

Convex-polygon estimation from support-line measurements and applications to target reconstruction from laser-radar data

A. S. Lele

Harvard Law School, Harvard University, Cambridge, Massachusetts 02138

S. R. Kulkarni

Department of Electrical Engineering, Princeton University, Princeton, New Jersey 08544

A. S. Willsky

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

Received November 15, 1990; revised manuscript received April 16, 1992; accepted May 8, 1992

We address techniques for the estimation of convex polygons from support-line measurements and introduce the application of these methods to laser-radar data. The algorithms developed use varying degrees of prior information. Quantitative assessments of their performance with respect to various parameters are provided. As expected, prior information concerning object shape and orientation greatly improves performance. It is interesting that nearly the same performance is obtained with and without prior information about object orientation, and this enables us to extract an estimate of orientation. These convex-polygon 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. Although 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 by using these methods are more robust to noisy 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 polygons 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 is any line that just grazes the boundary of the set, so that the set lies entirely in one of the half-planes defined by the support line. Clearly, a convex set is completely determined by its support lines at all orientations, and it can be obtained by simply intersecting the corresponding half-planes. 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 half-planes may not yield satisfactory results. Prince¹ and Prince and Willsky² formulated the problem of estimating a convex set from noisy support-line measurements and studied a variety of algorithms. (Greschak,³ Stark and Peng,^{4,5} and others have done related work.)

Here we introduce three new estimation algorithms that use varying degrees of prior information. The first is simply an extension of an algorithm from Refs. 1 and 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 that fits the measurements. The final algorithm is similar to the second algorithm except that rotation of the constellation angles is allowed. Hence this algorithm provides both an N -gon reconstruction and 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, as observed by Stark and Peng,^{4,5} Prince,¹ and Prince and Willsky,² in tomographic imaging the nonzero extent of each transmission projection provides support information on the underlying mass distribution. 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 about the extent of the target in space. For ex-




Fig. 1. Block diagram of reconstruction procedure.

ample, a range-resolved measurement indicates where the target begins in range along the radar line of sight (LOS). If a plane is drawn perpendicular to the line of sight at this range, the target lies completely to one side of this plane and, in fact, is just grazed by it. Range-resolved measurements from a number of aspects yield a set of support planes of the target that could conceivably be used to obtain a three-dimensional estimate of the target. In this paper we restrict our attention to the case in which all the aspects, i.e., all the 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 described in this paper may be decomposed in the manner indicated by Fig. 1. In the first module support-line information is extracted from the laser-radar data by an estimation procedure developed by others^{6,7} 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 of target shape.

Section 2 provides the background and formulation of the 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. In Section 5 the reconstruction algorithms are applied 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.⁸ Also, a method is introduced whereby the tomographic reconstructions from unregistered data may be greatly improved by 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.

A. 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 Fig. 2). The support value $h_S(\theta_0)$ is defined as the maximum projection onto $\omega(\theta_0)$ of all points in S :

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

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

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

From all this it follows that the set S lies in a particular one of the two half-planes 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 $\theta_1, \theta_2, \dots, \theta_M$ yields a support vector $h_S = [h_S(\theta_1) \ h_S(\theta_2) \ \dots \ h_S(\theta_M)]^T$. Also, from Fig. 2 it should be apparent that support lines provide no information about 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 if and only if the set is convex. For this reason we restrict our attention to convex sets in subsequent discussions 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.⁹⁻¹³ 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 if and only if $h''(\theta) + h(\theta) \geq 0$ (for example, see Ref. 12). 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 sections are facilitated by considering the




Fig. 2. Support line of a set.

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 Fig. 3(a). 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 Refs. 1 and 2, this leads to the consistency condition for a triplet of support values at adjacent and, in general, nonuniformly spaced angles, given by Ref. 14:

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

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

$$Ch \geq 0, \quad (4)$$

where

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

With such a consistency condition we can formulate several estimation algorithms, which are discussed in Section 3.

3. ESTIMATION ALGORITHMS

The estimation algorithms that we present in Sections 3–6 arise from increasingly general formulations of the problem of obtaining polygonal shape estimates from noisy support measurements. The most specific case, in which a polygon with faces at a fixed number of uniformly spaced measurement angles is estimated, was considered by Prince¹ and Prince and Willsky.² 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 measurement angles. Both sets of angles are nonuniformly spaced, in general. Fourth, we might permit 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.

A. Reconstruction with Sides at the Measurement Angles

In this problem we have a finite set of noisy support measurements $\{y_1, y_2, \dots, y_M\}$ at angles $\theta_1 < \theta_2 < \dots < \theta_M$. 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

$$\hat{h} = \underset{Ch \geq 0}{\operatorname{argmin}} \sum_{i=1}^M (y_i - h_i)^2, \quad (6)$$

where C is given by Eq. (5). We refer to this estimator as NUA, an acronym for nonuniform angles, 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 by using standard quadratic-programming (QP) techniques.^{15,16} 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 $(0, 2)$, $(-0.25, 0)$, 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 4(a) and 4(b) depict the results in both object space and support-function space by using the estimator NUA. Figure 4(a) shows the bold outline of the true object (the

standard triangle), the noisy support lines, and the shaded polygonal reconstruction produced by NUA. Correspondingly, Fig. 4(b) 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 $E = 1.56$.

B. Best N -gon Fitting M Measurements with Fixed Reconstruction Angles

In this section we exploit prior information about the angles of the object's sides to obtain reconstructions of higher quality than those we expect to obtain by using NUA, which uses 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, 20 noisy measurements of an object known *a priori* to be triangular.

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

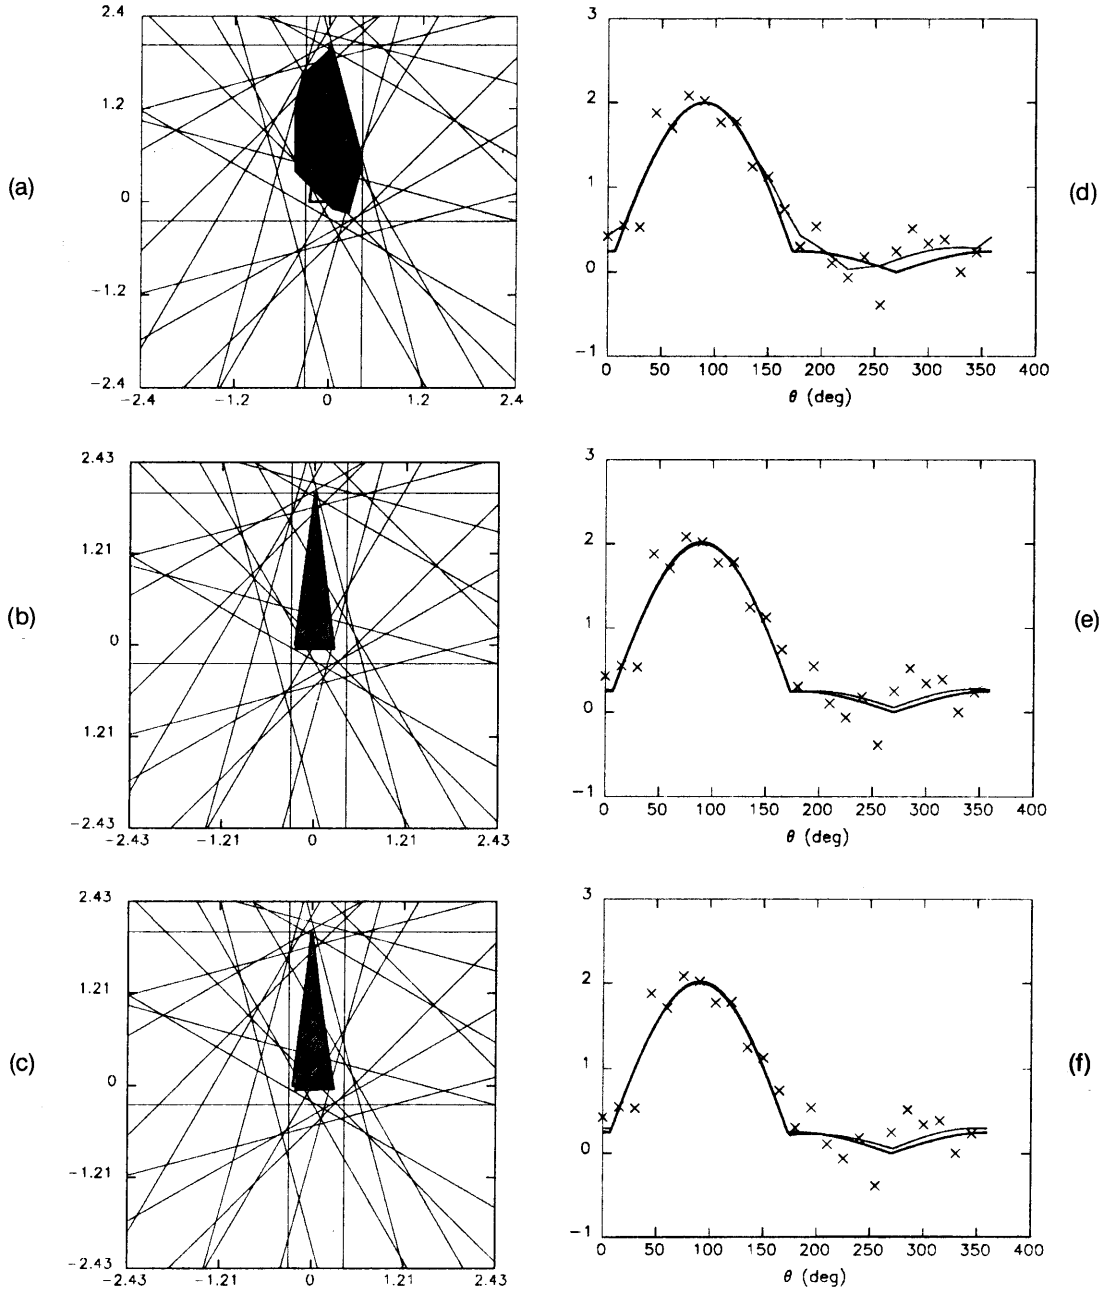


Fig. 4. Examples of the algorithms (a), (b) NUA; (c), (d) BNGON; and (e), (f) BNGONROT for the standard triangle in object space and support-function space for $M = 24$ uniformly spaced noisy ($\sigma = 0.25$) measurements. In (b), (d), and (f) the support function of the standard triangle is plotted 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.

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

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

where $h_\phi(\theta_i)$ denotes the value at θ_i of the support function $h_\phi(\cdot)$ of our estimated N -gon [see Eq. (8) and the associated explanation]. Equation (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

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

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 Fig. 3(a). 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 Refs. 1 and 2, this leads to the consistency condition for a triplet of support values at adjacent and, in general, nonuniformly spaced angles, given by Ref. 14:

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

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

$$Ch \geq 0, \quad (4)$$

where

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

With such a consistency condition we can formulate several estimation algorithms, which are discussed in Section 3.

3. ESTIMATION ALGORITHMS

The estimation algorithms that we present in Sections 3–6 arise from increasingly general formulations of the problem of obtaining polygonal shape estimates from noisy support measurements. The most specific case, in which a polygon with faces at a fixed number of uniformly spaced measurement angles is estimated, was considered by Prince¹ and Prince and Willsky.² 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 measurement angles. Both sets of angles are nonuniformly spaced, in general. Fourth, we might permit 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.

A. Reconstruction with Sides at the Measurement Angles
In this problem we have a finite set of noisy support measurements $\{y_1, y_2, \dots, y_M\}$ at angles $\theta_1 < \theta_2 < \dots < \theta_M$. 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

$$\hat{h} = \underset{Ch \geq 0}{\operatorname{argmin}} \sum_{i=1}^M (y_i - h_i)^2, \quad (6)$$

where C is given by Eq. (5). We refer to this estimator as NUA, an acronym for nonuniform angles, 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 by using standard quadratic-programming (QP) techniques.^{15,16} 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 $(0, 2)$, $(-0.25, 0)$, 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 4(a) and 4(b) depict the results in both object space and support-function space by using the estimator NUA. Figure 4(a) shows the bold outline of the true object (the

standard triangle), the noisy support lines, and the shaded polygonal reconstruction produced by NUA. Correspondingly, Fig. 4(b) 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 $E = 1.56$.

B. Best N -gon Fitting M Measurements with Fixed Reconstruction Angles

In this section we exploit prior information about the angles of the object's sides to obtain reconstructions of higher quality than those we expect to obtain by using NUA, which uses 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, 20 noisy measurements of an object known *a priori* to be triangular.

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

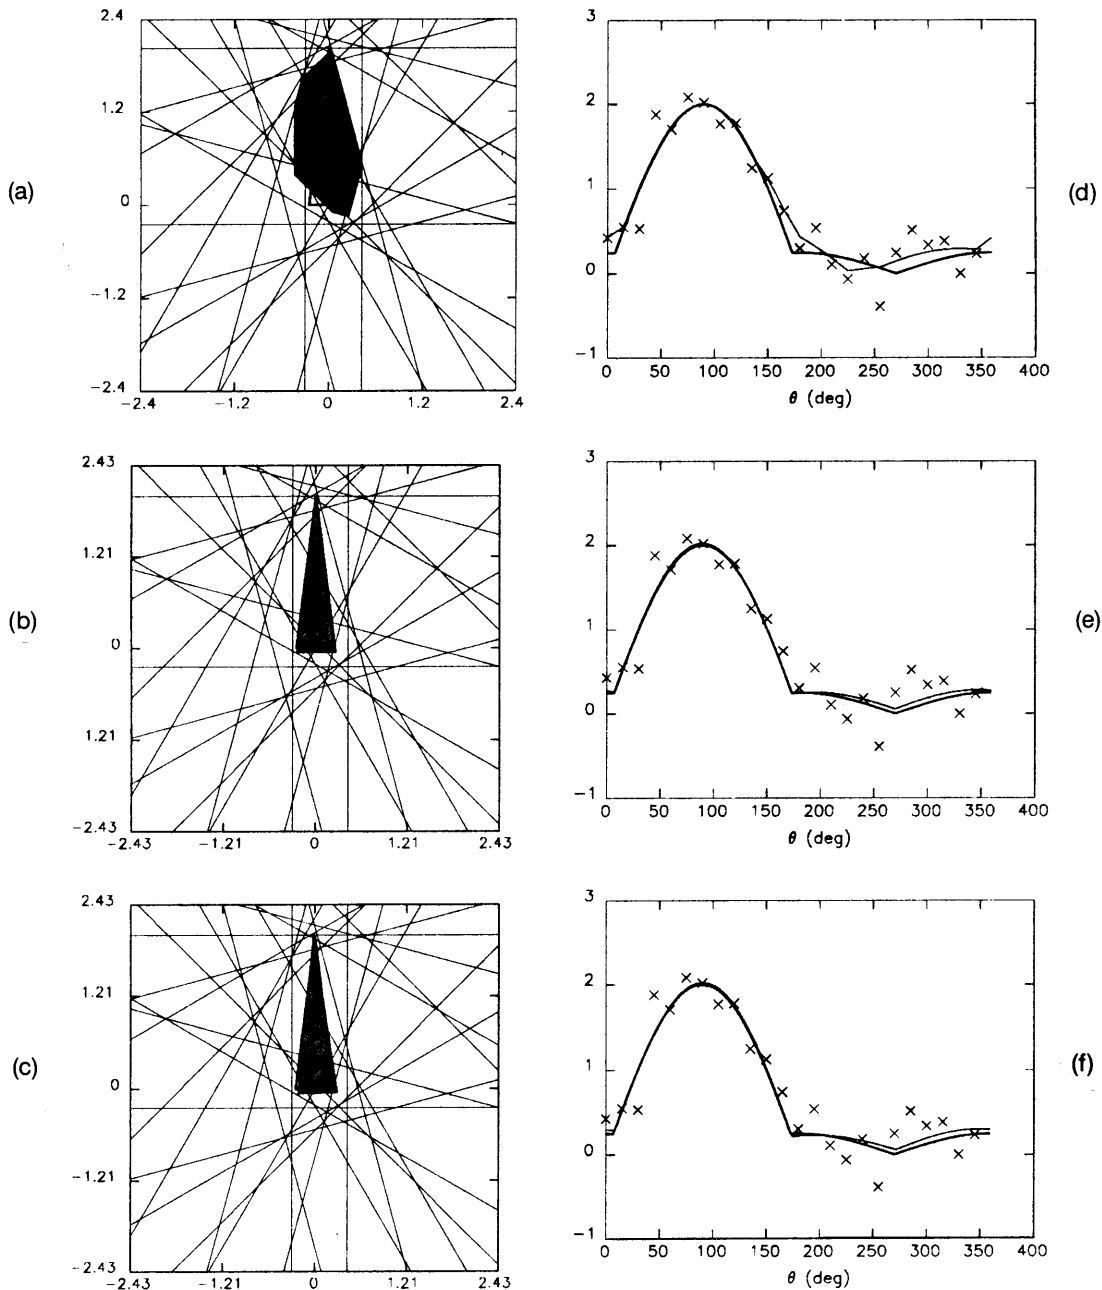


Fig. 4. Examples of the algorithms (a), (b) NUA; (c), (d) BNGON; and (e), (f) BNGONROT for the standard triangle in object space and support-function space for $M = 24$ uniformly spaced noisy ($\sigma = 0.25$) measurements. In (b), (d), and (f) the support function of the standard triangle is plotted 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.

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

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

where $h_\phi(\theta_i)$ denotes the value at θ_i of the support function $h_\phi(\cdot)$ of our estimated N -gon [see Eq. (8) and the associated explanation]. Equation (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

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

From Eqs. (7) and (8) our problem is formulated as

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

where $y = [y_1 y_2 \dots y_M]^T$ is the measurement vector, C is the consistency matrix of Eq. (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 $\{\theta_i\}$ by using Eq. (8). The i th row of the matrix A , corresponding to the i th measurement, has two adjacent (modulo N) nonzero entries, $\sin(\phi_{R_i} - \theta_i)/\sin(\phi_{R_i} - \phi_{L_i})$ and $\sin(\theta_i - \phi_{L_i})/\sin(\phi_{R_i} - \phi_{L_i})$, corresponding to the reconstruction angles ϕ_{L_i} and ϕ_{R_i} on either side of θ_i .

We refer to the estimator of Eq. (9) as BNGON. As before, since the cost function in Eq. (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 Ref. 14 for details.

An example of BNGON similar to that discussed above is shown in Figs. 4(c) and 4(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 Fig. 4 originates from the same set of measurements as the NUA reconstruction in the same figure (i.e., the same noise realization was used), permitting 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 $E_{\text{BNGON}} = 0.17$, whereas $E_{\text{NUA}} = 1.56$. However, prior information about 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 Subsection 3.D, 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, although 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 about the absolute orientation of the object. In Subsection 3.C we describe a generalization of the BNGON that addresses this problem.

C. 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 $\{\theta_1, \theta_2, \dots, \theta_M\}$ and $\{y(\theta_1), y(\theta_2), \dots, y(\theta_M)\}$ denote the M measurement angles and the measured support values at these angles, as before. How-

ever, unlike before, the reconstruction angles are given by $\{\phi_1 + \alpha, \phi_2 + \alpha, \dots, \phi_N + \alpha\}$, where $\{\phi_1, \phi_2, \dots, \phi_N\}$ are known and $\alpha \in [0, 2\pi)$ serves as an unknown offset parameter fixing the absolute locations of the reconstruction angles. Essentially, we wish to minimize the cost function in Eq. (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 estimate jointly values of α and $\{h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \dots, h_\phi(\phi_N + \alpha)\}$ that minimize

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

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

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

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

Note that although criterion (10) is quadratic in the values of h_ϕ at the reconstruction angles $\phi_1 + \alpha, \dots, \phi_N + \alpha$, it is certainly not quadratic in α [see Eq. (11)]. Thus the minimization of Eq. (11) is not a simple QP 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 :

$$J_{h_\phi}(\alpha) = \min_{\{h_\phi(\phi_i + \alpha)\}} J[\alpha, h_\phi(\phi_1 + \alpha), h_\phi(\phi_2 + \alpha), \dots, h_\phi(\phi_N + \alpha)]. \quad (12)$$

Note that solving Eq. (12) for any given value of α is simply a BNGON problem solved by means of a QP computation as described above. Note also that the minimization of Eq. (10) corresponds to choosing α to minimize $J_{h_\phi}(\alpha)$. Thus a brute-force approach to minimizing Eqs. (10) is to perform 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 values 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, note that

$$J_{h_\phi}(\alpha) = J[\alpha, h_\phi^*(\phi_1 + \alpha), h_\phi^*(\phi_2 + \alpha), \dots, h_\phi^*(\phi_N + \alpha)], \quad (13)$$

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

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

The difficulty here is that computing the sensitivity of the optimal support values $h_\phi^*(\phi_i + \alpha)$ with respect to α is not easy (since a QP optimization is involved). Thus, in-

stead of Eq. (14), we simply use

$$\begin{aligned} \frac{\partial J}{\partial \alpha} [\alpha, h_\phi^*(\phi_1 + \alpha), h_\phi^*(\phi_2 + \alpha), \dots, h_\phi^*(\phi_N + \alpha)] \\ = 2 \sum_{i=1}^M \frac{\partial h_\phi(\theta_i)}{\partial \alpha} [h_\phi^*(\theta_i) - y(\theta_i)], \quad (15) \end{aligned}$$

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

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

The second key point is that $J_{h_\phi}(\alpha)$ is, in general, a highly nonconvex function of α (see Ref. 14 for examples and discussion). Thus it is necessary to determine all the local minima of $J_{h_\phi}(\alpha)$ and choose the one with the 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 $\alpha = 0^\circ$ and solve the QP problem of Eq. (9). Using the estimated support values, we compute Eq. (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: (i) solve Eq. (9), (ii) compute the gradient, and (iii) 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 Eq. (9) and computing the gradient, determine whether our next series of steps will consist of gradient ascent or descent steps. Performing steps (i)–(iii) repeatedly, we reach our next maximum or minimum. We continue this traversal of the interval $[0^\circ, 360^\circ)$ until we have located all maxima and minima and then choose the global minimum $\hat{\alpha}$. Solving Eq. (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 or 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 Ref. 14). To deal with this we halve the step size λ of the gradient ascent or 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 (ensuring 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 suspect that its limitations are similar to those associated with the standard methods. Most important is the trade-off of speed versus accuracy as determined primarily by the choice of λ and the convergence criterion. For a given desired accuracy this algo-

rithm 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 the value that has the lowest cost.

An example of a reconstruction produced by BNGONROT is shown in Figs. 4(e) and 4(f). The true object and measurements are the same as before. The reconstruction forms an angle of $\alpha = 86.58^\circ$ 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 Figs. 4(a) and 4(b)]. Moreover, it is not much worse than the BNGON reconstruction [see Figs. 4(c) and 4(d)], indicating that not much is sacrificed in settling for a weaker prior knowledge, i.e., knowing relative rather than absolute reconstruction angles.

D. 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 Fig. 5(a) 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. Although 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 Ref. 14. However, even with this increased degradation, BNGONROT's performance is much better than that of NUA. Figure 5(b) 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. That is, if the quality of the prior information on which these algorithms are based—the number of faces and their absolute (for BNGON) or relative (for BNGONROT) angular locations—is in error, then it is possible for these algorithms to perform worse than the unconstrained NUA algorithm, especially for more severe errors in the assumed prior information. For example, let us examine the sensitivity of BNGON to errors in assumed face angles. 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 5(c) depicts plots of error E versus α for the two types of error in the prior information. A dotted 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

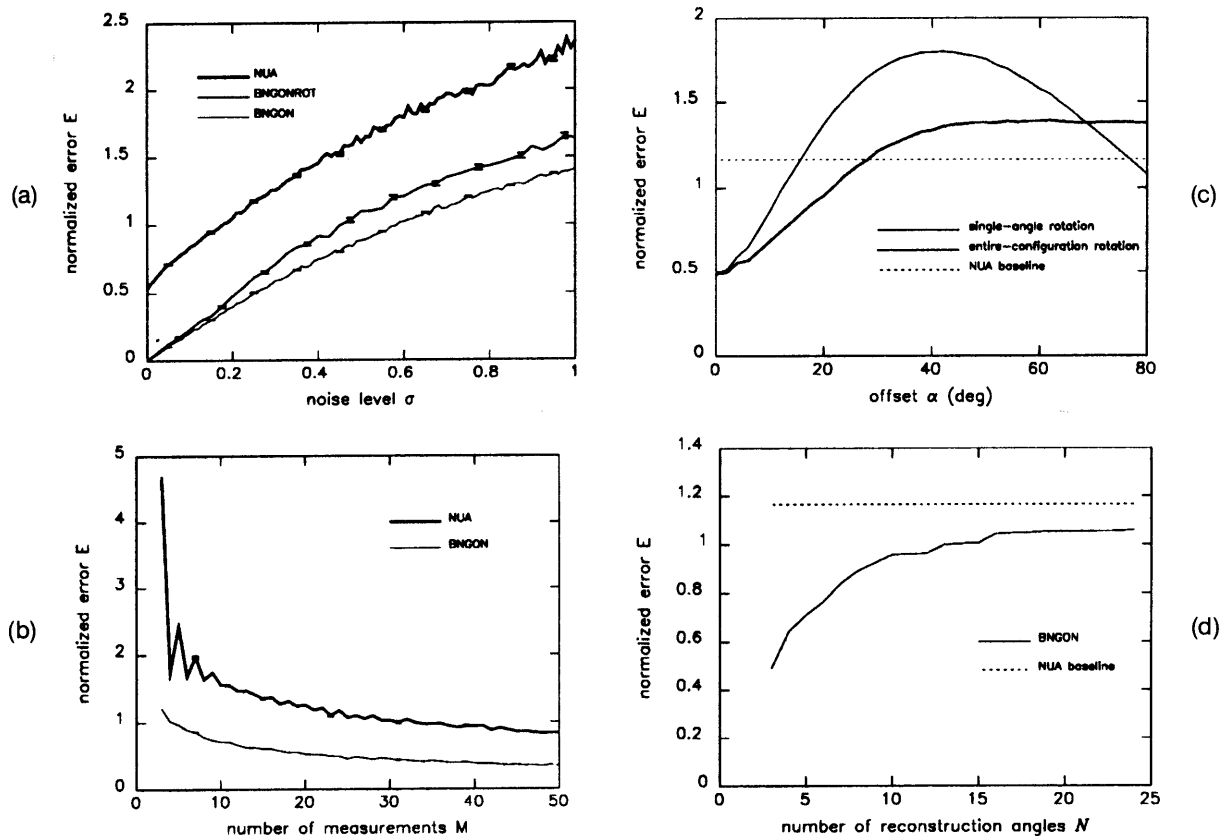


Fig. 5. Performance of estimators using the standard triangle: (a) E versus σ for NUA, BNGON, and BNGONROT, (b) E versus M for NUA and BNGON, (c) E versus α for BNGON with baseline NUA performance, and (d) E versus N for BNGON with baseline NUA performance.

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 $\approx 17^\circ$ and entire-configuration errors of up to $\approx 29^\circ$ 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 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 $\phi_1 = 7.125^\circ$, $\phi_2 = 82.875^\circ$, and $\phi_3 = 270^\circ$ for $N = 3$. For all values of $N > 3$ we choose ϕ_N such that it lies halfway between the most distant adjacent pair of the previous $N - 1$ reconstruction angles. For each value of N we solve the resulting BNGON problem in a Monte Carlo fashion to generate a data point in Fig. 5(d). 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 yields better performance. This is particularly ap-

parent 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 B we analytically determine the Cramer-Rao lower bound¹⁷ on the orientation-error variance as well as an approximate expression for the probability of obtaining anomalous orientation estimates (the threshold phenomenon mentioned above). Together, these expressions describe the orientation-error variance for a range of noise levels, which is given by

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

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

Monte Carlo simulations supplementing this analytical error analysis are 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 $\alpha_{\text{true}} = 90^\circ$. 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 $[\text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}})]^{1/2}$ 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 relation (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 $\text{Pr}(A)$ 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 we describe a technique to extract such support measurements from the data. Previous work in reconstructing targets from such laser-radar

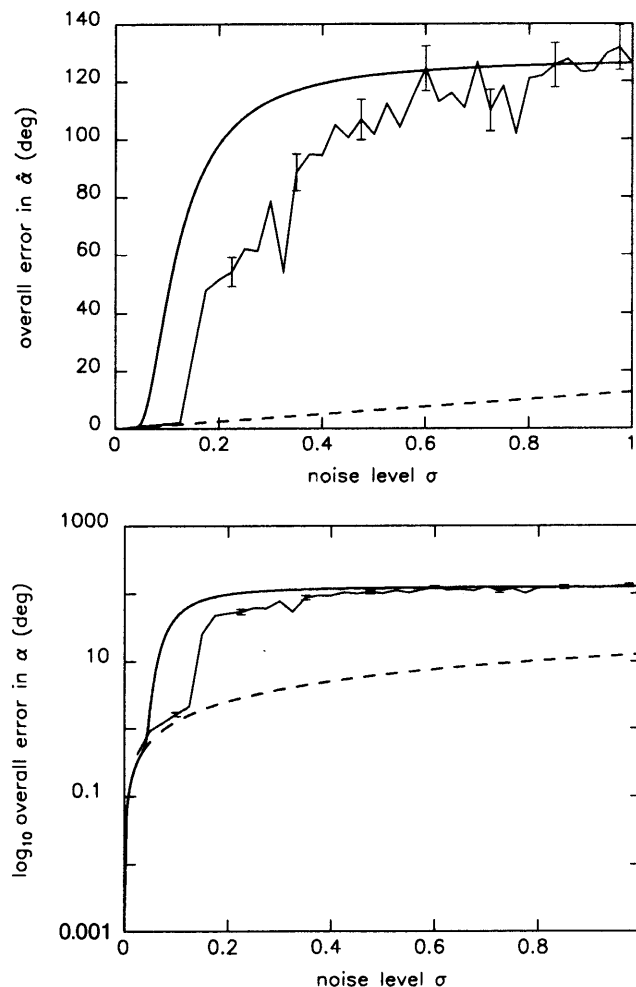


Fig. 6. Overall standard deviation $[\text{var}(\alpha - \hat{\alpha} | \alpha_{\text{true}})]^{1/2}$ of the orientation estimate versus σ , obtained analytically and through Monte Carlo simulation. Included for reference are the Cramer-Rao bound, shown dashed, and confidence intervals for the simulation results, indicated by vertical lines.

data have primarily employed techniques from transmissive tomography.^{8,18} These transmissive tomographic techniques are designed to provide reconstructions of an object's mass density, given measurements of line integrals of this mass density. As we discuss in Section 5, laser-radar measurements have precise interpretations as integrals of object reflectivity over that portion of object area falling into a particular range or Doppler bin (i.e., at a particular range from the sensor or with a particular Doppler shift). Indeed, as analyzed in detail in Ref. 19, although standard tomographic reconstruction (using, for example, filtered backprojection) does not exactly invert these measurements, it can, in principle, produce a reasonably well-focused image of the object, providing shape information and, less accurately, a map of reflectivity.

Our motivations for using convex-set reconstruction techniques to process laser-radar data are twofold. First, as we discuss in Subsection 4.B, laser-radar measurements do provide direct information about support lines to an object, and thus, if one is interested only in estimating object shape, use of our convex-set reconstruction methods provides this information directly, without the need for tomographic reconstruction. Second, as we discuss in Section 5, typical laser-radar situations involve relative motion of object and sensor to obtain data at different orientations. However, in contrast to medical imaging applications, the measurement process is far less tightly controlled and is subject to far more uncertainty. This can lead to registration errors that can severely degrade the quality of tomographic reconstructions. We will see later in this section that our convex-set reconstruction procedures provide an effective means for correcting such registration errors and thus can be of considerable value even if the ultimate goal is a full tomographic reconstruction.

A. 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.^{20,21} In this paper we restrict our 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 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. Such a measurement then actually corresponds to a set of measurements as a function of range.

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.²² For the case of a monostatic radar and a surface with isotropic reflectance properties, the bidirectional reflectance distribution function is given by $\rho(\psi)$, where ψ is the angle between the LOS and the local surface normal. In the case for which the wavelength of the illumination is large compared with surface aberrations of the target, the received intensity is proportional to

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

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 of the target. Hence for resolved data the intensity value associated with a particular bin is proportional to the laser-radar cross section 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 Section 3. Throughout we consider only the case in which the data are taken at aspects around a great circle, so that the LOS's all lie in a plane. With this restriction the entire situation is reduced to a two-dimensional problem in the plane containing the LOS's.

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, resolution of the target requires target motion, 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,^{8,18,23} we make several assumptions. Specifically, as we shall see, knowledge of the relative positions of the sensors is needed for reconstruction of 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 a 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 multisensor perspective.

Finally, note that, whether we have a multisensor system or a single-sensor system that uses relative motion between target and sensor to obtain multiple viewing angles, we must eventually deal with registration errors between

measurements taken at different views. For example, errors in the knowledge of the distances of two sensors to the target will cause an error manifesting itself as a relative shift in range-resolved data provided by the two sensors. Analogous registration errors also arise in the case of Doppler-resolved data (resulting, for example, from errors in knowledge of individual sensor velocity).

B. Support-Line Measurements from Laser-Radar Data

As illustrated in Fig. 7(a), 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 this 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 a plane, the problem is effectively reduced to a two-dimensional one, as mentioned above. That is, we need only to consider the projection of the target in the plane of LOS's. The support-plane information contained in the data corresponds to support-line




Fig. 7. Diagrams showing how (a) range-resolved and (b) Doppler-resolved measurements give rise to one and two support lines, respectively.

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 that is 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} , respectively, with nonzero return intensity correspond to the minimum and maximum cross range of any part of the target [see Fig. 7(b)]. Thus from a Doppler-resolved measurement we can extract two lines parallel to the LOS that graze the target that 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, we need a coordinate frame. This frame must serve as a common reference for all 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 are, 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 $\theta_i \pm 90^\circ$. Since target cross range is proportional to Doppler frequency after the Doppler spectrum is shifted by the Doppler frequency shift D_i produced by the target's translational velocity relative to the sensor, the support values are given by $(\lambda/2\omega)|D_{\min} - D_i|$ and $(\lambda/2\omega)|D_{\max} - D_i|$, where λ is the wavelength of the laser illumination.

Since a Doppler spectrum at aspect θ_i provides two support values at $\theta_i \pm 90^\circ$, the aspects θ_i and $\theta_i + 180^\circ$ yield duplicate support values if the support values are free of noise. For noisy data 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 because of two types of measurement error. One type of error arises from 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 Subsection 4.C. Second, 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.

C. Knot Location

Although determining r_{\min} or D_{\min} and D_{\max} is trivial if the data are noise free, such a determination for noisy data is a 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 Jones²⁴ for detecting abrupt changes in dynamic systems and later applied by Mier-Muth and Willsky⁷ to spline estimation. To cast our problem in the framework of Ref. 7, we model the range or Doppler spectrum as a linear spline, or piecewise-linear function. The points of discontinuity in the 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 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 maximum-likelihood estimates for each bin in the window, we perform a generalized-likelihood-ratio test for the knot-present and knot-absent hypotheses to determine whether a knot actually exists at the locations of any of the maximum-likelihood estimates. The first bin for which the generalized-likelihood ratio 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 backward along the spectrum. Details concerning the implementation and performance of this algorithm may be found in Refs. 6, 7, and 24. Note that the resulting support-line measurements are subject to two sources of error, namely, the error in knot location resulting from noise and registration error, which causes a shift in the entire range- or Doppler-resolved measurement. For range-resolved data this registration error enters as additive measurement uncertainty in one support line to be extracted. For Doppler-resolved data the registration error also enters additively—with opposite signs—in the two extracted support-line measurements.

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.²³ In particular, the values of the laser-radar return at ranges just higher than r_{\min} are determined by parts of the target whose surface normals roughly coincide with the LOS. As a result, $\psi \approx 0^\circ$, maximizing $\cos \psi$ in Eq. (18). Furthermore, since materials typically give high-intensity return at near-normal incidence and low-intensity return at near-grazing incidence, the bidirectional reflectance distribution function $\rho(\psi)$ is near maximum. Hence range spectra generally exhibit an abrupt increase in intensity




Fig. 8. (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, and (d) BNGONROT.

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

5. TARGET RECONSTRUCTIONS FROM LASER-RADAR DATA

In this section we apply the knot-location technique discussed in Subsection 4.C 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.

A. Convex-Set Reconstructions

The data for the first two examples are simulated²⁵ range-resolved and Doppler-resolved measurements of a cone

with a height of 200 cm, a radius of 25 cm, and Lambertian reflectance characteristics. The cone is positioned with the center of its base at the origin of a coordinate frame and is 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 Subsection 4.C and convert them to support values. If knot-location errors and registration errors for each aspect are modeled 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 $\sigma_{\text{eff}}^2 = \sigma_{kl}^2 + \sigma_{reg}^2$ for range-resolved data. However, for Doppler-resolved data at an even number of uniformly spaced aspects, (i) registration errors for aspects 180°

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

The support lines resulting from locating knots and corrupting the support values by measurement noise are shown in Fig. 8(a) for range-resolved data with noise level $\sigma_{eff} = 0.50$ m. The reconstructions produced by NUA, BNGON, and BNGONROT from this set of noisy support-line measurements are shown in Figs. 8(b)–8(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 Subsections 3.A–3.C. In particular, the prior knowledge of relative reconstruction angles allows BNGONROT to outperform NUA dramatically but does not cause it to underperform

BNGON significantly, since BNGON 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 ($\sigma_{kl} = 0.25$ m) and registration error ($\sigma_{reg} = 0.25$ m) are shown in Fig. 9.

The next example is one of reconstructing a triconic target with a height of 203 cm and a base radius of 39.5 cm (shown outlined in Fig. 10), given laboratory range-resolved measurements. The laboratory measurements were taken on a 10-m 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 Ref. 8 for details of the experimental setup. Support lines and reconstructions using the three estimators are shown in Fig. 10 for the uncorrupted laboratory data and in Fig. 11 for the laboratory data corrupted by adding additional knot-location errors with $\sigma_{eff} = 0.25$ m. Note that the uncorrupted laboratory data did contain some modest amount of noise, resulting in some errors in knot location. However, in a real system somewhat larger errors than are




Fig. 9. (a) Conical object and associated support-line measurements from simulated Doppler-resolved data with $\sigma_{kl} = \sigma_{reg} = 0.25$ m and reconstructions using (b) NUA, (c) BNGON, and (d) BNGONROT.




Fig. 10. (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, and (d) BNGONROT.

present in this controlled laboratory environment would be expected, and this provided the motivation for examining the effect of additional knot uncertainty in Fig. 11. (Note that the σ_{eff} value chosen is significant with respect to the target base radius). Also, 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 Fig. 12(a)], 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 by using a 10.6- μm CO₂ narrow-band 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 Ref. 8. Support lines and reconstructions produced by the three algorithms are shown in Fig. 12 for the uncorrupted data and in Fig. 13 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.

B. Comparison with and Improvements to Tomographic Imaging Methods

In previous work standard methods of tomographic image reconstruction²⁶ were applied to range-resolved and Doppler-resolved laser-radar data.^{8,18} In this subsection we compare the convex-set reconstructions of Subsection 5.A with reconstructions produced by using the tomographic methods. We then examine the effect of registration errors on both methods. As we shall see, the present algorithms are robust to registration errors, in contrast to tomographic reconstructions, which are rather sensitive to these errors. Finally, we show that the robustness of the present algorithms can be used to improve dramatically tomographic reconstructions from data with registration errors.

All the tomographic reconstructions in this subsection were obtained by using the standard method of filtered backprojection. (See Ref. 26 for methods of transmission tomography, and see Refs. 8 and 18 and the references contained therein for the application of

these methods to laser-radar reflective data.) Parts (a) of Figs. 14–17 show filtered backprojection reconstructions from the four data sets (free of registration errors) used in Subsection 5.A. It should be noted that we thresholded the Doppler data sets before backprojecting them 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 before 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 Figs. 8–13], the tomographic reconstructions contain intensity information within the outline of the target. Furthermore, the tomographic images differ from their convex-set counterparts in that they do not provide direct size or shape estimates of the target. Although, in principle, techniques for extracting 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 Subsection 4.C and

described and demonstrated in Ref. 23, reconstructed edges are not highlighted but are instead overwhelmed by the high intensities that are reconstructed in the interior of the target. Even if thresholding is used 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 Figs. 14–17 show the tomographic reconstructions resulting from shifting the data in each spectrum by an amount given by a zero-mean Gaussian random variable with standard deviation $\sigma_{\text{reg}} = 0.50, 0.25, 0.25$, and 0.10 m (with the shifts for the spectra being independent of one another except for the Doppler data sets, for which shifts for aspects 180° apart are negatives) and then using filtered back-projection. Clearly, one cannot expect any image-processing algorithm to extract shape information successfully 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 Figs. 8, 9, 11, and 13, which




Fig. 11. (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, and (d) BNGONROT.




Fig. 12. (a) Thor-Delta rocket body and associated support-line measurements from field Doppler-resolved data with $\sigma_{kl} = \sigma_{reg} = 0$ m and reconstructions using (b) NUA, (c) BNGON, and (d) BNGONROT.

were obtained from data suffering from registration errors identical to those used for the tomographic reconstructions (i.e., the same noise realizations were used) and from 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 by using implicit information about 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 about 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 back-projection.

Parts (c) and (d) of Figs. 14 and 15, parts (c)–(e) of Fig. 16, and part (c) of Fig. 17 show the tomographic reconstructions that result from 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 for which the reconstruction corresponding to BNGON was omitted, it could not be distinguished from that corresponding to BNGONROT. In the case for which reconstructions for both BNGON and BNGONROT were omitted, they were indistinguishable from that corresponding to NUA. Clearly, when the various images within each of Figs. 14–17 are compared, the improvement obtained by 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 using prior information, if available. The algorithms are computationally feasible, with two of the estimators resulting in QP problems and the third having a QP 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. The performance of BNGONROT is comparable with that of BNGON, and BNGONROT eliminates the need for prior orientation information. Its ability 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 by 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 with those produced by tomographic imaging methods, and this comparison resulted in the following observations. First, shape estimates are explicitly provided by our algorithms, whereas tomographic images can provide target-shape information only after the use of 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, in an effort to improve the quality of tomographic images, to exploit the tendency of our algorithms to correct unregistered data.

Our reconstruction algorithms might be extended in a variety of ways. One such extension may consist of developing more general formulations of the best N -gon algorithm, so that less stringent prior shape information could be used. 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




Fig. 13. (a) Thor-Delta rocket body and associated support-line measurements from field Doppler-resolved data with $\sigma_{kl} = \sigma_{reg} = 0.10$ m and reconstructions using (b) NUA, (c) BNGON, and (d) BNGONROT.




Fig. 14. Tomographic reconstructions from range-resolved measurements of a conical object with $\sigma_{\text{eff}} = 0.50$ m 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.

interesting to develop algorithms that provide smooth, rather than polygonal, shape estimates of objects. Incorporating the effects of noise into the measurement angles may be useful in some applications. Another useful generalization would be to extend the algorithms to three dimensions.

There are several possibly interesting extensions regarding the application to laser-radar data. After the restriction that aspects lie in a plane is relaxed (i.e., when general aspects are permitted in three dimensions), shape estimates of a target (rather than those of its projection onto a plane) will 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.

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 \mathcal{C} , then the NUA estimate is given simply by $\hat{h} = y$. On the other hand, if $y \notin \mathcal{C}$, then the estimation procedure consists of projecting y onto \mathcal{C} to yield the estimate $\hat{h} = P_1(y)$ [see Fig. 3(b)].

Some quantitative measure of the size of a reconstructed object is necessary for discussing bias in the estimator. One such measure is the magnitude of the projection of \hat{h} onto the vector $e = [1 \ 1 \dots 1]^T$, where the projection is denoted by $P_2(\hat{h})$. This quantity is proportional (with proportionality constant \sqrt{M}) to the perimeter of the reconstructed set, since the perimeter of a set having support function $\hat{h}(\theta)$ is given by $\int_0^{2\pi} \hat{h}(\theta) d\theta$, an expression whose discrete counterpart is $(1/M)\hat{h}^T e = (1/M)\sum_{i=1}^M \hat{h}_i = (1/\sqrt{M})P_2(\hat{h})$ 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 Fig. 18. Here the cone boundaries are given by the two heavier lines, and e is the center axis. The measurement is given by $y_0 = h + n_0$, where n_0 is zero-mean noise and 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 $A = P_2(h)$, let $p_n(n_0)$ denote the Gaussian-noise probability-density function, and let μ denote the expected value of $P_2[P_1(y_0)]$, then

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

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 Ref. 14 for a quantitative evaluation of the bias based on Monte Carlo simulations.

APPENDIX B: DERIVATION OF LOWER BOUND ON ERROR VARIANCE OF ORIENTATION ESTIMATE

As in all nonlinear estimation problems, the CRB¹⁷ provides a lower bound on estimation-error variances. In principle, in the problem of interest here we need to consider the joint estimation of all the support values $\{h_\phi(\phi_i)\}$ as well as the orientation angle α . However, since we wish to focus here on accuracy in estimating α , for simplicity let us look at the case in which the $\{h_\phi(\phi_i)\}$ are known. The CRB for this problem obviously provides a lower bound on orientation-error variance for the original problem in which the $\{h_\phi(\phi_i)\}$ must be estimated as well. For our simpler problem the observation vector is given by $y = h(\alpha) + n$, where $n \sim N(0, \sigma^2 I)$ and the i th component of $h(\alpha)$ is the true support value induced at θ_i when the object has an offset of α , as given by Eq. (11). The CRB on the error variance of any unbiased estimate of α is




Fig. 15. Tomographic reconstructions from Doppler-resolved measurements of a conical object with $\sigma_{kl} = \sigma_{reg} = 0.25$ m 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.




Fig. 17. Tomographic reconstructions from field Doppler-resolved measurements of a cone with $\sigma_{kl} = \sigma_{reg} = 0.10$ m from (a) data without registration errors, (b) data with registration errors, and (c) data with registration errors corrected by NUA.




Fig. 18. Support-cone geometry illustrating the fact that $P_2[P_1(y_0)]$ is farther along e than is $P_2(y_0)$, leading to bias in the estimator.

take the estimation-error variance when a global error occurs as the constant value $(180^\circ)^2$.

Thus we have characterized the estimation-error variance $\text{var}(\alpha - \hat{\alpha} | \alpha_{true})$ for the low-noise and high-noise regimes. For intermediate noise levels the error variance can be approximated, as in Ref. 17, 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 $\text{Pr}(A)$, 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 $\alpha = 90^\circ$ (zero error) and $\alpha = 270^\circ$ (180° global error). Using standard results on Gaussian binary hypothesis testing,¹⁷ we find that the resulting probability of anomaly is given by

$$\text{Pr}(A) = Q\left(\frac{\sqrt{H}}{2\sigma}\right) \leq \frac{1}{2} \exp\left(\frac{-H}{8\sigma^2}\right). \quad (21)$$

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,

$$H = [h(270^\circ) - h(90^\circ)]^T [h(270^\circ) - h(90^\circ)]. \quad (22)$$

Combining these results yields the orientation-error variance, formula (17).

ACKNOWLEDGMENTS

This research was conducted while A. S. Lele and S. R. Kulkarni were at MIT Lincoln Laboratory and the Laboratory for Information and Decision Systems, MIT. This research was supported by the U.S. Department of the Navy under U.S. Air Force contract F19628-90-C-0002, U.S. Army Research Office contract DAAL03-86-K-0171, National Science Foundation contract 9015281-MIP, and U.S. Office of Naval Research grant N00014-91-J-1004.

REFERENCES

1. J. L. Prince, "Geometric model-based estimation from projections," Ph.D. dissertation (Massachusetts Institute of Technology, Cambridge, Mass., 1988).
2. J. L. Prince and A. S. Willsky, "Estimating convex sets from noisy support line measurements," *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 377-389 (1990).
3. J. P. Gerschak, "Reconstructing convex sets," Ph.D. dissertation (Massachusetts Institute of Technology, Cambridge, Mass., 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, eds. (Elsevier Science, New York, 1988).
5. H. Stark and H. Peng, "Shape estimation in computer tomography from minimal data," *J. Opt. Soc. Am. A* **5**, 331-343 (1988).
6. A. M. Mier-Muth, "Adaptive knot location for spline approximation," S. M. thesis (Massachusetts Institute of Technology, Cambridge, Mass., 1976).
7. A. M. Mier-Muth and A. S. Willsky, "A sequential method for spline approximation with variable knots," *Int. J. Syst. Sci.* **9**, 1055-1067 (1978).
8. F. K. Knight, S. R. Kulkarni, R. M. Marino, and J. K. Parker, "Tomographic techniques applied to laser radar reflective measurements," *Lincoln Lab. J.* **2**, 143-160 (1989).
9. T. Bonnesen and W. Fenchel, *Theory of Convex Bodies* (BCS, Moscow, Idaho, 1987).
10. S. R. Lay, *Convex Sets and Their Applications* (Interscience, New York, 1982).

11. D. E. McClure and R. A. Vitale, "Polygonal approximation of plane convex bodies," *J. Math. Anal. Appl.* **51**, 326-358 (1975).
12. L. A. Santalo, *Integral Geometry and Geometric Probability*. Vol. 1 of Encyclopedia of Mathematics and Its Applications (Addison-Wesley, Reading, Mass., 1976).
13. R. A. Vitale, "A representation theorem for compact convex sets in the plane," *Div. Appl. Math. Rep.* (Brown University, Providence, R.I., 1974).
14. A. S. Lele, "Convex set reconstruction from support line measurements and its application to laser radar data," S. M. thesis (Massachusetts Institute of Technology, Cambridge, Mass., 1990).
15. A. H. Land and S. Powell, *FORTTRAN Codes for Mathematical Programming* (Interscience, London, 1973).
16. D. G. Luenberger, *Linear and Nonlinear Programming*, 2nd ed. (Addison-Wesley, Reading, Mass., 1984).
17. H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory* (Wiley, New York, 1968).
18. J. K. Parker, E. B. Craig, D. I. Klick, F. K. Knight, S. R. Kulkarni, R. M. Marino, J. R. Senning, and B. K. Tussey, "Reflective tomography: images from range-resolved laser radar measurements," *Appl. Opt.* **27**, 2642-2643 (1988).
19. B. T. Binder, "Laser radar tomography: the effects of speckle," Ph.D. dissertation (Massachusetts Institute of Technology, Cambridge, Mass., 1991).
20. C. G. Bachman, *Laser Radar Systems and Techniques* (Artech, Norwood, Mass., 1979).
21. A. W. Rihaczek, *Principles of High-Resolution Radar* (McGraw-Hill, New York, 1977).
22. F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, "Geometrical considerations and nomenclature for reflectance," *Natl. Bur. Stand. (U.S.) Monogr.* **160** (1977).
23. S. R. Kulkarni, A. S. Lele, M. F. Reiley, "On the qualitative differences between tomographic reconstructions from range-resolved versus Doppler-resolved data," Lincoln Lab. Proj. Memo. 52PM-ODT-0042 (MIT Lincoln Laboratory, Lexington, Mass., 1989).
24. A. S. Willsky and H. L. Jones, "A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems," *IEEE Trans. Autom. Control* **AC-21**, 108-112 (1976).
25. A. S. Lele, "TARMAN: a computer model for dynamic target simulation," S. B. thesis (Massachusetts Institute of Technology, Cambridge, Mass., 1987).
26. R. A. Brooks and G. DiChiro, "Principles of computer assisted tomography in radiographic and radioisotopic imaging," *Phys. Med. Biol.* **21**, 689-732 (1976).