

Convex Set Reconstruction from Support Line Measurements
and its Application to Laser Radar Data

by

Avinash Shreedhar Lele

S.B., Electrical Engineering and Computer Science
Massachusetts Institute of Technology
(1987)

Submitted to the Department of
Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements
for the Degree of

Master of Science

at the

Massachusetts Institute of Technology
February 1990

© Massachusetts Institute of Technology, 1990
All rights reserved

Signature of Author _____
Department of Electrical Engineering and Computer Science
October 31, 1989

Certified by _____
Alan S. Willsky
Thesis Supervisor

Accepted by _____
Arthur C. Smith
Chairman, Department Committee on Graduate Students

CONVEX SET RECONSTRUCTION FROM SUPPORT LINE MEASUREMENTS
AND ITS APPLICATION TO LASER RADAR DATA

by

Avinash Shreedhar Lele

Submitted to the Department of Electrical Engineering
and Computer Science on October 31, 1989 in partial
fulfillment of the requirements for the degree of
Master of Science

ABSTRACT

This thesis addresses the problem of estimating convex sets from noisy support line measurements. Several algorithms utilizing varying degrees of prior shape information are developed and studied. Quantitative assessments of their performance versus object parameters, measurement parameters, and quality of prior information are provided. The application of these algorithms to the problem of target reconstruction from resolved laser radar data is introduced and investigated. While target shape and size information may be obtained through other means (e.g., from reconstructed images using tomography), the present methods yield this information more directly. Furthermore, estimates obtained using these methods are more robust to noisy and/or sparse measurement data or data suffering from registration errors. Finally, a simple technique to improve tomographic images in the presence of registration errors is proposed and demonstrated.

Thesis Supervisor: Alan S. Willsky

Title: Professor of Electrical Engineering

Acknowledgements

Although many people at Lincoln Laboratory and M.I.T. have contributed in one way or another during my work on this thesis, I would like to make note of some special contributions. First, I thank Professor Alan Willsky for giving me the opportunity to work in his research group. I have learned a lot working with him, and deeply appreciate his guidance throughout the course of the research. His care and high standards always pushed me to learn more and substantially improved the quality of this thesis. I am also grateful to him for going on a 'reading binge' during the final weeks of my thesis research in spite of his particularly hectic schedule at the time.

Four years ago, I was very fortunate to have met Sanj Kulkarni, during our first summer at Lincoln Lab. The discussions we have had concerning this thesis research have been not only invaluable but were also particularly enjoyable due to our mutual tendency to digress onto any tangent, no matter how irrelevant! In general, I thank him for being sort of a 'big brother' to me, both technically and non-technically.

At Lincoln Lab, I am deeply indebted to Fred Knight and Alan Stein, for their extreme willingness to drop everything to answer any question that I had. In particular, I thank Fred for providing me with laboratory data, and I thank both for sharing their considerable computer expertise (and in Alan's case, his programming wizardry—the only cost being a good-natured laugh, at my expense!).

I gratefully acknowledge Bill Keicher for arranging financial support for my research and for having provided the opportunity for me to work at Lincoln every summer since my freshman year at M.I.T. He has been more than supportive in terms of allowing me to work on problems of interest to me.

Also, this thesis would have been more painful had it not been for Ellen Breau, Mike Reiley, George Sarafinas, and Wayne Schollenberger. I truly appreciate their help and am particularly grateful to George for going out of his way and driving in to Lincoln just to bring the VAX up for me at a critical stage in the research.

In addition, I wish to acknowledge the Defense Advanced Research Projects Agency for their support under Air Force Contract F19628-90-C-0002.

But my most special thanks must go to my family, for their contributions can never be surpassed. Whatever successes I have had truly belong to my mother and father, whose wisdom, encouragement, and understanding are unparalleled. I also thank my sisters, Aruna and Aparna, for being there whenever I need them. Surely I am luckier to have them as sisters than they are to have me as their brother.

*To my parents,
who sacrifice so much
for their children.*

Table of Contents

Abstract	ii
Acknowledgements	iii
List of Figures	vii
1 Introduction	1
1.1 Background and Motivation	1
1.2 Contributions and Organization	3
2 Problem Formulation and Previous Work	6
2.1 Background and Definitions	6
2.2 Formulation	10
2.3 Metrics on Convex Sets and Support Functions	12
2.4 Previous Work	15
2.5 Continuous-angle Framework	20
3 Reconstruction from Nonuniformly-spaced Measurements	24
3.1 The Estimator NUA	25
3.2 Examples	28
3.3 Performance Assessment	37
3.4 Perimeter Bias in the Estimator	44
4 Best N -gon Fitting M Measurements with Fixed Reconstruction Angles	54
4.1 The Estimator BNGON	54
4.2 Examples	58
4.3 Performance Assessment	60
4.4 Nonuniqueness of Solutions	71
5 Best N -gon with Fixed Relative Spacing of Reconstruction Angles	77
5.1 The Nature of the Cost Function	78
5.2 The Estimator BNGONROT	85
5.3 Examples	89
5.4 Performance Assessment	91
6 Target Reconstruction from Laser Radar Data	111
6.1 Laser Radar Data and Problem Scenario	111

6.2 Support Values from Resolved Laser Radar Data	115
6.3 Knot Location	119
6.4 Reconstructions from Simulated, Laboratory, and Field Laser Radar Data	120
6.5 Comparisons of Convex Set Reconstruction Methods with Tomographic Methods	130
7 Conclusion	143
7.1 Summary	143
7.2 Suggestions for Further Work	147
A Computation of the Area of the Symmetric Difference	152
B Calculation of the Probability of Anomalous Orientation Estimate $\Pr(A)$..	154
References	156

List of Figures

1.1	Block diagram of complete reconstruction process	3
2.1	Support line of a set	7
2.2	Members of the equivalence class of objects sharing the support vector $h = [2 \ 1 \ 2 \ 1]$	8
2.3	Support lines for an invalid support vector	9
2.4	Hausdorff distance between two sets	13
2.5	Equal Hausdorff distances for two cases	14
2.6	The symmetric difference between two sets	15
2.7	Projection of the invalid vector y onto the support cone \mathcal{C}	17
2.8	Consistency for a triplet of adjacent support lines spaced uniformly in angle	19
2.9	Support value for the point $(x_0, 0)$	22
2.10	A polygon and its support function $h(\theta)$	22
3.1	Consistency for a triplet of adjacent support lines spaced nonuniformly in angle	26
3.2	The standard triangle	29
3.3	NUA reconstruction from nonuniformly-spaced measurements of the standard triangle	31
3.4	NUA reconstruction from uniformly-spaced measurements of the standard triangle	32
3.5	The standard ellipse	33
3.6	NUA reconstruction from nonuniformly-spaced measurements of the standard ellipse	34
3.7	Geometry contributing to larger reconstruction errors for ellipses than for polygons	36
3.8	E versus σ for the standard ellipse	39
3.9	E versus M for the standard ellipse. Several values of eccentricity, uniformly-spaced noise-free measurements	40
3.10	E versus M for the standard ellipse. Single value of eccentricity,	

	uniformly-spaced measurements for several noise levels	41
3.11	<i>E</i> versus <i>M</i> for the standard ellipse. Two different configurations of measurement angles, several eccentricities	43
3.12	<i>E</i> versus <i>a</i> for an ellipse. Several values of <i>M</i> , noise-free uniformly-spaced measurements.....	44
3.13	Illustration of the benefit of having 90° and 270° in a measurement angle configuration for an ellipse	45
3.14	<i>E</i> versus <i>a</i> for an ellipse. Several values of <i>M</i> , radius-of-curvature measurement angle configuration, $\sigma = 0$	45
3.15	<i>E</i> versus <i>a</i> for an ellipse. Several noise levels, uniformly-spaced measurement angles, $M = 17$	46
3.16	Cross-section of support cone geometry giving rise to perimeter bias in NUA.....	48
3.17	Geometry illustrating the way in which invalid vectors cause perimeter bias.....	49
3.18	Non-Euclidean distance to inconsistency for sets of support lines of an ellipse.....	51
3.19	Normalized bias versus noise level and versus eccentricity.....	53
4.1	Support function picture of the BNGON problem	55
4.2	BNGON reconstruction from uniformly-spaced measurements of the standard triangle.....	59
4.3	Reference BNGON reconstruction from noise-free uniformly-spaced measurements of the standard ellipse.....	61
4.4	BNGON reconstruction from noisy uniformly-spaced measurements of the standard ellipse.....	62
4.5	<i>E</i> versus σ for the standard triangle, for BNGON and NUA	63
4.6	<i>E</i> versus <i>M</i> for the standard triangle, for BNGON and NUA	64
4.7	<i>E</i> versus σ for the standard triangle, for various incorrect priors (rotated reconstruction angles).....	65
4.8	<i>E</i> versus α for two types of rotation errors, and comparison with NUA baseline.....	67
4.9	<i>E</i> versus <i>N</i> , illustrating the effect of extraneous prior information, and comparison with NUA baseline	68

4.10	Reconstructions using extraneous prior information, and comparison to the NUA reconstruction	70
4.11	Underestimation versus overestimation of N	72
4.12	Nonuniqueness of solution: zero cost	73
4.13	Nonuniqueness of solution: non-zero cost	73
5.1	Sample plots of $J_{h_p}(\alpha)$ in the noise-free case, for uniformly-spaced measurement and reconstruction angles	81
5.2	Illustration of angles for which zero cost is attained, for a particular example	83
5.3	Deadspace extent for a particular example	84
5.4	BNGONROT reconstruction from uniformly-spaced measurements of the standard triangle	90
5.5	BNGONROT reconstruction from uniformly-spaced measurements of the standard ellipse	92
5.6	E versus σ for the standard triangle, for BNGONROT, BNGON, and NUA	93
5.7	Local and global errors	95
5.8	Cramer-Rao bound on orientation estimation error	98
5.9	Histograms of orientation estimation error for several noise levels	100
5.10	Probability of anomaly $\Pr(A)$	101
5.11	Orientation estimation error versus σ	103
5.12	Probability of anomaly $\Pr(A)$ (Monte Carlo)	104
5.13	Orientation estimation error versus σ for local errors (Monte Carlo)	105
5.14	Orientation estimation error versus σ for global errors (Monte Carlo)	106
5.15	Overall orientation estimation error versus σ (Monte Carlo)	107
5.16	Depth ratio γ versus σ for local errors and for global errors (Monte Carlo)	109
5.17	Overall depth ratio γ versus σ (Monte Carlo)	110
6.1	Range-resolved return from a target	112
6.2	Doppler-resolved return from a target	113
6.3	Support value identification from range-resolved data	118
6.4	Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.25$	123

6.5	Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.35$	124
6.6	Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.5$	125
6.7	Support lines and convex set reconstructions for simulated Doppler-resolved measurements of a cone. $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.25$	126
6.8	Support lines and convex set reconstructions for laboratory range-resolved measurements of the Mk-11 triconic. $\sigma_{\text{eff}} = 0$	128
6.9	Support lines and convex set reconstructions for laboratory range-resolved measurements of the Mk-11 triconic. $\sigma_{\text{eff}} = 0.25$	129
6.10	Support lines and convex set reconstructions for field Doppler-resolved measurements of the Thor Delta rocket body. $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0$	131
6.11	Support lines and convex set reconstructions for field Doppler-resolved measurements of the Thor Delta rocket body. $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.10$	132
6.12	Tomographic reconstructions from simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.25$	134
6.13	Tomographic reconstructions from simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.35$	135
6.14	Tomographic reconstructions from simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.5$	136
6.15	Tomographic reconstructions from simulated Doppler-resolved measurements of a cone. $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.25$	137
6.16	Tomographic reconstructions from laboratory range-resolved measurements of the triconic. $\sigma_{\text{eff}} = 0.25$	138
6.17	Tomographic reconstructions from field Doppler-resolved measurements of the Thor Delta rocket body. $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.10$	139
6.18	Block diagram of method by which convex set algorithms aid tomography in recovering from registration errors	141

Chapter 1 Introduction

1.1 Background and Motivation

In this thesis, we investigate the problem of reconstructing convex sets given noisy measurements of their support. The fundamental type of measurement that we consider is that of a support line to a set, defined as a line that ‘just grazes’ the set. From this definition, it follows that the set lies entirely in a particular one of the two halfplanes defined by each of its support lines. Clearly, a set has support lines at all orientations, or angles, and given the support lines at *all* angles, the halfplanes corresponding to these support lines may be intersected to yield the convex hull of the set.

Perhaps the simplest example of an application in which support line measurements are obtained is that of a parallel plate robot jaw obtaining information as to the shape of a flat object by tactile sensing [20]. Each time the jaw clamps down on the object, its two plates just graze the object and define two of its support lines. Repeated clamping from a number of different angles provides a set of support lines, and the associated halfplanes, when intersected, approximate the convex hull of the object. Another application arises in the field of medical imaging [19, 20], where pencil beam x-rays fired at the body from a number of angles provide projections of the density of body tissues. The projections provide support line information which can be used to determine the convex hull of, say, a tumor.

The application motivating our work is that of estimating the shape of a target given support line measurements which can be obtained from laser radar measurements of the target. Specifically, resolved laser radar measurements of a target provide

information as to the extent of the target in space. For example, a range-resolved measurement indicates where the target begins in range along the radar line of sight. Hence a line drawn at this range perpendicular to the line of sight just grazes the target and is a support line of the target. In a similar way, Doppler-resolved measurements of a spinning target contain support line information (see Chapter 6). Range-resolved or Doppler-resolved measurements from a number of aspects all lying in a plane yield a set of support lines, and given the support lines of a target at all aspects in a plane, the convex hull of the projection of the target onto this plane is determined.

However, in each of these applications the support line measurements are, in practice, noisy. This noise may arise from measurement errors in the tactile, x-ray, or laser radar sensor, or from incorrect spatial alignment, or *registration*, of the support lines obtained at a number of angles. Support lines that are corrupted by measurement noise may not correspond to the object that gave rise to them, and in fact may be *inconsistent* with one another, so that taken together there is *no* object having all of the measured lines as support lines. Our objective is to estimate the object which gave rise to a set of (noisy) support lines such as these. We do this by first extracting support line information from the raw data using a procedure known as *knot location* and then estimating the object which gave rise to these support line measurements, with both steps being optimal in some sense (see Figure 1.1). In this thesis, we focus our attention on the second step, using an algorithm developed by others [30, 16, 19] for the first step.

It appears that the estimation of convex sets from noisy support line measurements (the second module in the figure) was first formulated by Prince [19]. Prince studied the problem under the assumption of uniformly-spaced measurements in angle




Figure 1.1: Block diagram of complete reconstruction process.

and was motivated primarily by the potential for improvement in tomographic reconstructions from x-ray, or transmission, projections with the availability of prior shape information. Specifically, Prince used the projection data as a source for determining support line measurements. The shape estimate obtained from these measurements was then used as input to part of a hierarchical algorithm that ultimately produced a tomographic image through convolution backprojection.

The algorithms developed by Prince embody one basic approach, which consists of first characterizing those sets of support lines that are consistent, in that they are in fact support lines for some object. Given such a characterization, the estimation procedures have the effect of translating, or shifting, each support line in directions perpendicular to the support line, such that a cost function incorporating the sum of the squares of the shifts as well as prior information (if available) is minimized subject to the constraint that the set of shifted support lines is consistent. The shape estimate is given by the intersection of the halfplanes corresponding to these consistent support lines. We extend the work of Prince in several ways. Our extensions are described in some detail in the following section.

1.2 Contributions and Organization

The contributions of this thesis may be classified into two areas. In the first, we introduce and study several algorithms extending those of Prince, occasionally using

a continuous-angle framework in addition to Prince's discrete-angle one. The first extension consists of relaxing the assumption that the support line measurements be uniformly-spaced in angle, so that a polygon with faces at arbitrary angles is reconstructed. We point out that this estimator is biased in a specific sense, and provide explanations as to the origin of this bias in addition to quantifying it through a Monte Carlo approach. Next, we formulate a problem in which the faces of the polygonal reconstruction may be specified at a fixed set of angles independent of the measurement angles. This corresponds to the incorporation of prior information concerning the shape of the object—specifically, the number of sides and their absolute orientations. Finally, we allow rotations of this constellation of reconstruction angles (i.e., we assume knowledge of only *relative* rather than absolute angles of the sides of the object) in order to obtain joint orientation and shape estimates of objects. As we shall see, estimation of orientation leads to a nonlinear problem with the usual difficulties. For each of these three algorithms, we present quantitative performance analyses that consist of evaluating the area of the symmetric difference between the true objects and the corresponding reconstructions, for a range of object and measurement parameters. For the nonlinear estimator, we evaluate the Cramer-Rao bound on the error variance of the orientation estimate and an approximation to the probability of a large, anomalous orientation estimation error. This analytical error analysis is supported by Monte Carlo analysis.

The second area of contributions consists of introducing and investigating the application of convex set reconstruction techniques to laser radar data. In particular, we describe how range-resolved and Doppler-resolved laser radar data can serve to provide support line measurements of a target, and demonstrate the application of our reconstruction algorithms to resolved data obtained through a simulation model

and through laboratory and field measurements. Furthermore, we compare the reconstructions obtained by using the present methods with the reconstructed images produced by standard tomographic methods. In addition to not providing *direct* size and shape estimates of objects as do the convex set reconstruction methods, the tomographic methods are also not as robust to noisy and sparse measurement data. Furthermore, when we include registration errors, we find that the present methods are far more robust, due to the fact that they are based on consistency constraints that tend to reduce such errors. Finally, we introduce and demonstrate a simple but effective technique by which the tomographic reconstructions from unregistered data may be greatly improved using our estimation procedure as a preprocessor of the data.

The thesis is organized as follows. We begin Chapter 2 with background information and precise definitions of quantities of importance throughout the thesis, and then formulate the basic problem we consider. We also briefly describe the work of Prince as well as the continuous-angle framework that we occasionally use. Each of the following three chapters treats one of the reconstruction algorithms previously described, while Chapter 6 is concerned with the application of the three algorithms to laser radar data and the comparison of the convex set reconstructions to those produced by tomographic methods. Finally, we summarize the results of our work and propose further avenues of research.

Chapter 2 Problem Formulation and Previous Work

In this chapter we introduce the ideas of support lines, support functions, and convex sets, and observe that a convex set and its support function satisfy a one-to-one relationship. We also note that a function is the support function of some set iff it satisfies certain conditions. Next, we formulate the problem of estimating a valid support function (and hence, the associated convex set) given noisy support line measurements. Motivated by our desire to formulate estimation problems that are computationally tractable as well as our desire to minimize some geometrically appealing measure of error between the true object and the estimated object, we discuss measures of distance in support function space and object space. We then briefly review the work of Prince, and conclude the chapter by describing a continuous-angle framework that is useful in formulating and studying the estimation problems of Chapters 3–5.

2.1 Background and Definitions

We begin this section with several definitions. Minkowski [27] introduced the notion of a *support line* to a set. Specifically, using a coordinate frame fixed with respect to the set, we define the support line of the set S at angle θ_0 (denoted by $L_S(\theta_0)$) to be the line orthogonal to the vector $\omega(\theta_0) = [\cos \theta_0 \ \sin \theta_0]^T$ that ‘just grazes’ the set (see Figure 2.1). The *support value* $h_S(\theta_0)$ is defined as the maximum projection onto $\omega(\theta_0)$ of all points in S :

$$h_S(\theta_0) = \sup_{s \in S} s^T \omega(\theta_0), \quad (2.1)$$

and its magnitude, $|h_S(\theta_0)|$, is the minimum distance from $L_S(\theta_0)$ to the origin. In fact, since all points on $L_S(\theta_0)$ have the same projection onto $\omega(\theta_0)$ —namely, the support value $h_S(\theta_0)$ —the support line may be expressed precisely as

$$L_S(\theta_0) = \{x \in \mathbb{R}^2 \mid x^T \omega(\theta_0) = h_S(\theta_0)\}. \quad (2.2)$$

From all of this, it follows that the set S lies in a particular one of the two halfplanes defined by $L_S(\theta_0)$.




Figure 2.1: Support line of a set.

We will also refer to $h_S(\theta)$ as the *support function* of the set S . This function is continuous and periodic with period 2π . Sampling the support function at a finite number

of angles $\theta_1, \theta_2, \dots, \theta_M$ yields a *support vector* $h_S = [h_S(\theta_1) \ h_S(\theta_2) \ \dots \ h_S(\theta_M)]^T$.

From Figure 2.1, it should be apparent that support lines provide no information as to concavities in the set. Recall that a set S is *convex* if, for every pair of points $x, y \in S$, it is true that $\overline{xy} \subseteq S$, where \overline{xy} is the line segment joining x and y . Furthermore, the convex hull of a set S , denoted by $\text{conv}(S)$, is defined as the smallest convex set satisfying $S \subseteq \text{conv}(S)$. Since the convex hull and the support function of a set satisfy a one-to-one relationship, the support function of a set determines the set iff the set is convex. For this reason, we restrict attention to convex sets in subsequent discussion unless otherwise stated.

Note further that if the support function of a set is known for only a finite number of angles, the set is not uniquely determined, since an entire *equivalence class* of sets shares the same support vector. An example of this nonuniqueness is illustrated in Figure 2.2, where several sets that share the same set of support lines are shown. In this thesis, the set that we associate with any given support vector is the polygonal set bounded by the support lines, which is, of course, the largest set in the equivalence class. This set is the first one shown in the example in Figure 2.2.




Figure 2.2: Several members of the equivalence class of objects having support vector $h = [2 \ 1 \ 2 \ 1]^T$.

Although every convex set has a support function defined on $[0, 2\pi)$ that uniquely determines it, not every function defined on this domain is the support function of some set. Naturally, the same is true of support vectors. Consider the example of

the vector $[1\ 2\ 1\ 2\ 1\ 2\ 1\ 2]^T$ with support values at eight uniformly-spaced angles, whose support lines are shown in Figure 2.3. If the support lines corresponding to the support values equal to 1 (solid lines in the figure) do in fact support some object, then those support lines corresponding to values equal to two (shown dashed) *cannot* support the object. Hence, there is *no* object having $[1\ 2\ 1\ 2\ 1\ 2\ 1\ 2]^T$ as its support vector, and in such a case, the support values are said to be *inconsistent* (or, equivalently, the support vector is said to be *invalid*). Motivated by this example, we would like to characterize those functions and vectors that are in fact valid support functions and support vectors.




Figure 2.3: Support lines for the invalid support vector $h = [1\ 2\ 1\ 2\ 1\ 2\ 1\ 2]^T$

A number of necessary and sufficient conditions for a function to be a valid support function have been developed by others and are stated below without proof. In this thesis, we will be using a version of condition (1) suitable for support vectors. The second and third characterizations of a valid support function are included for completeness; they will not be used subsequently.

1. A twice-differentiable function $h(\theta)$ is a valid support function iff $h''(\theta) + h(\theta) \geq 0$ [24]. Roughly, the reasoning behind this condition is that for convex objects, the curvature of the boundary is given by $K(\theta) = h''(\theta) + h(\theta)$ and may never be negative.
2. A function $h(\theta)$ is a valid support function iff it can be written in the form

$$h(\theta) = a \cos \theta + b \sin \theta + \int_0^{2\pi} \sin(\theta - \lambda) dR(\lambda), \quad (2.3)$$

where a and b are constants and R is a finite measure on $[0, 2\pi)$ satisfying $\int_0^{2\pi} \cos \theta dR(\theta) = \int_0^{2\pi} \sin \theta dR(\theta) = 0$. This condition is part of a representation theorem proved by Vitale [29].

3. In contrast to our definition of the support function $h(\theta)$, the support function for a bounded set S is sometimes defined as $h_S(\mathbf{x}) = \sup_{s \in S} s^T \mathbf{x}$ for all $\mathbf{x} \in R^2$ [11].

The support function $h(\mathbf{x})$ is essentially a scaled version of $h(\theta)$ with scale factor $\|\mathbf{x}\|_2$ and reduces to $h(\theta)$ if its domain is restricted to the unit circle $\|\mathbf{x}\|_2 = 1$, with $\|\cdot\|_2$ denoting L^2 or Euclidean norm.

A function $h(\mathbf{x})$ is a valid support function iff it satisfies [2]

- (a) $h(0) = 0$
- (b) $h(\lambda \mathbf{x}) = \lambda h(\mathbf{x})$, for $\lambda > 0$
- (c) $h(\mathbf{x} + \mathbf{y}) \leq h(\mathbf{x}) + h(\mathbf{y})$

Incidentally, the continuity of $h(\mathbf{x})$ and $h(\theta)$ can be shown via (c).

2.2 Formulation

The fundamental problem we consider in this thesis is one in which a valid support vector is corrupted by noise, so that a finite number of noisy support measurements are given by

$$y_i = h_i + n_i, \quad i = 1, 2, \dots, M \quad (2.4)$$

at a set of angles $\{\theta_i\}$, where h_i replaces $h(\theta_i)$ for notational convenience and the $\{n_i\}$ are independent, identically distributed Gaussian random variables. We emphasize

that by 'noisy measurements' we mean uncertainty in the support values and not in the measurement angles. Because of the noise, the measurements $\{y_i\}$ may not correspond to the set S which gave rise to them and, in fact, may not correspond to *any* convex set. Our objective is to estimate S in some optimal fashion.

Each of the noisy support measurements gives rise to a support line and to the corresponding halfplane in which the measurement indicates that the set lies. A naïve method of reconstructing the set S might consist of intersecting these halfplanes. Unfortunately, as more and more of these halfplanes are intersected, more and more of the set may be whittled away due to measurement errors. In particular, those halfplanes corresponding to support values arising from negative noise samples cause this phenomenon.

There are two reasons that this *intersection method* is not a desirable method. First, we would like a *consistent* estimator, one which converges to the true value S as the number of measurements increases. Also, in using the intersection method, we implicitly place more confidence in certain support values than in others. That is, after all of the halfplanes have been intersected, certain support lines will support the estimated set while others will not. Producing such an estimate essentially indicates that the lines that *do* support the estimated set were considered by the estimator as being known perfectly, as opposed to those lines that do not support the set, which were thought to be in error.

In this thesis, we develop several estimation procedures that have more desirable properties for the fundamental problem formulated earlier in this section. The approach we take in estimating the set that gave rise to a set of noisy measurements follows that of Prince. Specifically, we first estimate the *valid* support vector that is in some sense 'closest' to the noisy support vector and then take as our estimate the

polygon bounded by the estimated support lines. In the next section, we investigate the idea of ‘closeness’ by discussing metrics on object space and support function space. Much of our work extends the work of Prince, which is described in some detail in Section 2.4.

2.3 Metrics on Convex Sets and Support Functions

In this section, we discuss measures of distance on the set of all convex sets (denoted by \mathcal{K}) and on the set of all valid support functions (denoted by \mathcal{C}), as well as relationships between them. In doing so, we would like to identify a metric on \mathcal{K} or on \mathcal{C} that reflects (with sufficient accuracy) the qualitative error between true objects and reconstructed objects as perceived by our eyes. Furthermore, the minimization of distance associated with our choice of metric must be amenable to computation.

Probably the most commonly used metric on \mathcal{K} is the *Hausdorff metric* [11], which has been referred to as the ‘standard metric’ on \mathcal{K} by Shephard and Webster [25]. To define the Hausdorff metric $d_{\mathcal{K}}^H(K_1, K_2)$ for two sets $K_1, K_2 \in \mathcal{K}$, we first define the Minkowski set sum of two sets K_1 and K_2 by $K_1 \oplus K_2 = \{k_1 + k_2 | k_1 \in K_1, k_2 \in K_2\}$, and denote the disk of radius δ by $B(\delta)$. Now let the quantity δ_1 be the minimum non-negative number for which $K_2 \subseteq K_1 \oplus B(\delta_1)$ and let δ_2 be the minimum non-negative number for which $K_1 \subseteq K_2 \oplus B(\delta_2)$. Then $d_{\mathcal{K}}^H(K_1, K_2) = \max(\delta_1, \delta_2)$. Moreover, it turns out that the Hausdorff metric on \mathcal{K} induces the L^∞ norm on \mathcal{C} , defined by

$$\|h_{K_1}(\theta) - h_{K_2}(\theta)\|_\infty = \sup_{\theta} |h_{K_1}(\theta) - h_{K_2}(\theta)|, \quad (2.5)$$

where the general L^p norm on \mathcal{C} is defined as

$$\|h_{K_1}(\theta) - h_{K_2}(\theta)\|_p = \left(\int_0^{2\pi} |h_{K_1}(\theta) - h_{K_2}(\theta)|^p d\theta \right)^{1/p}. \quad (2.6)$$

An example illustrating the Hausdorff distance between two sets is shown in Figure 2.4. To gain some geometrical intuition about Hausdorff distance, we can loosely think of δ_1 as being the amount by which K_1 would have to be ‘inflated’ or ‘expanded’ to include K_2 , and δ_2 as being the analogous amount after reversing the roles of K_1 and K_2 . The Hausdorff distance is the maximum of these amounts.




Figure 2.4: Illustration of Hausdorff distance between two sets K_1 and K_2

Although the Hausdorff metric on \mathcal{K} has a certain geometrical appeal, it is particularly deceiving in certain cases of interest. Consider the two examples in Figure 2.5, in which K_1 is the true object (a square) and K_2 is the reconstruction (a triangle in Figure 2.5a and a diagonal line in Figure 2.5b). In both examples the Hausdorff distance between K_1 and K_2 equals half of the length of the square’s diagonal, since

both the triangle and the line need to be ‘inflated’ by this amount to include the lower left corner of the square. But clearly the triangular reconstruction is ‘closer’ to the square than the line is, as perceived by our eyes.




Figure 2.5: Equal Hausdorff distances between K_1 and K_2 for two cases. Note that $\delta_1 = 0$ in both cases.

Consequently, we turn to another measure of distance on \mathcal{K} , which has more geometric appeal. This metric is the *symmetric difference metric* defined by

$$d_{\mathcal{K}}^S(K_1, K_2) = \mu(K_1 \Delta K_2) = \mu[(K_1 \cup K_2) \setminus (K_1 \cap K_2)]. \quad (2.7)$$

where $\mu(K)$ denotes the Lebesgue measure or area of K . In Figure 2.6, two sets are shown with their symmetric difference shaded. Unfortunately, this metric has no simple interpretation in support function space.

As we shall see subsequently, our reconstructions are obtained by minimizing L^2 distance in support vector space (which has no simple geometric interpretation as a




Figure 2.6: Two sets with their symmetric difference shown shaded

metric on \mathcal{K}) rather than by minimizing some suitable distance in object space, primarily because the former is computationally less difficult. Nevertheless, minimizing distance in \mathcal{K} is of direct interest in evaluating the quality of our reconstructions. As a result, we wish to assess the performance of our algorithms by some quantitative measure of the error between the true object and its reconstruction. Our choice of measure of distance between the true object S and the corresponding reconstructed object \hat{S} is the area of their symmetric difference $S\Delta\hat{S} = (S\cup\hat{S})\setminus(S\cap\hat{S})$, for reasons that include geometric appeal and ease of computation. Although the symmetric difference metric on \mathcal{K} does not correspond to the L^2 norm on \mathcal{C} , we expect that the distances in the two spaces are related in most cases of interest. Our method of computation of $\mu(S\Delta\hat{S})$ is outlined in Appendix A.

2.4 Previous Work

It appears that the estimation of convex sets from noisy support line measurements was first studied by Prince [19]. (Greschak [4], Stark and Peng [26], and others have done related work.) Prince's work was motivated primarily by the potential for improvement in tomographic reconstructions with the availability of prior shape

information consisting of convex set reconstructions from support measurements.

Prince considered the problem of estimating the support vector from noisy support line measurements $\{y_i\}$ taken at a set of uniformly-spaced angles $\{\theta_i\}$. In particular, the measurements are given by

$$y_i = h_i + n_i, \quad i = 1, 2, \dots, M$$

at the angles

$$\theta_i = \frac{2\pi}{M}(i - 1)$$

where the $\{n_i\}$ are samples from independent, identically distributed Gaussian random variables.

One of the estimation procedures considered by Prince essentially consists of finding the valid support vector \hat{h} minimizing the L^2 distance between h and y in \mathbb{R}^M . The space of valid support vectors in \mathbb{R}^M is a cone [27], which we shall refer to as the support cone \mathcal{C} following the notation of Prince. That is, if h is a valid support vector, then so too is λh , $\lambda \geq 0$. Hence, the problem of finding the closest valid support vector in the L^2 sense geometrically takes the form of projecting the noisy support vector onto the support cone (see Figure 2.7), and analytically takes the form

$$\hat{h} = \arg \min_{h \in \mathcal{C}} \sum_{i=1}^M (y_i - h_i)^2.$$

Prince derived the consistency constraint imposed upon candidate support vectors and proved its necessity and sufficiency in [19]. His approach is a geometric one, showing that a triplet of support lines at consecutive angles is consistent iff each line supports the object. In Figure 2.8, we see that if lines L_{i-1} and L_{i+1} support the object (shown shaded), then line L_i supports the object iff it passes through or to the




Figure 2.7: Projection of the noisy support vector y onto the support cone C to yield the estimate \hat{h} .

left of the intersection point of lines L_{i-1} and L_{i+1} . Mathematically, this translates to the inequality

$$\frac{1}{2 \cos(\frac{2\pi}{M})} h_{i-1} - h_i + \frac{1}{2 \cos(\frac{2\pi}{M})} h_{i+1} \geq 0 \quad (2.8)$$

which, when enforced for the support lines corresponding to all adjacent triplets of angles, yields the consistency condition for a candidate support vector

$$Ch \geq 0$$

where

$$C = \begin{pmatrix} -1 & k & 0 & \dots & 0 & k \\ k & -1 & k & 0 & \dots & 0 & 0 \\ 0 & k & -1 & k & 0 & \dots & 0 \\ \vdots & \dots & \dots & \dots & \dots & & \vdots \\ k & 0 & 0 & \dots & 0 & k & -1 \end{pmatrix},$$

and $k = 1/(2 \cos(\frac{2\pi}{M}))$. The vector inequality is an inequality for each component. As a result, the solution takes the form

$$\hat{h} = \arg \min_{Ch \geq 0} \sum_{i=1}^N (y_i - h_i)^2. \quad (2.9)$$

Since the cost function to be minimized is quadratic in the $\{h_i\}$ and the constraints are linear, the solution to Equation 2.9 may be obtained through standard quadratic programming (QP) methods [13, 10]. The estimated set is given by the polygon bounded by the estimated support lines.

Prince also investigated the geometry of the support cone C . In particular, he found that the consistency matrix C has a two-dimensional nullspace; consequently, the support cone C is not a proper cone, as it contains this nullspace. In terms of




Figure 2.8: Consistency for a triplet of adjacent support lines

object geometry, the component of h lying in the nullspace of C corresponds to a shift of the object in the plane. This leads to the notion of a centered object whose nullvector is zero. For the polygon bounded by the support lines $\{\hat{L}_i\}$ corresponding to the support vector \hat{h} , this corresponds to that shift for which the average of the support line intersection points $\{\hat{L}_i \cap \hat{L}_{i+1}\}$ is $(0,0)$.¹ Further, components of the support vector of the centered object that independently determine its size and shape were identified, leading to a size, shape, and shift decomposition of a valid support vector. Finally, Prince introduced the notion of a discrete radius of curvature for polygonal objects, thereby characterizing their smoothness. These ideas led to the development of several alternative estimation algorithms which incorporated available prior information as to object location, size, shape, or smoothness.

2.5 Continuous-angle Framework

The derivations of several of the results to be described in subsequent chapters are facilitated by considering the support function of a set, rather than its support vector. In this section, we describe this continuous-angle framework. The set S_1 consisting of a single point located on the x -axis at $(x_0, 0)$ is shown in Figure 2.9 to have the support function $h_{S_1}(\theta) = x_0 \cos \theta$. Similarly, the set S_2 located at $(0, y_0)$ can be seen to have the support function $h_{S_2}(\theta) = y_0 \sin \theta$. By the following property of support functions [11]

$$h_{A \oplus B} = h_A + h_B, \quad (2.10)$$

¹Note that this quantity is not necessarily the average of the polygonal vertices. This is because $\hat{L}_i \cap \hat{L}_{i+1}$ may equal $\hat{L}_j \cap \hat{L}_{j+1}$ for some i, j , so that several support lines intersect to form a single vertex.

where, as before, \oplus denotes the Minkowski set sum defined by $A \oplus B = \{a + b | a \in A, b \in B\}$, we have that the support function of a point (x_0, y_0) is given by

$$h(\theta) = x_0 \cos \theta + y_0 \sin \theta. \quad (2.11)$$

Now, by representing a polygon S_P as the convex hull of its vertices (x_i, y_i) :

$$S_P = \text{conv} \left[\bigcup_{i=1}^N (x_i, y_i) \right]$$

and from the property that [11]

$$h_{\text{conv}(A \cup B)} = \max(h_A, h_B) \quad (2.12)$$

we find that the support function of a polygon is given by

$$h_{S_P}(\theta) = \max_i (x_i \cos \theta + y_i \sin \theta), \quad (2.13)$$

An example polygon and its support function are shown in Figure 2.10. The local maxima in $h(\theta)$ correspond to vertices of the polygon, while the cusps correspond to the faces of the polygon, at which support is transferred from one vertex to another. If we denote as θ_{i-1} and θ_i the angles of the faces which intersect to form the i^{th} vertex, we may rewrite Equation 2.13 as

$$h_{S_P}(\theta) = x_i \cos \theta + y_i \sin \theta \quad \text{on } [\theta_{i-1}, \theta_i] \quad (2.14)$$

so as to bring out more clearly the trigonometric spline form that $h_{S_P}(\theta)$ takes.

Incidentally, through this framework we can characterize the shift of an object in a manner that is a generalization of Prince's notion of shift. Expressing $h_S(\theta)$ in




Figure 2.9: Support value for the point $(x_0, 0)$




Figure 2.10: A polygon and its support function $h(\theta)$

a Fourier series as $h_S(\theta) = \sum_{n=0}^{\infty} x_n \cos n\theta + y_n \sin n\theta$, we may rearrange terms and group the first harmonic terms apart from the others, labelling the two parts as $h_A(\theta)$ and $h_B(\theta)$:

$$h_S(\theta) = \underbrace{x_1 \cos \theta + y_1 \sin \theta}_{h_A(\theta)} + x_0 + \underbrace{\sum_{i=2}^{\infty} x_n \cos n\theta + y_n \sin n\theta}_{h_B(\theta)} \quad (2.15)$$

Now from Equation 2.11, $h_A(\theta)$ is the support function of the point (x_1, y_1) . Also, the Minkowski set sum of a point with a set simply translates the set by the coordinates of the point. *Defining* (x_1, y_1) as the coordinates of the center of the object S results in a notion of object center, which in fact reduces to that of Prince under the assumption of uniformly-spaced angles. Moreover, since it has no first harmonic, it follows that $h_B(\theta)$ is the support function of S shifted such that its center (by our definition) is located at $(0, 0)$.

Chapter 3 Reconstruction from Nonuniformly-spaced Measurements

This chapter begins a three chapter sequence in which we introduce algorithms to estimate convex sets from noisy support line measurements. Specifically, the estimation algorithms that we present arise from increasingly general formulations of the problem of obtaining polygonal shape estimates from noisy support measurements. The most specific case, considered by Prince, was reviewed in Section 2.4. In this case, a polygon with faces at a fixed number of uniformly-spaced measurement angles is estimated. A generalization of this algorithm results in relaxing the assumption of uniform spacing. A third formulation consists of estimating a polygon with faces at a set of prespecified reconstruction angles that are independent of the measurement angles. Both sets of angles are nonuniformly-spaced, in general. Fourth, we might allow rotations of the prespecified constellation of reconstruction angles in order to obtain joint orientation and shape estimates of objects. The fifth level of generality results in specifying only the number but not the values of the reconstruction angles. The most general formulation is one in which neither the number nor the values of the reconstruction angles is specified. Such an estimator would be essentially the same as the one just discussed, with the exception that larger numbers of reconstruction angles would be penalized in the cost function.

In Chapters 3–5, we consider the second, third, and fourth formulations, respectively. The fifth and sixth formulations involve angle estimation and appear to be difficult problems because of nonlinearities that enter into the cost functions and consistency constraints.

3.1 The Estimator NUA

In this chapter, we develop the first of several extensions to Prince's work, occasionally making use of the continuous framework from Section 2.5. In particular, we relax the requirement that the measurement angles be uniformly-spaced, motivated by limited-angle scenarios, in which it is difficult or impossible to obtain measurements at some angles. In this case, we have a finite set of noisy support measurements $\{y_1, y_2, \dots, y_M\}$ at angles $\{\theta_1, \theta_2, \dots, \theta_M\}$ about which no assumptions are made, and we seek the following solution:

$$\hat{h} = \arg \min_{h \text{ valid}} \sum_{i=1}^M (y_i - h_i)^2. \quad (3.1)$$

The only problem in solving Equation 3.1 lies in deriving a necessary and sufficient consistency condition on h for nonuniformly-spaced angles. Here, we show only sufficiency, since the proof of necessity is essentially the same as the somewhat involved proof of Prince for uniformly-spaced angles [19]. One approach to showing sufficiency is a geometric one, similar to Prince's. Specifically, consider the geometry in Figure 3.1, in which support lines L_{i-1} and L_{i+1} at angles θ_{i-1} and θ_{i+1} can be thought of as adjacent sides of a polygon intersecting to form a vertex at the point (v_x, v_y) . This vertex gives rise to the support function $g(\theta) = v_x \cos \theta + v_y \sin \theta$ on $[\theta_{i-1}, \theta_{i+1}]$, from Equation 2.14. Now, for a third support line L_i at angle θ_i , $\theta_{i-1} < \theta_i < \theta_{i+1}$, to be consistent with those at θ_{i-1} and θ_{i+1} , it is apparent from Figure 3.1 that we require

$$\begin{aligned} h_i &\leq g(\theta_i) \\ &= v_x \cos \theta_i + v_y \sin \theta_i. \end{aligned} \quad (3.2)$$

so that L_i does not lie in the invalid region, shown hatched and to the right of the dashed line.




Figure 3.1: Consistency for support lines at a triplet of nonuniformly-spaced angles

To determine the intersection point (v_x, v_y) , since

$$\begin{bmatrix} h_{i-1} \\ h_{i+1} \end{bmatrix} = \begin{bmatrix} g(\theta_{i-1}) \\ g(\theta_{i+1}) \end{bmatrix} = \begin{bmatrix} \cos \theta_{i-1} & \sin \theta_{i-1} \\ \cos \theta_{i+1} & \sin \theta_{i+1} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

we have that

$$\begin{bmatrix} v_x \\ v_y \end{bmatrix} = \frac{1}{\sin(\theta_{i+1} - \theta_{i-1})} \begin{bmatrix} h_{i-1} \sin \theta_{i+1} - h_{i+1} \sin \theta_{i-1} \\ h_{i+1} \cos \theta_{i-1} - h_{i-1} \cos \theta_{i+1} \end{bmatrix}. \quad (3.3)$$

By substituting Equation 3.3 into Equation 3.2 and using sum-of-angle formulas, we arrive at the consistency condition for a triplet of support values adjacent and nonuniformly-spaced in angle:

$$h_{i-1} \sin(\theta_{i+1} - \theta_i) - h_i \sin(\theta_{i+1} - \theta_{i-1}) + h_{i+1} \sin(\theta_i - \theta_{i-1}) \geq 0, \quad (3.4)$$

which is a generalization of Prince's consistency condition, reducing to Equation 2.8 for uniformly-spaced angles. Enforcing this for all adjacent triplets yields the matrix consistency condition

$$Ch \geq 0$$

where

$$C = \begin{pmatrix} -\sin(\theta_2 - \theta_M) & \sin(\theta_1 - \theta_M) & 0 & 0 & \sin(\theta_2 - \theta_1) \\ \sin(\theta_3 - \theta_2) & -\sin(\theta_3 - \theta_1) & \sin(\theta_2 - \theta_1) & 0 & 0 \\ 0 & \sin(\theta_4 - \theta_3) & -\sin(\theta_4 - \theta_2) & \sin(\theta_3 - \theta_2) & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots \\ \sin(\theta_M - \theta_{M-1}) & 0 & 0 & \sin(\theta_1 - \theta_M) & -\sin(\theta_1 - \theta_{M-1}) \end{pmatrix} \quad (3.5)$$

Another approach to deriving the consistency condition is based on the property that twice-differentiable valid support functions satisfy the inequality $K(\theta) = h''(\theta) + h(\theta) \geq 0$, as mentioned in Section 2.1. On substituting Equation 3.3 (with the index i replacing $i + 1$) into Equation 2.14, we have that the support function of a polygon is given by the trigonometric spline

$$h(\theta) = \frac{h_{i-1} \sin \theta_i - h_i \sin \theta_{i-1}}{\sin(\theta_i - \theta_{i-1})} \cos \theta + \frac{h_i \cos \theta_{i-1} - h_{i-1} \cos \theta_i}{\sin(\theta_i - \theta_{i-1})} \sin \theta \text{ on } [\theta_{i-1}, \theta_i]. \quad (3.6)$$

Although this support function is not twice-differentiable due to the cusps corresponding to the faces of the polygon, we interpret the inequality in a distributional sense. Specifically, $K(\theta)$ for a polygon is zero everywhere except at angles corresponding to faces, where Dirac-delta functions $\delta(\cdot)$ arise from the cusps in $h(\theta)$ (see Section 2.5). Requiring that these impulses have non-negative areas:

$$\left. \frac{dh(\theta)}{d\theta} \right|_{\theta=\theta_i^-} \leq \left. \frac{dh(\theta)}{d\theta} \right|_{\theta=\theta_i^+}, \quad \forall i$$

and using Equation 3.6 yields the same consistency condition as in Equation 3.4.

With the derived consistency condition, we can formulate an estimation algorithm similar to the one developed by Prince for uniformly-spaced angles. We will refer to this estimator as NUA:¹

Estimator NUA

$$\hat{h} = \arg \min_{Ch \geq 0} \sum_{i=1}^M (y_i - h_i)^2 \quad (3.7)$$

where

$$C = \begin{pmatrix} -\sin(\theta_2 - \theta_M) & \sin(\theta_1 - \theta_M) & 0 & 0 & \sin(\theta_2 - \theta_1) \\ \sin(\theta_3 - \theta_2) & -\sin(\theta_3 - \theta_1) & \sin(\theta_2 - \theta_1) & 0 & 0 \\ 0 & \sin(\theta_4 - \theta_3) & -\sin(\theta_4 - \theta_2) & \sin(\theta_3 - \theta_2) & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots \\ \sin(\theta_M - \theta_{M-1}) & 0 & 0 & \sin(\theta_1 - \theta_M) & -\sin(\theta_1 - \theta_{M-1}) \end{pmatrix} \quad (3.8)$$

3.2 Examples

To illustrate the behavior of NUA, we consider two examples. In the first, we examine the reconstruction arising from nonuniformly-spaced noisy support measurements of a triangle. The triangle used in this example is an isosceles triangle with vertices at $(2, 0)$, $(-0.25, 0)$, and $(0.25, 0)$ and is shown in Figure 3.2. We use this triangle throughout the thesis and refer to it as the ‘standard triangle.’ The measurement angles have spacings of $\Delta\theta = 30^\circ$ on $[-90^\circ, 90^\circ]$ and $\Delta\theta = 2^\circ$ on $[90^\circ, 270^\circ]$ yielding

¹NUA is an acronym for NonUniform Angles.

$M = 96$ total measurements with noise level $\sigma = 0.25$.




Figure 3.2: The standard triangle.

Figure 3.3 depicts the estimation process in object space and support function space. In Figure 3.3a are shown the bold outline of the true object (the standard triangle), the noisy support lines, and the shaded polygonal reconstruction produced by NUA. For this particular example, we also include a picture without support lines, since it is difficult to see the reconstruction in the picture with support lines. Correspondingly, Figure 3.3b shows the support function $h(\theta)$ of the true object, the noisy support values $\{y_i\}$, and the support function $\hat{h}(\theta)$ of the estimated object. The display conventions used in Figure 3.3 are also used throughout the thesis. Clearly, the reconstruction corresponding to the left half of the triangle is superior to that corresponding to the right half, a difference attributable to the different densities in measurement angles over the two halves, as seen in object space as well as in support function space. Moreover, one would expect that as the density of measurements increased on $[90^\circ, 270^\circ]$, the reconstruction would approximate the left half of the

triangle with greater fidelity.

The quantitative measure of reconstruction error that we use throughout the thesis consists of the area of the symmetric difference between the reconstructed object and the true object, normalized by the area of the true object. This error is denoted by E , and for the present example has the value $E = 1.17$. Although an error that is 117% of the area of the true object seems quite bad, on a visual evaluation the reconstruction does not appear so bad. As a result, the quantitative evaluation of error should not be considered in an absolute sense but rather in a relative sense. That is, a reconstruction with $E = 0.50$ should appear to be better than one with $E = 1.00$.

For the purpose of comparing reconstructions produced by NUA with those produced by estimators to be developed in subsequent chapters (given identical measurement data), we include here a reconstruction arising from $M = 24$ *uniformly-spaced* noisy ($\sigma = 0.25$) measurements (see Figure 3.4). This example, when compared with that shown in Figure 3.3, also shows the effect that measurement angle density has. Here, $E = 1.56$.

The second example is similar to the first except for the facts that the true object here is an ellipse and that the constellation of $M = 96$ measurement angles is rotated by -90° , so that the high and low density regions coincide with the intervals $[0^\circ, 180^\circ]$ and $[180^\circ, 360^\circ]$ respectively. Here, the noise level $\sigma = 0.5$. The ellipse, centered on the origin, has semimajor and semiminor axes aligned with the x - and y -axes, with lengths $a = 2$ and $b = 1$, respectively (see Figure 3.5). In this chapter and subsequent chapters, we will use this particular ellipse as a default when the true object is an ellipse, referring to it as the ‘standard ellipse.’

Figure 3.6 shows the estimation process in object space and support function


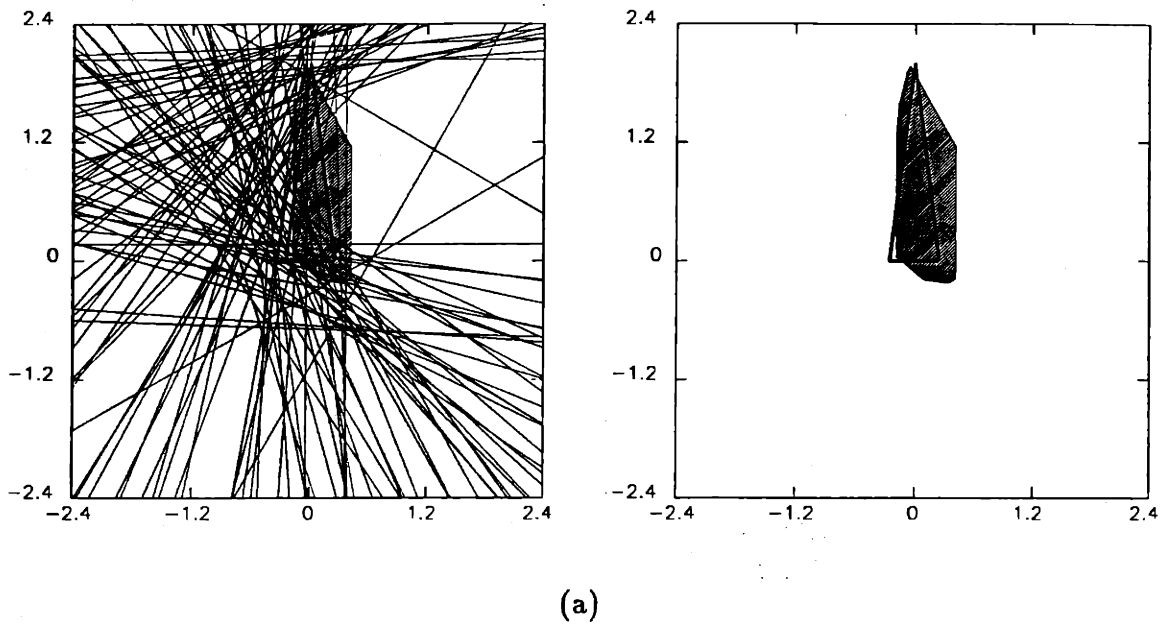


Figure 3.3: An example of NUA for the standard triangle in (a) object space and (b) support function space, for the case $M = 96$, $\sigma = 0.25$. Shown in (a) are the true object (bold outline), noisy support lines, and the reconstructed object (shaded). Shown in (b) are the true support function (thick curve), noisy support measurements (marked by \times 's), and the support function of the reconstruction (thin curve).





Figure 3.4: An example of NUA for the standard triangle in (a) object space and (b) support function space, for $M = 24$ uniformly-spaced noisy ($\sigma = 0.25$) measurements.




Figure 3.5: The standard ellipse.

space, where the support function $h(\theta)$ of the ellipse can be easily derived as

$$h(\theta) = \sqrt{a^2 \cos^2 \theta + b^2 \sin^2 \theta}. \quad (3.9)$$

The reconstructed object incurs an error of $E = 0.41$ with respect to the true object. As in the previous example, the different densities yield reconstructions of different quality over the two halves of the object. (This difference is particularly noticeable in support function space.) However, in object space, the reconstruction in this example does not approximate the top half of the ellipse as well as the Figure 3.3 reconstruction approximates the left half of the standard triangle.

There are several plausible explanations for this fact, and we digress momentarily to investigate these. First, because error minimization is done in \mathcal{C} , we expect the reconstruction to look better in support function space than in object space. Indeed, in Figures 3.3b and 3.6b both reconstructions are comparable on those intervals with dense measurements. However, reconstruction quality as perceived by our eyes is





Figure 3.6: An example of NUA for the standard ellipse in (a) object space and (b) support function space, for the case $M = 96$, $\sigma = 0.5$.

probably more closely related to distance in \mathcal{K} than in \mathcal{C} , accounting for the observed disparity in reconstruction quality in object space.

Another explanation is provided by examining in more detail the support cone geometry characterizing the estimation process in support vector space. In our examination below, we show that the support vector h_{tri} arising from noise-free measurements of a triangle (or polygon, in general) lies on a *vertex* of the cone while the support vector h_{ell} of an ellipse lies in its interior. This difference results in more accurate reconstructions for the triangle, on average.

First we make several observations. Given a valid support vector h , if consistency is satisfied with equality at a *single* triplet, then h lies on the boundary of \mathcal{C} (i.e., on a face of \mathcal{C}). We will refer to the point at which the three support lines intersect as a *triple point*. Furthermore, if a quadruplet of support lines intersect at a single point (a *quadruple point*) so that Equation 3.4 is satisfied at two consecutive triplets, then h lies on the intersection of two faces of \mathcal{C} , i.e., on a *vertex* of \mathcal{C} . For the general case in which k support lines intersect at a single point (a *k-tuple point*), h lies on a higher-order vertex of the cone.

If the true object is a triangle, all of the support lines will pass through a vertex of the triangle². For any $M > 3$, this results in k -tuple points so that h_{tri} lies on a vertex of \mathcal{C} . In our example, since $M = 96$, h_{tri} must lie on a high-order vertex of \mathcal{C} . In contrast, none of the support lines of an ellipse give rise to triple points. Hence, h_{ell} lies in the *interior* of \mathcal{C} .

Consider the geometry in Figure 3.7, which contains a cross-sectional view of \mathcal{C} obtained by slicing the cone perpendicular to $e = [1 \ 1 \ \dots \ 1]^T$, the central axis of \mathcal{C} . Due to the difficulty in visualizing objects in more than three dimensions, the

²not to be confused with a vertex of the support cone

figure shows a two-dimensional geometry. The true support vector h_{tri} is shown lying on a vertex of C , while h_{ell} lies in the interior. Since $y = h + n$, the ensemble of noisy vectors results from centering a Gaussian noise ‘sphere’ on h_{tri} or h_{ell} and considering points in the sphere. As we mentioned before, if $y \in C$, then $\hat{h} = y$; otherwise \hat{h} is the projection of y onto C . For the triangle, an entire sector of invalid noisy vectors (shown hatched in the figure) projects back to h_{tri} , leading to error-free reconstructions. Noisy vectors near this sector project to point on C that are near h_{tri} , leading to reconstructions with small error. In contrast, for the ellipse, *no* invalid noisy vectors are projected onto h_{ell} , since h_{ell} is in the interior of C . Hence, the reconstruction error for elliptical true objects is on average greater than that for triangular true objects.




Figure 3.7: The geometry in a cross-section of the support cone C which contributes to larger average reconstruction errors for ellipses than polygons.

The support cone picture also offers degrees-of-freedom interpretations in polyg-

onal estimation for the cases of polygons versus ellipses. Specifically, for a polygon, the projection of a sizable fraction of noisy vectors onto a high-order vertex (which is a low-dimensional subspace of the support cone \mathcal{C}) gives rise to fewer degrees of freedom in the solution space when compared to the number of degrees of freedom in the case of an ellipse, in which almost all noisy vectors project onto a face of the cone (a higher-dimensional subspace of the cone compared to a vertex). Hence, in the case of a triangle, we obtain more measurements per degree of freedom than in the case of an ellipse. In fact, in the next two chapters, we exploit the large number of measurements per degree of freedom in the case of a polygon. The estimation algorithms developed in these chapters utilize explicit prior information as to degrees of freedom expected in the solution.

3.3 Performance Assessment

The performance of our algorithms depends on a number of parameters of the measured data as well as on characteristics of the true object itself. In what follows, we evaluate NUA by computing the average normalized symmetric difference E between the true object and the reconstructions, for a range of values of relevant parameters. This evaluation is carried out by performing Monte Carlo analyses over a large number of noise samples.

In particular, we evaluate the performance of NUA against measurement noise level σ , number of measurements M , and object eccentricity (implicitly given by the length of the semimajor axis a , since we fix the length of the semiminor axis b at unity), for a true object which is an ellipse centered at the origin with semimajor and semiminor axes aligned with the coordinate axes. Our choice of an ellipse as the true object

is motivated by the desire to explore the behavior of NUA for a class of objects, and in particular those that vary continuously in shape. The ease with which the shape of an ellipse is parametrized leads us to this choice. However, we expect that the performance of NUA versus σ , M , and a for an elliptical true object is similar to its performance versus the same for a polygonal one, with polygonal eccentricity defined suitably. Moreover, Chapter 4 treats an estimation algorithm for which it is natural to use a polygonal true object. The performance assessment of that algorithm includes comparisons with NUA, for a polygonal true object, providing corroboration of the performance results presented here.

Figure 3.8 shows a plot of the normalized area of the symmetric difference E versus σ for the fixed values $M = 12$ and $a = 2$, with each data point representing the average of the areas of the symmetric difference arising from 1000 noise samples. As expected, E increases with σ .³ Note that $E \neq 0$ for $\sigma = 0$, since in this case we are inscribing the ellipse in a polygon.

We now consider the dependence of E on the number of measurements M , again for a large number of noise samples. In the noise-free case, we expect that E is a monotonically non-increasing function of M . However, this behavior is exhibited *provided that* the M measurement angles are arranged in what we call an ‘optimal’ manner, such that a circumscribing polygon with faces at those angles is of minimum area among all such polygons with the same number of faces. The monotonicity follows trivially from the fact that given an optimal configuration of M measurement angles, placing the $(M+1)$ st measurement at any angle cannot increase E and that the optimal configuration for $M+1$ measurements cannot be worse than this suboptimal

³Throughout the thesis, errorbars on plots extend one standard deviation from top to bottom. Occasionally, Monte Carlo plots will be shown without errorbars. In these cases, the number of noise samples used was so high that the errorbars would not be visible.




Figure 3.8: E versus σ for the standard ellipse. $M = 12$.

construction.

If the measurement angles are not optimally arranged, then monotone behavior cannot be expected, in general. One simple example that yields non-monotonic sequences in M is the set of uniformly-spaced measurements (see below). Throughout the present discussion, our focus will be on minimizing the error E rather than on achieving monotonicity. In fact, monotonicity can be achieved rather easily at the expense of incurring large errors. One such example consists of starting with three angles, and successively adding angles to the configuration without changing existing angles.

The problem of determining the optimal configuration of M measurement angles for an ellipse with some specified eccentricity appears to be a difficult one, requiring the solution of a system of quartic equations.⁴ Consequently, we examine E versus M

⁴Our formulation is the following. We wish to find M points on the ellipse such that the straight

for two suboptimal configurations. The first such configuration consists of uniformly-spaced angles. Figure 3.9 depicts plots of $\log_{10} E$ versus M for the three semimajor axis lengths $a = 3, 6,$ and 9 , in the noise-free case. For each of the plots, although the behavior at a fine scale is certainly not monotone, the trend at a somewhat coarser scale is. The local minima occurring at $M = 4k$, k an integer, are due to the fact that for an ellipse, measurements placed at 90° and 270° correspond to circumscribing polygons with much smaller areas than those arising from sets of measurements not including 90° and/or 270° , for comparable values of M .




Figure 3.9: E versus M for the standard ellipse, for several values of eccentricity. Uniformly-spaced measurements, $\sigma = 0$.

An interesting observation is that each plot for large enough M becomes monotonically non-increasing. We conjecture that, in the noise-free case with uniformly-spaced lines tangent to the ellipse at these points bound a polygon of minimum area. The method of Lagrange multipliers yields a system of quartic equations.

measurements, E becomes monotonically non-increasing after some value M^* which increases with eccentricity. (For the plots in Figure 3.10, values of M^* are 9, 33, and 54, corresponding to $a = 3, 6,$ and $9,$ respectively.) With the addition of measurement noise, the plots tend to smooth out somewhat. An example is shown in Figure 3.10, for the noise levels $\sigma = 1.0$ and $\sigma = 2.0$ and semimajor axis $a = 6.0$.




Figure 3.10: E versus M for the standard ellipse, for several noise levels. Uniformly-spaced measurements, $a = 6.0$.

The second configuration is also suboptimal but has some heuristic basis: a support line at an angle for which the ellipse has low curvature (e.g., 90° or 270°) remains close to the ellipse for a larger range of angles than a support line at a high-curvature angle (e.g., 0° or 180°), and therefore gives a smaller contribution to the symmetric difference. Hence, we believe that the angles in the optimal configuration are perhaps distributed with respect to the ellipse's radius of curvature, given by $K^{-1}(\theta) = (\cos^2 \theta + a^2 \sin^2 \theta)^{3/2} / a$ for $b = 1$. We construct our suboptimal configura-

tion by placing the first measurement at $\theta_1 = 90^\circ$ and then placing the other $M - 1$ measurements at angles such that the area under $K^{-1}(\theta)$ is divided into M equal⁵ parts by the M intervals $\{[\theta_i, \theta_{i+1}]\}$. We then take the measurement angle closest to 270° and set it equal to 270° , to gain the advantage that a measurement at this angle provides. This configuration shall be referred to as the ‘radius-of-curvature’ configuration.

Figure 3.11 depicts plots of $\log_{10} E$ versus M in the noise-free case, for $a = 3, 6,$ and 9 . Each plot compares $\log_{10} E$ for the uniformly-spaced and radius-of-curvature configurations. In each case, the latter configuration clearly outperforms the former.

In concluding this section, we observe the dependence of E on object eccentricity, as illustrated in Figure 3.12 for several values of M in the case of noise-free uniformly-spaced measurements. The value of M is seen to play a substantial role, as multiples of four yield asymptotically bounded behavior, while non-multiples yield unbounded behavior in that the ratio of the area of the circumscribing M -gon to the area of the ellipse increases without bound with increasing eccentricity. Again, this difference is due to the presence of the angles 90° and 270° in the constellation only for $M = 4k$. We illustrate the importance of not having these angles in the case of $M = 6$, in Figure 3.13. In fact, we conjecture that E is a monotonically non-decreasing function of eccentricity, if the constellation of measurement angles is held fixed. However, we conjecture that if, for each value of eccentricity, the measurement angles are placed in the optimal manner defined above, then E is a monotonically non-increasing function of eccentricity. We provide some support for this statement by replacing the uniformly-spaced measurements in the example of Figure 3.12 by

⁵The division into M parts actually results in approximately equal areas, in our implementation. As a result, each angle is up to 0.1° away from the desired value.



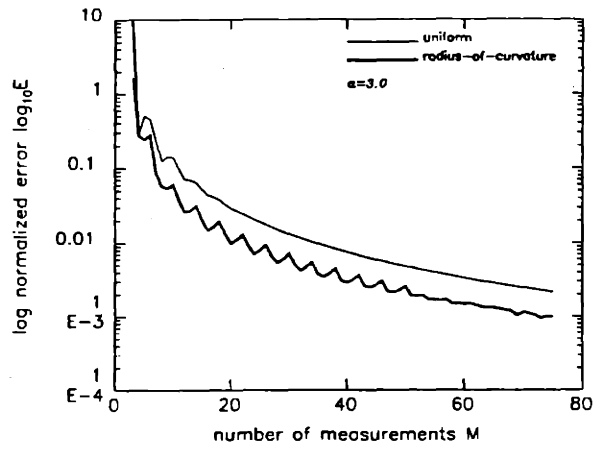


Figure 3.11: E versus M for the standard ellipse, for uniformly-spaced and radius-of-curvature configurations. (a) $a = 3.0$, (b) $a = 6.0$, and (c) $a = 9.0$. $\sigma = 0$.

measurements spaced according to the radius-of-curvature approach. The results are shown in Figure 3.14 and indicate that that much might be gained in judiciously choosing measurement angles for highly eccentric objects. (Note the difference in scale along the y -axis of the plots in Figures 3.12 and 3.14.) Finally, the addition of measurement noise does not affect qualitative behavior much, but essentially serves to shift the plots upward, as greater errors are incurred on average (see Figure 3.15).

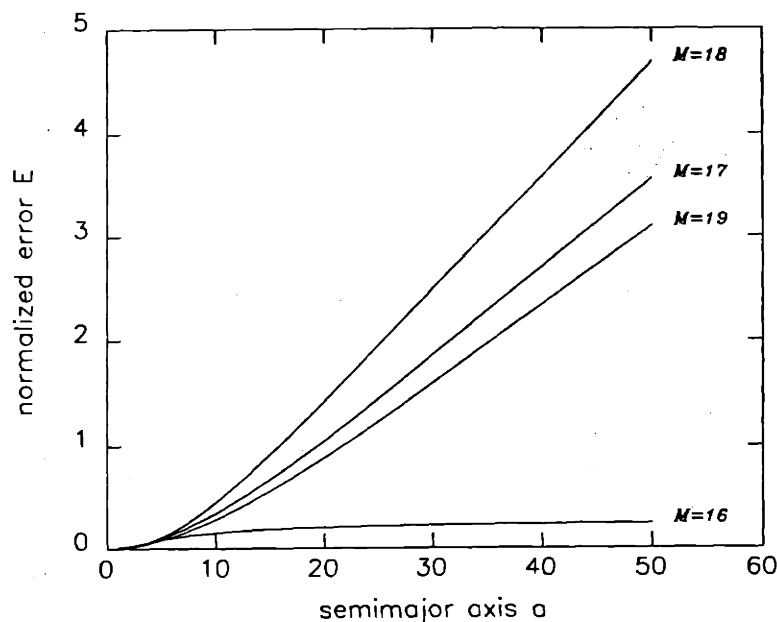


Figure 3.12: E versus a for an ellipse, for several values of M , and uniformly-spaced measurement angles. $b = 1.0$, $\sigma = 0$.

3.4 Perimeter Bias in the Estimator

The estimation procedure NUA does not appear to be an unbiased estimator in that it produces estimates on average larger than the true object. In this section, we point out the origin of this bias and quantify it through a Monte Carlo approach.




Figure 3.13: Sets of support lines for three ellipses of increasing eccentricity, illustrating the importance of including 90° and 270° in a measurement angle configuration. (a) $a/b = 1.0$ (b) $a/b = 5.0$ (c) $a/b = 10.0$, with ellipses normalized to have equal areas.




Figure 3.14: E versus a for an ellipse, for several values of M , and radius-of-curvature angles. $b = 1.0$, $\sigma = 0$.




Figure 3.15: E versus a for an ellipse, for uniformly-spaced measurement angles and several noise levels. $b = 1.0$, $M = 17$

Recall from earlier discussion that if the the noisy support vector $y = h + n$ (where n is a zero-mean Gaussian random vector) lies in the support cone \mathcal{C} , then the estimate is given by $\hat{h} = y$. On the other hand, if $y \notin \mathcal{C}$, then the estimation procedure consists of projecting y onto \mathcal{C} , to yield the estimate $\hat{h} = P_1(y)$ (see Figure 2.7).

In order to discuss bias in the estimator, we need some quantitative measure of the size of a reconstructed object. One such measure is the magnitude of the projection of \hat{h} onto the vector $e = [1 \ 1 \ \dots \ 1]^T$, where the projection is denoted by $P_2(\hat{h})$. This quantity is proportional (with proportionality constant \sqrt{M}) to the perimeter of the reconstructed set, since the perimeter of a set having support function $\hat{h}(\theta)$ is given by $\int_0^{2\pi} \hat{h}(\theta) d\theta$, an expression whose discrete counterpart is $\frac{1}{M} \hat{h}^T e = \frac{1}{M} \sum_{i=1}^M \hat{h}_i = \frac{1}{\sqrt{M}} P_2(\hat{h})$ for a set having support vector h with uniformly-spaced measurements.

Although the support cone is a cone in \mathbb{R}^M , we consider an analogous, lower dimensional problem due to the difficulty in visualizing objects in a high-dimensional space. We expect that the low-dimensional results reflect behavior in \mathbb{R}^M . In particular, we examine the problem in three dimensions, where a Gaussian noise ‘sphere’ is centered on the true support vector $h \in \mathcal{C}$. Slicing the cone along e to present a cross-sectional view (as seen in Figure 3.16), we mark the projection of h onto e as point $A = P_2(h)$. Now a particular invalid noisy support vector (i.e., one that does not lie in \mathcal{C}) such as y_0 (marked in the figure) first projects onto \mathcal{C} to yield $\hat{h}_0 = P_1(y_0)$. For this particular noisy vector, the subsequent projection of \hat{h}_0 onto e , denoted by $P_2(\hat{h}_0)$, is seen to be farther along e than point A , indicating that the reconstruction arising from y_0 is larger than⁶ the true object in the sense that $P_2(\hat{h}_0) > P_2(h) = A$.

In considering the noisy vectors $\{y\}$ arising from the entire noise sphere, we must distinguish between valid and invalid vectors, since a single projection is sufficient to determine perimeter for valid vectors, whereas the projection P_1 must precede P_2 for invalid ones. Denoting the Gaussian probability density function of the noise vector n by

$$p_n(n_0) = \frac{1}{(2\pi\sigma^2)^{M/2}} \exp\left(-\frac{1}{2\sigma^2} n_0^T n_0\right) \quad (3.10)$$

the mean perimeter of the estimated objects is given by

$$\begin{aligned} \mu &= \int_{h+n_0 \text{ valid}} P_2(h+n_0)p_n(n_0) dn_0 + \int_{h+n_0 \text{ invalid}} P_2(P_1(h+n_0))p_n(n_0) dn_0 \\ &> \int_{h+n_0 \text{ valid}} P_2(h+n_0)p_n(n_0) dn_0 + \int_{h+n_0 \text{ invalid}} P_2(h+n_0)p_n(n_0) dn_0 \end{aligned}$$

⁶Actually, this indicates that the reconstruction is larger than the polygonal object bounded by the support lines of the true support vector h , but for convenience we use the phrase “larger/smaller than the true object.”




Figure 3.16: A cross-sectional view of the support cone geometry giving rise to bias in the estimator for (a) a particular noisy support vector and (b) the entire noise sphere.

$$\begin{aligned}
&= \int_{\text{all } h+n_0} P_2(h+n_0)p_n(n_0) dn_0 \\
&= A,
\end{aligned} \tag{3.11}$$

where the inequality clearly follows from the fact that, for an invalid support vector $y_0 = h + n_0$, the double projection $P_2(P_1(y_0))$ is farther along the support cone axis e than the single projection $P_2(y_0)$ is (see Figure 3.17). The last two steps follow from linearity and symmetry, respectively. Hence the estimator is biased in perimeter, with the bias given by $\mu - A$.

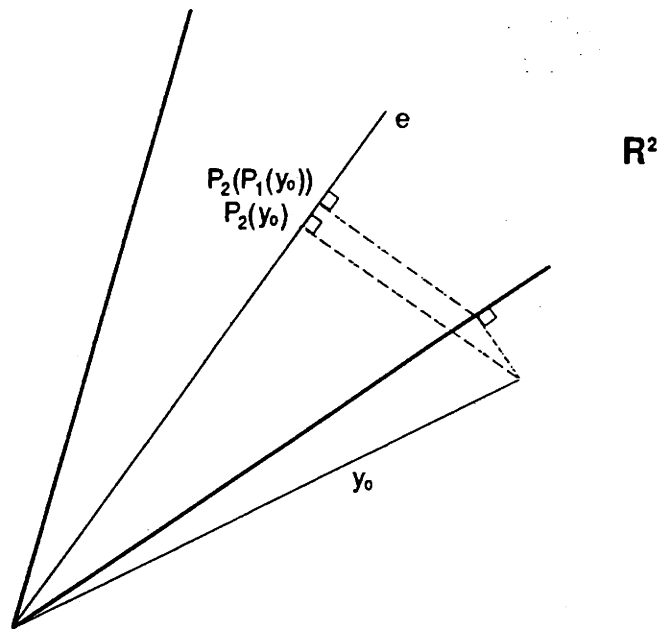


Figure 3.17: Support cone geometry illustrating the fact that $P_2(P_1(y_0))$ is farther along e than $P_2(y_0)$ is, leading to bias in the estimator.

We expect that the bias varies with noise level and object eccentricity. As the variance of the noise increases, a larger fraction of the noisy support vectors are invalid ones. Consequently, the difference between $\int_{h+n_0 \text{ invalid}} P_2(P_1(h+n_0))p_n(n_0) dn_0$ and $\int_{h+n_0 \text{ invalid}} P_2(h+n_0)p_n(n_0) dn_0$ increases, and from Equation 3.11, the bias increases.

To explain the dependence of bias on object eccentricity, we define the *distance to inconsistency* Δh^* for a *valid* support vector h as the Euclidean distance from h to the boundary \mathcal{B} of the support cone. Letting Δh denote a vector in \mathbb{R}^M , Δh^* can be expressed precisely as

$$\Delta h^* = \min_{(h+\Delta h) \text{ on } \mathcal{B}} \|\Delta h\|_2 \quad (3.12)$$

We conjecture that Δh^* is a monotone decreasing function of object eccentricity. What we will now show is that as eccentricity tends to infinity, Δh^* tends to zero. We do this by (1) defining an alternate, non-Euclidean distance to inconsistency Δh_{nE} ; (2) showing that Δh_{nE} decreases from some nonzero value to zero asymptotically with increasing object eccentricity; and (3) upper-bounding Δh^* by Δh_{nE} . First, define the non-Euclidean distance to inconsistency for a triplet of consistent support lines L_{i-1}, L_i, L_{i+1} (adjacent in angle) as being the minimum value $\Delta h_{nE, \text{triplet}}(L_{i-1}, L_i, L_{i+1})$ which would have to be added to *only* the central support value h_i to make the triplet inconsistent. Then, Δh_{nE} is defined as $\Delta h_{nE} = \min_i \Delta h_{nE, \text{triplet}}(L_{i-1}, L_i, L_{i+1})$.

Now consider the sets of noise-free support lines shown in Figure 3.18 for ellipses of several eccentricities. As will be the convention throughout the remainder of this section, we parametrize an ellipse by the ratio of the lengths of its semimajor and semiminor axes, a/b , with its area πab constrained to equal π . Although each set of support lines in the figure is consistent, the higher the value of a/b —and hence the higher the eccentricity—the lower the value of Δh_{nE} . (For each ellipse in the figure, we mark the triplet that is closest to being inconsistent by bold support lines, since it is this triplet for which $\Delta h_{nE} = \Delta h_{nE, \text{triplet}}(L_{i-1}, L_i, L_{i+1})$.) We also note the following two facts: (1) from Figure 3.18a where $a/b = 1.0$, we see that $\Delta h_{nE} \neq 0$ and

(2) it is apparent from observing the triplet of bold support lines in the three figures that Δh_{nE} asymptotically approaches zero. From all of this, we may conclude that Δh_{nE} decreases from some nonzero value to zero monotonically. Since the Euclidean distance Δh^* from the vector h to \mathcal{B} must be less than or equal to the distance Δh_{nE} from the vector h to \mathcal{B} achieved by changing only one support value in h , Δh^* must also decrease to zero (though not necessarily monotonically) with increasing object eccentricity. Hence, the Euclidean distance between the support vector h and the boundary of the support cone \mathcal{C} asymptotically approaches zero with increasing object eccentricity.




Figure 3.18: Three sets of support lines that illustrate the notion of non-Euclidean distance to inconsistency Δh_{nE} for a set. (a) $a/b = 1.0$ (b) $a/b = 2.0$ (c) $a/b = 3.0$.

For a support vector h that is closer to \mathcal{B} than some other support vector, a larger fraction of the noisy support vectors are invalid ones. Therefore, the difference between $\int_{h+n_0 \text{ invalid}} P_2(P_1(h+n_0))p_n(n_0) dn_0$ and $\int_{h+n_0 \text{ invalid}} P_2(h+n_0)p_n(n_0) dn_0$ increases, and from Equation 3.11, the bias increases. From all of this, we reach the conclusion that perimeter bias increases with object eccentricity.

In an effort to quantify the bias and observe the trends predicted above, let us describe a simple experiment. To be specific, we take an ellipse with some fixed

eccentricity and generate the NUA estimate from the noisy support vector arising from the sum of the ellipse's true support vector and a Gaussian random vector with standard deviation σ . The support vector is 12-dimensional, with support values at uniformly-spaced angles on $[0^\circ, 360^\circ)$. Given the NUA estimate, we determine the perimeter of the reconstruction. Repeating for a large number of noise samples (here, 5,000) yields an estimate of the average normalized bias β , defined as

$$\beta = \frac{\text{mean perimeter of estimate} - \text{perimeter of true object}}{\text{perimeter of true object}},$$

that one should expect at that eccentricity and noise level. In Figure 3.19a, β is plotted as a function of noise level σ for ellipses of three eccentricities ($a/b = 1.0, 2.0,$ and 5.0), while in Figure 3.19b it is plotted as a function of eccentricity for three noise levels ($\sigma = 0.10, 0.15,$ and 0.25).

The plots in Figure 3.19 reflect the behavior predicted by the preceding analysis. Specifically, the bias increases with increasing noise level, with greater bias for more eccentric ellipses. The bias also increases with increasing eccentricity, with greater bias for higher noise levels. However, the bias appears to very quickly reach some asymptotic level in the plots in Figure 3.19b. This is to be expected for the following reason. The true support vector h asymptotically approaches B as the semimajor axis a of the ellipse increases. As a result, there is very little change in bias for changes in a , at sufficiently high values of a .

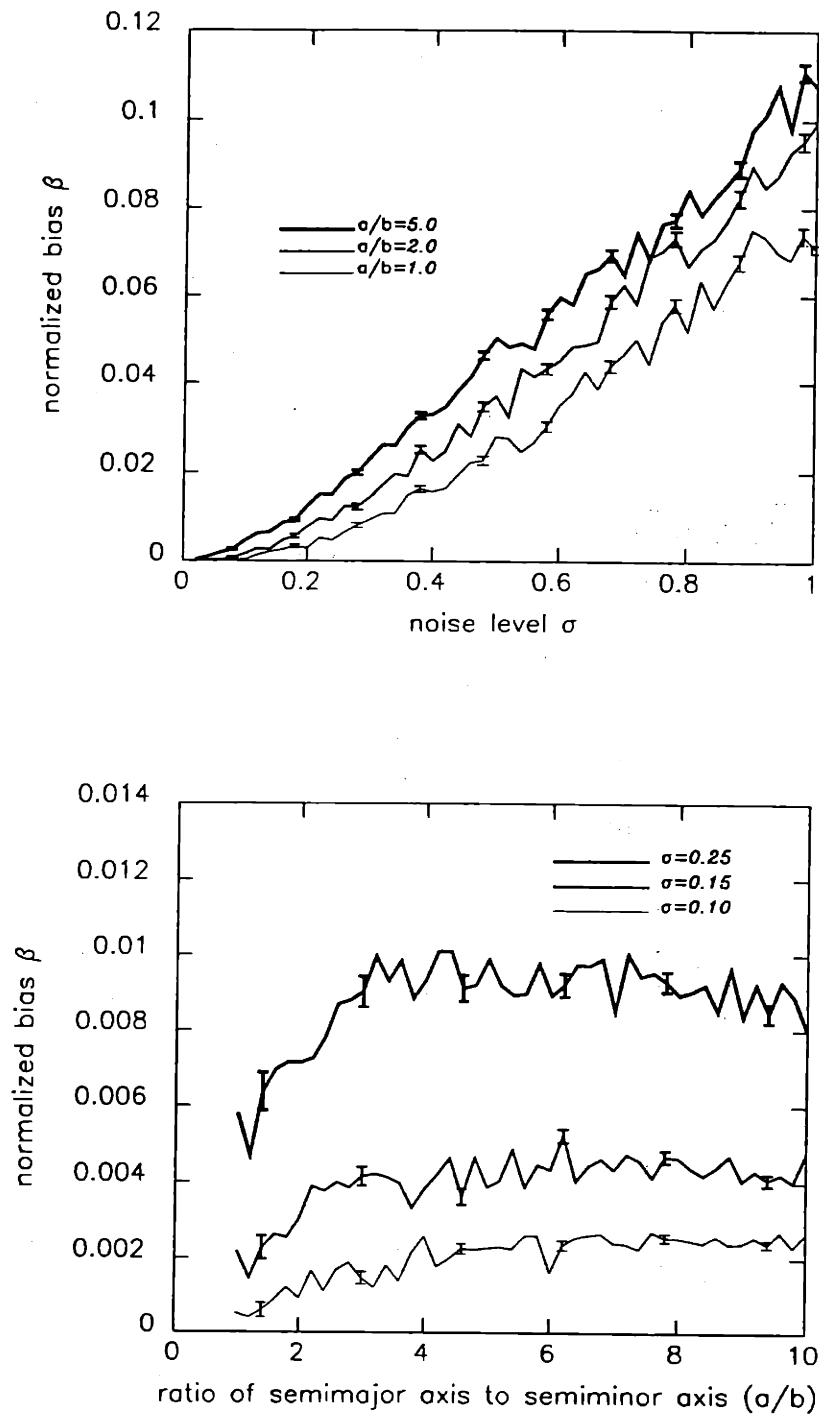


Figure 3.19: Normalized bias β (a) versus noise level σ , for three values of eccentricity and (b) versus a/b , for three noise levels.

Chapter 4 Best N -gon Fitting M Measurements with Fixed Reconstruction Angles

In the last chapter, we pointed out that there are more measurements per degree of freedom in the solution for a polygon than there are for an ellipse. Here we wish to exploit this fact by explicitly imposing the restriction that the solution space has N degrees of freedom if we know *a priori* that the true object is an N -sided polygon. With this restriction, many measurements will serve to collectively fix each degree of freedom (polygonal side) in the reconstruction.

In particular, we now consider the problem of determining the best N -sided polygon, with specified face angles, that fits a set of noisy support values at M measurement angles. For example, one might wish to reconstruct the best equilateral triangle having some specified orientation given a set of, say, twenty noisy measurements of an object known *a priori* to have this shape. Applications of this problem might arise in computer vision, tactile sensing, or target reconstruction scenarios, in which rough *a priori* shape data is used along with the measurements to obtain size and more accurate shape estimates of an object. Note that this problem is the third formulation of the general polygonal estimation problem that we outlined in the introduction to Chapter 3.

4.1 The Estimator BNGON

In formulating this problem, we let $\{\theta_1, \theta_2, \dots, \theta_M\}$, $\{y_1, y_2, \dots, y_M\}$, and $\{\phi_1, \phi_2, \dots, \phi_N\}$ denote the M measurement angles, the measured support values at these angles, and the N reconstruction angles, respectively. Given these quantities, we wish to estimate

an N -gon specified by the consistent set of support values $\{h_\phi(\phi_1), h_\phi(\phi_2), \dots, h_\phi(\phi_N)\}$ which minimizes

$$J(h_\phi(\phi_1), h_\phi(\phi_2), \dots, h_\phi(\phi_N)) = \sum_{i=1}^M (h_\phi(\theta_i) - y_i)^2, \quad (4.1)$$

where $h_\phi(\theta_i)$ denotes the value at θ_i of the support function $h_\phi(\cdot)$ of our estimated N -gon. As shown in Figure 4.1, Equation 4.1 corresponds to finding a set of support values at the reconstruction angles that minimizes the sum of the squared deviations between the measured support values and the piecewise sinusoidal curve given by Equation 2.14 that joins the reconstructed support values.




Figure 4.1: Support function picture of the best N -gon problem.

Letting ϕ_{L_i} and ϕ_{R_i} denote the reconstruction angles immediately to the left and right of the i^{th} measurement angle θ_i , and letting h_{L_i} and h_{R_i} denote the corresponding reconstructed support values, we have from Equation 2.14 that the support function

of the reconstructed object is given by $h_\phi(\theta) = v_x \cos \theta + v_y \sin \theta$, $\phi_{L_i} \leq \theta \leq \phi_{R_i}$, where (v_x, v_y) are the coordinates of the N -gon vertex arising from the intersection of the support lines at angles ϕ_{L_i} and ϕ_{R_i} . Hence the support value of the reconstructed N -gon at the measurement angle θ_i is given by

$$h_\phi(\theta_i) = v_x \cos \theta_i + v_y \sin \theta_i. \quad (4.2)$$

Replacing $\theta_{i-1}, \theta_{i+1}, h_{i-1}$, and h_{i+1} by $\phi_{L_i}, \phi_{R_i}, h_{L_i}$, and h_{R_i} in Equation 3.3 and substituting the result into Equation 4.2 yields, after some manipulation,

$$h_\phi(\theta_i) = \frac{\sin(\phi_{R_i} - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{L_i} + \frac{\sin(\theta_i - \phi_{L_i})}{\sin(\phi_{R_i} - \phi_{L_i})} h_{R_i}. \quad (4.3)$$

Our estimation algorithm, to be referred to as BNGON, arises on substituting Equation 4.3 into Equation 4.1:

Estimator BNGON

$$\hat{h}_\phi = \begin{bmatrix} \hat{h}_\phi(\phi_1) \\ \hat{h}_\phi(\phi_2) \\ \vdots \\ \hat{h}_\phi(\phi_N) \end{bmatrix} = \arg \min_{Ch_\phi \geq 0} (Ah_\phi - y)^T Ah_\phi, \quad (4.4)$$

where

$$y = [y_1 \ y_2 \ \cdots \ y_M]^T,$$

$$C = \begin{pmatrix} -\sin(\phi_2 - \phi_N) & \sin(\phi_1 - \phi_N) & 0 & 0 & \sin(\phi_2 - \phi_1) \\ \sin(\phi_3 - \phi_2) & -\sin(\phi_3 - \phi_1) & \sin(\phi_2 - \phi_1) & 0 & 0 \\ 0 & \sin(\phi_4 - \phi_3) & -\sin(\phi_4 - \phi_2) & \sin(\phi_3 - \phi_2) & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots \\ \sin(\phi_N - \phi_{N-1}) & 0 & 0 & \sin(\phi_1 - \phi_N) & -\sin(\phi_1 - \phi_{N-1}) \end{pmatrix}, \quad (4.5)$$

and A is an $M \times N$ matrix mapping the M support values at the $\{\theta_i\}$ to induced support values at the $\{\phi_i\}$. The i^{th} row of the matrix A , corresponding to the i^{th} measurement, has two adjacent (modulo N) non-zero entries, $l_i = \frac{\sin(\phi_{R_i} - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})}$ and $r_i = \frac{\sin(\theta_i - \phi_{L_i})}{\sin(\phi_{R_i} - \phi_{L_i})}$, corresponding to the reconstruction angles ϕ_{L_i} and ϕ_{R_i} on either side of θ_i . In the degenerate case in which a measurement angle has the same value as a reconstruction angle, the two entries become zero and unity.

Note that each measurement angle θ_i gives rise to a single nonzero row in the matrix A . However, an adjacent pair of reconstruction angles may give rise to many, one, or no nonzero columns in A if there are many, one, or no measurement angles between them. As a result, A has the following pseudobanded structure:

$$A = \begin{pmatrix} l_1 & r_1 & 0 & 0 & 0 \\ 0 & l_2 & r_2 & 0 & 0 \\ 0 & 0 & l_3 & r_3 & 0 \\ r_4 & 0 & 0 & 0 & l_4 \\ r_5 & 0 & 0 & 0 & l_5 \end{pmatrix}$$

for the following example configuration of angles:

$$0 < \phi_1 < \theta_1 < \phi_2 < \theta_2 < \phi_3 < \theta_3 < \phi_4 < \phi_5 < \theta_4 < \theta_5 < 2\pi.$$

Since the cost function in Equation 4.4 is quadratic in the reconstructed support values and the consistency constraint is linear, the problem can be solved by QP techniques.

4.2 Examples

Two examples of BNGON, similar to those discussed in Section 3.2, are now considered. The first example consists of reconstructing the best triangle with reconstruction angles at 7.125° , 82.875° , and 270° equal to those of the standard triangle (introduced in the previous chapter), given $M = 24$ uniformly-spaced noisy ($\sigma = 0.25$) support measurements. The pictures in object space and in support function space are shown in Figure 4.2, with the reconstructed object incurring an error $E = 0.17$ with respect to the true object.

The BNGON reconstruction in the figure originates from the same set of measurements as the NUA reconstruction of Figure 3.4 (i.e., the same noise realization was used), allowing us to compare the two. From a visual comparison, it is clear that the prior information that the true object is a triangle with known face angles allows BNGON to outperform NUA. This is also seen quantitatively by noting that $E_{\text{BNGON}} = 0.17 < E_{\text{NUA}} = 1.56$.

Our second example is one of estimating the rectangle that best fits a set of uniformly-spaced support measurements of an ellipse, where the sides of the rectangle are aligned with the axes of the ellipse. We use the standard ellipse introduced in Chapter 3 in our example. To see what such a fit looks like under noise-free





Figure 4.2: An example of BNGON for the standard triangle in (a) object space and (b) support function space, for the case $M = 24$, $\sigma = 0.25$.

conditions, we first consider the reference reconstruction shown in Figure 4.3, for the case $M = 24$. A rectangular reconstruction from $M = 48$ uniformly-spaced noisy ($\sigma = 0.5$) measurements is shown in Figure 4.4. Incidentally, this example can be considered as one of estimating the approximate size of an ellipse of known orientation.

4.3 Performance Assessment

In this section, we assess the performance of BNGON against measurement parameters and compare this performance to that of NUA. We then evaluate the robustness of BNGON to incorrect prior information.

Figure 4.5 contains plots of the reconstruction error E versus noise level σ for $M = 24$ uniformly-spaced measurements of the standard triangle, for the algorithms BNGON and NUA. The reconstruction angles used for BNGON are the correct face angles of the standard triangle. The value of E at each data point is given by the average error from 1,000 noise realizations.

While the error plot for BNGON increases with increasing σ as NUA does, there are two important differences. First, BNGON clearly outperforms NUA as we expected. Second, the error for BNGON equals zero at $\sigma = 0$ unlike that for NUA. This is due to the fact that BNGON is fitting a triangle with face angles equal to those of the standard triangle to a set of perfect measurements. This prior knowledge allows BNGON to produce an exact reconstruction. In contrast, since the 24 measurement angles do not include two face angles of the triangle (the congruent faces), NUA is not able to produce an exact reconstruction.

In Figure 4.6, we show plots of the reconstruction error E versus number of measurements M for the noise level $\sigma = 0.25$, for BNGON and NUA. For each value of M , the

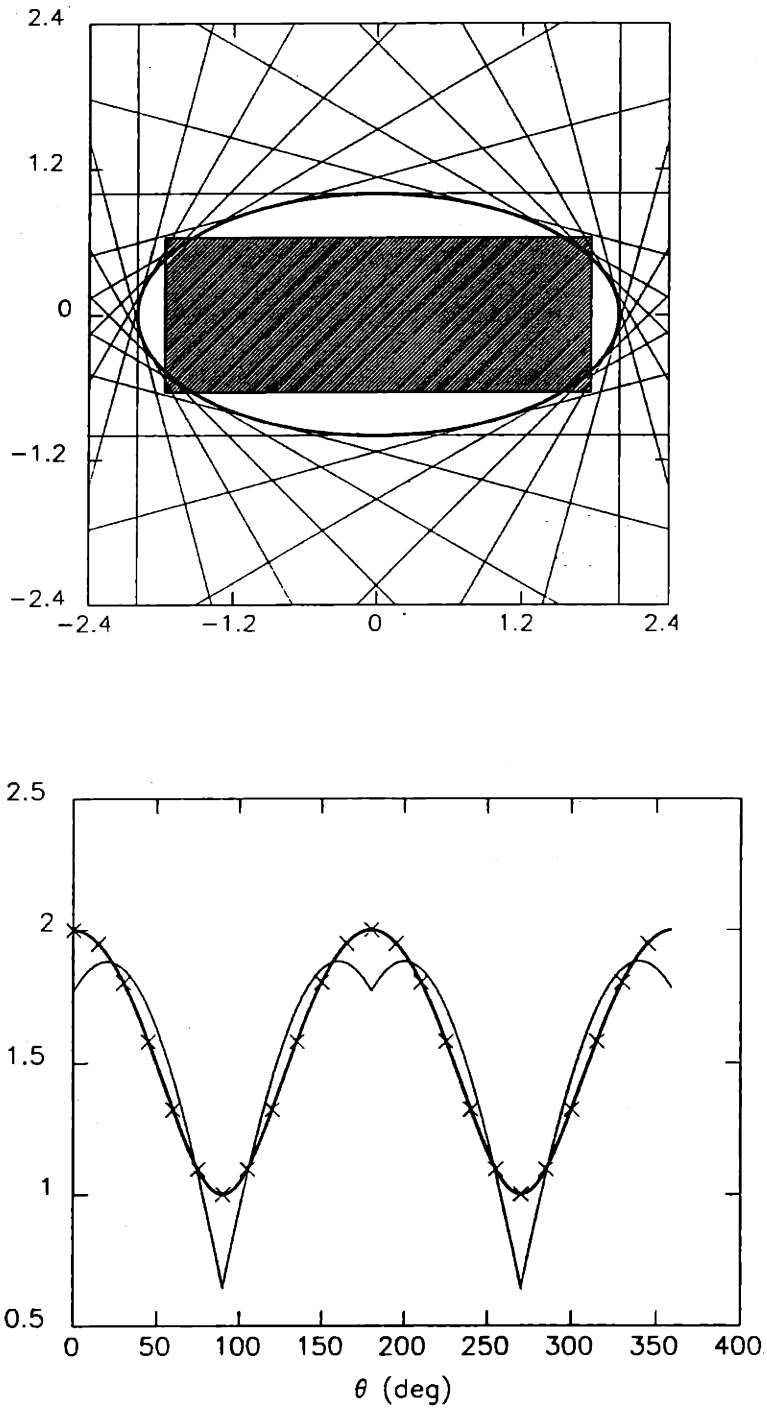


Figure 4.3: A reference example of BNGON for the standard ellipse in (a) object space and (b) support function space, for the case $M = 24$, $\sigma = 0$.





Figure 4.4: An example of BNGON for the standard ellipse in (a) object space and (b) support function space, for the case $M = 48$, $\sigma = 0.5$.




Figure 4.5: E versus σ for the standard triangle, for the algorithms BNGON and NUA. $M = 24$.

measurements are uniformly-spaced. For BNGON, 1,000 noise realizations were used to generate each data point. For NUA, 500 were used.

The error plots for BNGON and NUA exhibit the same general trend: both decrease with increasing M , as expected. Here also there are several notable differences. Not surprisingly, BNGON clearly outperforms NUA. Second, the difference in reconstruction error for the two algorithms decreases with M . This is because, for low M , the sparse measurement angles can be quite far from the face angles of the triangle, thereby causing large errors in the NUA reconstructions. However, since BNGON always uses reconstruction angles that are the triangle's face angles, the reconstruction error is not much larger for small values of M than for large values of M . Finally, the BNGON error plot is much smoother than the NUA plot, especially for small values of M , for much the same reasons as we just discussed. Namely, the choice of measurement




Figure 4.6: E versus M for the standard triangle, for the algorithms BNGON and NUA. $\sigma = 0.25$.

angles is much more critical for NUA than it is for BNGON.

Unlike NUA, the performance of BNGON also depends on the quality of the prior information as to numbers and values of reconstruction angles. Hence, in addition to the assessing the performance of BNGON against measurement parameters and comparing its performance to NUA's, we also focus our assessment on the robustness of the algorithm to incorrect priors. Throughout the remainder of this section, all reconstructions will be based on 24 uniformly-spaced noisy support measurements of the standard triangle.

Figure 4.7 depicts plots of normalized error E versus noise level σ for several types and degrees of inaccuracy in the prior. The dashed lines correspond to rotating the reconstruction angle at 7.125° by either $\alpha = 5^\circ$ or $\alpha = 10^\circ$ (so that $\phi'_1 = \phi_1 + \alpha$), while the solid lines correspond to rotating the entire configuration of reconstruction

angles by either 5° or 10° (so that $\phi'_i = \phi_i + \alpha, \forall i$). The dotted line corresponding to a correct prior is included for reference.




Figure 4.7: E versus σ for the standard triangle, for various rotations of the reconstruction angles.

Several observations can be made. Perhaps most importantly, rotating a single reconstruction angle by some amount is worse than rotating the entire configuration by the same amount. Another, somewhat interesting observation is that as σ increases to rather high levels, plots corresponding to entire-configuration rotations merge with the reference plot, while those arising from single-angle rotations do not. Finally and trivially, the error E increases with α .

A natural question to ask is “How much error in prior information should be tolerated before deciding to use NUA, which uses no prior information?” In Figure 4.8 we answer this question for a particular noise level ($\sigma = 0.25$), by plotting E versus offset α for single-angle and entire-configuration errors and comparing with the base-

line performance of NUA. For values of α below 69° , entire-configuration errors are less damaging than corresponding single-angle errors, as we saw earlier. On noting the intersections of the BNGON plots with the NUA baseline, we may conclude that for this particular noise level, one should tolerate single-angle errors of up to $\approx 17^\circ$ and entire-configuration errors of up to $\approx 29^\circ$ before abandoning the BNGON algorithm.

However, we should *not* abandon the prior information itself and resort to NUA. In the next chapter, we formulate an estimation problem identical to BNGON except for the fact that the estimator is free to rotate the constellation of reconstruction angles to achieve a minimum cost solution. The solution consists of an orientation estimate and a shape estimate produced by BNGON using the reconstruction angles rotated according to the orientation estimate. The resulting algorithm overcomes entire-configuration errors in the prior for moderate noise levels; for high noise levels, the orientation estimate can be quite poor (as we will see in Section 5.4) so that entire-configuration errors are not corrected.

Next, we observe how performance degrades as the number of reconstruction angles N increases, again with 24 uniformly-spaced noisy ($\sigma = 0.25$) measurements of the standard triangle. That is, we start with reconstruction angles at $\phi_1 = 7.125^\circ$, $\phi_2 = 82.875^\circ$, and $\phi_3 = 270^\circ$ for $N = 3$. For all values of $N > 3$, we choose ϕ_N such that it lies halfway between the most distant adjacent pair of the previous $N - 1$ reconstruction angles. For each value of N , we solve the resulting BNGON problem in a Monte Carlo fashion in order to generate a data point in Figure 4.9. Examples of BNGON reconstructions for the values $N = 6$ and $N = 18$ are shown in Figure 4.10, with their NUA counterpart that arises from the same noise realization. Through this constructive process, a constellation of 24 more or less uniformly-spaced reconstruction angles is built up. The performance of BNGON for the set of reconstruc-




Figure 4.8: E versus α for two types of rotation errors in the BNGON prior, compared with the baseline E associated with NUA.

tion angles constructed in this manner is compared with the baseline performance of NUA, which uses the set of 24 uniformly-spaced measurement angles. From the plot, we may conclude that for a polygon of N sides, as long as the N reconstruction angles are known, adding extraneous reconstruction angles degrades performance but not to the extent that switching to NUA is better. This is particularly apparent from the fact that near $N = 24$, BNGON produces estimates which incur on average somewhat less reconstruction error E than those produced by NUA, indicating that the original three reconstruction angles that are not available to NUA are quite helpful to BNGON.

Incidentally, the magnitude of the difference in the error E at $N = 24$ for the two algorithms depends on the number and values of the measurement angles relative to the face angles of the triangle. For example, if the values of the measurement angles were such that three of the measurement angles coincided with the triangle's face






Figure 4.9: E versus N , illustrating the effect of extraneous information in the BNGON prior, compared with the baseline E associated with NUA.

angles, then the NUA baseline would be somewhat lower. In fact, its height would be roughly the error associated with BNGON at $N = 24$. Since this is not the case with uniformly-spaced measurements, the errors at $N = 24$ differ for the two algorithms. Also, as the *number* of measurement angles M increased, the NUA baseline would be lowered. This is because with large numbers of (uniformly-spaced) measurement angles, the three measurement angles nearest to the three face angles of the triangle would be nearer to these face angles than the three nearest would be for a small number of measurement angles.

In the preceding discussion, we investigated the effect of *overestimating* the number of reconstruction angles N . If, in fact, we *underestimate* the number of reconstruction angles by omitting one or more of them, the performance of BNGON degrades to a larger extent than it does with overestimation. A particularly illustrative example consists






Figure 4.10: Reconstructions illustrating the effect of extraneous information in the BNGON prior, with comparisons to the corresponding NUA reconstruction. All reconstructions arise from an identical set of measurements. (a) $N = 6$, BNGON: $E = 0.69$. (b) $N = 18$, BNGON: $E = 1.4$. (c) NUA: $E = 1.6$.

of taking the trapezoid that results from chopping off the top half of the standard triangle and producing reconstructions based on 3 and 6 reconstruction angles. In this case, the trapezoid has a face angle at 90° in addition to the face angles of the standard triangle. In the example of underestimation, we omit the reconstruction angle at 90° . In the example of overestimation, we add two extraneous reconstruction angles, at 45° and 135° . The reconstructions that result (in the noise-free case) are shown in Figure 4.11. While the $N = 3$ reconstruction has a error of $E = .67$, the $N = 6$ reconstruction is perfect ($E = 0$). Clearly, extraneous information is more tolerable than insufficient information in this example. We expect that this is true in general.

4.4 Nonuniqueness of Solutions

Under certain conditions, given a set of measurements and reconstruction angles, the estimator BNGON yields a nonunique solution. That is, an entire family of solutions achieves minimum cost. An example of nonuniqueness occurs for the set of measurements $\{1, 1, 1, 1\}$ at the angles $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$, with reconstruction angles at $\{45^\circ, 135^\circ, 225^\circ, 315^\circ\}$. Figure 4.12 shows the measured support lines and the (inscribed) rectangular reconstruction achieving zero cost, for several members of the family of solutions. Nonuniqueness for the case in which minimum cost is nonzero is illustrated in Figure 4.13, where the measured support values are $\{2, 1, 2, 1\}$ and the measurement and reconstruction angles are the same as before.

Whether or not a particular solution is nonunique depends on the relative numbers and values of the measurement and reconstruction angles, and on the measured support values. Consequently, we analyze the problem for several separate cases: (1)





Figure 4.11: Examples of BNGON when the number of reconstruction angles is (a) underestimated and (b) overestimated.




Figure 4.12: Nonuniqueness of solution arising in a case with zero cost. Shown are the measured support lines as well as several inscribed rectangular reconstructions (each of which is drawn with a different line thickness) that achieve zero cost.




Figure 4.13: Nonuniqueness of solution arising in a case with non-zero cost. Shown are the measured support lines as well as several rectangular reconstructions (each of which is drawn with a different line thickness) that achieve minimum cost.

$N > M$, so that there are too many reconstruction angles for the measurement angles; (2) $N \leq M$, but with too many reconstruction angles between some pair(s) of measurement angles; (3) $N = M$, with interleaved measurement and reconstruction angles. The analysis is not meant to be exhaustive but is instead intended to provide some insight as to when nonunique solutions may occur.

We investigate nonuniqueness by characterizing the nullspace of $Ah_\phi = y$ in Equation 4.4. Whether this equation has a single solution, many solutions, or no solutions depends on the rank of the matrix A as well as on whether or not y lies in the range of A . For equal numbers of measurement and reconstruction angles ($M = N$), checking whether the determinant of A is nonzero is equivalent to checking if its rank is full.

In case 1, where $N > M$, the matrix A has a rank of at most M . Hence, we have $N - M$ degrees of freedom in the solution which correspond to free variables appearing in the solution to $Ah_\phi = y$. However, this case is not an important one, since we are attempting to obtain a reconstruction with more degrees of freedom than we have measurements.

The same degree-of-freedom argument applies to case 2, where $N \leq M$ and where the reconstruction angles sufficiently outnumber the measurement angles near some θ_i . Again, we expect nonunique solutions. A simple example illustrating this degree-of-freedom argument is that in which three reconstruction angles $\phi_{j-1}, \phi_j, \phi_{j+1}$ lie between the adjacent measurement angles θ_i and θ_{i+1} , so that we have $\phi_{j-2} < \theta_i < \phi_{j-1} < \phi_j < \phi_{j+1} < \theta_{i+1} < \phi_{j+2}$. Here, the support values at ϕ_{j-1} and ϕ_{j+1} , along with those at the reconstruction angles ϕ_{j-2} and ϕ_{j+2} , completely determine the sinusoidal curves on the intervals $[\phi_{j-2}, \phi_{j-1}]$ and $[\phi_{j+1}, \phi_{j+2}]$, which in turn determine the cost incurred at θ_i and θ_{i+1} . The reconstruction angle ϕ_j is thereby 'insulated' from the two measurement angles by ϕ_{j-1} and ϕ_{j+1} , and as a consequence the support value

$h_\phi(\phi_j)$ can take on a range of values, subject only to the consistency constraint in Equation 4.5.

In the case in which we have equal numbers of measurement and reconstruction angles ($M = N$, case 3), it appears that nonuniqueness is a rather rare phenomenon. We restrict our analysis to the situation in which the $M = N$ measurement angles and reconstruction angles are interleaved (so that $\phi_i \leq \theta_i < \phi_{i+1}$ and $\theta_i \leq \phi_{i+1} < \theta_{i+1} \forall i$, where we define $\theta_{N+1} = \theta_1$ and $\phi_{N+1} = \phi_1$ for convenience). This restriction is motivated by our discussion concerning cases 1 and 2, in which there were too many reconstruction angles between at least one pair of adjacent measurement angles. With interleaving, the matrix A takes the form

$$A = \begin{pmatrix} l_1 & r_1 & 0 & \cdots & 0 \\ 0 & l_2 & r_2 & 0 & \vdots \\ \vdots & 0 & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & l_{N-1} & r_{N-1} \\ r_N & 0 & \cdots & 0 & l_N \end{pmatrix}, \quad (4.6)$$

where $l_i = \frac{\sin(\phi_{i+1} - \theta_i)}{\sin(\phi_{i+1} - \phi_i)}$ and $r_i = \frac{\sin(\theta_i - \phi_i)}{\sin(\phi_{i+1} - \phi_i)}$, and has determinant

$$\det A = \prod_{i=1}^N \frac{\sin(\phi_{i+1} - \theta_i)}{\sin(\phi_{i+1} - \phi_i)} + (-1)^{N+1} \prod_{i=1}^N \frac{\sin(\theta_i - \phi_i)}{\sin(\phi_{i+1} - \phi_i)} \quad (4.7)$$

In order to make more concrete conclusions, we further restrict attention to the special case of what we call *perfectly-interleaved* measurement and reconstruction angles, in which the interleaved measurement and reconstruction angles also satisfy $\theta_i = \frac{\phi_i + \phi_{i+1}}{2}$ and $\phi_i = \frac{\theta_{i-1} + \theta_i}{2}$. With perfect interleaving, Equation 4.7 reduces to

$$\det A = [1 + (-1)^{N+1}] \frac{1}{(2 \cos \Delta)^N}, \quad (4.8)$$

where $\Delta = \theta_i - \phi_i = \phi_{i+1} - \theta_i = \frac{\theta_{i+1} - \theta_i}{2} = \frac{\phi_{i+1} - \phi_i}{2}$. Since the face angles $\{\phi_1, \phi_2, \dots, \phi_N\}$

of a polygon must satisfy $\max_i(\phi_{i+1} - \phi_i) < \pi$, we have that $\Delta < \pi/2$. Hence, for N odd ($N \geq 3$), $\det A > 0$ so that A is invertible, and for any measurement vector y , there exists a *unique* solution $h_\phi = A^{-1}y$. However, for N even ($N \geq 4$), $\det A = 0$. In this case, $\text{rank}(A) < N$ so that for any y , there exists a *subspace* of solutions to $Ah_\phi = y$.

For angles that are interleaved but not necessarily perfectly interleaved, when N is odd, the solution is unique since both terms in Equation 4.7 must be positive (again, because $\max_i(\phi_{i+1} - \phi_i) < \pi$). When N is even, nonuniqueness results if the spacings between the angles satisfy

$$\prod_{i=1}^N \sin(\phi_{i+1} - \theta_i) = \prod_{i=1}^N \sin(\theta_i - \phi_i),$$

which is not generically true for interleaved measurement and reconstruction angles.

Chapter 5 Best N -gon with Fixed Relative Spacing of Reconstruction Angles

In the previous chapter, we saw that prior knowledge of the face angles of an object that is roughly polygonal can dramatically improve reconstruction quality. However, the utility of the algorithm BNGON is somewhat limited in that it requires knowledge of the object's face angles *and* its orientation. Since it is more reasonable to assume the availability of rough shape information without having orientation information, in this chapter we formulate a problem in which the *relative* (rather than the absolute) angles of the object's sides are known. Recall that this formulation constitutes the fourth level of generality in the general polygonal shape estimation problem that was outlined at the beginning of Chapter 3.

In formulating this problem, we let $\{\theta_1, \theta_2, \dots, \theta_M\}$ and $\{y(\theta_1), y(\theta_2), \dots, y(\theta_M)\}$ denote the M measurement angles and the measured support values at these angles, as before. The reconstruction angles are given by $\{\phi_1 + \alpha, \phi_2 + \alpha, \dots, \phi_N + \alpha\}$, where $\{\phi_1, \phi_2, \dots, \phi_N\}$ are known and $\alpha \in [0, 2\pi)$ serves as an unknown offset parameter fixing the absolute locations of the reconstruction angles or equivalently the orientation of the object. We wish to jointly estimate values of α and $\{h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \dots, h_\phi(\phi_N + \alpha)\}$ by choosing the values that minimize

$$\sum_{i=1}^M (h_\phi(\theta_i) - y(\theta_i))^2 \stackrel{\text{def}}{=} J(\alpha, \{h_\phi(\phi_i + \alpha)\}) \quad (5.1)$$

where $h_\phi(\theta_i)$ is given by Equation 4.3 and is repeated here for convenience

$$h_\phi(\theta_i) = \frac{\sin(\phi_{R_i} + \alpha - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{L_i} + \frac{\sin(\theta_i - \phi_{L_i} - \alpha)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{R_i} \quad (5.2)$$

and the $\{h_\phi(\phi_i + \alpha)\}$ are constrained to be a set of consistent support values. Unfortunately, the cost in Equation 5.1 is nonlinear in α , and QP techniques cannot directly be used.

However, for a *fixed* value of α , the cost function and consistency constraint *are* quadratic and linear (respectively) in the support values $\{h_\phi(\phi_i)\}$, reducing the problem to a QP problem. Hence, to minimize the cost in Equation 5.1, we can use an outer loop that minimizes over α with an inner step consisting of solving a QP problem. As we will see, the estimation algorithm that we introduce in Section 5.2 is based on standard gradient descent methods for the outer loop α -minimization. Prior to developing the algorithm, we characterize the cost function in Section 5.1, demonstrating its highly complicated nature. This characterization will be useful in developing the details of the algorithm.

5.1 The Nature of the Cost Function

The cost function $J_{h_\phi}(\alpha)$ is a function of $N + 1$ variables, namely the support values at the N reconstruction angles and the orientation parameter α . Since locating the global minimum of a multivariable nonlinear function is generally difficult, we decompose the minimization into a pair of successive minimizations as follows:

$$\begin{aligned} \min_{\alpha, \{h_\phi(\phi_i + \alpha)\}} J(\alpha, \{h_\phi(\phi_i + \alpha)\}) &= \min_{\alpha} \min_{\{h_\phi(\phi_i + \alpha)\}} J(\alpha, \{h_\phi(\phi_i + \alpha)\}) \\ &\stackrel{\text{def}}{=} \min_{\alpha} J_{h_\phi}(\alpha). \end{aligned} \quad (5.3)$$


so that the inner minimization consists of solving the BNGON QP problem of Equa-

tion 4.4 for a fixed value of α , and the outer minimization consists of choosing that solution to Equation 4.4 which yields the minimum cost for all values of $\alpha \in [0, 2\pi)$.


Equation 5.3 suggests the following ‘brute-force’ approach to minimizing $J_{h_\phi}(\alpha)$. We simply choose many values of α , solve the QP problem of Equation 4.4 for each, and then choose that value of α which yields minimum cost. The major drawback of this approach is a desired-accuracy versus computational-requirement. Since $J_{h_\phi}(\alpha)$ is sampled at finitely many values of α (say p values), our estimate $\hat{\alpha}$ will be somewhat inaccurate, being on average $(360/(4p))^\circ$ away from the true minimum α_{true} . So to achieve a reasonably accurate estimate, a prohibitively large number of QP problems would have to be solved. For these reasons, we abandon this as a viable algorithm.

In this section, we describe qualitative features of the optimization problem by examining the nature of the cost function $J_{h_\phi}(\alpha)$, to gain some insight concerning our problem and its solution. We restrict ourselves to the case in which the measurement and reconstruction angles are uniformly-spaced and the measured support values are *noise-free*. On a look at several sample plots of $J_{h_\phi}(\alpha)$ (see Figure 5.1), it is clear that the cost function can be highly complex, and we expect that behavior in the presence of noise is at least as complicated.


Two interesting features present in the sample plots are worth discussing: (1) regions in which cost is zero and (2) local extrema, and in particular maxima that take the form of cusps. The regions of zero cost, corresponding to values of α at which reconstructions that perfectly fit the measurements can be found, will be referred to as ‘deadspaces.’ To understand why these deadspaces occur as well as what determines their extent, we examine the example of Figure 5.1b. Here, $J_{h_\phi}(\alpha)$ contains four identical deadspaces, each extending $2 \tan^{-1} \frac{1}{2} \approx 53.13^\circ$. Due to the symmetry of $J_{h_\phi}(\alpha)$ on $[0^\circ, 360^\circ)$, we restrict discussion to the interval $[0^\circ, 45^\circ]$, where 45° is the




(a)



(b)



(c)



(d)

Figure 5.1: Sample plots of $J_{h_p}(\alpha)$ for noise-free uniformly-spaced measurement angles and reconstruction angles. Note that the plot in (c) extends only 90° .

angle at which the first cusp is located.

The example in Figure 5.1b corresponds to $M = 4$ measurements of an ellipse with $N = 4$ reconstruction angles, with both sets of angles being multiples of 90° . The four measurements give rise to a set of support lines that bound a rectangle having sides aligned with the coordinate axes. We consider the reconstruction that results from rotating the constellation of reconstruction angles by $\alpha = 15^\circ$. Figure 5.2c shows the measured support lines as well as the rectangular reconstruction that results from using the (rotated) reconstruction angles $\phi_1 = 15^\circ, \phi_2 = 105^\circ, \phi_3 = 195^\circ$, and $\phi_4 = 285^\circ$. Since the reconstruction is inscribed in the support line rectangle and touches each of the support lines, the solution incurs zero cost.

In fact, *any* value of $\alpha \in [0^\circ, \tan^{-1} \frac{1}{2}]$, giving rise to the reconstruction angles $\phi_1 = \alpha, \phi_2 = 90^\circ + \alpha, \phi_3 = 180^\circ + \alpha$, and $\phi_4 = 270^\circ + \alpha$, corresponds to a rectangular reconstruction that achieves zero cost. Several such rectangular reconstructions are shown in Figure 5.2a–e for values of $\alpha \in [0^\circ, \tan^{-1} \frac{1}{2}]$. In contrast, for a value of $\alpha \in [\tan^{-1} \frac{1}{2}, 45^\circ]$, a zero-cost reconstruction does not exist. As an example, the minimum-cost reconstruction for $\alpha = 45^\circ$ is shown in Figure 5.2f.

That the extent of the deadspace on $[0^\circ, 45^\circ]$ is given by $\tan^{-1} \frac{1}{2}$ is explained by Figure 5.3. As we saw in Figure 5.2a–e, zero-cost solutions exist for those values of α for which the reconstructed rectangle is inscribed in the measurement rectangle. With our restricted interval $[0^\circ, 45^\circ]$, this occurs only for $\alpha \in [0^\circ, \tan^{-1} \frac{b}{a}]$, since for $\alpha > \tan^{-1} \frac{b}{a}$, the ‘southeast’ side of the reconstructed rectangle shown in Figure 5.3 would no longer touch the rightmost vertical side of the measurement rectangle. For our problem, the values $a = 2$ and $b = 1$ lead to a deadspace extent of $\tan^{-1} \frac{1}{2}$. Moreover, since $\tan^{-1} \frac{b}{a} \rightarrow 0^\circ$ as $\frac{b}{a} \rightarrow 0$, we conclude that deadspace extent decreases as object eccentricity increases.

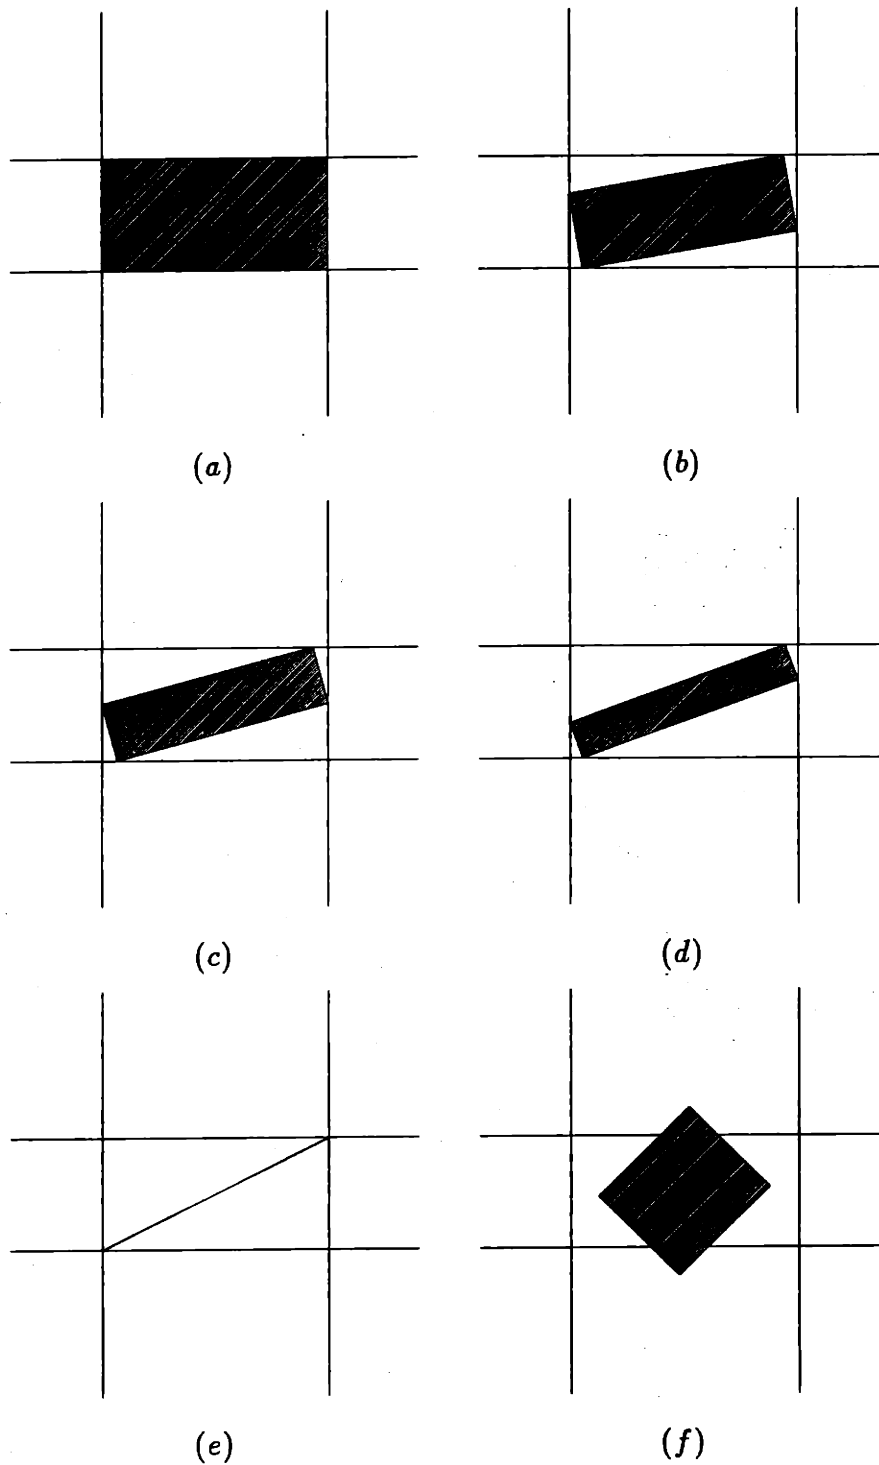


Figure 5.2: Measured support lines and reconstructed object for (a) $\alpha = 0^\circ$, (b) $\alpha = 10^\circ$, (c) $\alpha = 15^\circ$, (d) $\alpha = 20^\circ$, (e) $\alpha = \tan^{-1} \frac{1}{2}$, and (f) $\alpha = 45^\circ$.




Figure 5.3: Deadspace extent on the interval $[0^\circ, 45^\circ]$.

We now make some observations concerning the local maxima and minima in $J_{h_\phi}(\alpha)$. Probably the simplest case to consider is that of a circle (see Figure 5.1a for an example), for which we conjecture that the number of maxima (and hence minima) is given by $\text{LCM}(M, N)$, $M \neq N$, where $\text{LCM}(\cdot, \cdot)$ denotes the least common multiple of its arguments. (For $M = N$, $J_{h_\phi}(\alpha) \equiv 0$.) For the case in which the true object is an ellipse (see Figure 5.1b-c) or in which it is a polygon (Figure 5.1d), we observe maxima that take the form of cusps. Also, the example of Figure 5.1d suggests that extrema can occur arbitrarily close together. Although we have not investigated the numbers and locations of extrema or the origin of cusp-like behavior in the cost function in any significant detail, we believe that such an investigation would be rather difficult (especially on relaxing our restriction that the measurement and reconstruction angles be uniformly-spaced).

5.2 The Estimator BNGONROT

In the previous section, we found that the cost function $J_{h_\phi}(\alpha)$ is somewhat complicated even in the noise-free case, characterized by zero-cost regions, local extrema that can be close together, and maxima that take the form of cusps. In this section, we present an algorithm to solve Equation 5.3 given *noisy* measurements. The algorithm is based on standard gradient descent methods modified to accomodate cusps and deadspaces.

Standard optimization algorithms often require knowledge of the gradient of the function to be minimized, but since the functional form of $J_{h_\phi}(\alpha)$ is not defined in closed form but rather implicitly through a QP optimization, we do not have access to the derivative of $J_{h_\phi}(\alpha)$ with respect to α . However, using the chain rule we can express $\frac{dJ_{h_\phi}(\alpha)}{d\alpha}$ as

$$\frac{dJ_{h_\phi}(\alpha)}{d\alpha} = \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} + \frac{\partial J_{h_\phi}(\alpha)}{\partial h^*(\alpha)} \frac{\partial h^*(\alpha)}{\partial \alpha} \quad (5.4)$$

where $h^*(\alpha)$ is the solution to the BNGON QP problem resulting from rotating the reconstruction angles by α , and $\frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha}$ results from taking the partial derivative of J with respect to α in Equation 5.1 using $h^*(\alpha)$ for the $\{h_{L_i}\}$ and $\{h_{R_i}\}$ in Equation 5.2. We may then *approximate* the total derivative $\frac{dJ_{h_\phi}(\alpha)}{d\alpha}$ by the partial derivative $\frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha}$, which we can compute easily.

Using the partial derivative in place of the total derivative, we can then use gradient methods. However, since the value of $\frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha}$ for any particular of α depends on $h^*(\alpha)$, each gradient step, given by

$$\alpha^{(j+1)} = \alpha^{(j)} \pm \lambda \left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j)}}, \quad (5.5)$$

where λ is the gradient step size, must follow a QP step. Our estimation algorithm, to be referred to as BNGONROT, essentially consists of performing interleaved gradient ascent/descent steps and QP steps on $J_{h_\phi}(\alpha)$ from $\alpha = 0^\circ$ to $\alpha = 360^\circ$ to locate all local extrema. The solution is obtained by choosing that local minimum that yields minimum cost.

Specifically, we begin at $\alpha = 0^\circ$ and solve the BNGON QP problem. Using the estimated support values, we compute $\frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha}$ and perform a gradient ascent or descent step depending on whether its sign is positive or negative, to obtain a new value of α . We are then committed to performing gradient ascent until we reach the first maximum or gradient descent until we reach the first minimum. We then perform the following steps repeatedly: (1) solve the BNGON QP step; (2) compute the gradient based on the estimated support values; and (3) perform a gradient step. Once an appropriate convergence criterion has been met (as discussed below), indicating that a local minimum or maximum has been found, we store this value of α as well as the associated cost $J_{h_\phi}(\alpha)$. We then advance by some small amount in α , and by solving the BNGON QP problem and computing the gradient, determine from its sign whether our next series of interleaved steps will consist of gradient ascent or descent steps. Performing steps (1)–(3) repeatedly, we reach our next maximum or minimum. We continue this traversal of the interval $[0^\circ, 360^\circ)$ until we have located all maxima and minima, and then choose the global minimum $\hat{\alpha}$. Solving the BNGON QP problem with $\alpha = \hat{\alpha}$ yields the solution to our problem.

The criterion for convergence is met when either of two conditions is satisfied. The first condition is the usual termination rule for standard gradient ascent/descent, i.e., convergence is said to have occurred when the magnitude of the gradient falls below some prespecified value. The need for a second convergence condition is due to

the inability of standard gradient ascent/descent algorithms (and their convergence criteria) to deal with cusps (discontinuities in slope) that can occur in the cost function $J_{h_\phi}(\alpha)$. To deal with this, we halve the step size λ of the gradient ascent/descent every time the sign of the derivative changes (indicating that a maximum or minimum has been crossed) provided that the magnitude of the derivative is sufficiently large (assuring that we are near a discontinuity in slope rather than a smooth maximum or minimum). The second convergence condition is met when λ falls below some prespecified value.

Because the algorithm is based on standard gradient ascent/descent methods, modified to obtain precise solutions near cusps, we expect that its limitations are similar to those associated with the standard methods. Most important is the trade-off of speed versus accuracy as determined primarily by the choice of λ and the convergence criterion. For a given desired accuracy this algorithm is generally much more efficient than the 'brute-force' approach of solving a QP problem at each of many independently chosen values of α and choosing that value having lowest cost. A typical BNGONROT problem requires on the order of 1,000 QP steps to yield an estimate of α that is within roughly 0.001 degrees of α_{true} . To achieve such accuracy using the brute-force method would require the solution of 90,000 QP problems.

The following is a complete description of the BNGONROT estimation algorithm:

Estimator BNGONROT

We wish to find the global minimum $\hat{\alpha}$ of $J_{h_\phi}(\alpha)$, as well as the set of reconstructed support values $\{h_\phi(\phi_1 + \hat{\alpha}), h_\phi(\phi_2 + \hat{\alpha}), \dots, h_\phi(\phi_N + \hat{\alpha})\}$ arising from $\hat{\alpha}$, assuming no particular spacings for the measurement and reconstruction angles.

1. Initialize by setting the minima counter $i = 0$ and by setting the offset $\alpha^{(0)} = 0^\circ$.
2. Solve the BNGON QP problem using the reconstruction angles arising from $\alpha = \alpha^{(j)}$ and evaluate the local slope $\left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j)}}$. If the derivative is $\begin{cases} \text{positive} \\ \text{negative} \end{cases}$, then 'gradient step' in (3) refers to a gradient $\begin{cases} \text{ascent} \\ \text{descent} \end{cases}$ step. If the derivative is zero, then increment the minima counter $i = i + 1$, let $\alpha_i = \alpha^{(j)}$, and go to (5).
3. Repeat steps (a)–(c) until either of the following two convergence criteria is met: (i) $\left| \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right| \leq \epsilon$ for some small prespecified $\epsilon > 0$ or (ii) $\lambda < \lambda_{\text{threshold}}$, where $\lambda_{\text{threshold}}$ is prespecified.
 - (a) Evaluate the local slope $\left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j)}}$ and perform a single gradient step to arrive at a new offset value: $\alpha^{(j+1)} = \alpha^{(j)} \pm \lambda \left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j)}}$. If $\left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j)}}$ and $\left. \frac{\partial J_{h_\phi}(\alpha)}{\partial \alpha} \right|_{\alpha=\alpha^{(j-1)}}$ differ in sign and if either has magnitude greater than some prespecified threshold, then $\lambda = \lambda/2$.
 - (b) Solve the BNGON QP problem using the reconstruction angles arising from $\alpha = \alpha^{(j+1)}$.
 - (c) Increment the iteration index: $j = j + 1$.
4. If gradient descent was being performed in (3), increment the minima counter $i = i + 1$ and let $\alpha_i = \alpha^{(j)}$.
5. Set $\alpha^{(0)} = \alpha^{(j)} + f(\lambda)$, where f is a strictly positive function that satisfies $f(\lambda) \ll$ expected distance between the two nearest stationary points. If $\alpha^{(0)} < 360^\circ$, go to (2). Otherwise, $\hat{\alpha} = \arg \min_{\alpha_i} J_{h_\phi}(\alpha_i)$, and \hat{h} is given by solving the BNGON QP problem using $\hat{\alpha}$.

It should be noted that BNGONROT has been developed without much emphasis on optimizing the speed of the algorithm. Several simple modifications to increase speed, desirable for the Monte Carlo performance analysis to be described in Section 5.4, were made, however. These include tolerating crude estimates of those quantities that are of not much interest (e.g., maxima as well as those minima which clearly have higher cost than previously found minima) by relaxing appropriate convergence criteria.

5.3 Examples

We now present two examples of reconstructions produced by BNGONROT, the first of which is that of the standard triangle introduced earlier. In Figure 5.4 are shown the true object, $M = 24$ uniformly-spaced noisy support lines with noise level $\sigma = .25$, and the reconstructed object. The reconstruction forms an angle of $\alpha = 91.63^\circ$ with the positive x -axis and incurs an error $E = 0.19$ with respect to the true object.

Not surprisingly, the reconstruction is qualitatively and quantitatively far better than that corresponding to NUA, as seen on a comparison with Figure 3.4 and from the fact that $E_{\text{BNGONROT}} = 0.19 < E_{\text{NUA}} = 1.56$. Moreover, on a comparison with Figure 4.2 we see that the BNGONROT reconstruction is not much worse than the BNGON reconstruction ($E_{\text{BNGONROT}} = 0.19 \approx E_{\text{BNGON}} = 0.17$), indicating that not much is sacrificed in settling for a weaker prior, i.e., knowing relative rather than absolute reconstruction angles. The performance assessment in the next section illustrates this attribute of BNGONROT for a range of noise levels.

The second example is similar to one considered in Section 4.2, in which a rough size estimate of an ellipse was formed by fitting the best rectangle (with sides at





Figure 5.4: Example of reconstruction of the standard triangle in (a) object space and (b) support function space. $M = 24$, $N = 3$, $\sigma = 0.25$.

0°, 90°, 180°, and 270°) to noisy support measurements of an ellipse. Here, we introduce a generalization by estimating both the orientation and rough size of an ellipse.

Figure 5.5 depicts the standard ellipse of Chapters 3 and 4, $M = 24$ uniformly-spaced noisy support lines with noise level $\sigma = 0.5$, and the rectangular reconstruction. The estimated and actual orientations, defined by the moment axes of the rectangle and ellipse respectively, differ by 1.82°. The lengths of the sides of the rectangle are estimated to be 1.36 and 3.58 units, and roughly indicate the size of the ellipse. Again, on comparing Figures 4.4 and 5.5, it should be noted that less information (a relative angle prior as opposed to an absolute angle prior) available to BNGONROT does not cause it to underperform BNGON significantly.

5.4 Performance Assessment

In this section, we evaluate the performance of BNGONROT versus noise level using Monte Carlo techniques, and compare this performance to that of NUA and BNGON. In addition, since BNGONROT provides nonlinear orientation estimates, we determine an analytical expression approximating the orientation error variance as a function of measurement noise variance and corroborate it with Monte Carlo simulation.

In Figure 5.6, we show plots of the reconstruction error E versus noise level σ for $M = 24$ uniformly-spaced measurements of the standard triangle, for the algorithms BNGONROT, BNGON, and NUA. The reconstruction angles used for BNGON are the correct face angles of the triangle, while those used for BNGONROT are the correct relative face angles. Each data point in the plots corresponding to the NUA and BNGON algorithms arises from 1000 noise realizations. Each point in the BNGONROT plot arises from 500





Figure 5.5: Example of reconstruction of the standard ellipse in (a) object space and (b) support function space. $a = 2.0$, $M = 48$, $N = 4$, $\sigma = 0.50$.

noise realizations. Note that the plots corresponding to NUA and BNGON are repeated from Figure 4.5.




Figure 5.6: E versus σ for the standard triangle, for the algorithms BNGONROT, BNGON, and NUA. $M = 24$.

As we expect, the reconstruction error increases with increasing noise level for all three algorithms. However, there are several more important observations to make. First, the performance of BNGONROT is much better than that of NUA but only slightly worse than BNGON for relatively low noise levels. As noise level increases, the performance of BNGONROT degrades slowly relative to that of BNGON until noise levels near $\sigma = 0.17$, where the degradation becomes much more pronounced. For noise levels beyond this threshold level, the BNGONROT algorithm is somewhat worse than BNGON, but is still much better than NUA. We investigate this threshold effect in more detail in the remainder of this section, in our analysis of orientation estimation error.

To further characterize the performance of BNGONROT, we note that once the value

of α has been estimated, the problem is reduced to the BNGON problem of Equation 4.4 with the absolute spacings of the reconstruction angles determined by $\hat{\alpha}$. Since the performance of BNGON versus α was investigated in Section 4.3 for the case in which the entire configuration of reconstruction angles is rotated by an incorrect non-zero offset, it suffices to examine the performance of BNGONROT only in estimating the orientation parameter α .

In what follows, we analytically determine the Cramer-Rao lower bound [28] on orientation error variance as well as an approximate expression for the probability of obtaining anomalous orientation estimates. Together, these expressions describe the orientation error variance for a range of noise levels. We supplement this analytical error analysis with results from Monte Carlo simulations. Throughout, our discussion is based on $M = 24$ uniformly-spaced measurements of the standard triangle, so that the true orientation, or offset parameter is given by $\alpha_{\text{true}} = 90^\circ$.

We first note that errors in estimating α can be classified into two distinct categories [28]. Consider the cost function $J_{h_\phi}(\alpha)$ that arises from taking a particular noise realization for each of the noise levels $\sigma = 0.25$ and $\sigma = 0.75$, as shown in Figure 5.7. On looking at $J_{h_\phi}(\alpha)$ for $\sigma = 0.25$, one finds that the global minimum at $\hat{\alpha}$ lies in the same *basin* as does the true value of 90° . That is, both share the same pair of adjacent maxima. However, in our example plot for $\sigma = 0.75$, the global minimum lies in a basin different from that belonging to the true value of 90° . We shall refer to the former error in estimating α as a *local error* and to the latter as a *global error*.

The approach that we take in our analysis is characteristic of that for a standard nonlinear estimation problem. For low noise levels the nonlinear estimator exhibits behavior characteristic of a linear one, and indeed can be thought of as being linearized





Figure 5.7: Examples of $J_{h_g}(\alpha)$ for (a) a local error ($\sigma = 0.25$) and (b) a global error ($\sigma = 0.75$).

about the true value α_{true} . Namely, errors of large magnitude are not very likely in a linear estimator unless noise levels are extraordinarily high, and hence local errors predominate. As a result, the error variance of the estimate is well-approximated by the Cramer-Rao bound (CRB) for low levels of noise.

However for high noise levels, the observations are not a strong function of the quantity to be estimated (in our case α). In this case, noise values are more likely to lower the cost $J_{h_\phi}(\alpha)$ for some α_0 far from α_{true} such that a basin whose depth is greater than that of α_{true} 's basin is created, thereby giving rise to a global error. In this case, the CRB is an overly optimistic descriptor of achievable performance.

For convenience, in what follows we use notation that is slightly different from that used in the beginning of this chapter. In particular, let us denote by $\mathbf{x} = [\alpha \ h_\phi(\phi_1) \ h_\phi(\phi_2) \ \dots \ h_\phi(\phi_N)]^T$ the parameter vector to be estimated, so that the observation, or measurement vector is given by $\mathbf{y} = g(\mathbf{x}) + \mathbf{n}$, where the i^{th} component of $g(\mathbf{x})$ is the scalar quantity in Equation 5.2. With this notation and if \mathbf{x} were completely unconstrained, the Cramer-Rao bound on the error covariance of any unbiased estimate of \mathbf{x} would be given by

$$\begin{aligned} \text{cov}(\mathbf{x} - \hat{\mathbf{x}} | \alpha = \alpha_{\text{true}}) &\geq I_y^{-1}(\mathbf{x})|_{\alpha=\alpha_{\text{true}}} \\ &= \left\{ E \left[\left(\frac{\partial \ln p_{\mathbf{y}|\mathbf{x}}(\mathbf{y}_0|\mathbf{x})}{\partial \mathbf{x}} \right)^2 \right] \right\}^{-1} \Big|_{\alpha=\alpha_{\text{true}}}, \end{aligned} \quad (5.6)$$

where $I_y(\mathbf{x})$ is referred to as the *Fisher information* in \mathbf{y} about \mathbf{x} . Since the measurement noise vector \mathbf{n} is a Gaussian random vector with zero mean and covariance matrix $\sigma^2 I$, Equation 5.6 may be rewritten as

$$\text{cov}(\mathbf{x} - \hat{\mathbf{x}} | \alpha = \alpha_{\text{true}}) \geq \sigma^2 \left[\frac{\partial g^T(\mathbf{x})}{\partial \mathbf{x}} \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right]^{-1} \Big|_{\alpha=\alpha_{\text{true}}}, \quad (5.7)$$

where the (i, j) element of the $M \times (N + 1)$ matrix $\frac{\partial g(\mathbf{x})}{\partial \mathbf{x}}$ is given by $\frac{\partial g_i}{\partial x_j}$. The error variance of α is lower-bounded by the $(1, 1)$ element of $I_y^{-1}(\mathbf{x})$.

In our problem, however, N of the parameters to be estimated—namely, the $\{h_\phi(\phi_i)\}$ —are constrained by the consistency condition of Equation 4.5, so that the CRB of Equation 5.7 is not strictly applicable. As a result, we *approximate* the bound on the error variance of the orientation estimate. We do this by replacing the matrix CRB given by Equation 5.7 by the scalar CRB on α , obtained by treating α along as the parameter to be estimated and fixing the $\{h_\phi(\phi_i)\}$ at their true values. In this case, the observation vector is given by $y = f(\alpha) + n$, where $n \sim N(0, \sigma^2 I)$. The CRB on the error variance of any unbiased estimate of α is given by

$$\begin{aligned} \text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}}) &\geq I_y^{-1}(\alpha_{\text{true}}) \\ &= \left\{ E \left[\left(\frac{\partial \ln p_{y|x}(y_0 | \mathbf{x})}{\partial \mathbf{x}} \right)^2 \right] \right\}^{-1} \Big|_{\alpha = \alpha_{\text{true}}} \end{aligned} \quad (5.8)$$

$$= \frac{\sigma^2}{\frac{\partial h^T(\alpha)}{\partial \alpha} \frac{\partial h(\alpha)}{\partial \alpha}} \Big|_{\alpha = \alpha_{\text{true}}}, \quad (5.9)$$

where the i^{th} component of $\frac{\partial f(\alpha)}{\partial \alpha}$ is given by $\frac{\partial f_i(\alpha)}{\partial \alpha}$. Taking the square root of the Cramer-Rao bound on error variance given by Equation 5.9 yields a bound on the standard deviation of the estimation error (in degrees), plotted as a function of noise level σ in Figure 5.8. As we mentioned above, this bound is an accurate predictor of performance for sufficiently low noise levels, where our definition of ‘sufficiently low’ will be based on our subsequent calculation of the probability of a global error or anomaly.

For some nonlinear estimation problems, such as those arising in communication problems [28], under high noise levels the anomalous minima that result should occur with equal likelihood anywhere except for values of α within the basin corresponding




Figure 5.8: Cramer-Rao bound on $\sqrt{\text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}})}$ versus σ .

to α_{true} , and as a consequence are modeled as being distributed uniformly over the interval of interest. The error variance at these noise levels is therefore approximated by the variance of a uniformly distributed random variable. However, in Monte Carlo simulations that we have performed, the locations of anomalous global minima were observed to cluster near 270° (see the histograms of estimation error for several noise levels, shown in Figure 5.9). This behavior is attributable to the fact that the underlying triangular object is highly eccentric (i.e., it has a high moment) and hence gives rise to a dominant secondary minimum in $J_{h_*}(\alpha)$ corresponding to a 180° rotation of the object, under noise-free conditions. Under high noise conditions, the noise often lowers the cost near 270° so that the secondary minimum becomes the global minimum, and this is by far the most likely global error. Hence, in our particular example, we take the estimation error variance when a global error occurs as being the constant value $(180^\circ)^2$.

We have characterized the estimation error variance $\text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}})$ for the low noise and high noise regimes. For intermediate noise levels, the error variance can be approximated, as in [28], by a linear combination of the error variances in the low σ and high σ cases, where the weighting function is the probability of anomaly, or global error $\Pr(A)$ as a function of σ . We may approximate this probability by formulating a binary hypothesis testing problem in which the two hypotheses are $\alpha = 90^\circ$ (zero error) and $\alpha = 270^\circ$ (180° global error), determining the means and variances of the sufficient statistic, which is conditionally Gaussian under each of the two hypotheses, and integrating the Gaussian density. (Details are included in Appendix B.) The resulting probability of anomaly is given by

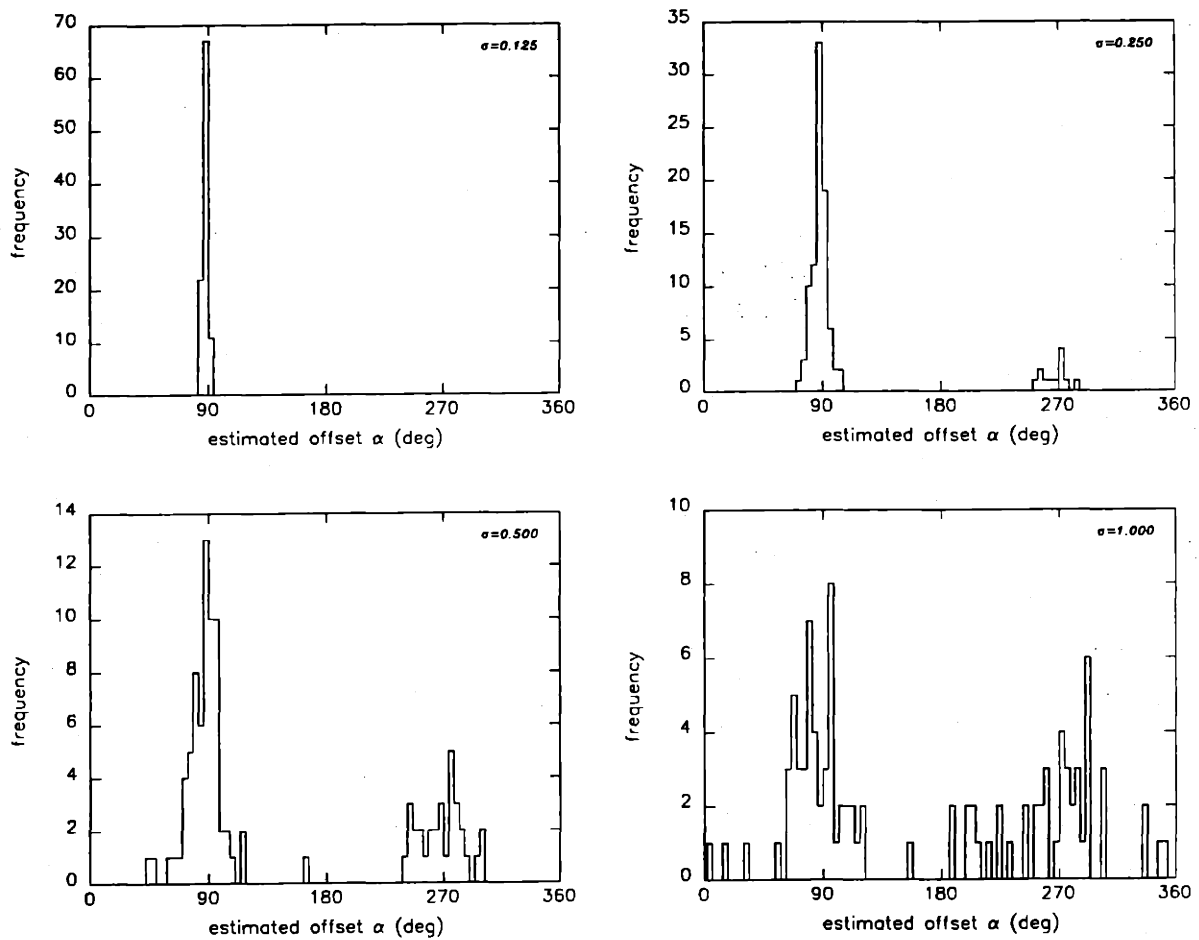


Figure 5.9: Histograms of error in α for several noise levels σ .

$$\Pr(A) = Q\left(\frac{\sqrt{H}}{2\sigma}\right) \quad (5.10)$$

$$\leq \frac{1}{2} \exp\left(\frac{-H}{8\sigma^2}\right), \quad (5.11)$$

where $H = [h(270^\circ) - h(90^\circ)]^T [h(270^\circ) - h(90^\circ)]$ can be interpreted as the energy in the difference of the support vectors under the two hypotheses. Figure 5.10 contains a plot of $\Pr(A)$ versus noise level σ . We note that $\Pr(A)$ increases monotonically with noise level, asymptotically approaching 0.5.




Figure 5.10: Probability of anomaly $\Pr(A)$ versus σ .

Using Equations 5.9 and 5.11 and that fact that the estimation error variance is approximately $(180^\circ)^2$ under high noise conditions, we have that $\text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}})$ for all noise levels is approximately given by

$$\begin{aligned}
\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}}) &\approx [1 - \Pr(A)](\text{CRB}) + \Pr(A)(180^\circ)^2 \\
&\approx \left[1 - \frac{1}{2} \exp\left(\frac{-H}{8\sigma^2}\right)\right] \frac{\sigma^2}{\frac{\partial h^T(\alpha)}{\partial \alpha} \frac{\partial h(\alpha)}{\partial \alpha}} \Big|_{\alpha=\alpha_{\text{true}}} \\
&\quad + \frac{1}{2} \exp\left(\frac{-H}{8\sigma^2}\right) (180^\circ)^2.
\end{aligned} \tag{5.12}$$

Figures 5.11 contain plots of $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ versus σ on normal and logarithmic scales. As we expect for a nonlinear estimator, BNGONROT exhibits dramatic threshold behavior, in which error abruptly increases near $\sigma = 0.05$.

As we mentioned in the discussion accompanying Figure 5.9, we have performed a Monte Carlo performance assessment in support of the preceding analytical error analysis. In each of the following plots, each data point is represented by 200 noise realizations. Also, the plots corresponding to the analytical analysis are overlaid in the appropriate Monte Carlo plots for comparison. We do not explicitly identify the analytical and Monte Carlo plots in each figure. The fact that the former ones are smooth while the latter ones are irregular allows identification.

Figure 5.12 illustrates the percentage of errors that are classified as being global errors (or equivalently, the probability of anomaly $\Pr(A)$) versus σ . Figure 5.13 shows the estimation error $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ for those noise realizations giving rise to local errors, while Figure 5.14 shows the estimation error for global error-producing realizations. The local error plot exhibits a linear trend, roughly matching the analytical results. Also, the global error Monte Carlo plot matches our intuition that global error magnitude is essentially independent of σ . Finally, Figure 5.15 depicts the overall estimation error (i.e., for *all* noise realizations). The simulation results exhibit the same threshold effect as the analytical results, the only major difference being in the noise level at which the threshold is present. This difference is probably due to the




Figure 5.11: Standard deviation $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ of the orientation estimate versus σ . Included for reference is the Cramer-Rao bound, shown dashed.

fact that *many* more noise realizations are needed to generate each data point in the Monte Carlo analysis, especially for small values of $\Pr(A)$.




Figure 5.12: Probability of anomaly $\Pr(A)$ versus noise level σ , obtained analytically and through Monte Carlo simulation.

In addition to knowing the probability of anomaly $\Pr(A)$, which indicates the percentage of errors that are global errors, we would like to know the degree to which they are local or global. That is, we would like information as to how close local errors are to being global errors, and vice-versa, as a function of noise level. One heuristical way to quantify such information is to define a *depth ratio* γ . For local errors we define the depth ratio by

$$\gamma = \frac{J_1}{J_2} \quad (5.13)$$

where

J_1 = the cost at the minimum of the second lowest basin

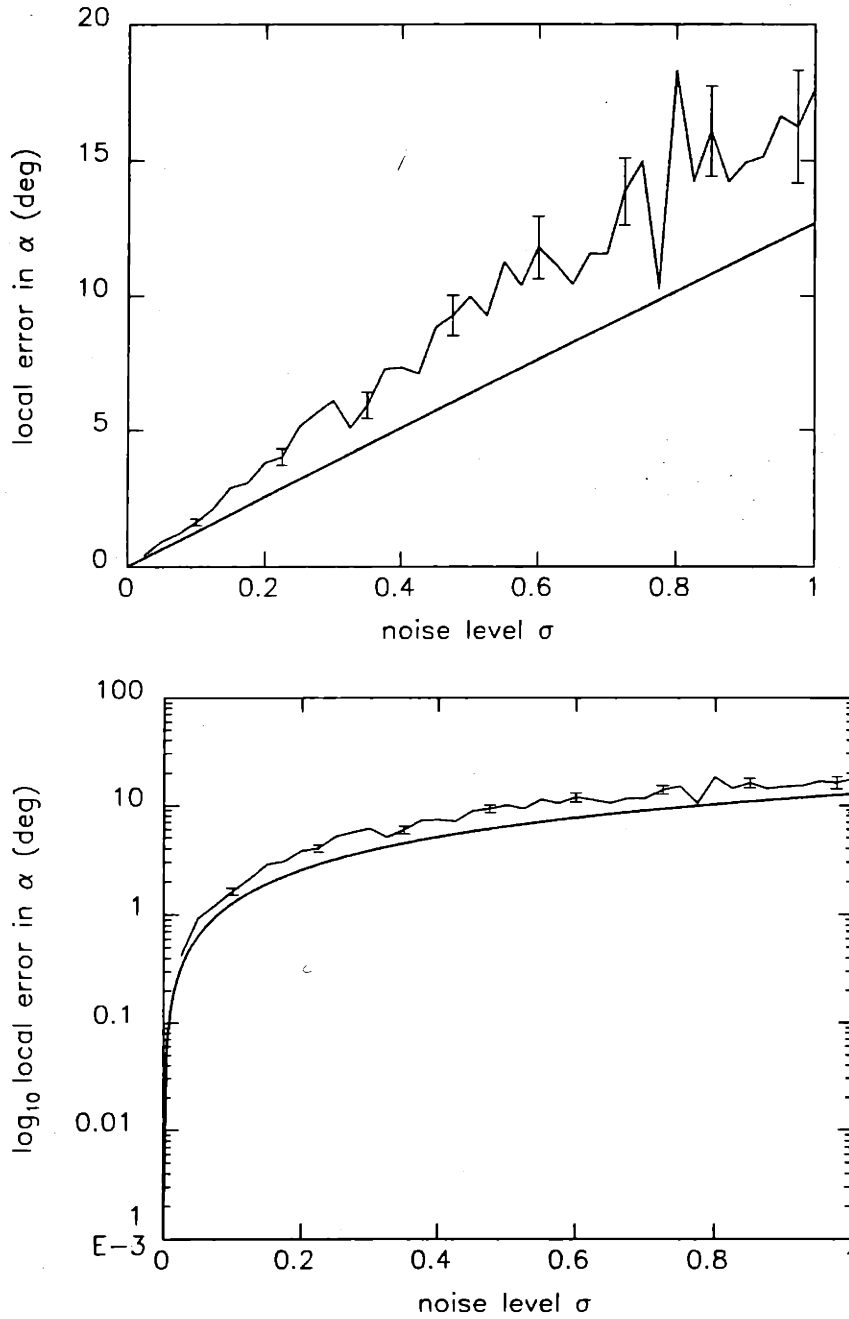


Figure 5.13: Standard deviation $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ of the orientation estimate for local errors versus σ , obtained analytically and through Monte Carlo simulation.




Figure 5.14: Standard deviation $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ of the orientation estimate for global errors versus σ , obtained analytically and through Monte Carlo simulation.

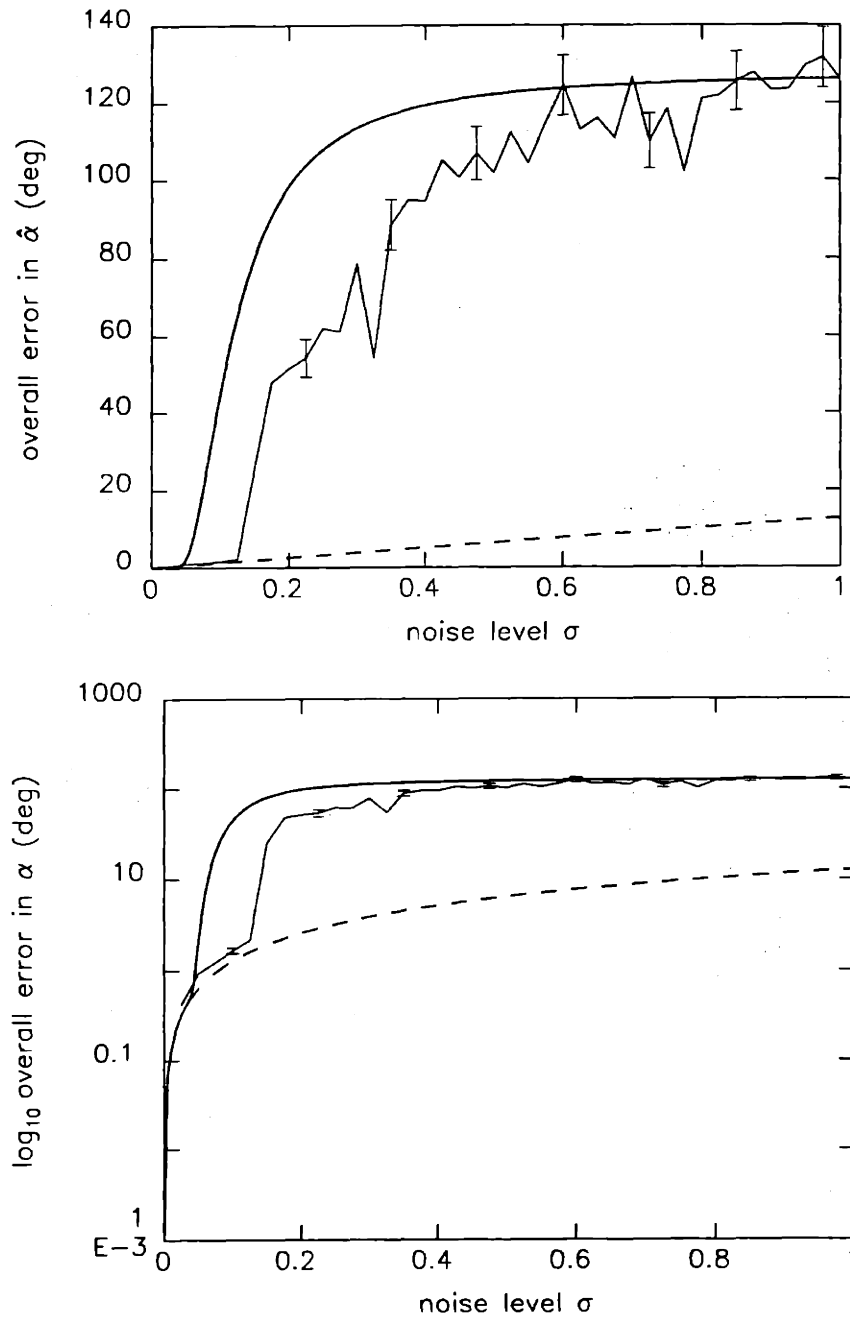


Figure 5.15: Overall standard deviation $\sqrt{\text{var}(\alpha - \hat{\alpha}|\alpha_{\text{true}})}$ of the orientation estimate versus σ , obtained analytically and through Monte Carlo simulation. Included for reference is the Cramer-Rao bound, shown dashed.

J_2 = the cost at the minimum of the lowest basin, i.e., the one containing the true value of α (or, equivalently, the estimated value $\hat{\alpha}$)

as illustrated by the values of J_1 and J_2 shown in Figure 5.7a. Similarly, for global errors γ is defined by Equation 5.13 with

J_1 = the cost at the minimum of the lowest basin, i.e., the one containing the estimated value $\hat{\alpha}$

J_2 = the cost at the minimum of the basin containing the true value of α

as illustrated in Figure 5.7b. By definition, then, $\gamma > 1$ for local errors and $\gamma < 1$ for global errors, with crossover occurring at $\gamma = 1$.

We expect that the depth ratio is closely related to the probability of anomaly $\text{Pr}(A)$, so that low $\text{Pr}(A)$ corresponds to high γ and high $\text{Pr}(A)$ (near 0.5) corresponds to low γ (below 1.0). Figure 5.16 contains plots of depth ratio γ versus σ for those noise realizations giving rise to local and global errors, respectively. The local error plot decreases with σ and asymptotically approaches $\gamma = 1$, confirming our expectation that the lower the noise level, the deeper the global minimum is relative to all other minima. The plot also exhibits a distinct knee at noise levels near the threshold that characterizes the nonlinear estimator. In contrast, the depth ratio for global errors is essentially constant at $\gamma = 0.9$ regardless of noise level, indicating that for global errors, the lowest basin is on average not much deeper than the true basin. Finally, Figure 5.17 is a weighted average of the plots in Figure 5.16 (where the weighting function is given by the plot in Figure 5.12), depicting the depth ratio for all (local and global) errors.





Figure 5.16: Depth ratio γ versus noise level σ for noise realizations giving rise to (a) local errors and (b) global errors. The dashed line at $\gamma = 1$ is included for reference and delineates the 'local error' and 'global error' regions.




Figure 5.17: Depth ratio γ versus noise level σ , based on all noise realizations.

Chapter 6 Target Reconstruction from Laser Radar Data

In this chapter, we apply the convex set reconstruction algorithms NUA, BNGON, and BNGONROT to data consisting of laser radar measurements of various objects, in order to produce shape estimates of these objects. We begin by describing laser radar data and the way in which such data can serve to provide support line measurements, and then describe a technique to estimate support values from the laser radar data. In Section 6.4, our techniques are applied to data obtained through a simulation model as well as to laboratory and field measurements. Finally, we compare our reconstructions with those produced by standard tomographic methods, and present a method by which tomographic images from unregistered data may be substantially improved using convex set reconstruction techniques.

6.1 Laser Radar Data and Problem Scenario

By illuminating a target with radiation at optical frequencies and receiving the reflected signal, laser radars provide information about the surface characteristics of the target [1, 23]. Laser radars can be designed to resolve the return from the target with respect to various quantities. In this thesis, we restrict attention to range-resolved and Doppler-resolved laser radar data. Furthermore, we consider only the case of a monostatic radar, in which the transmitter and receiver are at the same location.

A *range-resolved* measurement (also called a range spectrum) is one in which the return is distributed in range along the line of sight (LOS) of the laser radar (see Figure 6.1). That is, only those parts of the target that are a distance r_0 away from

the laser radar (with distance measured along the LOS) may contribute to the value of the range spectrum at range r_0 . Although a range-resolved measurement is ideally a continuous function of range, in practice it takes the form of a histogram with bins of finite range extent, where each bin is referred to as a range bin.




Figure 6.1: A laser radar illuminating a target along some LOS and receiving a range-resolved return.

Alternatively, for a target undergoing motion, different parts of the target may have different components of velocity along the LOS. A *Doppler-resolved* measurement (also called a Doppler spectrum) is one in which the return is distributed with respect to these variations in velocity (see Figure 6.2). As with a range spectrum, the Doppler spectrum takes the form of a histogram. The value associated with a particular Doppler bin arises from the return of all illuminated parts of the target having the

corresponding component of velocity along the LOS.

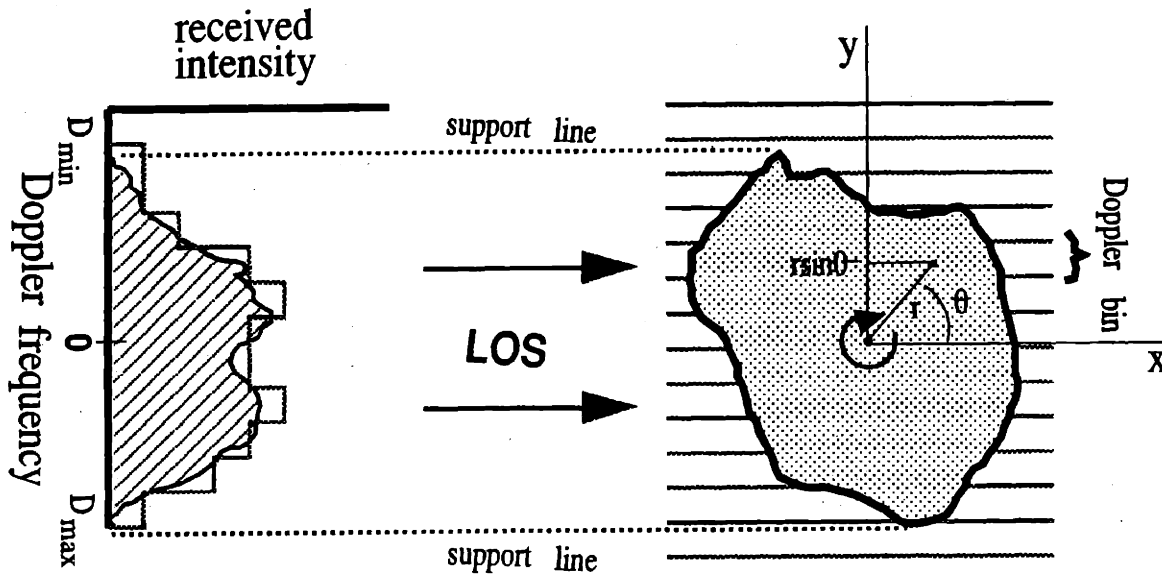


Figure 6.2: A laser radar illuminating a target along some LOS and receiving a Doppler-resolved return.

The received intensity from a target illuminated by a laser radar is proportional to¹ what is known as the *laser radar cross section* (LRCS) of the target. The LRCS σ is defined to be the area of a non-absorbing planar surface with Lambertian reflectance properties which, when viewed normally, produces a scattered intensity at the receiver equal to that of the target. Specifically,

$$\sigma = \frac{\pi I}{E}, \quad (6.1)$$

where I is the apparent radiant intensity of the target in units of W/sr and E is the irradiance at the target in W/cm² [31]. For a more general surface, the LRCS is dependent on the geometry and reflectance properties of the surface. That is, the

¹The proportionality constant depends on parameters of the laser radar's optics and parameters of the atmosphere.

reflectance properties are usually characterized by a function known as the bidirectional reflectance distribution function (BRDF) [17]. For the case of a monostatic radar and a surface with isotropic reflectance properties, the BRDF is denoted by $\rho(\psi)$ where ψ is the angle between the LOS and the local surface normal. In the case where the wavelength of the illumination is large compared to surface aberrations of the target, the LRCS is given by

$$\sigma = 2\pi \iint_S \rho(\psi) \cos^2 \psi \, dA, \quad (6.2)$$

where the integration is performed over the visible (illuminated) part of the surface, denoted by S . Hence, for resolved data, the intensity value associated with a particular bin is proportional to the LRCS arising from those portions of the target that contribute to that bin.

In this chapter, we apply the convex set algorithms to reconstruct a target from a series of range-resolved or Doppler-resolved measurements. Throughout, we consider only the case in which the data is taken at aspects around a great circle, so that the lines of sight all lie in a plane. With this restriction, the entire scenario is reduced to a two-dimensional problem in the plane containing the lines of sight, so that our algorithms may be used.

We can view the set of measurements at a number of aspects as being obtained by the sensor in either of two ways: (1) by a sensor fixed in space, observing a target rotating with its rotation axis normal to the sensor-target LOS for the duration of an entire rotation of the target, or (2) by a number of sensors distributed about the target, each taking a 'snapshot' of the target at a single instant in time. Both scenarios are equivalent in that they yield the same set of laser radar measurements. However, in terms of processing the data using our algorithms, different assumptions

must be made in each. In the single-sensor scenario, we assume that the target's rotation axis and rotation rate are fixed and known. In the multi-sensor scenario, we assume that the sensors are all observing, or tracking the target and that the relative positions of the sensors are known. Because target motion is required to resolve the target for Doppler-resolved measurements, in this case we also assume that the target is rotating with a fixed and known rotation rate and that the Doppler velocity of the target relative to each sensor is known. In the remainder of this chapter, we view the problem from the multi-sensor perspective.

6.2 Support Values from Resolved Laser Radar Data

Our interest in resolved laser radar measurements of a target lies in the fact that range-resolved and Doppler-resolved data contain support line information for the target. Given a range-resolved measurement, the minimum range r_{\min} with nonzero return intensity indicates that the distance from the sensor to *any* part of the target is at least r_{\min} (see Figure 6.1). Under far-field assumptions, the above indicates that the target lies completely on one side of the line perpendicular to the LOS at range r_{\min} . Moreover, since some part of the target is at range r_{\min} , this line actually grazes the target and is therefore a support line of the target.²

A Doppler-resolved measurement contains similar information. To see this, consider the two-dimensional coordinate frame defined in Figure 6.2 for a target undergoing simple rotation with known rate ω . For the moment, we assume that the target's translational velocity in the direction of the LOS is zero. The coordinate frame lies in

²Note that the maximum range r_{\max} with nonzero return intensity *does not necessarily* provide another support line. This is because parts of the target at ranges greater than r_{\max} may be blocked by parts of the target at lower range. Consequently, they will not be visible to the radar and will not contribute to the target's range spectrum.

a plane perpendicular to the rotation axis, with its origin located at the rotation axis and the x -axis aligned with the LOS. The y -axis is often referred to as the *cross-range* axis. Letting r denote the distance from a point on the target's surface to the origin of this frame, we have that the cross-range distance of the point is given by $r \sin \theta$. It follows that the target's extreme points in cross-range are given by $\min r \sin \theta$ and $\max r \sin \theta$, where the minimization and maximization are over all parts of the target's surface. Moreover, the Doppler frequency due to a point on the target is given by

$$D = \frac{2\omega r \sin \theta}{\lambda}, \quad (6.3)$$

with r as defined above. Hence, the Doppler frequency of a point is proportional to its cross-range distance, and in particular, the minimum and maximum Doppler frequencies, D_{\min} and D_{\max} , with nonzero return intensity correspond to the minimum and maximum cross-range of any part of the target. Thus, from a Doppler-resolved measurement we can extract two lines parallel to the LOS that graze the target which lies between them. These two lines are support lines of the target at angles 90° away from the aspect angle.

To identify the support value(s) associated with the support line(s) provided by a set of range or Doppler spectra, a coordinate frame such as that in Figure 2.1 is needed. This frame must serve as a common reference for all of the aspects, so that the sets of data may be spatially aligned, or *registered*.

For range-resolved measurements, the assumption that the positions of the laser radars are known relative to one another allows us to establish such a frame, say, with origin at the average of the laser radar position coordinates and 0° aspect defined by the LOS of the first laser radar (see Figure 6.3). The resulting position and orientation

of this coordinate frame is, of course, arbitrary; an alternate choice of frame would yield support values of a shifted and/or rotated target. Given such a coordinate frame, the support value corresponding to the i^{th} laser radar's range spectrum is equal to the minimum nonzero range r_{\min} subtracted from the distance from the laser radar to the origin along the i^{th} LOS. The set of support values obtained in this manner for the set of laser radars forms a support vector y .

A coordinate frame for Doppler-resolved measurements is established in the same way as for range-resolved measurements. As we discussed above, the i^{th} sensor (at aspect θ_i) gives rise to support values at $\theta_i \pm 90^\circ$. Since target cross-range is proportional to the Doppler frequency after shifting the Doppler spectrum by the Doppler frequency shift D_i produced by the target's translational velocity relative to the sensor, the support values are given by $\frac{\lambda}{2\omega} |D_{\min} - D_i|$ and $\frac{\lambda}{2\omega} |D_{\max} - D_i|$, where λ is the wavelength of the laser illumination. The support values arising from the Doppler spectra of the set of laser radars form a support vector y .

Since a Doppler spectrum at aspect θ_i provides two support values, at $\theta_i \pm 90^\circ$, the aspects θ_i and $\theta_i + 180^\circ$ yield duplicate support values, if the support values are free of noise. For noisy data (discussed below), the duplicate values may be averaged, thereby reducing the noise in the support measurements.

In general, this support vector y arising from range or Doppler data is noisy and may be invalid, due to two types of measurement errors. One type of error arises in incorrectly estimating the values of r_{\min} or D_{\min} and D_{\max} because of noise in the range or Doppler spectra. Secondly, incorrect knowledge of the relative laser radar positions (and for Doppler data, incorrect knowledge of the Doppler velocity of the target's center of gravity relative to each sensor) leads to registration errors. The convex set reconstruction algorithms produce shape estimates of targets given




Figure 6.3: Support value identification from range-resolved data.

support vectors having these measurement errors.³

6.3 Knot Location

In Sections 6.1 and 6.2, we described how the minimum range r_{\min} and minimum and maximum Doppler frequencies D_{\min} and D_{\max} having nonzero return intensity give rise to support values of the target. Although determining these minimum and maximum values is simple if the laser radar data is noise-free, doing so for noisy data is a quite difficult problem in general. The most obvious method—thresholding the data—suffers greatly from its nonrobustness to noise ‘spikes’ in the data. As a result, we turn to a method based on a technique developed by Willsky and Jones [30] for detecting abrupt changes in dynamic systems, and later applied by Mier-Muth and Willsky [16] to spline estimation and by Prince [19] to support line extraction from transmission projections. To cast our problem in the framework of [16], we model the range or Doppler spectrum as a linear *spline*, or piecewise linear function. The points of discontinuity in derivative are referred to as *knots*. Our goal is to determine the first knot in a range spectrum and the first and last knots in a Doppler spectrum.

Since [30], [16], and [19] contain complete, well-presented descriptions of the knot location algorithm and analyses of its performance, we provide only a brief description of the algorithm here. The basic approach consists of using a Kalman filter based on a linear ramp model for the range or Doppler spectrum. Initializing the filter with zero slope, we run the filter along the spectrum. At each bin, we use the innovations sequence to determine a set of maximum-likelihood (ML) estimates of the slope of the ramp at the current bin assuming that a knot was located at each of the previous

³Errors in knowing the laser radar positions may also cause angular errors in the measurements. However, our algorithms ignore these.

bins in some finite window. Using the ML estimates for each bin in the window, we perform a generalized likelihood ratio (GLR) test for the two hypotheses ‘knot present’ and ‘knot absent’ in order to determine whether a knot actually exists at the locations of any of the ML estimates. The first bin for which the GLR exceeds a prespecified threshold corresponds to the first knot in the spectrum. For a Doppler spectrum, to locate the last knot, we repeat the above process running the Kalman filter backwards along the spectrum.

In concluding this section, we note that it is in general more difficult to locate knots in a Doppler spectrum than in a range spectrum. This difference is due to the properties of typical target materials combined with the viewing geometries associated with the two data types [8]. In particular, the values of the laser radar return at ranges just higher than r_{\min} are determined by parts of the target whose surface normals roughly coincide with the LOS. As a result, $\psi \approx 0^\circ$, maximizing $\cos \psi$ in Equation 6.2. Furthermore, since materials typically give high intensity return at near-normal incidence and low intensity return at near-grazing incidence, the BRDF $\rho(\psi)$ is near maximum. Hence, range spectra generally exhibit an abrupt increase in intensity at the knot having range r_{\min} .

In contrast, the values of the laser radar return at Doppler velocities just greater than D_{\min} and just less than D_{\max} are determined by parts of the target having surface normals that are nearly perpendicular to the LOS. Consequently, $\psi \approx 90^\circ$ giving rise to values of $\cos \psi$ and $\rho(\psi)$ that are nearly zero. Hence, Doppler spectra generally vary slowly in intensity near the two knots.

6.4 Reconstructions from Simulated, Laboratory, and Field Laser Radar Data

In this section, we apply the reconstruction algorithms NUA, BNGON, and BNGONROT to laser radar measurements of several targets, in order to obtain shape estimates of the targets. The examples presented are those of reconstructions from sets of range and Doppler spectra obtained through a simulation computer program[9, 12] as well as spectra obtained through laboratory and field measurements.

The data for the first two examples are simulated range-resolved and Doppler-resolved measurements of a cone of height 200 cm and radius 25 cm with Lambertian reflectance characteristics. Note that projecting this cone onto a plane containing its axis of symmetry yields the standard triangle used in earlier chapters. The cone is positioned with the center of its base at the origin of a coordinate frame and oriented with its axis of symmetry lying in the xy -plane. In order to be resolved in Doppler, the cone rotates in the xy -plane about the z -axis at one revolution per second, in a manner resembling end-over-end tumble. Measurements are taken at an instant in time when the cone's axis is aligned with the frame's x -axis, at 72 aspects uniformly-spaced around the great circle of radius 10,000 m in the xy -plane, and with a resolution of 2 cm for the range data and a resolution of 3.750 KHz for the Doppler data. The values of these measurement parameters were chosen to reflect actual system parameters.

To reconstruct the targets, we first locate the knots by the Kalman filtering technique described in Section 6.3 and convert them to support values. Modeling knot location errors and registration errors for each aspect by statistically-independent samples from Gaussian distributions⁴ with variances σ_{kl}^2 and σ_{reg}^2 , the effective measurement error is Gaussian, with variance $\sigma_{eff}^2 = \sigma_{kl}^2 + \sigma_{reg}^2$ for range-resolved data. However, for Doppler-resolved data at an even number of uniformly-spaced aspects,

⁴In this section and the next, all Gaussian distributions are assumed to have zero mean.

(1) registration errors for aspects 180° apart are negatives of each other, and (2) the duplicate support value measurements provided by aspects 180° apart are averaged together. As a result, the knot location error may be obtained by drawing samples from a Gaussian distribution with variance $\sigma_{kl}^2/2$ for each aspect. The registration error may be obtained by drawing samples from a Gaussian distribution with variance $\sigma_{reg}^2/2$ for aspects $\theta_1, \theta_2, \dots, \theta_{M/2}$, and using the negatives of these samples for the aspects $\theta_{M/2+1}, \theta_{M/2+2}, \dots, \theta_M$. The effective measurement error is given by the sum of these two errors, for each aspect.

The support lines resulting from locating knots and corrupting the support values by measurement noise are shown in Figures 6.4a, 6.5a, and 6.6a for range-resolved data, with noise levels $\sigma_{eff} = 0.25, 0.35$, and 0.50 respectively. The reconstructions produced by NUA, BNGON, and BNGONROT from these three sets of noisy support line measurements are shown in Figures 6.4b-d, 6.5b-d, and 6.6b-d. The display conventions used in these figures will be used throughout this section. The reconstructions exhibit behavior similar to that seen for the standard triangle reconstructions of Chapters 3–5. In particular, the prior knowledge of relative reconstruction angles allows BNGONROT to dramatically outperform NUA, but does not cause it to significantly underperform BNGON, which has absolute angle information. Also, the quality of the reconstructions is rather impressive in light of the fact that the noise levels are so high, having standard deviations ranging from half of the target's width to its full width.

The corresponding results for the Doppler-resolved measurements arising from knot location error ($\sigma_{kl} = 0.25$) and registration error ($\sigma_{reg} = 0.25$) are shown in Figure 6.7. Conclusions similar to those for the range-resolved measurements may be drawn.




Figure 6.4: Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.25$. The display convention used in this figure and all subsequent figures in this section is $\frac{a}{c} \frac{b}{d}$, with (a) noisy support lines, (b) NUA reconstruction, $E_{\text{NUA}} = 0.50$, (c) BNGON reconstruction, $E_{\text{BNGON}} = 0.16$, and (d) BNGONROT reconstruction, $E_{\text{BNGONROT}} = 0.16$. Orientation estimate for BNGONROT: $\hat{\alpha} = -1.2 \times 10^{-6}$ degrees.




Figure 6.5: Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.35$. Reconstruction errors: 0.67, 0.24, 0.24. Orientation estimate: $\hat{\alpha} = -0.05^\circ$.




Figure 6.6: Support lines and convex set reconstructions for simulated range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.50$. Reconstruction errors: 0.91, 0.36, 0.36. Orientation estimate: $\hat{\alpha} = -0.24^\circ$.




Figure 6.7: Support lines and convex set reconstructions for simulated Doppler-resolved measurements of a cone. $\sigma_{kl} = \sigma_{reg} = 0.25$. Reconstruction errors: 0.60, 0.27, 0.27. Orientation estimate: $\hat{\alpha} = -1.51^\circ$.

The third example is one of reconstructing a triconic⁵ target of height 203 cm and base radius 39.5 cm (shown outlined in Figure 6.8a) given laboratory range-resolved measurements. The laboratory measurements were taken on a ten-meter indoor range at 72 uniformly-spaced aspects in the horizontal plane containing the target's axis of symmetry, with a range resolution of 1 cm. See [7] for details of the experimental set-up. Support lines and reconstructions using the three algorithms are shown in Figure 6.8 for the uncorrupted laboratory data and in Figure 6.9 for the laboratory data corrupted with measurement noise ($\sigma_{\text{eff}} = 0.25$).

Two comments concerning these examples should be made. First, the support lines shown in Figure 6.8a indicate that the knot location algorithm performed quite well. The aspects for which the knot location errors are largest correspond to views of the target from near nose-on lines of sight. Since the half-angle of the nose cone is small, the directions of the local surface normals relative to the LOS for these aspects coupled with the target material's reflectance properties gave rise to shallow increases in intensity near r_{min} . Second, the quality of the BNGON and BNGONROT reconstructions is remarkable in that they are convex reconstructions of a non-convex target. Not much of the quality is lost in introducing registration errors.

Finally, we present reconstructions from Doppler-resolved field measurements. The target, a scaled aluminum model of the Thor Delta rocket body (shown outlined in Figure 6.10a), was rotated at approximately 1 rpm about an axis normal to its axis of symmetry. The measurements, taken at 72 aspects in a plane normal to the rotation axis, were made using a $10.6\mu\text{m}$ CO_2 narrowband laser radar on a 5.4

⁵The target is composed of three conic sections stacked on top of each other (a cone on top of a cylinder on top of another cone).




Figure 6.8: Support lines and convex set reconstructions for laboratory range-resolved measurements of the triconic. $\sigma_{\text{eff}} = 0$. Orientation estimate: $\hat{\alpha} = -0.43^\circ$.




Figure 6.9: Support lines and convex set reconstructions for laboratory range-resolved measurements of the triconic. $\sigma_{\text{eff}} = 0.25$. Orientation estimate: $\hat{\alpha} = -5.06^\circ$.

km ground range, and have a Doppler resolution of approximately 200 Hz. Details of the experiment may be found in [7]. Support lines and reconstructions produced by the three algorithms are shown in Figure 6.10 for the uncorrupted field data and in Figure 6.11 for the field data corrupted with measurement noise ($\sigma_{kl} = 0.10$ and $\sigma_{reg} = 0.10$).

On looking at the support lines in Figure 6.10a, we note that the knots were detected reasonably well, except for those aspects near the tail of the target. Here, knots seem to cluster around two values. In fact, the laser radar data contain a spurious peak, probably an artifact of the field measurement set-up. These knot location errors give rise to reconstruction error near the tail that is somewhat larger than the error around the rest of the target. Very little reconstruction quality is sacrificed in introducing registration errors, as seen on a comparison of Figures 6.10 and 6.11.

6.5 Comparisons of Convex Set Reconstruction Methods with Tomographic Methods

In previous work, standard methods of tomographic image reconstruction [3, 6] have been applied to range-resolved and Doppler-resolved laser radar data [7, 18]. In this section, we compare the convex set reconstructions of the previous section with reconstructions produced using the tomographic methods. We then examine the effect of registration errors on both methods. As we shall see, the present algorithms are quite robust to registration errors, in contrast to tomographic reconstructions, which are rather sensitive to these errors. Finally, we show that the robustness of the present algorithms can be used to dramatically improve tomographic reconstructions




Figure 6.10: Support lines and convex set reconstructions for field Doppler-resolved measurements of the Thor Delta rocket body. $\sigma_{kl} = \sigma_{reg} = 0$. Orientation estimate: $\hat{\alpha} = -0.318^\circ$.




Figure 6.11: Support lines and convex set reconstructions for field Doppler-resolved measurements of the Thor Delta rocket body. $\sigma_{kl} = \sigma_{reg} = 0.10$. Orientation estimate: $\hat{\alpha} = -1.26^\circ$.

from data with registration errors.

All of the tomographic reconstructions in this section were obtained using the standard method of filtered backprojection. (See the references cited above for methods of transmission tomography and the application of these methods to laser radar reflective data.) Parts (a) of Figures 6.14–6.17 show filtered backprojection reconstructions from the four data sets (free of registration errors) used in Section 6.4. It should be noted that the Doppler data sets were thresholded prior to being backprojected in order to improve the tomographic reconstructions. This is necessary since the high intensities that are typically near the center of a Doppler spectrum tend to give rise to a dominant high intensity region in the center of the reconstruction. Incidentally, we threshold the data sets prior to backprojecting rather than thresholding the reconstructed images themselves, since the former approach appears to yield better results.

Unlike the convex set reconstructions (shown in parts (b)–(d) of Figures 6.4–6.11) which contain only binary information about the target (i.e., where the target is and where it is not), the tomographic reconstructions contain intensity information about the target. However, exactly what information the intensity values convey about the target's surface is not well understood. Furthermore, the tomographic images differ from their convex set counterparts in that they do not provide direct size or shape estimates of the target. While in principle techniques to extract edge and shape information could be used, the usual difficulties associated with image processing would be faced. In fact, since conventional image processing techniques assume white noise in an image, they would face additional difficulties in dealing with the correlated noise that results from backprojection. We also point out that processing a tomographic image arising from Doppler data would be more difficult than simi-




Figure 6.12: Tomographic reconstructions from range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.25$.




Figure 6.13: Tomographic reconstructions from range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.35$.




Figure 6.14: Tomographic reconstructions from range-resolved measurements of a cone. $\sigma_{\text{eff}} = 0.5$.




Figure 6.15: Tomographic reconstructions from Doppler-resolved measurements of a cone. $\sigma_{kl} = \sigma_{reg} = 0.25$.




Figure 6.16: Tomographic reconstructions from range-resolved measurements of the triconic. The top reconstruction is a ; the lower four are arranged as $\frac{b}{c}$. $\sigma_{\text{eff}} = 0.25$.




Figure 6.17: Tomographic reconstructions from Doppler-resolved measurements of the Thor Delta rocket body. The reconstructions are arranged as $\frac{a}{b/c}$. $\sigma_{kl} = \sigma_{reg} = 0.10$.

lar processing from one arising from range data. The reasons for this difference are suggested in Section 6.3 and described and demonstrated in [8]. Essentially, edges in Doppler-based reconstructions are not highlighted but are instead overwhelmed by the high intensities that are reconstructed in the interior of the target.

Like the convex set algorithms, tomographic techniques require knowledge of a common reference point, without which registration errors occur. The introduction of registration errors in the data has disastrous effects on the tomographic reconstructions that result. Parts (b) of Figures 6.12–6.17 show the tomographic reconstructions resulting from shifting the data in each spectrum by an amount given by a zero-mean Gaussian random variable with standard deviation $\sigma_{\text{reg}} = 0.25, 0.35, 0.50, 0.25, 0.25,$ and 0.10 (with the shifts for the spectra being independent of one another, except for the Doppler data sets, where shifts for aspects 180° apart are negatives), and then using filtered backprojection. Clearly, one cannot expect any image processing algorithm to successfully extract shape information from the tomographic images in these figures.

In contrast, the convex set algorithms are rather robust to registration errors. This is seen by the reconstructions shown in parts (b)–(d) of Figures 6.4–6.7, 6.9, and 6.11, obtained from data suffering from the *identical* registration errors as those used for the tomographic reconstructions (i.e., the same noise realizations were used), as well as knot location errors with the same standard deviations as above.

The difference in the robustness of tomographic and convex set methods to registration errors is due to the fact that the convex set algorithms attempt to register the data in the reconstruction process using implicit information as to the consistency of the measurements. That is, in adjusting the support values to achieve consistency, the algorithms are essentially shifting each range or Doppler spectrum such that the

sum of the squares of the shifts is minimal and such that the set of shifted laser radar data is *registered* data for *some* target.

In fact, we may exploit this registering property of the convex set algorithms as an aid to tomography, for data sets with registration errors. Specifically, we start with a possibly inconsistent set of measured support values $\{y_i\}$, which are estimated from the laser radar data by knot location. If we have no prior information as to the target's shape, we use NUA to obtain a consistent set of support values $\{\hat{h}_i\}$. If we have prior shape information, we use BNGON or BNGONROT to estimate a consistent set of support values at the reconstruction angles, and then sample the support function of the reconstruction, given by Equation 2.14, at the measurement angles to yield a consistent set of support values $\{\hat{h}_i\}$. Then, given the $\{\hat{h}_i\}$ and $\{y_i\}$, we shift the i^{th} range or Doppler spectrum by an amount $\hat{h}_i - y_i$, for all values of i . The resulting registered data set is then processed tomographically by the usual methods. This process is outlined in the block diagram of Figure 6.18.




Figure 6.18: Block diagram illustrating the process by which data containing registration errors are corrected by a convex set estimation algorithm and then processed tomographically.

Recall that parts (a) Figures 6.12–6.15, 6.16, and 6.17 show the tomographic reconstructions from laser radar data sets free of registration errors. Also recall that parts (b) of the same figures show the tomographic reconstructions from the data sets suffering from registration errors. Parts (c) and (d) of Figures 6.12–6.15 show the reconstructions arising with the aid of registration correction using NUA and

BNGONROT, for the standard cone. The reconstruction corresponding to BNGON is not shown since it appears identical to the BNGONROT reconstruction. Similarly, parts (c)–(e) of Figure 6.16 and (c) of Figure 6.17 depict the reconstructions obtained from registration-corrected data of the *Triconic* and the Thor Delta rocket body. In Figure 6.17, both the BNGON and BNGONROT reconstructions are omitted as they looked the same as the NUA reconstruction. On comparing these reconstructions with those in parts (b) of Figures 6.12–6.17 that arise from unregistered data, the dramatic improvement in the tomographic images is readily apparent.

Chapter 7 Conclusion

In this thesis, we have investigated the estimation of convex sets from noisy support line measurements. Specifically, we have developed several estimation algorithms that extended the work of Prince [19, 20, 21], analyzed their performance, and introduced and demonstrated their application to the problem of target reconstruction from laser radar data. In this chapter, we highlight the main results of our work and suggest potential topics for further research.

7.1 Summary

In Chapters 3, 4, and 5, we have described algorithms to estimate a convex set from a set of noisy support lines, using prior shape information, if available. The basic approach taken is was one of minimizing a cost function that incorporates the measurements and prior information, if available, subject to constraining the estimated support lines to be consistent with one another. The minimization is performed in support vector space, with the solution being given by that valid support vector closest in Euclidean distance to the measured support vector. The estimate of the convex set then consists of the (polygonal) intersection of the halfplanes associated with the set of estimated support lines.

The three algorithms we have developed were of increasing generality. The first algorithm, NUA, produces the maximum likelihood estimate of the convex set given measured support lines at a set of arbitrary angles. The resulting polygonal reconstruction has faces at the measurement angles. Often, however, prior shape information concerning the true object is available and should be utilized to improve the

quality of the reconstructions. If the object is a polygon (or is well-approximated by one), this prior may take the form of knowledge of the polygonal face angles. In this case, we would like to estimate a polygon whose faces are at a set of specified angles, given support line measurements at some arbitrary set of measurement angles that is independent of the set of reconstruction angles. Formulating such a problem results in our second algorithm, BNGON. However, since it is often unreasonable to assume the availability of such stringent prior information, we relaxed the prior to include knowledge of the shape of the object, but not its orientation. The resulting algorithm, BNGONROT, produces a polygonal reconstruction with known *relative* face angles by jointly estimating orientation and support values.

Since the first two algorithms, NUA and BNGON, are characterized by a quadratic cost function and linear consistency constraints, they can be implemented using standard quadratic programming techniques. The BNGONROT algorithm, however, is characterized by a nonquadratic cost function and nonlinear consistency constraints, which led us to an iterative implementation comprised of successive QP and gradient ascent/descent steps.

In analyzing our algorithms, we have made several observations concerning their properties. First, we noted that they are biased in that they produce reconstructions with average perimeter larger than that of the true object. Analyzing the geometry of the support cone for NUA, we concluded that the bias increases with measurement noise level and object eccentricity. Moreover, we quantified the bias through Monte Carlo simulation. We also pointed out the nonuniqueness that may arise in minimum-cost solutions of the BNGON, and hence BNGONROT, problems. This nonuniqueness is caused by having too few measurements for a given set of reconstruction angles.

Finally, we have analyzed the performance of the three algorithms versus a number

of object and measurement parameters and versus the quality of the prior information. We chose as our measure of error the area of the symmetric difference between the true object and the reconstruction, normalized by the area of the true object. This choice was motivated by our desire to obtain reconstructions that are in some sense 'close' to the true object, despite the fact that our algorithms perform the minimization of distance in support vector space for computational reasons.

In our assessment, we arrived at several important conclusions. First, as we expected, the reconstruction error increases with noise level and decreases with number of measurements for all three algorithms. Our plots of error versus object eccentricity lead us to believe that the choice of measurement angles becomes increasingly important as the underlying object becomes more eccentric. On explicitly comparing the performances of the three algorithms given correct prior information for BNGON and BNGONROT, we found that the performance of BNGON is far better than that of NUA, due to the prior information. However, the more important observation that was made is that BNGONROT substantially outperforms NUA, but is *not much worse* than BNGON, indicating that sacrificing prior orientation information does not lead to a sizable sacrifice in performance.

Further, we analyzed the performance of BNGON in the presence of incorrect prior information. We found that extraneous reconstruction angles degrade performance, but not to the extent that it would be better to abandon the prior and switch to NUA. Incorrect prior information may also consist of incorrectly knowing the relative or absolute reconstruction angles. In investigating the effect of rotating a single reconstruction angle or the entire configuration of these angles, we found that both yield average errors worse than those associated with NUA, for sufficiently large rotations. However, the single-angle rotations caused a greater degradation in perfor-

mance than did entire-configuration rotations of the same amount. Moreover, by switching to BNGONROT, the reconstruction error associated with entire-configuration errors in the prior would be substantially reduced, as this is precisely the uncertainty which BNGONROT is designed to eliminate.

Our assessment of BNGONROT's performance also included the classical error analysis for a nonlinear estimation problem. Specifically, formulating and solving a hypothesis testing problem enabled us to predict the probability of a global, anomalous orientation error as a function of noise level. In combination with the evaluation of the Cramer-Rao bound on the error variance of the orientation estimate, this provided us with an analytical expression describing the error variance as a function of noise level. As expected, BNGONROT experiences the dramatic threshold effect characteristic of a nonlinear estimator. These calculations were supported by Monte Carlo results.

In Chapter 6, we introduced the application of convex set reconstruction methods to laser radar data. In particular, we showed that range-resolved and Doppler-resolved measurements of a target contain support line information for the target. By using the knot location technique developed by Willsky and Jones [30], Mier-Muth and Willsky [16], and Prince [19] to extract the support line information from range or Doppler spectra and then using one of our three reconstruction algorithms, a direct size and shape estimate of the convex hull of the target's projection onto the plane of aspect angles is obtained. We then demonstrated this process by reconstructing several targets from simulation, laboratory, and field measurements.

There are several critical points we have observed in contrasting the reconstructions obtained through use of the present algorithms with those produced by tomographic imaging methods. First, shape estimates are explicitly provided by our algorithms, as opposed to the images produced by tomographic techniques, which

can provide target shape information only after image processing techniques—which are often somewhat unreliable—have been used. In fact, standard tomographic techniques such as convolution backprojection act to *correlate* the noise, causing unwanted artifacts in the reconstructions that may severely affect the performance of edge-detection algorithms. This is especially true in the case of sparse measurement data. We then investigated the effects of registration error on both methods and found that the reconstructions produced by tomographic methods experience substantial degradation, unlike those arising from the present methods, which are rather robust. These observations motivated us to exploit the ability of our algorithms to correct unregistered data, in an effort to improve the quality of tomographic images. Although the technique to do so is rather simple, its effectiveness was seen to be rather remarkable.

The fundamental difference between the two methods is due to the different approaches in which they deal with inadequacies in the measurement data. Given noise-free data at *all* angles, both produce accurate reconstructions. However, as noise level increases or number of measurements decreases, their performance differs: while the convex set algorithms *estimate* solutions given a set of measurements, the tomographic methods have no such mechanism to overcome noise, other than relying on the averaging that takes place during the convolution backprojection.

7.2 Suggestions for Further Work

We propose the following potential avenues for further research:

- *Extensions to three dimensions*: In many applications (e.g., medical imaging, computer vision, tactile sensing, and target reconstruction), the underlying scenario is a three-dimensional one, in which it is more natural to think of support *plane* rather than support line measurements. Upon relaxing the restriction

that aspects lie in a plane and allowing general aspects in three dimensions, our algorithms would produce shape estimates of three-dimensional objects rather than their projections onto a plane. The main problem in extending the algorithms lies in developing a necessary and sufficient consistency condition on the support planes. In two dimensions, we define consistency for a triplet of adjacent angles, and enforce this consistency for all triplets of adjacent angles. However, in three dimensions, there is no natural ordering of the aspect angles. Instead, the notion of a particular aspect angle lying 'between' three others on the unit sphere is useful. Karl has derived a consistency condition for such quadruplets of aspect angles. Moreover, on enforcing the quadruplet condition for all such quadruplets to yield a consistency condition on the set of support planes, he has proven its necessity.

- *More general laser radar scenarios:* The extension of the algorithms to three dimensions is related to useful extensions in the laser radar application. First, the general viewing geometry for laser radars is bistatic, where the transmitter and receiver have, in general, different lines of sight to the target. It may be possible to extract support plane information from bistatic laser radar data. Second, it has been suggested in [7] that the application of tomographic techniques to two-dimensional laser radar data resolved in both range and Doppler may provide three-dimensional tomographic reconstructions. In a similar way, it might be possible to obtain three-dimensional target shape estimates using the convex set reconstruction methods.
- *Noisy measurement angles:* In general, the measurement angles are not error-free. Modelling the angles as random variables with some known distribution

would lead to a formulation in which the angles and the support values at these angles would be jointly estimated. Solving the resulting problem appears to be very difficult, however, due to the fact that the entries of the consistency matrix C (see Chapter 3) are nonlinear functions of the angles.

- *Generalizations of the best N -gon problem:* It would be useful to develop algorithms that are generalizations of the two best N -gon problems considered in this thesis, so as to use lesser degrees of prior information. For example, one may want to reconstruct a triangle of unknown shape given some set of measurement data. The first level of generality would be achieved by specifying only the number of reconstruction angles, and not their (relative) values. Like BNGONROT, the resulting algorithm would be nonlinear. However, it would undoubtedly be far more complex, since N reconstruction angles (rather than a single orientation parameter) would have to be estimated. The second level of generality would consist of specifying neither the number nor values of the reconstruction angles. This problem would be formulated in essentially the same way as the one just discussed, with the exception that the cost function to be minimized would penalize a large number of angles.
- *Investigating the support cone geometry for the best N -gon problems:* It may be useful to investigate the support cone geometry for the best N -gon problems, in order to gain an understanding of conditions under which nonunique solutions occur. This might be done by considering two cones, in \mathbb{R}^M and \mathbb{R}^N , and exploring mappings between them.
- *Further improvements to tomographic reconstructions:* In Section 6.5, we introduced a technique by which tomographic images from unregistered data are

improved by using the convex set reconstruction algorithms to pre-register the data. Using this registered data as the input to Prince's hierarchical reconstruction algorithm [19, 22] would probably yield further improvements in the tomographic images.

- *Smooth shape estimates:* The present algorithms provide polygonal reconstructions of objects given support line measurements. It may be possible to obtain *smooth* shape estimates by developing estimation procedures that use a smooth representation of the object's support function. One such example is the use of a bandlimited Fourier series representation, where the choice of the bandwidth would determine the degree of smoothness of the reconstruction. However, enforcing consistency for this example appears to lead to highly nonlinear consistency constraints.
- *Minimization in curvature space:* Our algorithms minimize Euclidean distance between the measured support vector y and the estimated support vector \hat{h} subject to $\hat{h} \in \mathcal{C}$. Although the curvature function $K(\theta) = y''(\theta) + y(\theta)$ computed directly from the measurements may have had negative impulses, the consistency constraint ensures that only nonnegative impulses are present in $\hat{K}(\theta) = \hat{h}''(\theta) + \hat{h}(\theta)$. An alternative estimation procedure might consist of minimizing distance in curvature space, i.e., by simply setting any negative impulses to zero. It may be useful to investigate the consequence of such a procedure in object space.
- *Joint knot location and set reconstruction:* At present, the process of reconstructing an object from raw data consists of *first* locating knots in the data and *then* estimating the object using the measured support values. We feel that

jointly estimating the knots and the object should provide reconstructions that are at least as good as those provided by the present process. One possible implementation might be an iterative one, in which knots are first located from the raw data, then an object estimate is formed using the present algorithms, then knots are again located from the raw data, this time with prior information provided by the object estimate, and so on.

- *Weighted cost functions*: It seems natural to weight the measurement terms in the cost functions by quantities inversely proportional to the knot location variances. This has been studied in some detail by Prince [19], but more thorough investigations are probably warranted.

Appendix A Computation of the Area of the Symmetric Difference

In assessing the performance of our algorithms, we compute the area of the symmetric difference between an underlying object S that is either elliptical or polygonal and the polygonal estimate \hat{S} arising from noisy support line measurements of S . For all but two degenerate cases, in which the area of the symmetric difference is either the sum (when $S \cap \hat{S} = \emptyset$) or the positive difference (when $S \subseteq \hat{S}$ or $\hat{S} \subseteq S$) of the areas of S and \hat{S} , our algorithm consists of first determining points of intersection between S and \hat{S} and considering those parts of S and \hat{S} that lie outside the polygon formed by these intersection points. For an underlying object that is an ellipse, each line segment having endpoints at a pair of adjacent intersection points serves, in general, as an edge of a segment of the ellipse and an edge of a segment of the polygon (which of course is itself a polygon); see Figure A.1. Taking the magnitude of the difference between the areas of the elliptical segment and the polygon, and summing over all pairs of adjacent intersection points yields the area of the symmetric difference. The area of the symmetric difference between two polygons is computed in essentially the same manner. It should be noted that the present algorithm is not necessarily optimal in terms of, say, speed but is simply a straightforward method of performing the calculation.




Figure A.1: Segments of the ellipse and the polygon, whose areas contribute to the symmetric difference reconstruction error E

Appendix B Calculation of the Probability of Anomalous Orientation Estimate $\Pr(A)$

Our Monte Carlo simulations indicate that anomalous global minima tend to cluster at $\alpha = 270^\circ$, corresponding to a 180° rotation of the object. Motivated by this observation, we formulate a binary hypothesis testing problem in which the two hypotheses are $\alpha = 90$ (zero error) and $\alpha = 180$ (global error) in order to determine an approximation to the probability of anomaly $\Pr(A)$. For notational convenience, we define $h(\alpha) = [h_\phi(\phi_1) h_\phi(\phi_2) \cdots h_\phi(\phi_M)]^T$.

Under hypothesis H_0 , which corresponds to $\alpha = 90$, the measurement model is given by $y = h(90^\circ) + n$, where $n \sim N(0, \sigma^2 I)$. Likewise, under hypothesis H_1 we have $y = h(270^\circ) + n$, with the same probability distribution on n . With a minimum error probability cost assignment and equiprobable hypotheses, the decision rule is characterized by the Likelihood-Ratio Test (LRT)

$$\mathcal{L}(Y) \equiv \frac{p_{y|H_1}(Y|H_1)}{p_{y|H_0}(Y|H_0)} \begin{array}{c} \text{say } H_1 \\ > \\ < \\ \text{say } H_0 \end{array} 1. \quad (\text{B.1})$$

Since $p_{y|H_1}(Y|H_1)$ and $p_{y|H_0}(Y|H_0)$ are normally distributed with means $h(90^\circ)$ and $h(270^\circ)$ and covariance $\sigma^2 I$, the LRT reduces to

$$\ell \equiv (h(270^\circ) - h(90^\circ))^T y \begin{array}{c} \text{say } H_1 \\ > \\ < \\ \text{say } H_0 \end{array} \frac{h(270^\circ)^T h(270^\circ) - h(90^\circ)^T h(90^\circ)}{2} \quad (\text{B.2})$$

after some manipulation. The sufficient statistic ℓ is conditionally Gaussian under each of the hypotheses, with means and variances given by

$$\begin{aligned}
E(\ell|H_0) &= E[(h(270^\circ) - h(90^\circ))^T(h(90^\circ) + n)] \\
&\quad h(270^\circ)^T h(90^\circ) - h(90^\circ)^T h(90^\circ) \\
E(\ell|H_1) &= E[(h(270^\circ) - h(90^\circ))^T(h(270^\circ) + n)] \\
&\quad h(270^\circ)^T h(270^\circ) - h(270^\circ)^T h(90^\circ) \\
\text{var}(\ell|H_0) &= [h(270^\circ) - h(90^\circ)]^T [h(270^\circ) - h(90^\circ)] \\
\text{var}(\ell|H_1) &= [h(270^\circ) - h(90^\circ)]^T [h(270^\circ) - h(90^\circ)].
\end{aligned}$$

Because we have assumed equally likely hypotheses and a minimum error probability cost assignment, the probability of error (which, in our problem, is $\text{Pr}(A)$) is in fact the probability of detection or probability of false alarm. Using the means and variances of the conditional Gaussian distribution on ℓ from above, we have that

$$\text{Pr}(A) = Q\left(\frac{\sqrt{H}}{2\sigma}\right) \tag{B.3}$$

$$\leq \frac{1}{2} \exp\left(\frac{-H}{8\sigma^2}\right). \tag{B.4}$$

with H denoting $[h(270^\circ) - h(90^\circ)]^T [h(270^\circ) - h(90^\circ)]$.

References

- [1] C.G. Backman, *Laser Radar Systems and Techniques*. Artech House, Inc., 1979.
- [2] T. Bonnesen and W. Fenchel, *Theory of Convex Bodies*. BCS Associates, Moscow, Idaho, 1987. Originally published as *Theorie der Konvexen Körper* in German, Berlin, 1934.
- [3] R.A. Brooks and G. DiChiro, "Principles of Computer Assisted Tomography in Radiographic and Radioisotopic Imaging." *Phys. Med. Biol.* **21**, pp. 689-732, 1976.
- [4] J.P. Greschak, *Reconstructing Convex Sets*. Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 1985.
- [5] P.M. Gruber, "Approximation of Convex Bodies." From *Convexity and Its Applications*, edited by Gruber and Wills, 1983.
- [6] G.T. Herman (ed.), *Image Reconstruction from Projections*. Topics in Applied Physics, **32**, Springer-Verlag, 1979.
- [7] F.K. Knight, *et al.* "Tomographic Techniques Applied to Laser Radar Reflective Measurements." *The Lincoln Laboratory Journal*, **2**, pp.143-160, 1989.
- [8] S.R. Kulkarni, A.S. Lele, M.F. Reiley, "On the Qualitative Differences Between Tomographic Reconstructions from Range-resolved Versus Doppler-resolved Data." Lincoln Laboratory Project Memorandum 52PM-ODT-0042, 1989.
- [9] S.R. Kulkarni and A.S. Lele, *TARMAN: A Model for the Prediction of Laser Radar Signatures of Multiple Complex Dynamic Targets*. Lincoln Laboratory Project Report ODT-7, 1987.
- [10] A.H. Land and S. Powell, *Fortran Codes for Mathematical Programming*. Wiley-Interscience, London, 1973.
- [11] S.R. Lay, *Convex Sets and their Applications*. Wiley-Interscience, New York, 1982.
- [12] A.S. Lele, *TARMAN: A Computer Model for Dynamic Target Simulation*. S.B. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 1987.

- [13] D.G. Luenberger, *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA, second edition, 1984.
- [14] D.E. McClure and R.A. Vitale, "Polygonal Approximation of Plane Convex Bodies." *Journal of Mathematical Analysis and Applications*, **51**, pp. 326-358, 1975.
- [15] A.M. Mier-Muth, *Adaptive Knot Location for Spline Approximation*. S.M. Thesis, Massachusetts Institute of Technology, Department of Electrical Engineering, 1976.
- [16] A.M. Mier-Muth and A.S. Willsky, *A Sequential Method for Spline Approximation with Variable Knots*. Technical Report ESL-P-759, M.I.T. Electronic Systems Laboratory, 1977.
- [17] F.E. Nicodemus, *et al.*, *Geometrical Considerations and Nomenclature for Reflectance*. NBS Monograph 160, National Bureau of Standards, U.S. Department of Commerce, Washington, DC, 1977.
- [18] J.K. Parker, *et al.*, "Reflective Tomography: Images from Range-resolved Laser Radar Measurements." *Applied Optics*, **27**, pg. 2642, 1988.
- [19] J.L. Prince, *Geometric Model-Based Estimation from Projections*. Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 1988.
- [20] J.L. Prince and A.S. Willsky, *Estimation Algorithms for Reconstructing a Convex Set Given Noisy Measurements of its Support Lines*. Technical Report LIDS-P-1638, M.I.T. Laboratory for Information and Decision Systems, 1987.
- [21] J.L. Prince and A.S. Willsky, *Estimating Convex Shapes from Support Line Measurements Using Prior Geometric Information*. Technical Report LIDS-P-1823, M.I.T. Laboratory for Information and Decision Systems, 1988.
- [22] J.L. Prince and A.S. Willsky, *A Hierarchical Algorithm for Limited-Angle Reconstruction*. Technical Report LIDS-P-1843, M.I.T. Laboratory for Information and Decision Systems, 1989.
- [23] A.W. Rihaczek, *Principles of High-resolution Radar*. Mark Resources, Inc., 1977.
- [24] L.A. Santalo, *Integral Geometry and Geometric Probability*. Volume 1 of *Encyclopedia of Mathematics and its Applications*, Addison-Wesley, Reading, MA, 1976.
- [25] G.C. Shephard and R.J. Webster, "Metrics for Sets of Convex Bodies." *Mathematika*, **12**, pp. 73-88, 1965.
- [26] H. Stark and H. Peng, "Shape Estimation in Computer Tomography from Minimal Data," in *Pattern Recognition and Artificial Intelligence*. E.S. Gelsema and L.N. Kanal, ed. Elsevier Science Publishers B.V., North Holland, 1988.
- [27] F.A. Valentine, *Convex Sets*. McGraw-Hill, New York, 1964.

- [28] H.L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory*. Wiley and Sons, Inc., New York, 1968.
- [29] R.A. Vitale, *A representation theorem for compact convex sets in the plane*. Division of Applied Mathematics Report, Brown University, 1974.
- [30] A.S. Willsky and H.L. Jones, "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems." *IEEE Trans. Auto. Control*, Vol. AC-21, pp. 108-112, 1976.
- [31] P.W. Wyman, *Applied Optics* 7, pg. 207, 1968.