November 1990 LIDS-P-2005

Convex Set Estimation from Support Line Measurements and Applications to Target Reconstruction from Laser Radar Data*

A.S. Lele[1,2] S.R. Kulkarni[1,2] A.S. Willsky[2]

[1]Massachusetts Institute of Technology

Lincoln Laboratory

244 Wood Street

Lexington, Massachusetts 02173

[2]Massachusetts Institute of Technology

Center for Intelligent Control Systems and

Laboratory for Information and Decision Systems

77 Massachusetts Avenue

Cambridge, MA 02139

Abstract

This paper addresses techniques of convex set estimation from support line measurements and introduces the application of these methods to laser radar data. The algorithms developed

*This research was conducted at M.I.T. Lincoln Laboratory, sponsored by the Department of the Navy under Air Force contract F19628-90-C-0002, and at M.I.T. under U.S. Army Research Office contract DAAL03-86-K-0171 and National Science Foundation contract ECS-8700903.

utilize varying degrees of prior information. Quantitative assessments of their performance with respect to various parameters are provided. As expected, prior information as to object shape and orientation greatly improves performance. Interestingly, nearly the same performance is obtained with and without using the prior information about object orientation, enabling us to extract an estimate of orientation.

These convex set estimation techniques are applied to the problem of target reconstruction from range-resolved and Doppler-resolved laser radar data. The resulting reconstructions provide size and shape estimates of the targets under observation. While such information can be obtained by 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 and are much more robust to data suffering from registration errors. Finally, the present methods are used to improve tomographic images in the presence of registration errors.

1. INTRODUCTION

In this paper, we develop techniques for estimating convex sets from support line measurements and introduce the application of these methods to target reconstruction from resolved laser radar measurements. A support line of a two-dimensional convex set refers to any line which just grazes the boundary of the set, so that the set lies entirely in one of the halfplanes defined by the support line. Clearly, a convex set is completely determined by its support lines at all orientations, and can be obtained by simply intersecting the corresponding halfplanes. From a finite number of support lines an approximation to the convex set may be obtained in the same manner. However, any physical measurements giving rise to support lines are in general noisy, and simply intersecting the halfplanes may not yield satisfactory results. Prince and Willsky [1, 2] formulated the problem of estimating a convex set from noisy support line measurements and studied a variety of algorithms. (Greschak [3], Stark and Peng [4], and others have done related work.)

Here, we introduce three new estimation algorithms which utilize varying degrees of prior information. The first is simply an extension of an algorithm from [1, 2], which allows the measurements to be spaced nonuniformly in angle. The reconstructed polygon has sides at the measurement angles. The second algorithm allows the reconstruction angles to be specified independently of the measurement angles. This corresponds to the incorporation of prior information concerning the shape of the object, in that the reconstruction is the best N-gon with sides at fixed angles fitting the measurements. The final algorithm is similar to the second algorithm except that rotation of the constellation of reconstruction angles is allowed. Hence, this algorithm provides both an N-gon reconstruction as well as an orientation estimate.

There are a number of applications in which support line information can be extracted from physical measurements of an object. For example, in tomographic imaging the nonzero extent of each transmission projection provides support information of the underlying mass distribution

[1, 2]. Another possible application arises in tactile sensing, in which the support information can be obtained by a robot jaw repeatedly grasping an object [1, 2]. In these applications, convex set estimation algorithms can be used to provide reconstructions of the object, either independently or in conjunction with other algorithms.

An application introduced and studied in this paper is that of target reconstruction from laser radar data. 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 (LOS). If a plane is drawn at this range perpendicular to the line of sight, the target lies completely to one side of this plane and in fact is just grazed by this plane. Range-resolved measurements from a number of aspects yield a set of support planes of the target which could conceivably be used to obtain a three-dimensional estimate of the target. In this paper, we restrict attention to the case in which all the aspects, i.e. LOS's, lie in a plane, so that the problem is reduced to two dimensions. In a similar manner, Doppler-resolved measurements of a spinning target contain support plane information, which reduces to support line information under the restriction that all aspects lie in a plane.

The basic reconstruction procedure of this paper may be decomposed in the manner indicated by Figure 1. In the first module, support line information is extracted from the laser radar data using an estimation procedure developed by others [5, 6] for the location of knots in spline approximations. The support line measurements obtained are in general noisy for various reasons. In the second module, the convex set estimation techniques developed in this paper provide an estimate of the target given these noisy support line measurements and prior knowledge as to target shape.

Section 2 provides the background, formulation, and basic approach to the estimation problem. The three specific estimators are presented in Section 3, along with an assessment of their performance. In Section 4, we describe the laser radar data, the way in which the data contain support

line information, and the technique used to extract this information. The reconstruction algorithms are applied in Section 5 to range-resolved measurements and Doppler-resolved measurements obtained through simulated, laboratory, and field measurements. Reconstructions obtained by using the present methods are compared with reconstructed images produced by standard tomographic methods [7]. Also, a a method is introduced whereby the tomographic reconstructions from unregistered data may be greatly improved using our reconstruction algorithm as a preprocessor of the data. Finally, in Section 6 our results are summarized and possible directions for further work are suggested.

2. CONVEX SET ESTIMATION FROM SUPPORT LINE MEASUREMENTS

In this section we first discuss the ideas of support lines and support functions of convex sets. Although the exact support values at all angles characterize a convex set, in many applications only a finite number of noisy measurements are available. Accordingly, we formulate and discuss the basic approach to the problem of estimating a convex set from such measurements.

2.1 Background and Definitions

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). The support value $h_S(\theta_0)$ is defined as the maximum projection onto $\omega(\theta_0)$ of all points in S:

$$h_S(\theta_0) = \sup_{s \in S} s^T \omega(\theta_0), \tag{1}$$

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

$$L_S(\theta_0) = \{x \in \mathbf{R}^2 | x^T \omega(\theta_0) = h_S(\theta_0)\}. \tag{2}$$

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

We will 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 $\theta_1, \theta_2, \ldots, \theta_M$ yields a support vector $h_S = [h_S(\theta_1) \quad h_S(\theta_2) \quad \cdots \quad h_S(\theta_M)]^T$. Also, from Figure 2, it should be apparent that support lines provide no information as to concavities in the set, so that support lines at all orientations determine only the convex hull of the set. In fact, the support function of a set uniquely 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. In this paper, 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.

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. A number of necessary and sufficient conditions for a function to be a valid support function have been developed [8, 9, 10, 11, 12]. In this paper, we will be using a version of the following condition suitable for support vectors: A twice-differentiable function $h(\theta)$ is a valid support function iff $h''(\theta) + h(\theta) \geq 0$ (for example, see [11]). 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.

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 continuous-angle framework, support functions possess several useful properties which reduce in a natural way to corresponding properties for support vectors. The set S_1 consisting of a single point located on the x-axis at $(x_0, 0)$ has support function $h_{S_1}(\theta) = x_0 \cos\theta$. Similarly, the set S_2 located at $(0, y_0)$ has the support function $h_{S_2}(\theta) = y_0 \sin\theta$. By the property of support functions $h_{A\oplus B} = h_A + h_B$ [9] where $\oplus$ 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 $h(\theta) = x_0 \cos\theta + y_0 \sin\theta$. Now, by representing a polygon S_P as the convex hull of its vertices (x_i, y_i) and from the property that [9] $h_{conv(A\cup B)} = \max(h_A, h_B)$ we find that the support function of a polygon is given by $h_{S_P}(\theta) = \max_i(x_i \cos\theta + y_i \sin\theta)$. For a polygon, the cusps of $h(\theta)$ correspond to the sides of the polygon, at which support is transferred from one vertex to another. Denoting θ_{i-1} and θ_i as the angles of the faces which intersect to form the i^{th} vertex, we may rewrite the support function of a polygon as $h_{S_P}(\theta) = x_i \cos\theta + y_i \sin\theta$ on $[\theta_{i-1}, \theta_i]$ to more clearly reveal the trigonometric spline form that $h_{S_P}(\theta)$ takes.

2.2 Formulation and Basic Approach

We model our support measurements $y_1, y_2, \ldots, y_M$ at measurement angles $\theta_1, \theta_2, \ldots, \theta_M$ as consisting of the true support values of the set $h_i = h(\theta_i)$ corrupted by noise. That is, $y_i = h_i + n_i$ for $i = 1, 2, \ldots, M$ where the $\{n_i\}$ are samples from some noise distribution. We emphasize that by noisy measurements we mean uncertainty in the support values and not in the measurement angles. The measured support values y_i constitute the elements of what we refer to as the measured support vector y.

Due to the presence of the noise, the measurements $\{y_i\}$ may not correspond to the set which

gave rise to them and, in fact, may not correspond to *any* set. We adopt the approach introduced in [1, 2] for estimating a convex set by finding the valid support vector $\hat{h}$ that is closest in some sense to y. A natural choice is to minimize Euclidean distance in support vector space. This choice results in an estimation problem for which there exist efficient computational algorithms.

As previously mentioned, there are many convex sets sharing the same (valid) estimated support vector $\hat{h}$. The particular convex set that we associate with $\hat{h}$ is the polygon bounded by the support lines corresponding to $\hat{h}$. Since we are more directly interested in minimizing some measure of distance in object space, we evaluate the quality of our reconstructions by using a quantitative measure of the error between the true object S and its reconstruction $\hat{S}$. The measure of error we use here is the area of their symmetric difference $S \Delta \hat{S} = (S \cup \hat{S}) \setminus (S \cap \hat{S})$.

The idea of choosing $\hat{h}$ as the valid support vector closest to y can be given a nice geometric interpretation [1, 2]. Specifically, the set of valid M-dimensional support vectors forms a cone C in $\mathbf{R}^M$, and the estimation procedure can be regarded as projecting the measurement vector y onto C (see Figure 3b). To carry out this projection, a characterization of the support cone is required. Prince and Willsky gave such a characterization by proving a necessary and sufficient condition for a vector h to be a valid support vector in the case of uniformly spaced measurement angles. Our estimators require an extension of this consistency condition to the case of angles which are in general nonuniformly-spaced. This condition is a discrete version of the consistency condition $h''(\theta) + h(\theta) \geq 0$ for twice-differentiable support functions. In fact, the discrete version can be obtained by using the trigonometric spline form of support function of a polygon and interpreting the derivatives in a distributional sense. The discrete condition can also be obtained from a geometric approach as in Figure 3a. Given support lines L_{i-1} and L_{i+1} at θ_{i-1} and θ_{i+1}, a third support line at θ_i is consistent only if it lies to the left of the intersection point of L_{i-1} and L_{i+1}. Together with sufficiency as shown in [1, 2], this leads to the consistency condition for

a triplet of support values at adjacent and in general nonuniformly-spaced angles, given by [13]:

$$h_{i-1}\sin(\theta_{i+1}-\theta_i) - h_i\sin(\theta_{i+1}-\theta_{i-1}) + h_{i+1}\sin(\theta_i-\theta_{i-1}) \geq 0. \tag{3}$$

Enforcing this inequality for all adjacent triplets yields a necessary and sufficient condition for a vector to be a valid support vector. Namely,

$$Ch \geq 0 \tag{4}$$

where

$$C = \begin{pmatrix} -\sin(\theta_2-\theta_M) & \sin(\theta_1-\theta_M) & 0 & 0 & \sin(\theta_2-\theta_1) \\ \sin(\theta_3-\theta_2) & -\sin(\theta_3-\theta_1) & \sin(\theta_2-\theta_1) & 0 & 0 \\ 0 & \sin(\theta_4-\theta_3) & -\sin(\theta_4-\theta_2) & \sin(\theta_3-\theta_2) & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots \\ \sin(\theta_M-\theta_{M-1}) & 0 & 0 & \sin(\theta_1-\theta_M) & -\sin(\theta_1-\theta_{M-1}) \end{pmatrix} \tag{5}$$

With such a consistency condition, we can formulate several estimation algorithms which are discussed in the next section.

3. ESTIMATION ALGORITHMS

The estimation algorithms that we present in the following three sections arise from increasingly general formulations of the problem of obtaining polygonal shape estimates from noisy support measurements. The most specific case was considered by Prince and Willsky [1, 2], in which 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 not necessarily the same as the set of of 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 following subsections treat the three generalizations.

3.1 Reconstruction with sides at the measurement angles

In this problem, we have a finite set of noisy support measurements $\{y_1, y_2, \ldots, y_M\}$ at angles $\theta_1 < \theta_2 < \cdots < \theta_M$. We wish to reconstruct a convex polygon, or equivalently the valid support vector $\hat{h}$, with sides at the measurement angles, that is closest to the measurement vector y. The solution is obtained by solving

$$\hat{h} = \text{argmin}_{Ch \geq 0} \sum_{i=1}^{M} (y_i - h_i)^2 \tag{6}$$

where C is given by (5). We refer to this estimator as **NUA**, an acronym for **N**on**U**niform **A**ngles, since the present algorithm is an extension of one developed by Prince and Willsky for uniformly-spaced angles. Since the cost function is quadratic and the constraints are linear, the solution to this problem can be obtained using standard quadratic programming techniques [14, 15]. We note that **NUA**, as well as the algorithms of the following subsections result in biased estimators (see Appendix A).

To illustrate the behavior of **NUA**, we consider the following example. The original object used in this example is an isosceles triangle with vertices at $(2, 0)$, $(-0.25, 0)$, and $(0.25, 0)$. We use this triangle throughout the paper and refer to it as the 'standard triangle.' The data consist of $M = 24$ uniformly-spaced noisy measurements ($\sigma = 0.25$). Figures 4a,b depict the results in both object space and support function space using the estimator **NUA**. Figure 4a shows the bold outline of the true object (the standard triangle), the noisy support lines, and the shaded polygonal reconstruction produced by **NUA**. Correspondingly, Figure 4b 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 quantitative measure of reconstruction error that we use throughout the paper 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 $E = 1.56$.

3.2 Best N-gon fitting M measurements with fixed reconstruction angles

In this section, we exploit prior information as to the angles of the object's sides, in order to obtain reconstructions of higher quality than those we expect to obtain using **NUA**, which utilizes no prior information. Specifically, we consider the problem of determining the N-sided polygon with prespecified face angles that best fits a set of noisy support values at M measurement angles. For example, one might wish to reconstruct the best equilateral triangle given a set of, say, twenty noisy measurements of an object known *a priori* to be triangular.

In formulating this problem, we let $\{\theta_1, \theta_2, \ldots, \theta_M\}$, $\{y_1, y_2, \ldots, y_M\}$, and $\{\phi_1, \phi_2, \ldots, \phi_N\}$ 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 $\{h_\phi(\phi_1), h_\phi(\phi_2), \ldots, h_\phi(\phi_N)\}$ at the specified N-gon face angles which minimizes

$$J(h_\phi(\phi_1), h_\phi(\phi_2), \ldots, h_\phi(\phi_N)) = \sum_{i=1}^{M} (h_\phi(\theta_i) - y_i)^2, \tag{7}$$

where $h_\phi(\theta_i)$ denotes the value at θ_i of the support function $h_\phi(\cdot)$ of our estimated N-gon (see (8) and the associated explanation). Eq. (7) corresponds to finding a set of support values (i.e. finding $h_\phi(\phi_i)$ for all i) at the reconstruction angles that minimizes the sum of the squared deviations between the measured support values and the values of the piecewise sinusoidal support function of the reconstructed polygon at the measurement angles.

Let 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 let h_{L_i} and h_{R_i} denote the corresponding reconstructed support values. Since the support function of a polygon is piecewise sinusoidal with cusps at the face angles, we can obtain the entire support function from its values at the face angles by simply determining the appropriate sinusoid in each interval. That is, the support function of the reconstructed object evaluated at θ_i is given by

$$h_\phi(\theta_i) = \frac{\sin(\phi_{R_i} - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{L_i} + \frac{\sin(\theta_i - \phi_{L_i})}{\sin(\phi_{R_i} - \phi_{L_i})} h_{R_i}. \tag{8}$$

From (7) and (8), our problem is formulated as

$$\hat{h}_\phi = \begin{bmatrix} \hat{h}_\phi(\phi_1) \\ \hat{h}_\phi(\phi_2) \\ \vdots \\ \hat{h}_\phi(\phi_N) \end{bmatrix} = \text{argmin}_{Ch_\phi \geq 0} (Ah_\phi - y)^T Ah_\phi, \tag{9}$$

where $y = \begin{bmatrix} y_1 & y_2 & \cdots & y_M \end{bmatrix}^T$ is the measurement vector, C is the consistency matrix of (5) with the θ_i's replaced by ϕ_i's, and A is an $M \times N$ matrix mapping the N support values at the $\{\phi_i\}$ to the corresponding N-gon support values at the $\{\theta_i\}$ using (8). The i^{th} row of the matrix A, corresponding to the i^{th} measurement, has two adjacent (modulo N) non-zero entries, $\frac{\sin(\phi_{R_i} - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})}$ and $\frac{\sin(\theta_i - \phi_{L_i})}{\sin(\phi_{R_i} - \phi_{L_i})}$, corresponding to the reconstruction angles ϕ_{L_i} and ϕ_{R_i} on either side of θ_i.

We refer to the estimator of (9) as **BNGON**. As before, since the cost function in (9) is quadratic in the reconstructed support values and the consistency constraint is linear, the problem can be solved by QP techniques. Incidentally, under certain conditions there may be nonunique solutions. However, this is not the generic case, and we refer the reader to Appendix B and [13] for more details.

An example of **BNGON**, similar to that discussed in Section 3.1, is shown in Figure 4c,d. The example consists of reconstructing the best triangle with reconstruction angles at $7.125°, 82.875°$, and $270°$ equal to those of the standard triangle, given $M = 24$ uniformly-spaced noisy ($\sigma = 0.25$) support measurements. The pictures in both object space and support function space are shown, with the reconstructed object incurring an error of $E = 0.17$ with respect to the true object.

The **BNGON** reconstruction in the figure originates from the same set of measurements as the **NUA** reconstruction in the same figure (i.e., the same noise realization was used), allowing comparison of the two. Visually, 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 $E_{\text{BNGON}} = 0.17$ while $E_{\text{NUA}} = 1.56$. However, prior information as to the number of faces and the face angles of the true object may not be known precisely, and in this case one would expect some degradation in performance. Nevertheless, as we will see in Section 3.4, **BNGON** still outperforms **NUA** even in the presence of a broad range of errors in the prior information. Furthermore, one important source of errors leads to a natural generalization of the **BNGON** algorithm. Specifically, while in many cases it may be reasonable to assume that one has prior information about number of faces and their relative angles, one would typically not expect to have prior information on the absolute orientation of the object. In the next subsection we describe a generalization of **BNGON** that addresses this problem.

3.3 Best N-gon with fixed relative spacing of reconstruction angles

In this subsection, we assume somewhat less prior information than in **BNGON** by formulating a problem in which the relative (rather than the absolute) angles of the object's sides are known. Specifically, let $\{\theta_1, \theta_2, \ldots, \theta_M\}$ and $\{y(\theta_1), y(\theta_2), \ldots, y(\theta_M)\}$ denote the M measurement angles and the measured support values at these angles, as before. However, unlike before, the recon-

struction angles are given by $\{\phi_1 + \alpha, \phi_2 + \alpha, \ldots, \phi_N + \alpha\}$, where $\{\phi_1, \phi_2, \ldots, \phi_N\}$ are known and $\alpha \in [0, 2\pi)$ serves as an unknown offset parameter fixing the absolute locations of the reconstruction angles. Essentially, we wish to minimize the cost function in (9), with the exception that the estimator here is free to rotate the constellation of reconstruction angles in order to achieve minimum cost in the estimate. That is, we wish to jointly estimate values of α and $\{h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \ldots, h_\phi(\phi_N + \alpha)\}$ that minimize

$$J(\alpha, h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \ldots, h_\phi(\phi_N + \alpha)) = \sum_{i=1}^{M} (h_\phi(\theta_i) - y(\theta_i))^2 \tag{10}$$

where $h_\phi(\theta_i)$, using (8), is given by

$$h_\phi(\theta_i) = \frac{\sin(\phi_{R_i} + \alpha - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{L_i} + \frac{\sin(\theta_i - \phi_{L_i} - \alpha)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{R_i} \tag{11}$$

and the $\{h_\phi(\phi_i + \alpha)\}$ are constrained to be a set of consistent support values.

Note that while the criterion (10) is quadratic in the values of h_ϕ at the reconstruction angles $\phi_1 + \alpha, \ldots, \phi_N + \alpha$, it is certainly not quadratic in α (see (11)). Thus, the minimization of (11) is not a simple quadratic programming problem. Nevertheless the structure of this criterion does allow us to obtain a reasonably efficient QP-based optimization algorithm, which we refer to as **BNGONROT**. Specifically, let $J_{h_\phi}(\alpha)$ denote the cost resulting from a best choice of support vector for a fixed value of α:

$$J_{h_\phi}(\alpha) = \min_{\{h_\phi(\phi_i + \alpha)\}} J(\alpha, h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \ldots, h_\phi(\phi_N + \alpha)) \tag{12}$$

Note that solving (12) for any given value of α is simply a **BNGON** problem solved via a QP computation as described previously. Note also that the minimization of (10) corresponds to choosing α to minimize $J_{h_\phi}(\alpha)$. Thus a brute force approach to minimizing (10) is to perfrom an

exhaustive search over the values of $J_{h_\phi}(\alpha)$, where each evaluation of this function involves a QP computation. Our more efficient method involves a gradient-like search for the optimum value of α. However, the nature of the problem is such that there are two important distinctions between our algorithm and a standard gradient descent algorithm. First of all, note that

$$J_{h_\phi}(\alpha) = J(\alpha, h_\phi^*(\phi_1 + \alpha), h_\phi^*(\phi_2 + \alpha), \ldots, h_\phi^*(\phi_N + \alpha)) \tag{13}$$

where $h_\phi^*(\phi_i + \alpha),\ i = 1, \ldots, N$ are the optimal values of the reconstructed support vector for the given value of α. Thus,

$$\frac{dJ_{h_\phi}(\alpha)}{d\alpha} = \frac{\partial J}{\partial \alpha} + \sum_{i=1}^{N} \frac{\partial J}{\partial h_\phi(\phi_i + \alpha)} \frac{\partial h_\phi^*(\phi_i + \alpha)}{\partial \alpha} \tag{14}$$

The difficulty here is that computing the sensitivity of the optimal support values $h_\phi^*(\phi_i + \alpha)$ with respect to α is not easily obtained (since a QP optimization is involved). Thus, instead of (14), we simply use

$$\frac{\partial J}{\partial \alpha}(\alpha, h_\phi^*(\phi_1 + \alpha), h_\phi^*(\phi_2 + \alpha), \ldots, h_\phi^*(\phi_N + \alpha)) = 2 \sum_{i=1}^{M} \frac{\partial h_\phi(\theta_i)}{\partial \alpha} (h_\phi^*(\theta_i) - y(\theta_i)) \tag{15}$$

where $h_\phi^*(\theta_i)$ is obtained from (11) with the h_{L_i}, h_{R_i} values corresponding to the $h_\phi^*(\phi_i + \alpha)$, and

$$\frac{\partial h_\phi(\theta_i)}{\partial a} = \frac{\cos(\phi_{R_i} + \alpha - \theta_i)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{L_i} - \frac{\cos(\theta_i - \phi_{L_i} - \alpha)}{\sin(\phi_{R_i} - \phi_{L_i})} h_{R_i} \tag{16}$$

The second key point is that $J_{h_\phi}(\alpha)$ is, in general, a highly nonconvex function of α (see [13] for examples and discussion). Thus it is necessary to determine *all* of the local minima of $J_{h_\phi}(\alpha)$ and choose the one with lowest cost. Since α is a scalar, constrained to lie between 0° and 360°, we can do this in the following manner.

We begin at $\alpha = 0°$ and solve the QP problem of (9). Using the estimated support values, we compute (14) 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 (9), (2) compute the gradient, 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. We then advance by some small amount in α, and by solving (9) and computing the the gradient, determine whether our next series of 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 $[0°, 360°)$ until we have located all maxima and minima, and then choose the global minimum $\hat{\alpha}$. Solving (9) 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. 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 $J_{h_\phi}(\alpha)$ (see [13]). 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 specified 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 tradeoff 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.

An example of a reconstruction produced by **BNGONROT** is shown in Figure 4e,f. The true object and measurements are the same as before. The reconstruction forms an angle of $\alpha = 86.58°$ with the positive x-axis. The error E equals 0.42. Not surprisingly, the reconstruction is qualitatively and quantitatively far better than that corresponding to **NUA** (see Figure 4a,b). Moreover, it is not much worse than the **BNGON** reconstruction (see Figure 4c,d), indicating that not much is sacrificed in settling for a weaker prior, i.e., knowing relative rather than absolute reconstruction angles.

3.4 Performance assessment of the estimation algorithms

We have evaluated our algorithms by computing the average normalized symmetric difference area E for a range of values of relevant parameters. In particular, this Monte Carlo analysis is carried out versus measurement parameters and quality of the prior information.

In Figure 5a, we show a plot of E versus measurement noise level σ for each of the three algorithms, with $M = 24$ uniformly-spaced measurements of the standard triangle. While the error for all three algorithms increases with σ as expected, the performance of **BNGONROT** is much better than that of **NUA** but only slightly worse than that of **BNGON** for relatively low noise levels. The difference in performance between **BNGON** and **BNGONROT** becomes more pronounced near $\sigma = 0.17$. This threshold effect is exactly that characterizing standard nonlinear estimators, and is analyzed below and in more detail in [13]. However, even with this increased degradation, **BNGONROT**'s performance is much better than that of **NUA**. Figure 5b shows a plot of E versus number of measurements M for **NUA** and **BNGON**, with noisy ($\sigma = 0.25$) measurements of the standard triangle. Again, **BNGON** outperforms **NUA**, where both yield decreasing values of E with increasing M.

The performance of **BNGON** and **BNGONROT** is also dependent on the quality of the prior information. For example, let us examine the sensitivity of **BNGON**. Specifically, we take $M = 24$ noisy ($\sigma = 0.25$) measurements of the standard triangle. However, we reconstruct a triangle whose face angles are not the same as those of the standard triangle. Errors in the reconstruction angles considered are entire-configuration errors in which all face angles are rotated by the same amount, and single-angle errors in which only the first reconstruction angle, originally at 7.125°, is in error. The angular error is denoted by α. Figure 5c depicts plots of error E versus α for the two types of errors in the prior. A dashed line denoting the baseline performance level of **NUA** is included for comparison. The figure indicates that for values of α less than 69°, entire-configuration errors are less damaging than corresponding single-angle errors. Moreover, 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 $\approx 17°$ and entire-configuration errors of up to $\approx 29°$ before abandoning the **BNGON** algorithm and resorting to either **BNGONROT** or **NUA**.

Obviously, knowledge of the precise number of sides of the desired reconstruction is a very powerful and, in some sense, unrealistic piece of prior information. Thus, it is of interest to see how **BNGON** performance degrades as the number N of reconstruction angles is increased. To do this, we again use 24 uniformly-spaced noisy ($\sigma = 0.25$) measurements of the standard triangle. We start with the correct triple of reconstruction angles at $\phi_1 = 7.125°$, $\phi_2 = 82.875°$, and $\phi_3 = 270°$ 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 5d. 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 reconstruction 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 **BNGON** performs significantly better than **NUA** near $N = 24$, indicating that the original three reconstruction angles that are not available to **NUA** are quite helpful to **BNGON**.

Finally, we investigate the performance of **BNGONROT** in estimating the orientation parameter α. In Appendix C, we analytically determine the Cramer-Rao lower bound [16] on the orientation error variance as well as an approximate expression for the probability of obtaining anomalous orientation estimates (the 'threshold phenomenon' mentioned previously). Together, these expressions describe the orientation error variance for a range of noise levels, and is given by

$$\begin{aligned} \mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}}) &\approx [1 - \Pr(A)](\mathrm{CRB}) + \Pr(A)(180°)^2 \\ &\approx \left[1 - \frac{1}{2}\exp\left(\frac{-H}{8\sigma^2}\right)\right] \frac{\sigma^2}{\frac{\partial h^T(\alpha)}{\partial \alpha}\frac{\partial h(\alpha)}{\partial \alpha}}\Bigg|_{\alpha=\alpha_{\mathrm{true}}} \\ &\quad + \frac{1}{2}\exp\left(\frac{-H}{8\sigma^2}\right)(180°)^2. \end{aligned} \tag{17}$$

where the probability $\Pr(A)$ of anomalous orientation estimate, the Cramer-Rao bound CRB, and H are given in Appendix C.

Monte Carlo simulations supplementing this analytical error analysis were also performed. As before, we use $M = 24$ uniformly-spaced measurements of the standard triangle so that the true orientation, or offset parameter is given by $\alpha_{\mathrm{true}} = 90°$. Figure 6 compares the analytical expression for the standard deviation of the orientation estimate with the Monte Carlo results, for a range of noise levels. Plots of $\sqrt{\mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}})}$ versus σ are shown on both normal and logarithmic scales, with each Monte Carlo data point representing 200 noise realizations. The Monte Carlo results agree reasonably well with (17) and, as expected for a nonlinear estimator, exhibit dramatic

threshold behavior as the noise variance increases from low values, where the CRB dominates, to higher values, where $\Pr(A)$ is significant.

4. LASER RADAR DATA AND EXTRACTION OF SUPPORT LINE MEASUREMENTS

In this section, we describe the laser radar data to be used as input to the reconstruction algorithms (namely, range-resolved and Doppler-resolved data). We then discuss the way in which the laser radar data can serve to provide support line information, and describe a technique to extract such support measurements from the data. Previous work in reconstructing targets from such laser radar data have primarily employed techniques from transmissive tomography [7, 17]. These transmissive tomographic techniques are designed to provide reconstructions of an object's mass density given measurements of line integrals of this mass density. Although the application of these techniques to laser radar data provide reconstructions containing some geometric qualities of the target, it is not at all clear what the reconstructed intensities represent. Our motivation in using convex set reconstruction techniques stems from the fact that although the laser radar data do not convey line integral information, they do in fact provide target support information.

4.1 Laser radar data and problem scenario

By illuminating a target and receiving the reflected signal, laser radars provide information about the surface characteristics of the target. Laser radars can be designed to resolve the return from the target with respect to various quantities [18, 19]. In this paper, 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. 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 ideally has perfect range resolution, 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.

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. 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 with the corresponding component of velocity along the LOS.

The received intensity from a surface illuminated by a laser radar is dependent on the geometry and reflectance properties of the surface. The reflectance properties are commonly characterized by a function known as the bidirectional reflectance distribution function (BRDF) [20]. For the case of a monostatic radar and a surface with isotropic reflectance properties, the BRDF is given 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 received intensity is proportional to

$$\sigma = 2\pi \iint_S \rho(\psi) \cos^2 \psi \, dA, \tag{18}$$

where the integration is performed over the visible (illuminated) part of the surface, denoted by S. The quantity σ is referred to as the laser radar cross section (LRCS) of the target. 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 paper, we investigate some methods to reconstruct a target from a series of range-resolved or Doppler-resolved measurements using the algorithms described in the preceding section.

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.

For range-resolved measurements, we can consider the data as being obtained either with a single sensor revolving around a stationary target, or with the sensor fixed and the target rotating, with known rotation rate, about an axis perpendicular to the plane in which the measurements are taken. For Doppler-resolved measurements, target motion is required to resolve the target, and so in this case we assume that the target is rotating as above with a fixed sensor.

Alternatively, in either case we may think of the data as being obtained simultaneously by a number of sensors distributed about the target. Also, as in previous work using tomographic techniques [7, 21, 17], we make several assumptions. Specifically, as we shall see, knowledge of the relative positions of the sensors is needed to reconstruct targets from both range-resolved and Doppler-resolved measurements. In addition, for Doppler-resolved measurements, we assume that the target is rotating about an axis perpendicular to the plane of aspects, with known rotation rate. Moreover, if the target is translating, the Doppler velocity of the target's center of gravity relative to each sensor must be known. Since each sensor is presumably tracking the target, we assume knowledge of the necessary quantities. In what follows, we view the problem from this multi-sensor perspective.

4.2 Support line measurements from laser radar data

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 $r_{\min}$. Under far-field assumptions, the above indicates that the target lies completely on one side of the plane perpendicular to the LOS at range $r_{\min}$. Moreover, since some part of the target is at range

$r_{\min}$, this plane actually grazes the target. Hence, this plane is precisely a support plane of the target. (Note that the maximum range $r_{\max}$ with nonzero return intensity does not necessarily provide another support line, since parts of the target at ranges greater than $r_{\max}$ may not be visible to the radar.) However, under our restriction that the LOS's all lie in plane, the problem is effectively reduced to a two-dimensional one, as mentioned above. That is, we need only consider the projection of the target in the plane of LOS's. The support plane information contained in the data corresponds to support line information of the target's projection. Hereafter, use of the word 'target' refers to the two-dimensional projection of the actual target.

A Doppler-resolved measurement contains similar information. For a target undergoing simple rotation with known rate ω, the Doppler frequency due to a point on the target is proportional to the distance from the point to the rotation axis in a direction perpendicular to the LOS (also called the cross-range distance of the point). 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. Hence, two support lines of the target are obtained in this case.

To identify the support value(s) associated with the support line(s) provided by a range or Doppler spectrum, a coordinate frame 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. The resulting position and orientation of this coordinate frame is, of course, arbitrary. 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. From above, the i^{th} sensor (at aspect θ_i) gives rise to support values at $\theta_i \pm 90°$. Since target cross-range is proportional to 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.

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

In general, the resulting 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 $D_{\max}$ amid noise in the range or Doppler spectra. The technique used to estimate $r_{\min}$, $D_{\min}$, and $D_{\max}$ from the laser radar data is briefly described in Section 4.3. 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. Errors in knowing the laser radar positions may also cause angular errors (i.e., errors in knowing the aspects). However, in this paper we ignore angular errors and assume throughout that the aspects of the measurements are known perfectly.

4.3 Knot location

Although determining $r_{\min}$ or $D_{\min}$ and $D_{\max}$ is trivial if the 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 [22] for detecting abrupt changes in dynamic systems, and later applied by Mier-Muth and Willsky [6] to spline estimation. To cast our problem in the framework of [6], 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.

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 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. Details concerning the implementation and performance of this algorithm may be found in [22], [6], and [5].

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 [21]. 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, $\psi \approx 0°$,

maximizing $\cos\psi$ in (18). 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, $\psi \approx 90°$ 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.

5. TARGET RECONSTRUCTIONS FROM LASER RADAR DATA

In this section, we apply the knot location technique discussed in Section 4.3 and the convex set estimation algorithms of Section 3 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 simulated, laboratory, and field measurements.

The data for the first two examples are simulated [23] range-resolved and Doppler-resolved measurements of a cone of height 200 cm and radius 25 cm with Lambertian reflectance characteristics. The cone is positioned with the center of its base at the origin of a coordinate frame and oriented such that its axis of symmetry lies 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.

To reconstruct the targets, we first locate the knots by the Kalman filtering technique described in Section 4.3 and convert them to support values. Modelling knot location errors and registration

errors for each aspect by statistically-independent samples from zero mean Gaussian distributions with variances σ_{kl}^2 and σ_{reg}^2, the effective measurement error is Gaussian, with variance $\sigma_{\text{eff}}^2 = \sigma_{\text{kl}}^2 + \sigma_{\text{reg}}^2$ for range-resolved data. However, for Doppler-resolved data at an even number of uniformly-spaced aspects, (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 modelled by drawing samples from a Gaussian distribution with variance $\sigma_{\text{kl}}^2/2$ for each aspect. The registration error may be obtained by drawing samples from a Gaussian distribution with variance $\sigma_{\text{reg}}^2/2$ for aspects $\theta_1, \theta_2, \ldots, \theta_{M/2}$, and using the negatives of these samples for the aspects $\theta_{M/2+1}, \theta_{M/2+2}, \ldots, \theta_M$. 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 Figure 7a for range-resolved data, with noise level $\sigma_{\text{eff}} = 0.50$ m. The reconstructions produced by **NUA**, **BNGON**, and **BNGONROT** from this set of noisy support line measurements are shown in Figures 7b–d. The display conventions of this figure will be used throughout this section. The reconstructions exhibit behavior similar to that seen for the standard triangle reconstructions of Sections 3.1–3.3. In particular, the prior knowledge of relative reconstruction angles allows **BNGONROT** to outperform **NUA** dramatically, but does not cause it to underperform **BNGON** significantly, which uses absolute angle information. Also, the quality of the reconstructions is rather impressive in light of the fact that the noise level is so high, having a standard deviation equal to the full width of the target. The corresponding results for the Doppler-resolved measurements arising from knot location error ($\sigma_{\text{kl}} = 0.25$ m) and registration error ($\sigma_{\text{reg}} = 0.25$ m) are shown in Figure 8.

The next example is one of reconstructing a triconic target of height 203 cm and base radius 39.5 cm (shown outlined in Figure 9) given laboratory range-resolved measurements. The labora-

tory 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 estimators are shown in Figure 9 for the uncorrupted laboratory data and in Figure 10 for the laboratory data corrupted with measurement noise ($\sigma_{\text{eff}} = 0.25$ m). Note that since the target is not convex, support values using **BNGON** and **BNGONROT** are reconstructed at five angles corresponding to the sides of the convex hull of the target.

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 11a), 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 10.6μm CO_2 narrowband laser radar on a 5.4 km ground range, and had 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 11 for the uncorrupted field data and in Figure 12 for the field data corrupted with measurement noise ($\sigma_{\text{kl}} = 0.10$ and $\sigma_{\text{reg}} = 0.10$). Again, since the target is not convex, support values using **BNGON** and **BNGONROT** are reconstructed at eight angles corresponding to the sides of the convex hull of the target.

5.2 Comparison with and Improvements to Tomographic Imaging Methods

In previous work, standard methods of tomographic image reconstruction [24] were applied to range-resolved and Doppler-resolved laser radar data [7, 17]. 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 recon-

structions, 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 from data with registration errors.

All of the tomographic reconstructions in this section were obtained using the standard method of filtered backprojection. (See [24] for methods of transmission tomography, and [7, 17] and references contained therein for the application of these methods to laser radar reflective data.) Parts (a) of Figures 13–16 show filtered backprojection reconstructions from the four data sets (free of registration errors) used in Section 5.1. 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 typically the high intensities are near the center of a Doppler spectrum and 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 7–12), the tomographic reconstructions contain intensity information within the outline of 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. This is especially true of reconstructions arising from Doppler data, where for reasons suggested in Section 4.3 and described and demonstrated in [21], reconstructed edges are not highlighted but are instead overwhelmed by the high intensities that are reconstructed in the interior of the target. Even when using thresholding as mentioned above, edges in the reconstructions from Doppler data are not sufficiently highlighted.

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 13–16 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 $\sigma_{\text{reg}} = 0.50, 0.25, 0.25,$ and 0.10 m (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 from the reconstructions shown in parts (b)–(d) of Figures 7,8,10, and 12, 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 (piecewise-sinusoidal support function of the reconstructed polygon 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 filtered backprojection.

Parts (c) and (d) of Figures 13 and 14, parts (c)-(e) of Figure 15, and part (c) of Figure 16 show the tomographic reconstructions that result using this process. The tomographic reconstructions resulting from preprocessing by each of the three convex set algorithms are not included in some of the figures. In the cases that the reconstruction corresponding to **BNGON** was omitted, it could not be distinguished from that corresponding to **BNGONROT**. In the case where reconstructions for both **BNGON** and **BNGONROT** were omitted, they were indistinguishable from that corresponding to **NUA**. Quite clearly, on comparing the various images within each of the Figures 13-16, the improvement obtained using the registration correction method is dramatic.

6. SUMMARY AND SUGGESTIONS FOR FURTHER WORK

In this paper, we first developed and studied several techniques for estimating convex sets from a set of noisy support line measurements. The basic approach to these methods involves reconstructing a polygon close to the measurements while enforcing a consistency condition and utilizing prior information, if available. The algorithms are computationally feasible with two of the estimators resulting in quadratic programming problems and the third having a quadratic programming core. The performance of the algorithms was assessed with respect to several parameters. As expected, if accurate prior information is available then **BNGON** and **BNGONROT** substantially outperform **NUA**. A useful observation is that the performance of **BNGONROT** is comparable to that of **BNGON** while elim-

inating the need for prior orientation information. The ability of BNGONROT to provide orientation estimates may be useful in certain applications.

We also introduced the use of these methods for reconstructing targets from resolved laser radar data. The reconstruction process consists of first extracting support line measurements from the data, and then producing a shape estimate using the convex set estimation techniques. The application of these techniques to laser radar data obtained through simulated, laboratory, and field measurements was demonstrated. The reconstructions obtained were compared to those produced by tomographic imaging methods, resulting in the following observations. First, shape estimates are explicitly provided by our algorithms, as opposed to tomographic images, which can provide target shape information only after using image processing techniques with their attendant difficulties. Second, we investigated the effects of registration error on both methods and found that the tomographic methods experience substantial degradation, unlike the present methods which are rather robust. These observations motivated us to exploit the tendency of our algorithms to correct unregistered data, in an effort to improve the quality of tomographic images.

A variety of extensions to our reconstruction algorithms might be made. One such extension may consist of developing more general formulations of the best N-gon algorithm, so that the use of less stringent prior shape information could be made. For example, one might consider a formulation in which only the number but not the values of the reconstruction angles are specified. A more general formulation might leave both the number and the values of the reconstruction angles unspecified, but would penalize larger numbers of reconstruction angles. Also, it may be interesting to develop algorithms that provide smooth shape estimates of objects, as opposed to polygonal estimates. Incorporating the effects of noise in the measurement angles may be useful in some applications. Another useful generalization would be to extend the algorithms to three dimensions.

There are several extensions regarding the application to laser radar data that may be interesting. Upon relaxing the restriction that aspects lie in a plane (i.e., allowing general aspects in three dimensions), shape estimates of a target (rather than those of its projection onto a plane) would be obtained. In addition, the application of our methods to two-dimensional laser radar data resolved in both range and Doppler may provide three-dimensional target shape estimates. Extending the viewing geometry to allow bistatic observations might prove useful as well.

APPENDICES

Appendix A: Perimeter Bias in the Estimator

The convex set reconstruction algorithms described in this paper produce estimates that appear on average larger than the true object. Here we point out the origin of this bias for **NUA** and quantify it through a Monte Carlo approach when the true objects are ellipses.

If the noisy support vector $y = h + n$ lies in the support cone C, then the **NUA** estimate is simply given by $\hat{h} = y$. On the other hand, if $y \notin C$, then the estimation procedure consists of projecting y onto C, to yield the estimate $\hat{h} = P_1(y)$ (see Figure 3b).

To discuss bias in the estimator, some quantitative measure of the size of a reconstructed object is necessary. One such measure is the magnitude of the projection of $\hat{h}$ onto the vector $e = [1\ 1\ \cdots\ 1]^T$, 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.

Since it is not possible to illustrate the M-dimensional support cone for meaningful values of M (i.e. $M \geq 5$), we use a two-dimensional illustration to provide the basic idea. Note that the vector

e defines the center axis of the support cone. Thus, consider the geometry of Figure 17. Here, the cone boundaries are given by the two bold lines, and e is its center axis. The measurement is given by $y_0 = h + n_0$, where n_0 is zero-mean noise and where h is some point in the cone. In the case illustrated, y_0 is not in the cone, and the **NUA** algorithm yields the estimate $\hat{h}_0 = P_1(y_0)$. Note that $P_2(P_1(y_0)) > P_2(y_0)$, i.e., the estimate in this case has a larger projection onto e. Thus, if we let $A = P_2(h)$, $p_n(n_0)$ denote the Gaussian noise probability density function, and μ denote the expected value of $P_2(P_1(y_0))$, then

$$\begin{aligned}
\mu &= \int_{h+n_0 \ valid} P_2(h+n_0)p_n(n_0)\,dn_0 + \int_{h+n_0 \ invalid} P_2(P_1(h+n_0))p_n(n_0)\,dn_0 \\
&> \int_{h+n_0 \ valid} P_2(h+n_0)p_n(n_0)\,dn_0 + \int_{h+n_0 \ invalid} P_2(h+n_0)p_n(n_0)\,dn_0 \\
&= \int_{all \ h+n_0} P_2(h+n_0)p_n(n_0)\,dn_0 \\
&= A. \qquad (19)
\end{aligned}$$

where the last equality follows from n_0 being zero mean. Also, as the variance of the noise increases, a larger fraction of the noisy support vectors are invalid ones. Consequently, the bias is an increasing function of noise variance. We refer the reader to [13] for a quantitative evaluation of the bias based on Monte Carlo simulations.

Appendix B: Nonuniqueness of Solutions in Best N-gon Problems

Under certain conditions, given a set of measurements and reconstruction angles, the estimators **BNGON** and **BNGONROT** yield nonunique solutions. That is, an entire family of solutions achieves minimum cost. For example, nonuniqueness occurs for the set of measurements $\{1,1,1,1\}$ at the angles $\{0°, 90°, 180°, 270°\}$, with reconstruction angles at $\{45°, 135°, 225°, 315°\}$. In this case, any rectangle with sides at the reconstruction angles inscribed within the measured support lines

corresponds to a solution achieving zero cost. Nonuniqueness for the case in which minimum cost is nonzero can also occur [13].

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. 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 (9). 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. Note first that if $N > M$, i.e. if we are attempting to obtain a reconstruction with more degrees of freedom than we have measurements, then A has a nullspace of dimension at least $N - M$.

The same degree-of-freedom argument applies the case where $N \leq M$ but where the reconstruction angles sufficiently outnumber the measurement angles near some θ_i. 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 (5).

In the case in which we have equal numbers of measurement and reconstruction angles ($M = N$), 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 $\theta_i \leq \phi_{i+1} < \theta_{i+1} \quad \forall i$, where we define $\theta_{N+1} = \theta_1$ and $\phi_{N+1} = \phi_1$ for convenience). This restriction is motivated by the preceding discussion, 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

$$A = \begin{pmatrix} l_1 & r_1 & 0 & \cdots & 0 \\ 0 & l_2 & r_2 & 0 & \vdots \\ \vdots & 0 & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & l_{N-1} & r_{N-1} \\ r_N & 0 & \cdots & 0 & l_N \end{pmatrix}, \tag{20}$$

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

$$\det A = \prod_{i=1}^{N} \frac{\sin(\phi_{i+1} - \theta_i)}{\sin(\phi_{i+1} - \phi_i)} + (-1)^{N+1} \prod_{i=1}^{N} \frac{\sin(\theta_i - \phi_i)}{\sin(\phi_{i+1} - \phi_i)} \tag{21}$$

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, (21) reduces to

$$\det A = [1 + (-1)^{N+1}] \frac{1}{(2\cos\Delta)^N}, \tag{22}$$

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 $\{\phi_1, \phi_2, \ldots, \phi_N\}$ 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 (21) 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.

Appendix C: Derivation of Lower Bound on Error Variance of Orientation Estimate

In this appendix, we analytically determine the Cramer-Rao lower bound [16] 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 first note that errors in estimating α can be classified into local errors and global errors [16]. For a local error the global minimum at $\hat{\alpha}$ lies in the same cost function basin as does the true value of α, i.e. both share the same pair of adjacent maxima in the cost function. A global error refers to the case in which the global minimum lies in a basin different from that belonging to the true value of α.

The present 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 about the true value of α. That is, 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, in nonlinear problems noise values are more likely to lower the cost $J_{h_\phi}(\alpha)$ for some α_0 far from α_{true} so 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 $\mathbf{x} = [\alpha \; h_\phi(\phi_1) \; h_\phi(\phi_2) \cdots h_\phi(\phi_N)]^T$ the parameter vector to be estimated, so that the observation, or measurement vector is given by $y = g(\mathbf{x}) + n$, where the i^{th} component of $g(\mathbf{x})$ is the scalar quantity in (8). With this notation and if $\mathbf{x}$ were completely unconstrained, the Cramer-Rao bound [16] on the error covariance of any unbiased estimate of $\mathbf{x}$ would be given by

$$\begin{aligned} \text{cov}(\mathbf{x} - \hat{\mathbf{x}} | \alpha = \alpha_{\text{true}}) &\geq I_y^{-1}(\mathbf{x})|_{\alpha=\alpha_{\text{true}}} \\ &= \left\{ E\left[\left(\frac{\partial \ln p_{y|\mathbf{x}}(y_0|\mathbf{x})}{\partial \mathbf{x}} \right)^2 \right] \right\}^{-1} \Bigg|_{\alpha=\alpha_{\text{true}}}, \end{aligned} \tag{23}$$

where $I_y(\mathbf{x})$ is referred to as the *Fisher information* in y about $\mathbf{x}$. Since the measurement noise vector n is a Gaussian random vector with zero mean and covariance matrix $\sigma^2 I$, (23) may be rewritten as

$$\text{cov}(\mathbf{x} - \hat{\mathbf{x}} | \alpha = \alpha_{\text{true}}) \geq \sigma^2 \left[\frac{\partial g^T(\mathbf{x})}{\partial \mathbf{x}} \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right]^{-1} \Bigg|_{\alpha=\alpha_{\text{true}}}, \tag{24}$$

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 (5) with the θ_i replaced by ϕ_i, so that the CRB of (24) 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 (24) by the scalar CRB on α, obtained by treating α alone 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 $n \sim N(0, \sigma^2 I)$ and the i^{th} component of $f(\alpha)$ is the true support value induced at θ_i when the object has an offset of α, as given by (11). The CRB on the error variance of any unbiased estimate of α is given by

$$\begin{aligned} \mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}}) &\geq I_y^{-1}(\alpha_{\mathrm{true}}) \\ &= \left. \frac{\sigma^2}{\frac{\partial f^T(\alpha)}{\partial \alpha}\frac{\partial f(\alpha)}{\partial \alpha}} \right|_{\alpha=\alpha_{\mathrm{true}}}, \end{aligned} \tag{25}$$

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 (25) yields a bound on the standard deviation of the estimation error (in degrees). As we have mentioned, the CRB 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.

Specifically, as discussed in [16], in many nonlinear estimation problems, such as those arising in communication problems, under high noise levels the anomalous minima that result occur with equal likelihood anywhere except for values of α within the basin corresponding 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 for the standard triangle, the locations of anomalous global minima were observed to cluster near 270° (see [13]). 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 $J_{h_\phi}(\alpha)$ 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, for this particular example, we take the estimation error variance when a global error occurs as being the constant value $(180°)^2$.

We have characterized the estimation error variance $\mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}})$ for the low noise and high noise regimes. For intermediate noise levels, the error variance can be approximated, as in [16], 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. For the standard triangle, we may approximate this probability by formulating a binary hypothesis testing problem in which the two hypotheses are $\alpha = 90°$ (zero error) and $\alpha = 270°$ (180° global error). Using standard results on Gaussian binary hypothesis testing [16], we find that the resulting probability of anomaly is given by

$$\Pr(A) = Q\left(\frac{\sqrt{H}}{2\sigma}\right) \tag{26}$$

$$\leq \frac{1}{2}\exp\left(\frac{-H}{8\sigma^2}\right), \tag{27}$$

Here $Q(x) = \int_{-\infty}^{x} p_n(n_0)dn_0$ and H is the norm-square of the difference in the means of the support vector measurements under the two hypotheses. That is,

$$H = [f(270°) - f(90°)]^T[f(270°) - f(90°)] \tag{28}$$

Using (25) and (27) and the fact that the estimation error variance is approximately $(180°)^2$ under high noise conditions, we have that $\mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}})$ for all noise levels is approximately given by

$$
\begin{aligned}
\mathbf{var}(\alpha - \hat{\alpha}|\alpha_{\mathbf{true}}) &\approx [1 - \Pr(A)](\mathrm{CRB}) + \Pr(A)(180^\circ)^2 \\
&\approx \left[1 - \frac{1}{2}\exp\left(\frac{-H}{8\sigma^2}\right)\right] \frac{\sigma^2}{\frac{\partial f^T(\alpha)}{\partial \alpha}\frac{\partial f(\alpha)}{\partial \alpha}}\Bigg|_{\alpha=\alpha_{\mathbf{true}}} \\
&\quad + \frac{1}{2}\exp\left(\frac{-H}{8\sigma^2}\right)(180^\circ)^2. \qquad (29)
\end{aligned}
$$

References

[1] J.L. Prince, *Geometric Model-Based Estimation from Projections.* Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science (1988).

[2] J.L. Prince and A.S. Willsky, "Estimating Convex Sets from Noisy Support Line Measurements," *IEEE Trans. PAMI*, **12**, pp. 377-389, April (1990).

[3] J.P. Greschak, *Reconstructing Convex Sets.* Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science (1985).

[4] 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).

[5] A.M. Mier-Muth, *Adaptive Knot Location for Spline Approximation.* S.M. Thesis, Massachusetts Institute of Technology, Dept. of Elec. Eng. (1976).

[6] 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).

[7] F.K. Knight, *et. al.*. "Tomographic Techniques Applied to Laser Radar Reflective Measurements," *The Lincoln Laboratory Journal,* **2**, pp.143-160 (1989).

[8] 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).

[9] S.R. Lay, *Convex Sets and their Applications.* Wiley-Interscience, New York (1982).

[10] D.E. McClure and R.A. Vitale, "Polygonal Approximation of Plane Convex Bodies." *Journal of Mathematical Analysis and Applications*, **51**, pp. 326-358, (1975).

[11] L.A. Santalo, *Integral Geometry and Geometric Probability.* Volume 1 of *Encyclopedia of Mathematics and its Applications,* Addison-Wesley, Reading, MA (1976).

[12] R.A. Vitale, *A representation theorem for compact convex sets in the plane.* Division of Applied Mathematics Report, Brown University, (1974).

[13] A.S. Lele, *Convex Set Reconstruction from Support Line Measurements and its Application to Laser Radar Data.* S.M. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science (1990).

[14] A.H. Land and S. Powell, *Fortran Codes for Mathematical Programming.* Wiley-Interscience, London, (1973).

[15] D.G. Luenberger, *Linear and Nonlinear Programming.* Addison-Wesley, Reading, MA, second edition (1984).

[16] H.L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory,* Wiley and Sons, Inc., New York, (1968).

[17] J.K. Parker, *et al.*, "Reflective Tomography: Images from Range-resolved Laser Radar Measurements," *Applied Optics*, **27**, 2642 (1988).

[18] C.G. Bachman, *Laser Radar Systems and Techniques*, Artech House, Inc. (1979).

[19] A.W. Rihaczek, *Principles of High-resolution Radar*, Mark Resources, Inc. (1977).

[20] 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, October (1977).

[21] 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, August (1989).

[22] 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, February (1976).

[23] 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).

[24] R.A. Brooks and G. DiChiro, "Principles of Computer Assisted Tomography in Radiographic and Radioisotopic Imaging." *Phys. Med. Biol.* **21**, pp. 689-732 (1976).

Figure Captions

Figure 1: Block diagram of reconstruction procedure.

Figure 2: Support line of a set.

Figure 3: (a) Support line consistency is achieved only if L_i does not lie in the invalid region, shown hatched and to the right of the dashed line at angle θ_i. (b) Geometry illustrating the projection of the invalid support vector y onto the support cone $\mathcal{C}$.

Figure 4: Examples of (a,b) **NUA**, (c,d) **BNGON**, and (e,f) **BNGONROT** for the standard triangle in object space and support function space, for $M = 24$ uniformly-spaced noisy ($\sigma = 0.25$) measurements. In (b,d, and f), the support function of the standard triangle is plotted with a thick curve, the noisy support values $\{y_i\}$ are marked by $\times$'s, and the support function $\hat{h}(\theta)$ of the estimated object is plotted with a thin curve.

Figure 5: Performance of estimators using the standard triangle: (a) E vs. σ for **NUA**, **BNGON**, and **BNGONROT**; (b) E vs. M for **NUA** and **BNGON**; (c) E vs. α for **BNGON** with baseline **NUA** performance; and (d) E vs. N for **BNGON** with baseline **NUA** performance.

Figure 6: Overall standard deviation $\sqrt{\mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}})}$ of the orientation estimate versus σ, obtained analytically and through Monte Carlo simulation. Included for reference is the Cramer-Rao bound, shown dashed.

Figure 7: (a) A conical object and associated support line measurements from simulated range-resolved data with $\sigma_{\mathrm{eff}} = 0.50$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 8: (a) A conical object and associated support line measurements from simulated Doppler-resolved data with $\sigma_{\mathrm{kl}} = \sigma_{\mathrm{reg}} = 0.25$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 9: (a) A triconic object and associated support line measurements from laboratory

range-resolved data with $\sigma_{\text{eff}} = 0$ m and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 10: (a) A triconic object and associated support line measurements from laboratory range-resolved data with $\sigma_{\text{eff}} = 0.25$ m and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 11: (a) Thor Delta rocket body and associated support line measurements from field Doppler-resolved data with $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0$ m and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 12: (a) Thor Delta rocket body and associated support line measurements from field Doppler-resolved data with $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.10$ m and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

Figure 13: Tomographic reconstructions from range-resolved measurements of a conical object with $\sigma_{\text{eff}} = 0.50$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by **NUA**; and (d) data with registration errors corrected by **BNGONROT**.

Figure 14: Tomographic reconstructions from Doppler-resolved measurements of a conical object with $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.25$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by **NUA**; and (d) data with registration errors corrected by **BNGONROT**.

Figure 15: Tomographic reconstructions from laboratory range-resolved measurements of a triconic object with $\sigma_{\text{eff}} = 0.25$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by **NUA**; (d) data with registration errors corrected by **BNGON**; and (e) data with registration errors corrected by **BNGONROT**.

Figure 16: Tomographic reconstructions from field Doppler-resolved measurements of a cone

with $\sigma_{kl} = \sigma_{reg} = 0.10$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; and (c) data with registration errors corrected by **NUA**.

Figure 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.

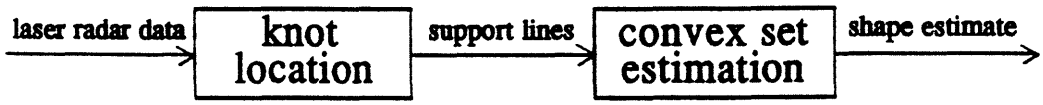

Figure 1: Block diagram of reconstruction procedure.

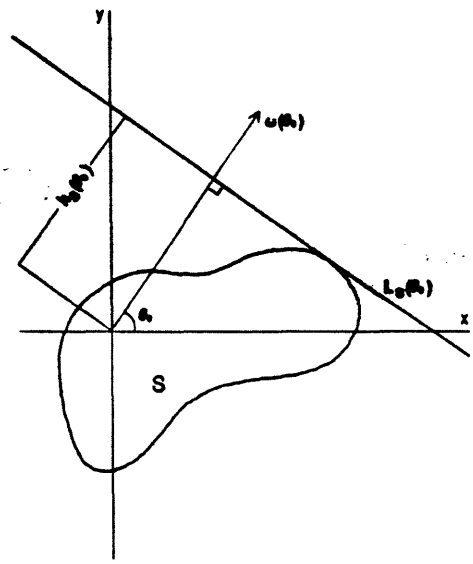

Figure 2: Support line of a set.

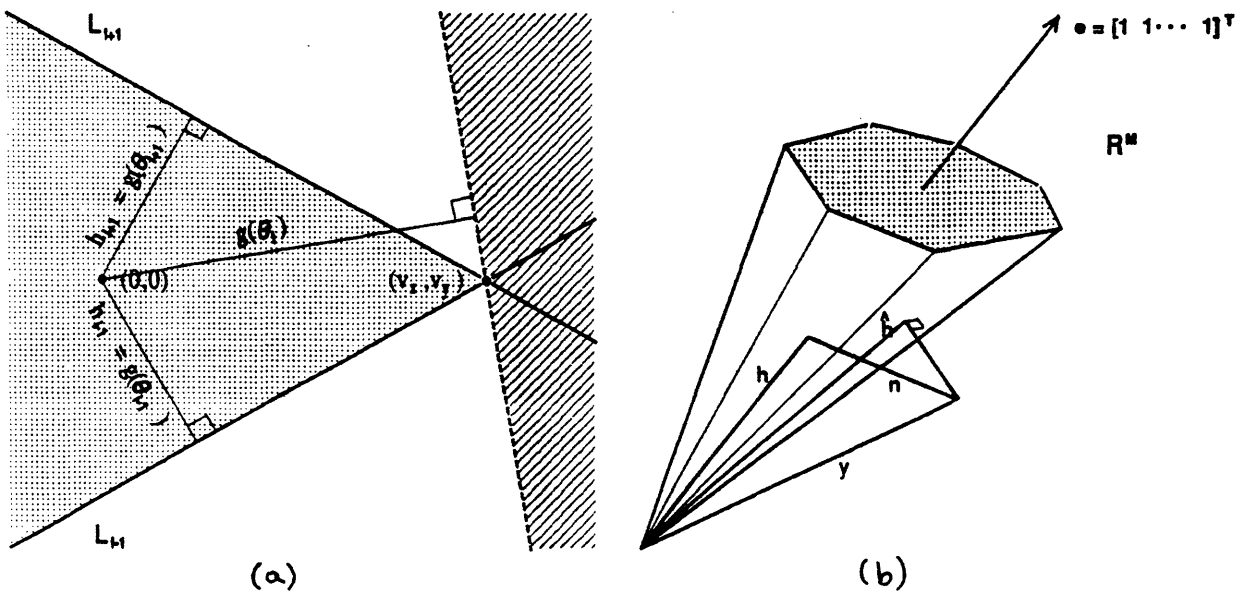

Figure 3: (a) Support line consistency is achieved only if L_i does not lie in the invalid region, shown hatched and to the right of the dashed line at angle θ_i. (b) Geometry illustrating the projection of the invalid support vector y onto the support cone $\mathcal{C}$.

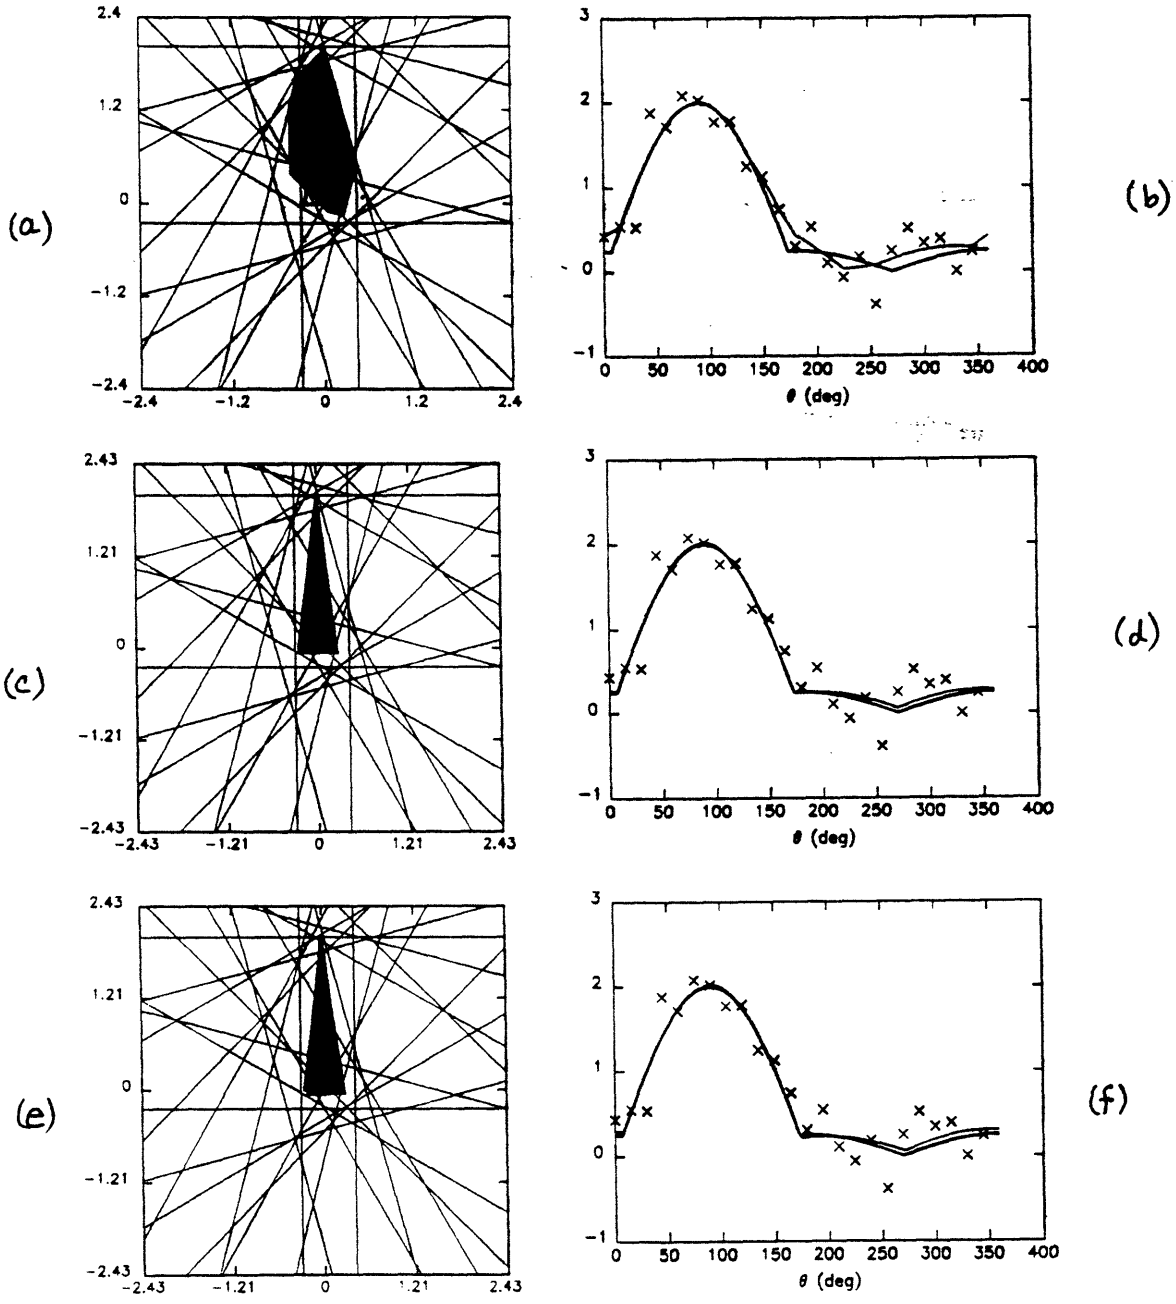

Figure 4: Examples of (a,b) **NUA**, (c,d) **BNGON**, and (e,f) **BNGONROT** for the standard triangle in object space and support function space, for $M = 24$ uniformly-spaced noisy ($\sigma = 0.25$) measurements. In (b,d, and f), the support function of the standard triangle is plotted with a thick curve, the noisy support values $\{y_i\}$ are marked by $\times$'s, and the support function $\hat{h}(\theta)$ of the estimated object is plotted with a thin curve.

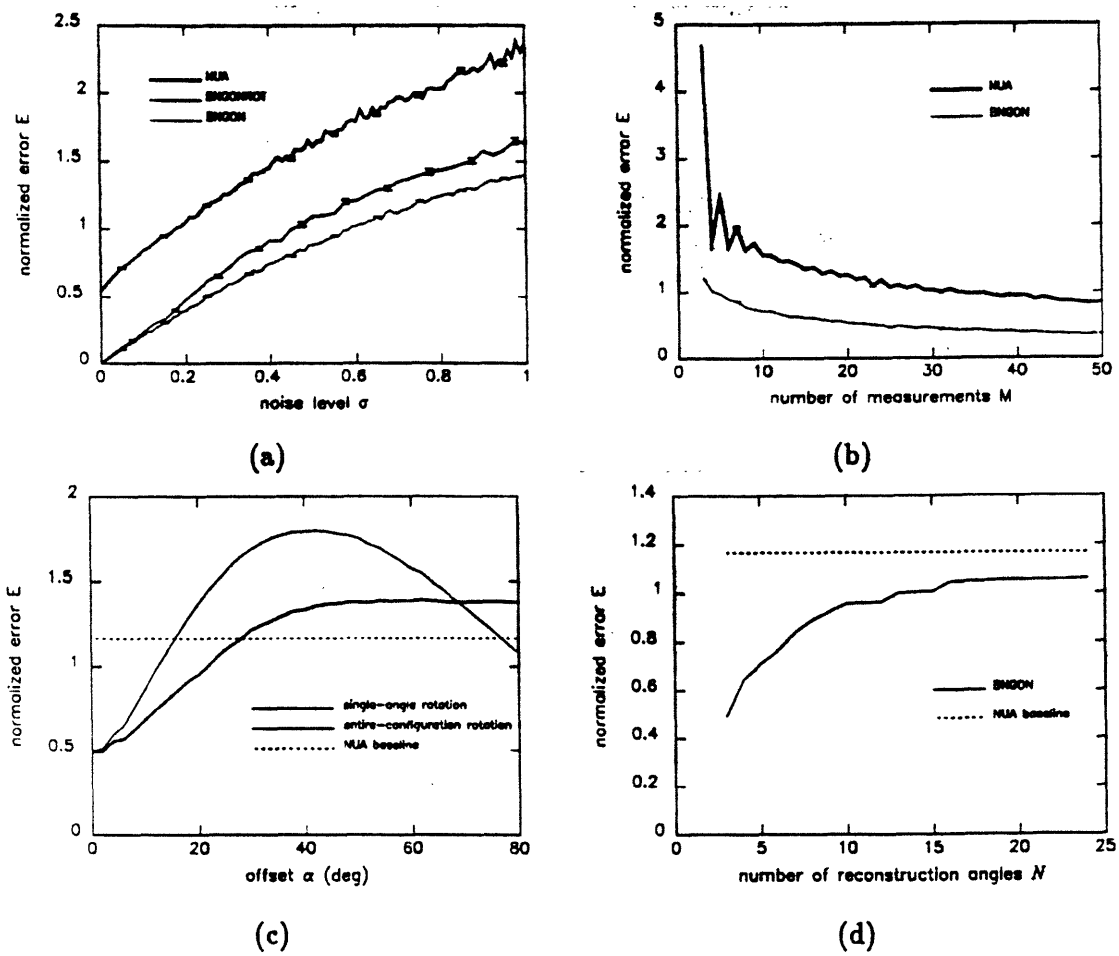

Figure 5: Performance of estimators using the standard triangle: (a) E vs. σ for **NUA**, **BNGON**, and **BNGONROT**; (b) E vs. M for **NUA** and **BNGON**; (c) E vs. α for **BNGON** with baseline **NUA** performance; and (d) E vs. N for **BNGON** with baseline **NUA** performance.

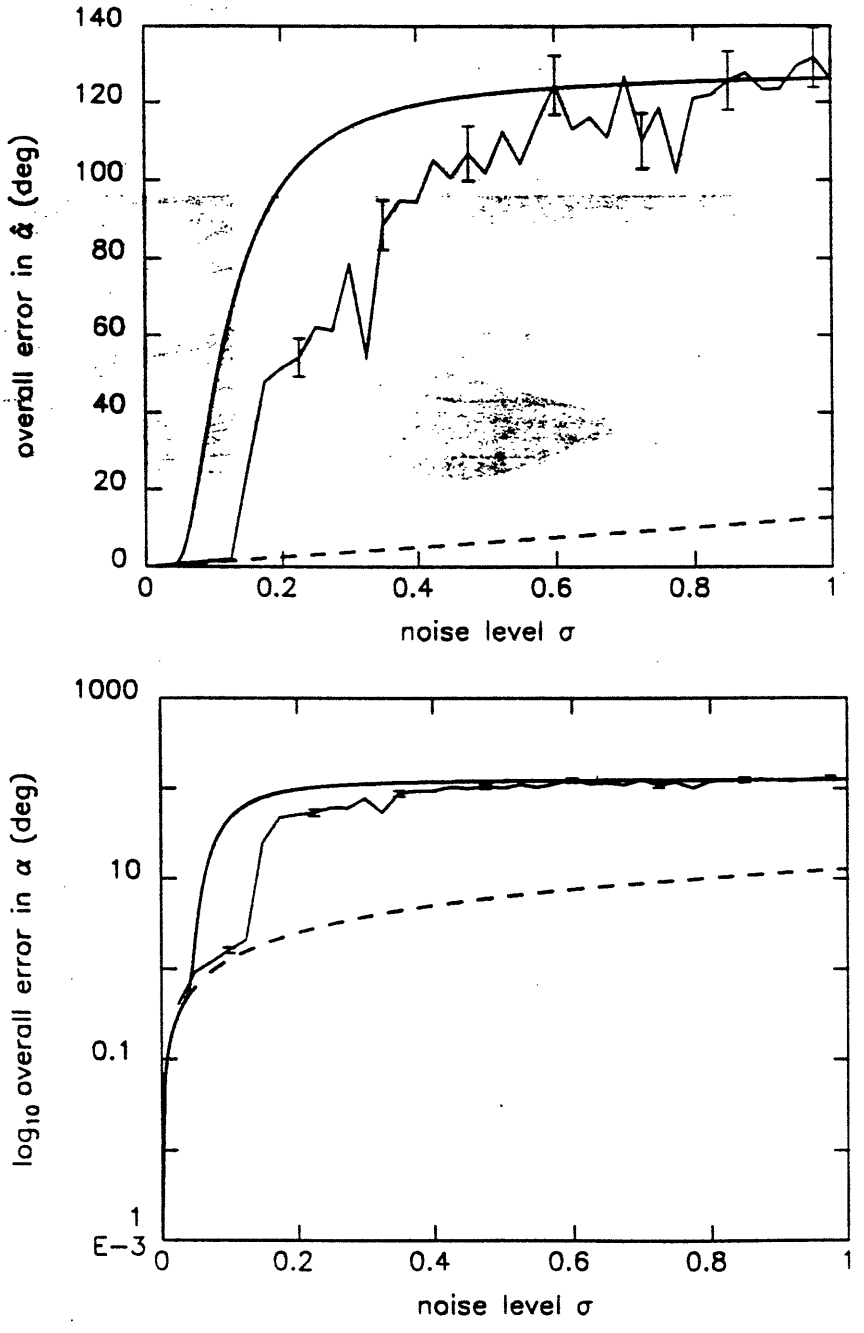

Figure 6: Overall standard deviation $\sqrt{\mathrm{var}(\alpha - \hat{\alpha}|\alpha_{\mathrm{true}})}$ of the orientation estimate versus σ, obtained analytically and through Monte Carlo simulation. Included for reference is the Cramer-Rao bound, shown dashed.

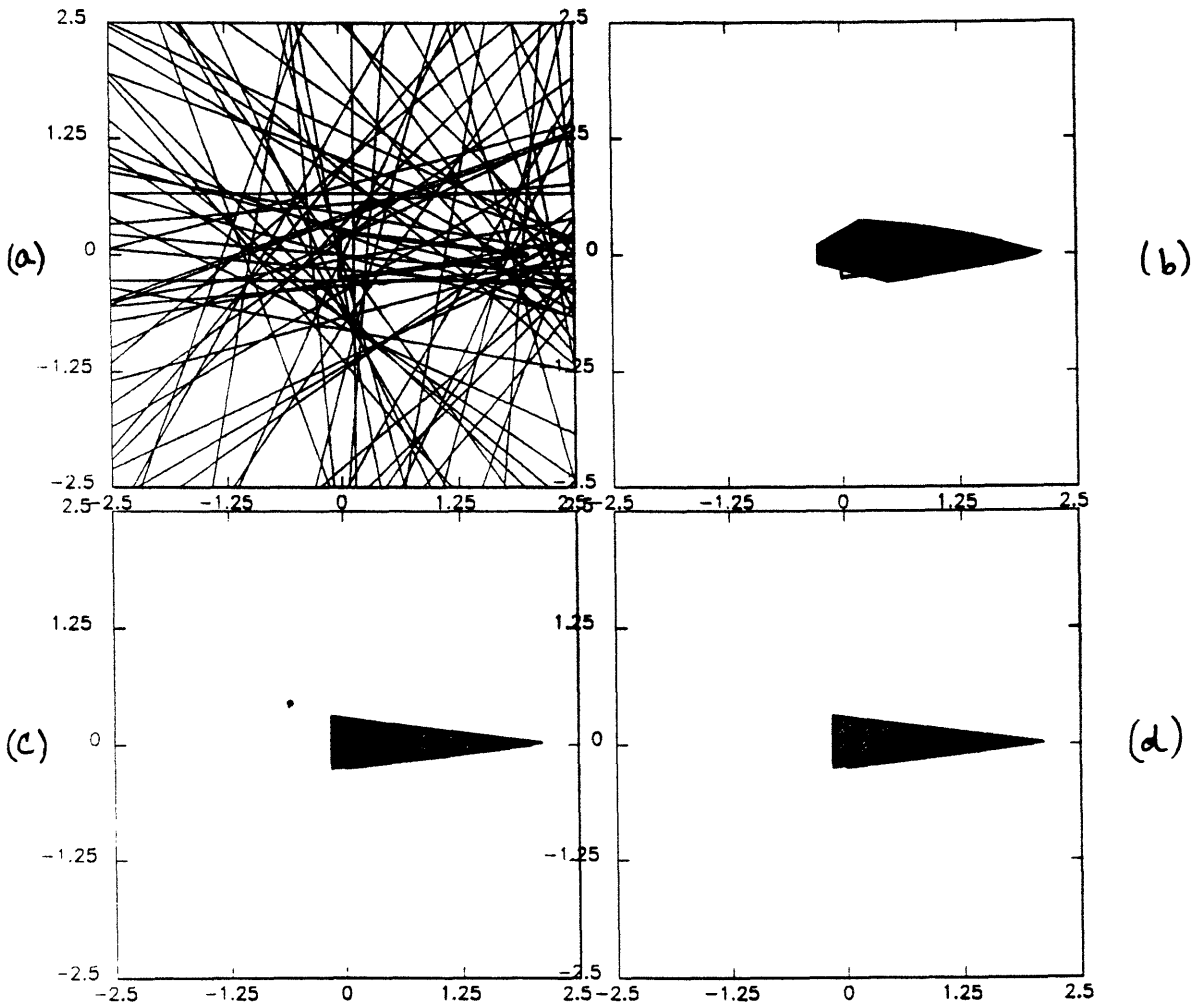

Figure 7: (a) A conical object and associated support line measurements from simulated range-resolved data with $\sigma_{\text{eff}} = 0.50$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

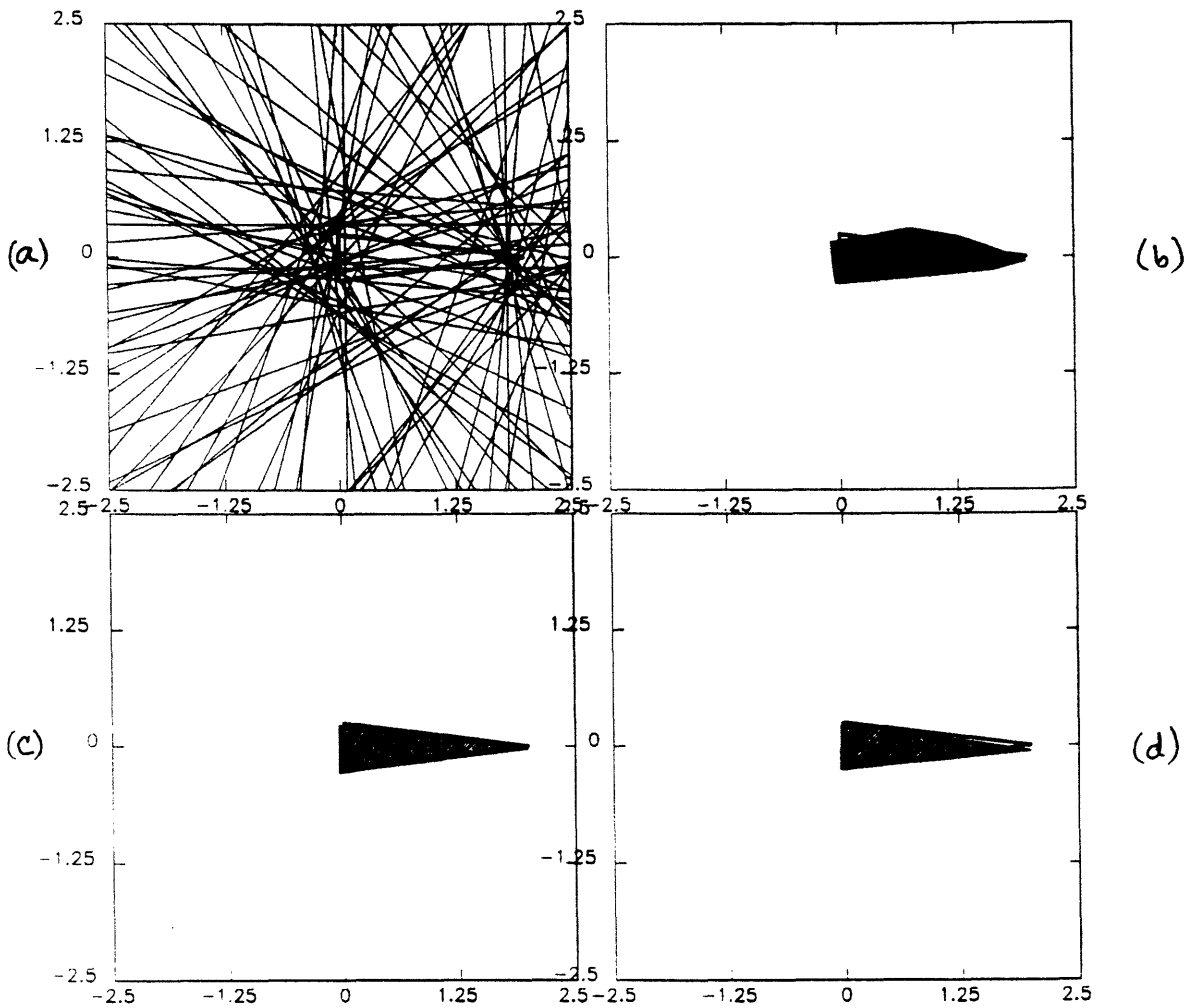

Figure 8: (a) A conical object and associated support line measurements from simulated Doppler-resolved data with $\sigma_{\mathrm{kl}} = \sigma_{\mathrm{reg}} = 0.25$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

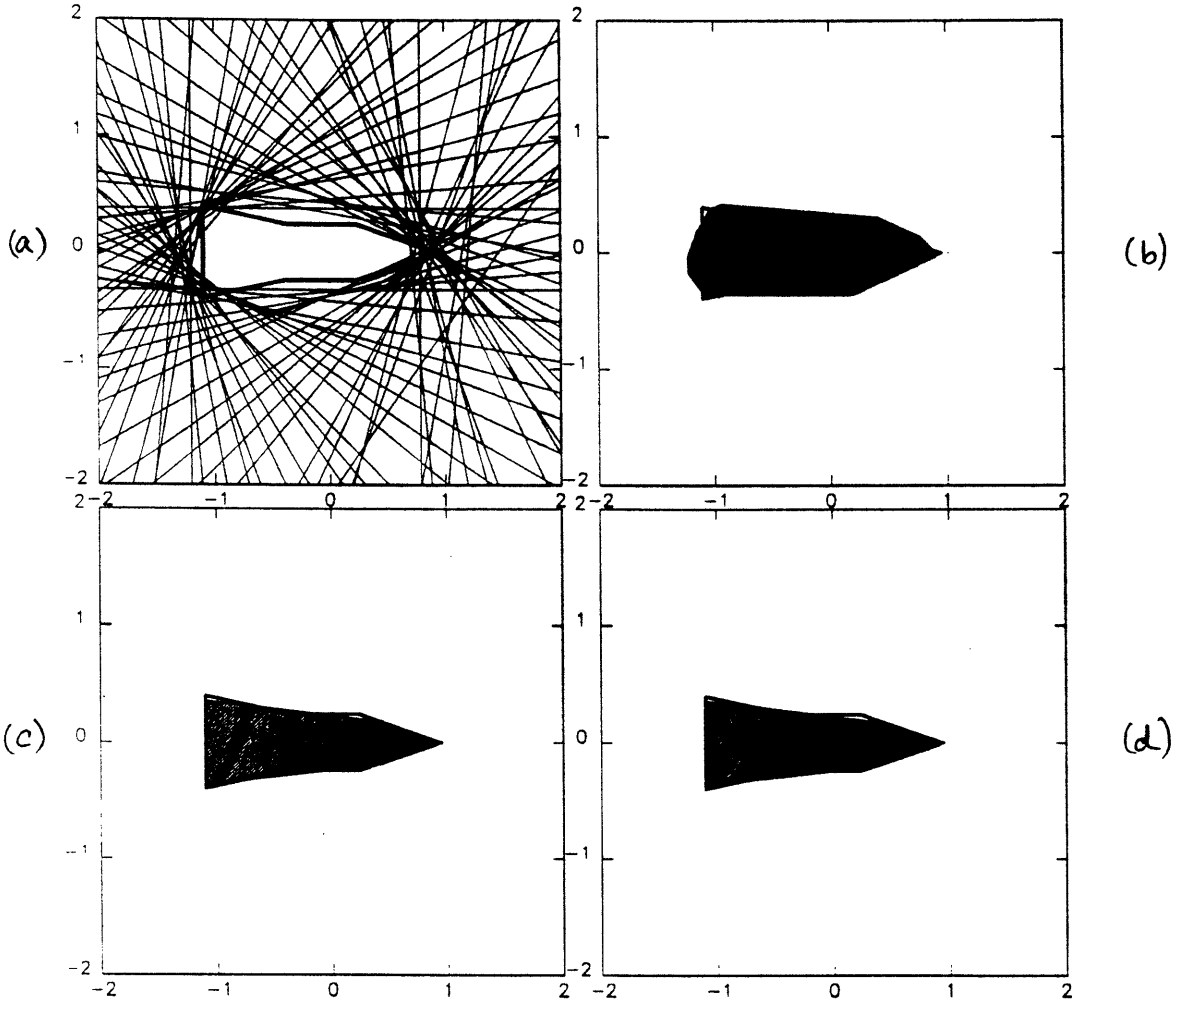

Figure 9: (a) A triconic object and associated support line measurements from laboratory range-resolved data with $\sigma_{\text{eff}} = 0$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

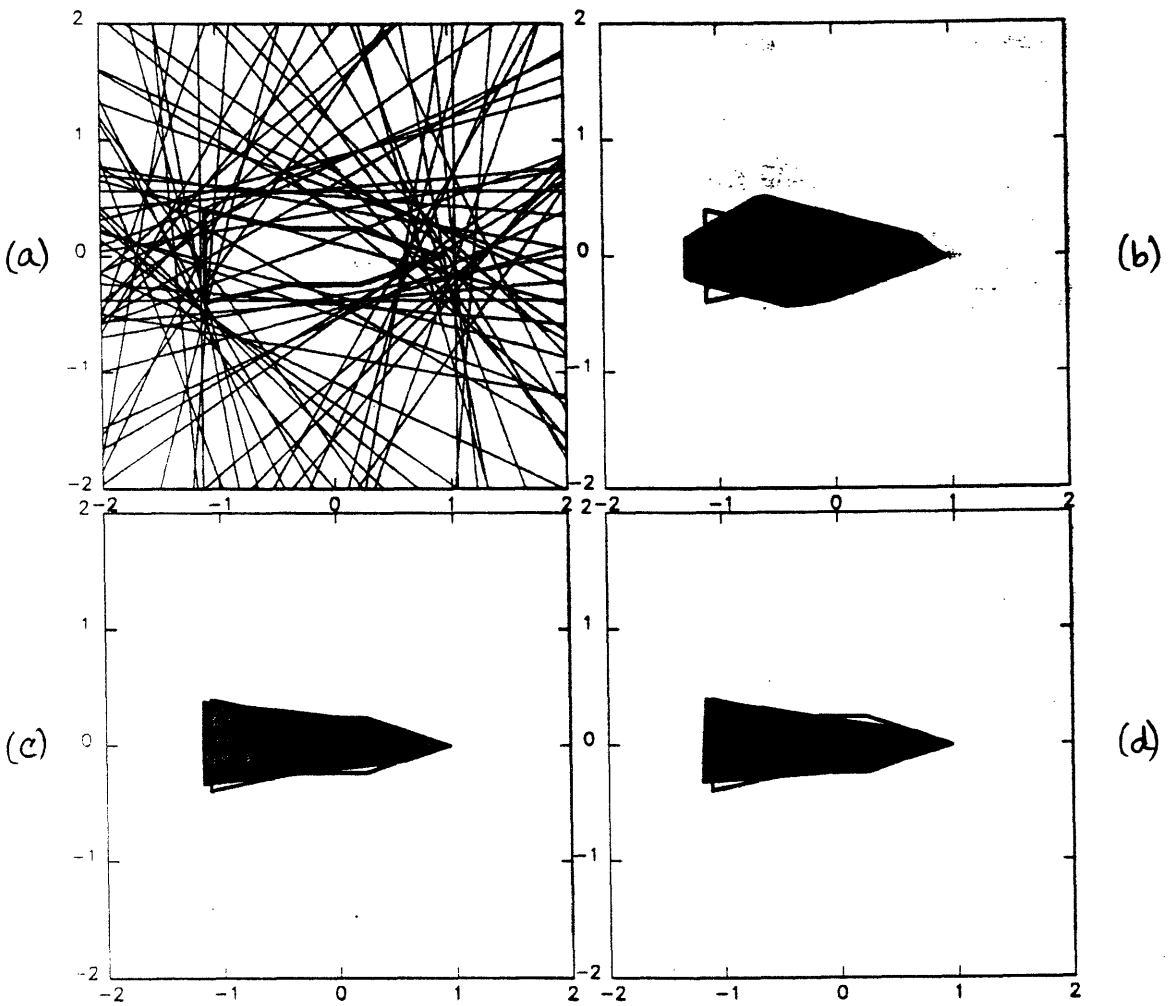

Figure 10: (a) A triconic object and associated support line measurements from laboratory range-resolved data with $\sigma_{\text{eff}} = 0.25$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

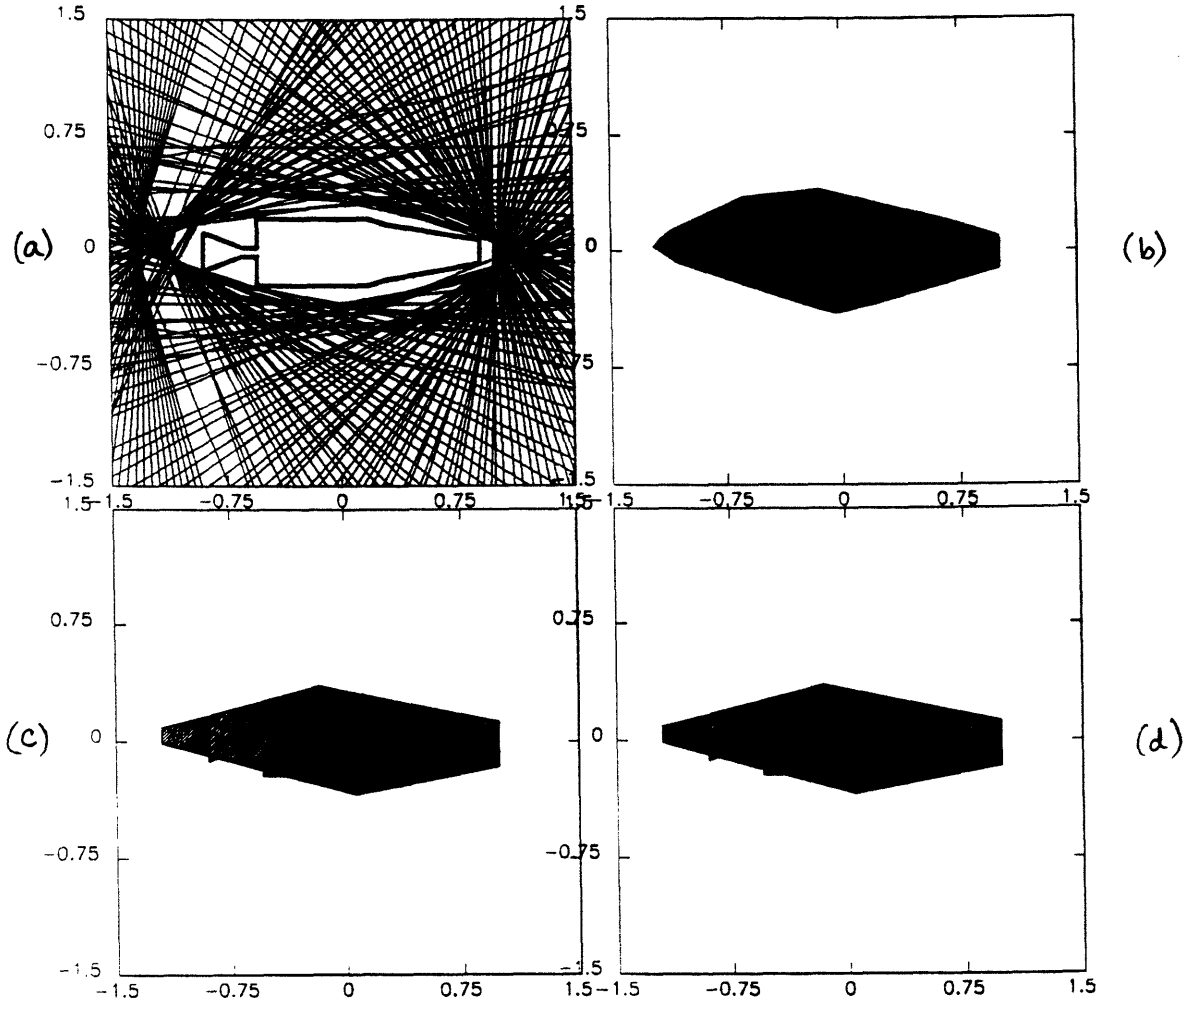

Figure 11: (a) Thor Delta rocket body and associated support line measurements from field Doppler-resolved data with $\sigma_{\mathbf{kl}} = \sigma_{\mathbf{reg}} = 0$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

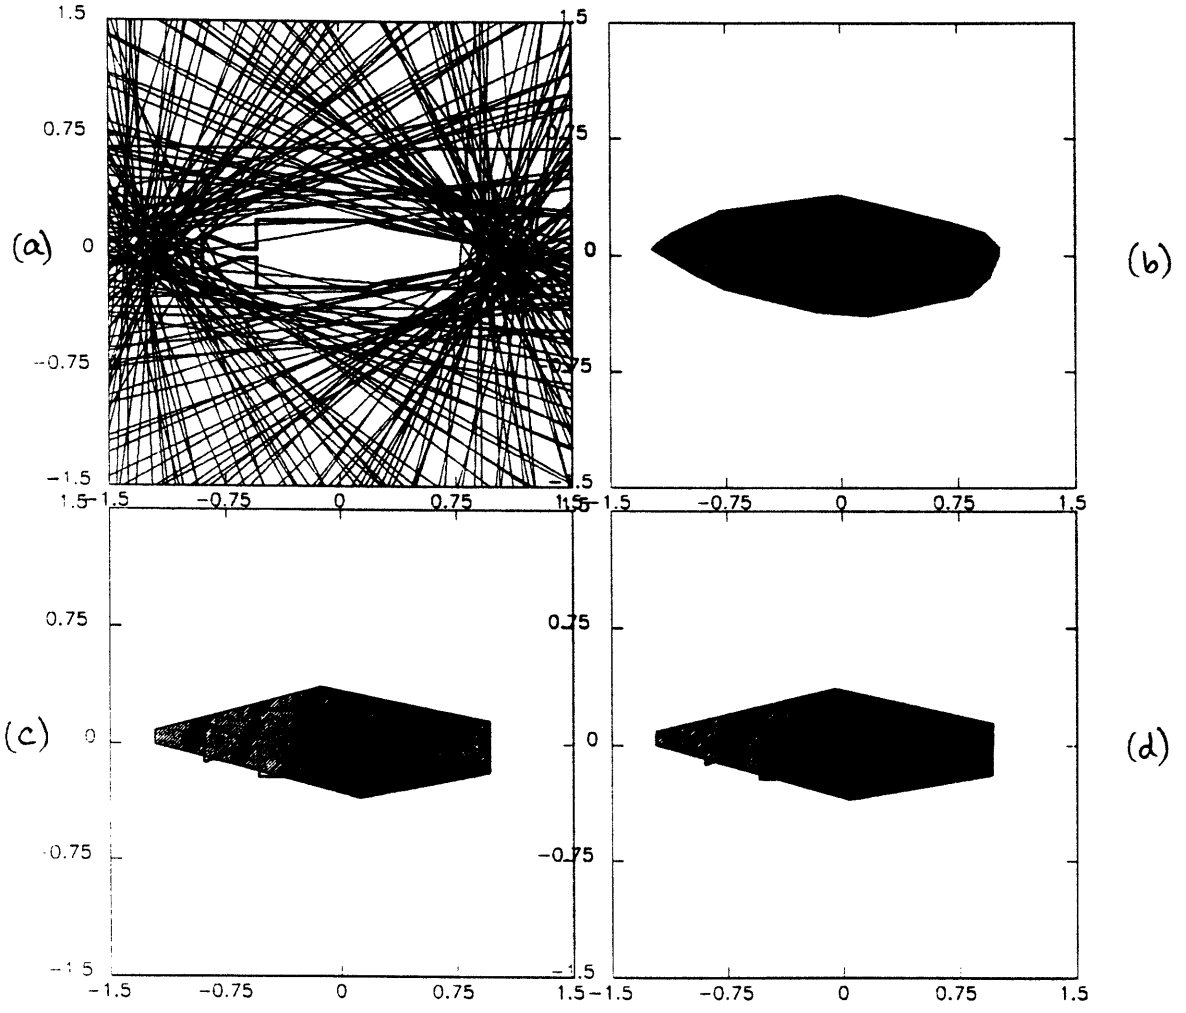

Figure 12: (a) Thor Delta rocket body and associated support line measurements from field Doppler-resolved data with $\sigma_{\mathrm{kl}} = \sigma_{\mathrm{reg}} = 0.10$ m, and reconstructions using (b) **NUA**; (c) **BNGON**; (d) **BNGONROT**.

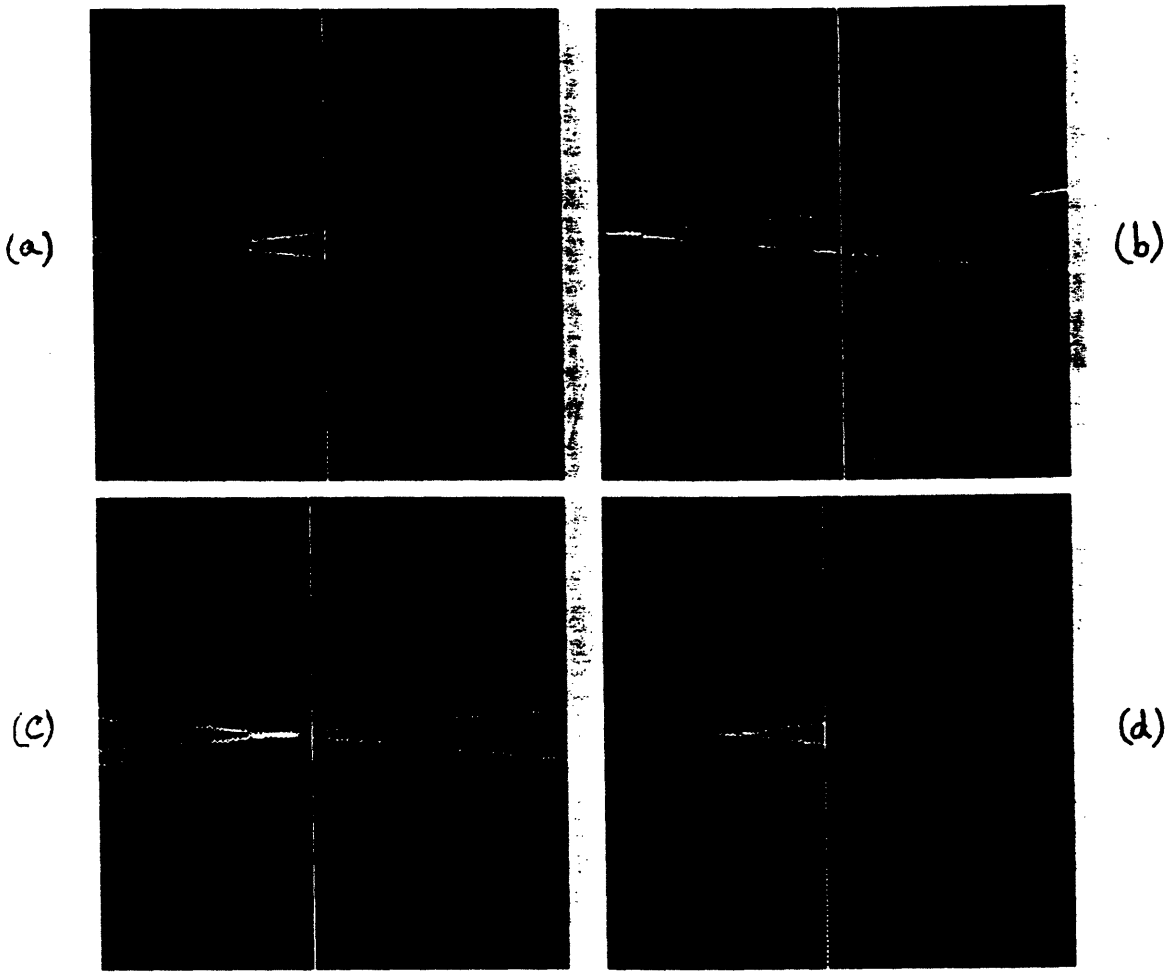

Figure 13: Tomographic reconstructions from simulated range-resolved measurements of a conical object with $\sigma_{\text{eff}} = 0.50$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by `NUA`; and (d) data with registration errors corrected by `BNGONROT`.

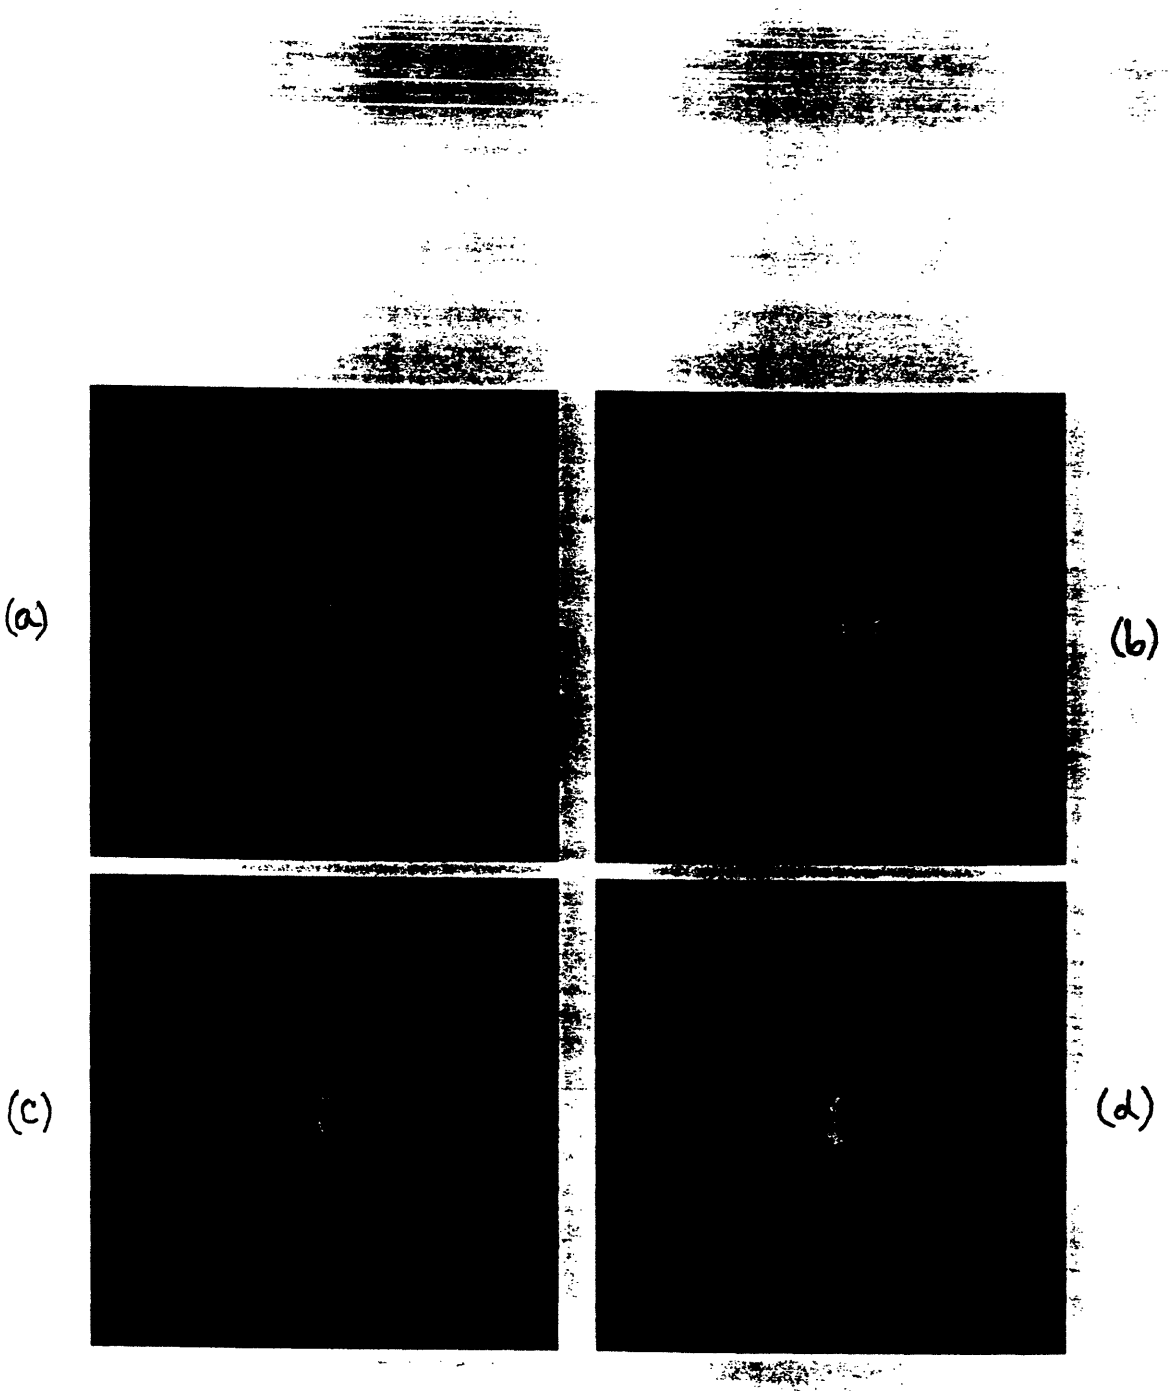

Figure 14: Tomographic reconstructions from simulated Doppler-resolved measurements of a conical object with $\sigma_{\text{kl}} = \sigma_{\text{reg}} = 0.25$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by `NUA`; and (d) data with registration errors corrected by `BNGONROT`.

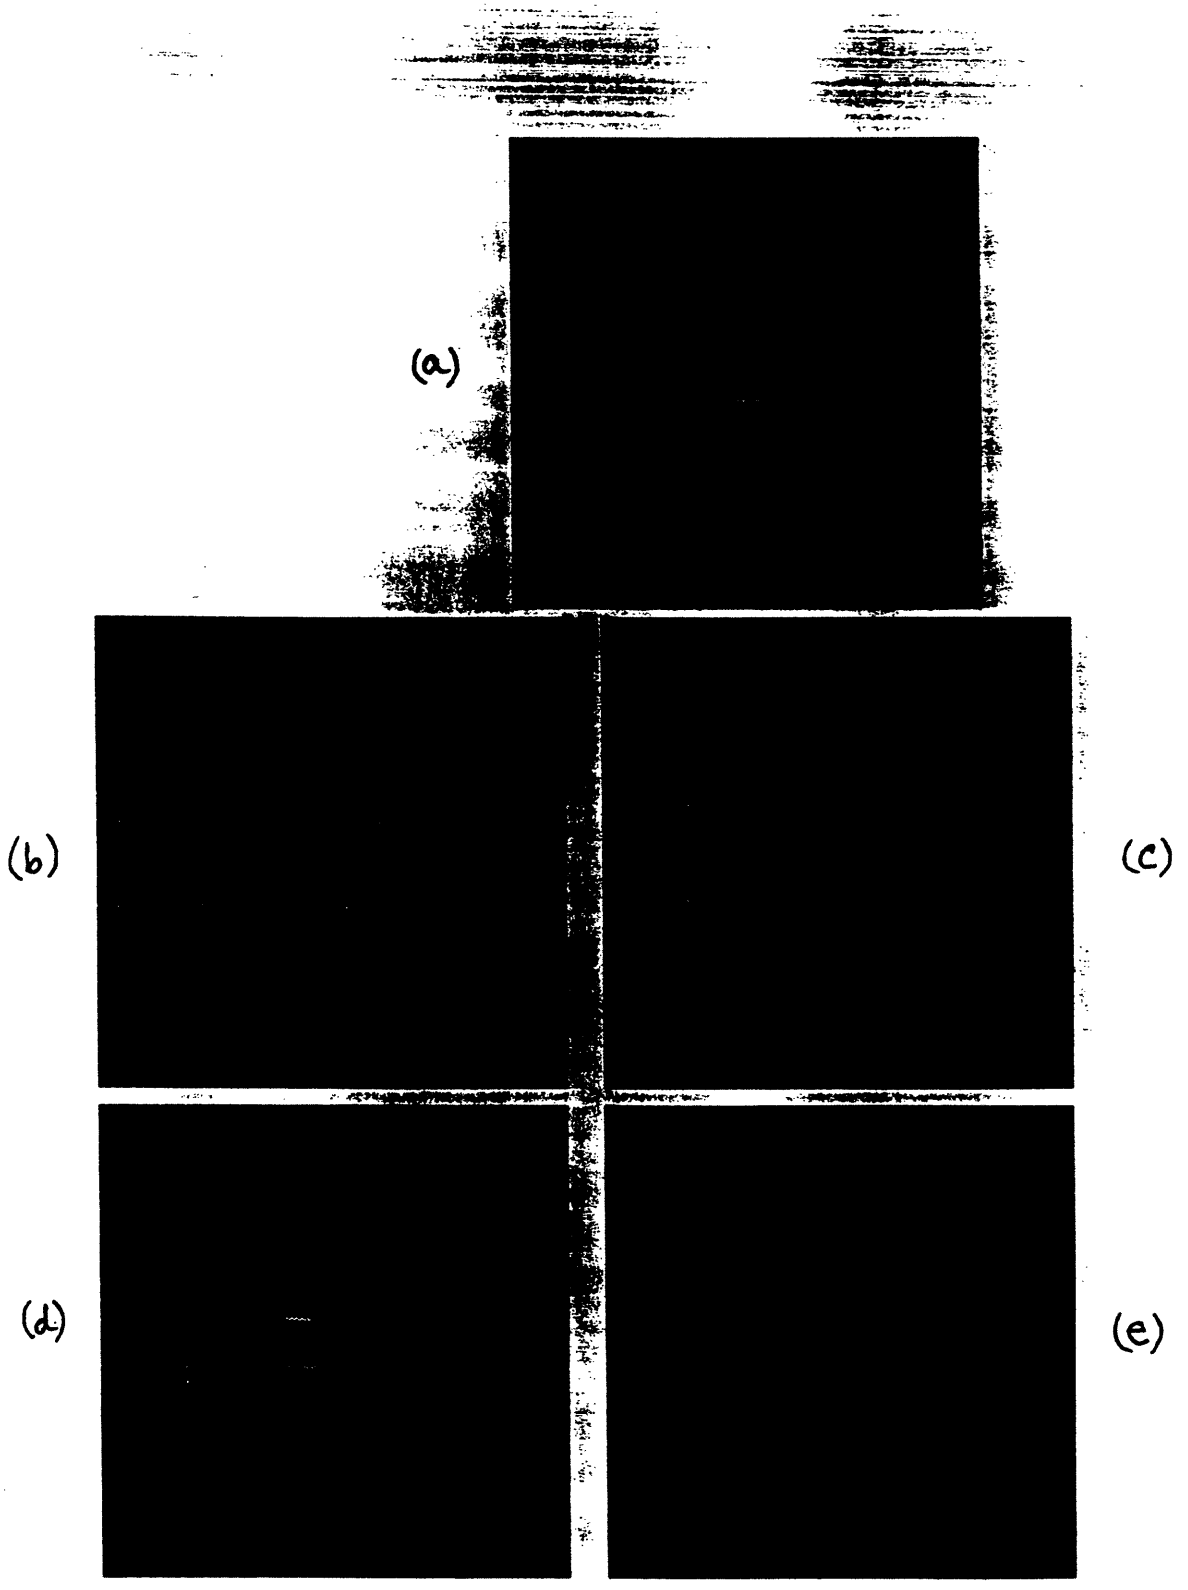

Figure 15: Tomographic reconstructions from laboratory range-resolved measurements of a triconic object with $\sigma_{\text{eff}} = 0.25$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; (c) data with registration errors corrected by **NUA**; (d) data with registration errors corrected by **BNGON**; and (e) data with registration errors corrected by **BNGONROT**.

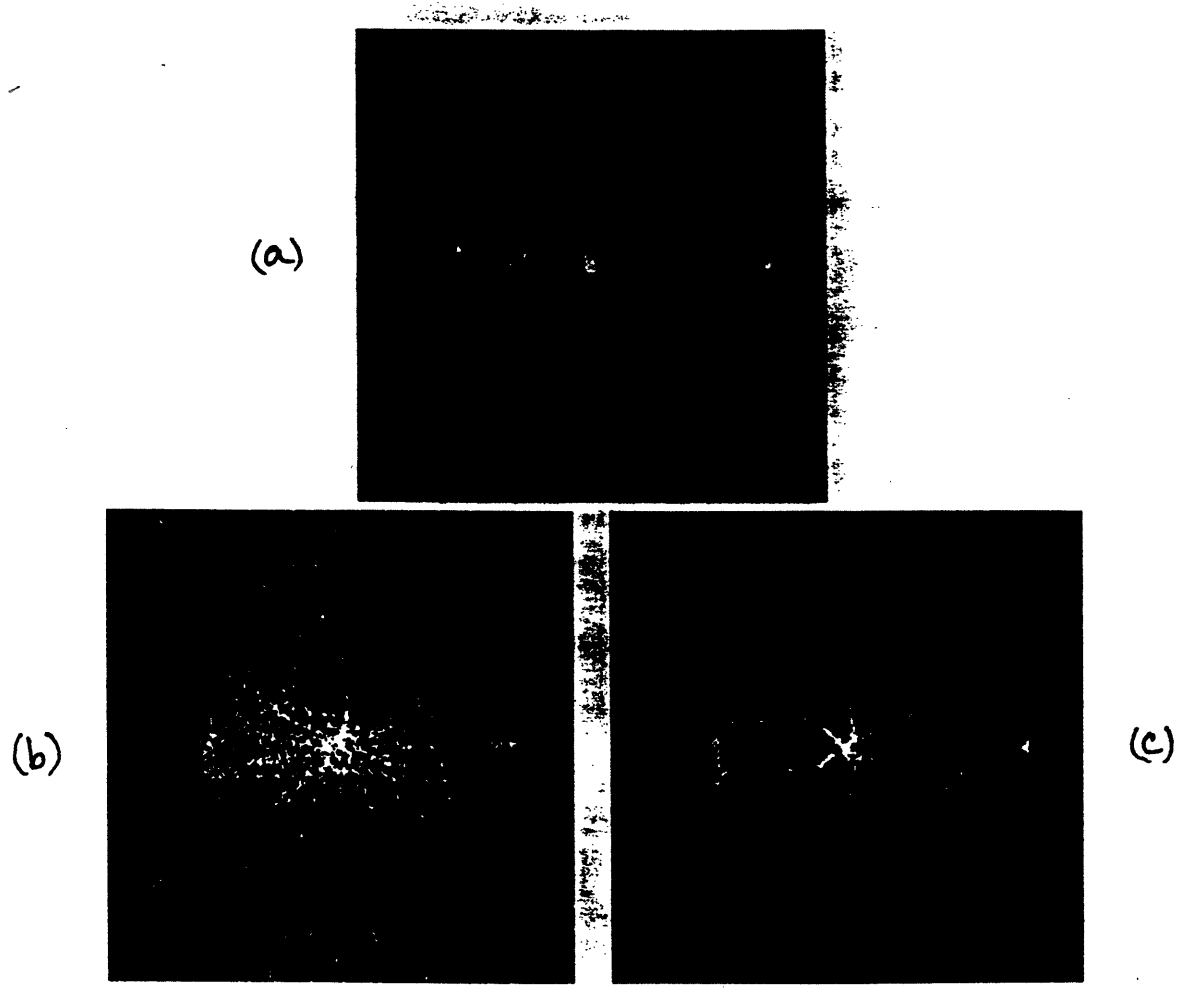

Figure 16: Tomographic reconstructions from field Doppler-resolved measurements of a Thor Delta rocket body with $\sigma_{kl} = \sigma_{reg} = 0.10$ m. The reconstructions shown are from (a) data without registration errors; (b) data with registration errors; and (c) data with registration errors corrected by NUA.

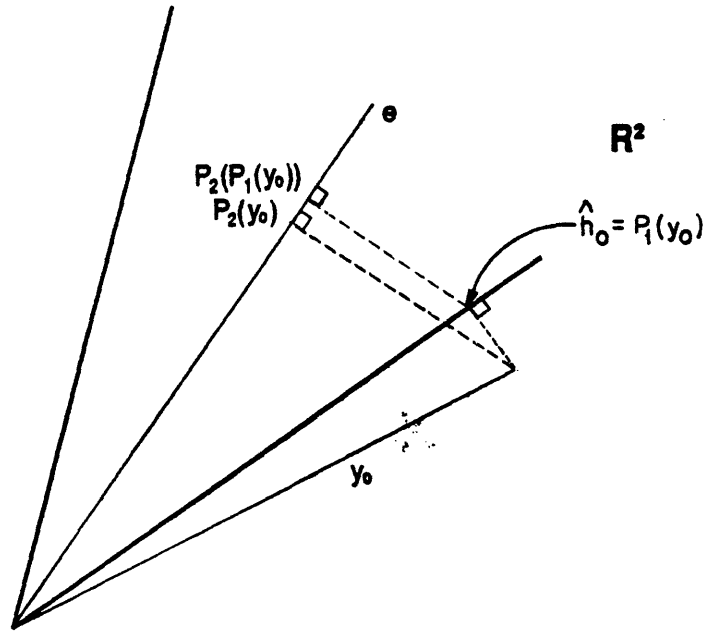

Figure 5: Support cone geometry illustrating the fact that $P_2(P_1(y_0))$ is farther along e than $P_2(y_0)$ is, leading to bias in the estimator.