



Shape Estimation from Support and Diameter Functions*

AMYN POONAWALA

Department of Computer Engineering, University of California, Santa Cruz, CA 95064, USA

PEYMAN MILANFAR

Department of Electrical Engineering, University of California, Santa Cruz, CA 95064, USA

RICHARD J. GARDNER

Department of Mathematics, Western Washington University, Bellingham, WA 98825, USA

Published online: 9 January 2006

Abstract. We address the problem of reconstructing a planar shape from a finite number of noisy measurements of its support function or its diameter function. New linear and non-linear algorithms are proposed, based on the parametrization of the shape by its Extended Gaussian Image. This parametrization facilitates a systematic statistical analysis of the problem via the Cramér-Rao lower bound (CRLB), which provides a fundamental lower bound on the performance of estimation algorithms. Using CRLB, we also generate confidence regions which conveniently display the effect of parameters like eccentricity, scale, noise, and measurement direction set, on the quality of the estimated shapes, as well as allow a performance analysis of the algorithms.

Keywords: support function, diameter function, brightness function, Extended Gaussian Image, parametric shape estimation, constrained Cramér-Rao lower bound, confidence regions.

1. Introduction

This paper is motivated by the problem of reconstructing an unknown planar shape from a finite number of noisy measurements of its support function or its diameter function. Given a measurement direction (i.e., a unit vector, or angle), the corresponding *support function* measurement gives the (signed) distance from some fixed reference point (usually taken to be the origin) to the support line to the shape orthogonal to the direction (see Fig. 2). The corresponding *diameter function* measurement provides the distance between the two support lines parallel to this direction (see Fig. 3). We shall refer to support and diameter functions collectively as *support-type functions*. In view of the data, it is natural to focus on convex bodies.

Support function data arise in a variety of physical experiments and therefore have been studied by researchers with diverse interests. Prince and Willsky [22] used such data in computerized tomography as a prior to improve performance, particularly when only limited data is available. Lele et al. [13] applied support function measurements to target reconstruction from range-resolved and Doppler-resolved laser-radar data. The general approach in these papers is to fit a polygon to the data, in contrast to that of Fisher et al. [2], who use spline interpolation and the so-called von Mises kernel to fit a smooth curve to the data. This method was taken up in [9] and [18], the former dealing with convex bodies with corners and the latter giving an example to show that the algorithm in [2] may fail for a given data set. Support function data has also featured in robotics via the notion of a line probe; however, the focus here has been on algorithmic complexity issues under the assump-

*Supported in part by U.S. National Science Foundation grants CCR-9984246 and DMS-0203527.




Figure 1. Two-step procedure for shape from support-type measurements.

tion of exact rather than noisy data (see, for example, [15] and [24]).

Diameter measurements can be obtained using an instrumented parallel-jaw gripper and shape from diameter has been studied by researchers in the robotics community. Rao and Goldberg [23], observing that shape cannot be uniquely recovered even from exact diameter measurements (see Section 2.1), used them to recognize a shape among a known finite set of shapes. Li [14] found the precise number of (exact) diameter measurements which together with location information allow one to reconstruct a polygon with a known number of edges.

There is a more general and much-studied form of data called the brightness function, the 2-D case of which is just the diameter function. For an n -dimensional body, brightness function measurements give the $(n - 1)$ -dimensional volumes of its orthogonal projections onto hyperplanes (i.e., areas of shadows). The problem of reconstruction from brightness function data is important in geometric tomography, the area of mathematics concerning the retrieval of information about a geometric object from data about its sections or projections (see [3]). Algorithms for reconstructing an object from noisy measurements of its brightness function were introduced by Gardner and Milanfar [5, 6]. Moreover, convergence of the principal algorithms has been established in this setting (see [4] for details).

One contribution of the present paper (see Sections 2 and 3) is a new approach to support function measurements that employs the Extended Gaussian Image (EGI). This encodes a convex polygon that approximates the shape in terms of the lengths a_k of its edges and their outer normal angles θ_k . The EGI is a suitable parametric model enabling a systematic statistical analysis of shape estimation from support functions. Moreover, for the diameter function, the EGI is the most natural and convenient parametrization (see Section 2.2). There is therefore the added advantage of a uniform treatment of both support and diameter functions.

Regardless of the underlying true shape, our method always produces a convex polygon that approximates it. The reconstruction problem is solved in two steps (see Fig. 1). In the first step, the noisy data are used to estimate the EGI of the unknown shape; we obtain, in fact, the EGI of the approximating convex polygon. The second step produces a Cartesian representation of this polygon from its EGI.

Another contribution of this paper is a systematic statistical analysis (see Section 4) of the problem of reconstructing a planar shape from noisy support or diameter function measurements. The approach involves the derivation of the constrained Cramér-Rao lower bound (CCRLB) on the estimated parameters. Using the CCRLB and following the approach employed in [27] and [28], local and global confidence regions can be calculated corresponding to any preassigned confidence level. These form uncertainty regions around points in the boundary of the underlying object, or around its whole boundary. Such confidence regions are tremendously powerful in displaying visually the dependence of measurement direction set, noise power, and the eccentricity, scale, and orientation of the underlying true shape on the quality of the estimated shape. They also allow a performance analysis of our algorithms, carried out in Section 5, where experimental results can also be found.

Support and diameter functions are typically measured using electronic sensor devices such as a camera, robot tool, grippers, etc., so the noise corrupting the measurements is generally the electrical noise coming from the sensors. This noise arises during data acquisition (readout noise) and can be approximated by a Gaussian distribution. Throughout the paper, therefore, the noise is modelled as Gaussian white noise. Two remarks about this are appropriate. Firstly, this model makes natural the use of least squares algorithms (see Section 3). However, since such algorithms are sensitive to outliers, one could, if the data were thought to contain many outliers, use instead L_p norm minimization for an appropriate $p \in [1, 2]$. Secondly, the CRLB analysis could




Figure 2. The support function.

if necessary be carried out for different noise statistics.

2. Support-Type Functions and the EGI

2.1. Background

If K is a convex body in the plane \mathfrak{R}^2 and \mathbf{v} is a unit vector, the *support function* $h_K(\mathbf{v})$ of K is

$$h_K(\mathbf{v}) = \sup_{\mathbf{x} \in K} \mathbf{x}^T \mathbf{v}. \quad (1)$$

Then $h_K(\mathbf{v})$ is the (signed) distance from the origin $\mathbf{0}$ to the *support line*

$$L(\mathbf{v}) = \{\mathbf{x} \in \mathfrak{R}^2 \mid \mathbf{x}^T \mathbf{v} = h_K(\mathbf{v})\}$$

to K orthogonal to \mathbf{v} . A convex body is completely determined by its support function; see [25, p. 38].

The most useful metric for calculating the distance between convex bodies can be defined by means of the support function. The *Hausdorff distance* between two convex bodies K and M is

$$\delta_H(K, M) = \|h_K - h_M\|_\infty, \quad (2)$$

where $\|\cdot\|_\infty$ is the infinity norm.

It is often convenient to write the unit vector $\mathbf{v} = [\cos \alpha, \sin \alpha]^T$. Then the support function $h_K(\mathbf{v}) = h_K(\alpha)$ is continuous and periodic with period 2π . Closely related to the support function is the *diameter function* $b_K(\mathbf{v})$ of K , given by

$$b_K(\mathbf{v}) = b_K(\alpha) = h_K(\alpha + \pi/2) + h_K(\alpha - \pi/2). \quad (3)$$

The diameter function is continuous and periodic with a period of π , and $b_K(\mathbf{v}) = b_K(\alpha)$ is the distance between the two supporting lines $L(\alpha \pm \pi/2)$ parallel to \mathbf{v} (see Fig. 3).




Figure 3. The diameter function.




Figure 4. Sets of convex bodies with equal diameter functions.

The diameter function does not determine a convex body, even up to translation, and even if the body is a polygon (see Fig. 4, left) or a disk (see Fig. 4, right). In fact, for each convex body there is an origin-symmetric convex body with the same diameter function. (An *origin-symmetric* body is one equal to its reflection in the origin.) However, any two origin-symmetric convex bodies having the same diameter function must be equal. This is a very special case of *Aleksandrov's projection theorem* [3, Theorem 3.3.6]. For detailed discussions, see [3, Chap. 3] and [6].

Underlying our whole approach is the idea of parameterizing a shape using its *Extended Gaussian Image* (EGI), which effectively encodes the shape in terms of the curvature as a function of the normal to its boundary. When K is a convex polygon, its EGI is simple to describe. For an N -sided polygon whose k th edge has length a_k and outer unit normal vector \mathbf{u}_k , the EGI can be represented by the N vectors $a_k \mathbf{u}_k$ (see the left two pictures in Fig. 5), or, if $\mathbf{u}_k = [\cos \theta_k, \sin \theta_k]^T$, by the set $\{a_1, \dots, a_N, \theta_1, \dots, \theta_N\}$. This provides a convenient set of parameters we shall often use later. The EGI has the desirable property that it determines a convex body, up to translation. We refer the reader to [3, 10, 25] for more details.

The important question naturally arises as to which sets $\{a_1, \dots, a_N, \theta_1, \dots, \theta_N\}$ of $2N$ real numbers correspond to the EGI of a convex polygon. The follow-

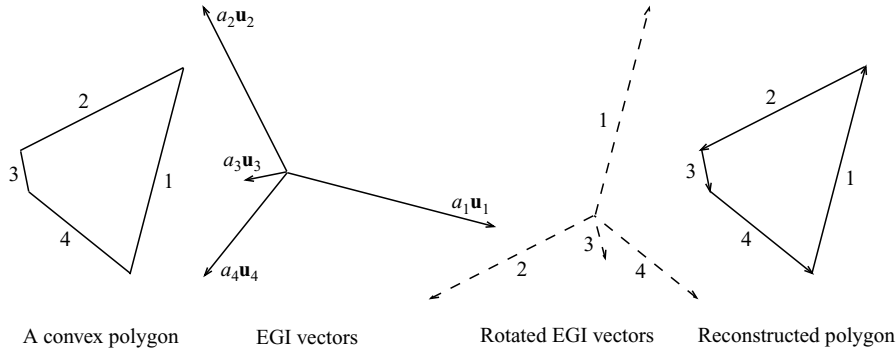


Figure 5. Obtaining the Cartesian coordinate representation from the EGI.

ing necessary and sufficient conditions are intuitively obvious:

$$a_k \geq 0 \quad \text{for } k = 1, \dots, N \quad (4)$$

and the EGI vectors sum to zero, that is,

$$\sum_{k=1}^N a_k \mathbf{u}_k = \left[\sum_{k=1}^N a_k \cos \theta_k, \sum_{k=1}^N a_k \sin \theta_k \right]^T = [0, 0]^T. \quad (5)$$

This observation is a very special case of *Minkowski's existence theorem* (see [3, p. 356] or [25, p. 390]).

2.2. Geometry from the EGI

Suppose that we have the EGI $\{a_1, \dots, a_N, \theta_1, \dots, \theta_N\}$ of a convex polygon K . Clearly K is uniquely determined, up to translation. It is then possible to obtain its diameter function, and, if the position of K is specified, its Cartesian representation and its support function.

We now obtain the corresponding formulas, beginning with the diameter function. By Cauchy's projection formula (see [3, p. 361]), the diameter of K in the direction $\mathbf{v} = [\cos \alpha, \sin \alpha]^T$ is

$$b_K(\alpha) = \frac{1}{2} \sum_{k=1}^N a_k |\cos(\alpha - \theta_k)|. \quad (6)$$

The formula becomes transparent on observing that $a_k |\cos(\alpha - \theta_k)|$ is the diameter in the direction \mathbf{v} of the k th edge of the polygon, a line segment of length a_k with normal angle θ_k .

The simplest way to reconstruct a translate of K from its EGI is to arrange the vectors in counterclockwise order, rotate each vector counterclockwise by $\pi/2$, and place them so that the tail of each vector lies at the head of the preceding vector (see Fig. 5 and [16]). Let

$\mathbf{r} = [r_1, \dots, r_N]^T = [a_1 e^{i(\theta_1 + \pi/2)}, \dots, a_N e^{i(\theta_N + \pi/2)}]^T$ list the rotated EGI vectors in complex form. Then, if the tail of the first vector is at $\mathbf{0}$, the vertices of this translate of K are represented by the complex numbers

$$0, r_1, r_1 + r_2, \dots, r_1 + r_2 + \dots + r_{N-1}. \quad (7)$$

In matrix-vector notation, the vertices in (7) are given by \mathbf{Pr} , where

$$\mathbf{P} = \begin{bmatrix} 0 & 0 & 0 & \dots & 0 & 0 \\ 1 & 0 & 0 & \dots & 0 & 0 \\ 1 & 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & 1 & 1 & \dots & 1 & 0 \end{bmatrix}. \quad (8)$$

For our purposes, however, this position of the reconstructed polygon is not optimal, since it causes any error in vertex positions to accumulate as the vertex index t increases from $t = 1$ to N . This can be avoided by positioning the reconstructed polygon so that its vertices have their barycenter at the origin (see Section 4.3.2). To obtain a formula for the latter, note that the barycenter of the vertices in (7) is represented by the complex number

$$\begin{aligned} & (0 + r_1 + (r_1 + r_2) + \dots + (r_1 + r_2 + \dots \\ & + r_{N-1}))/N = \frac{N-1}{N} r_1 + \frac{N-2}{N} r_2 + \dots \\ & + r_{N-1} = (1/N) \mathbf{e} \mathbf{e}^T \mathbf{Pr}, \end{aligned} \quad (9)$$

where $\mathbf{e} = [1, 1, \dots, 1]^T$ and where we used the fact that $\sum_{j=1}^N r_j = 0$, by (5). Then the vertices $\mathbf{s}_t = [x_t, y_t]^T$ of the translate of K whose vertices

have barycenter at $\mathbf{0}$ are given by $\mathbf{s} = \mathbf{D}\mathbf{r}$, where

$$\mathbf{D} = \mathbf{P} - (1/N)\mathbf{e}\mathbf{e}^T\mathbf{P}; \quad (10)$$

specifically,

$$\mathbf{s}_t = \begin{bmatrix} \sum_{k=1}^N d_{tk}a_k \cos(\theta_k + \pi/2), \\ \sum_{k=1}^N d_{tk}a_k \sin(\theta_k + \pi/2) \end{bmatrix}^T, \quad (11)$$

where $d_{tk} = \mathbf{D}_{t,k}$ is the (t, k) entry of the matrix \mathbf{D} .

Having established this formula for the Cartesian representation of the translate of K whose vertices have their barycenter at $\mathbf{0}$, we can also obtain its support function at $\mathbf{v} = [\cos \alpha, \sin \alpha]^T$. Indeed, choose t so that $\theta_t \leq \alpha \leq \theta_{t+1}$. Then, from (1) we have

$$h_K(\alpha) = h_{\{\mathbf{s}_t\}}(\alpha) = \mathbf{s}_t^T \mathbf{v} = x_t \cos \alpha + y_t \sin \alpha, \quad (12)$$

so from this and (11), we obtain

$$\begin{aligned} h_K(\alpha) &= \sum_{k=1}^N d_{tk}a_k \cos(\theta_k + \pi/2) \cos \alpha \\ &\quad + \sum_{k=1}^N d_{tk}a_k \sin(\theta_k + \pi/2) \sin \alpha \\ &= \sum_{k=1}^N d_{tk}a_k \sin(\alpha - \theta_k). \end{aligned} \quad (13)$$

Note that (13) holds only for polygons whose vertices have barycenter at $\mathbf{0}$ and when $\theta_t \leq \alpha \leq \theta_{t+1}$.

3. Algorithms for Shape from Support-type Functions

3.1. Algorithms Using Support Function Measurements

Throughout this section we consider an unknown planar convex body K with a finite number of noisy support function measurements modelled by $y(\alpha_m) = h_K(\alpha_m) + n(\alpha_m)$ in fixed measurement directions $\mathbf{v}_m = [\cos \alpha_m, \sin \alpha_m]^T$, $m = 1, \dots, M$. The noise $n(\alpha_m)$ is assumed to be white Gaussian noise with variance σ^2 . In vector form, the measurements are

$$\mathbf{y} = \mathbf{h}_K + \mathbf{n}, \quad (14)$$

where $\mathbf{h}_K = [h_K(\alpha_1), \dots, h_K(\alpha_M)]^T$ and $\mathbf{n} = [n(\alpha_1), \dots, n(\alpha_M)]^T$.

3.1.1. Some Previous Algorithms. Following Prince and Willsky [22], who dealt with the special case when the measurement directions are equally spaced, Lele et al. [13] proposed an algorithm for reconstructing a shape from noisy support function measurements. This algorithm solves the constrained linear least squares problem

$$\hat{\mathbf{h}} = \arg \min_{\mathbf{h}} \|\mathbf{y} - \mathbf{h}\|^2, \quad (15)$$

subject to the constraint

$$\mathbf{M}(\boldsymbol{\Omega})\mathbf{h} \geq \mathbf{0}, \quad (16)$$

where $\boldsymbol{\Omega} = [\alpha_1, \dots, \alpha_M]^T$ and $\mathbf{M}(\boldsymbol{\Omega})$ is a certain matrix depending only on $\boldsymbol{\Omega}$. The constraint (16) is a necessary and sufficient condition, essentially due to Rademacher (see [25, p. 47]), for \mathbf{h} to be consistent with a support function of a closed convex set.

There will in general be many convex polygons P with support function vector

$$\mathbf{h}_P(\boldsymbol{\Omega}) = [h_P(\alpha_1), h_P(\alpha_2), \dots, h_P(\alpha_M)]^T = \hat{\mathbf{h}}.$$

A natural choice among these for an output polygon is the largest such polygon \hat{P} , obtained by intersecting the half-planes defined by the corresponding M support lines. If $\hat{\mathbf{h}} = [\hat{h}_1, \dots, \hat{h}_M]^T$, this is given by

$$\hat{P} = \{\mathbf{x} \in \mathfrak{R}^2 \mid \mathbf{x}^T \mathbf{v}_m \leq \hat{h}_m, m = 1, \dots, M\}. \quad (17)$$

Lele et al. [13] also address the problem of reconstructing an N -sided polygon K with prescribed outer normal angles $\theta_1, \dots, \theta_N$ from the M noisy measurements (14).

3.1.2. Proposed Algorithms Using EGI Parametrization.

As described in the introduction, we propose to use the noisy support function measurements in (14) to estimate the EGI of the true shape K by finding the EGI of an approximating convex polygon. Once this is done, the method described in Section 2.2 can be used to construct this approximating polygon.

If we assume that K is an N -sided polygon whose vertices have barycenter at $\mathbf{0}$ and which has prescribed outer normal angles $\theta_k, k = 1, \dots, N$, then we proceed as follows. For each measurement angle α_m , we find

t_m so that $\theta_{t_m} \leq \alpha_m < \theta_{t_m+1}$ and then use (13) and (14) to obtain

$$y(\alpha_m) = \sum_{k=1}^N d_{t_m k} a_k \sin(\alpha_m - \theta_k) + n(\alpha_m), \quad (18)$$

where the unknowns are the edge lengths a_k , $k = 1, \dots, N$. In vector form this becomes

$$\mathbf{y} = \mathbf{S}(\Theta)\mathbf{a} + \mathbf{n}, \quad (19)$$

where \mathbf{y} and \mathbf{n} are as before, $\mathbf{a} = [a_1, \dots, a_N]^T$, $\Theta = [\theta_1, \dots, \theta_N]^T$, and the $M \times N$ matrix $\mathbf{S}(\Theta)$ has entries $\mathbf{S}(\Theta)_{m,k} = d_{t_m k} \sin(\alpha_m - \theta_k)$. We can then solve the constrained linear least squares problem

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a}} \|\mathbf{y} - \mathbf{S}(\Theta)\mathbf{a}\|^2, \quad (20)$$

subject to (4) and (5). Note that these are linear constraints. Note also that since these constraints must be satisfied by the output edge lengths and prescribed outer normal angles, the resulting $\hat{\mathbf{a}}$ is “valid” in the sense that it corresponds exactly to the EGI of a convex (approximating) polygon, which can then be reconstructed by the method of Section 2.2.

The linear algorithm based on (20) is equivalent to that of Lele et al. [13] described in Section 3.1.1; that is, for a given input vector \mathbf{y} and prescribed outer normal angle vector Θ , the outputs of the two algorithms are the same (provided both algorithms construct polygons whose vertices have barycenter at $\mathbf{0}$). If we take $N = M$ and replace Θ by Ω , the algorithm based on (20) is equivalent to that of Prince and Willsky [22].

In fact our approach allows treatment of a somewhat more general problem. Suppose that the outer normal angles are not known or prescribed, but that it is only known, for each measurement angle α_m , the value of t_m for which $\theta_{t_m} \leq \alpha_m < \theta_{t_m+1}$. This comprises weaker prior information than that assumed by the algorithms above. We may then solve the constrained *non-linear* least squares problem

$$(\hat{\mathbf{a}}, \hat{\Theta}) = \arg \min_{(\mathbf{a}, \Theta)} \|\mathbf{y} - \mathbf{S}(\Theta)\mathbf{a}\|^2, \quad (21)$$

again subject to (4) and (5). This algorithm is actually the maximum likelihood estimator (MLE) of the unknown parameters \mathbf{a} and Θ , due to the Gaussian noise assumption (see [12, pp. 223–6 and 254–5] for further details).

3.2. Algorithms Using Diameter Function Measurements

The algorithms described in this section were introduced by Gardner and Milanfar [5, 6]. In principle, suitable versions of the algorithms operate equally well in higher dimensions, when the diameter function is replaced by the brightness function.

In view of the remarks on the lack of uniqueness in reconstructing a convex body from its diameter function (see Section 2.1), there is no loss in restricting attention to origin-symmetric convex bodies. Therefore throughout this section we consider an unknown planar origin-symmetric convex body K with a finite number of noisy diameter function measurements modelled by $y(\alpha_m) = b_K(\alpha_m) + n(\alpha_m)$ in fixed measurement directions $\mathbf{v}_m = [\cos \alpha_m, \sin \alpha_m]^T$, $m = 1, \dots, M$. The noise $n(\alpha_m)$ is assumed to be white Gaussian noise with variance σ^2 . In vector form, the measurements are

$$\mathbf{y} = \mathbf{b}_K + \mathbf{n}, \quad (22)$$

where $\mathbf{b}_K = [b_K(\alpha_1), \dots, b_K(\alpha_M)]^T$ and $\mathbf{n} = [n(\alpha_1), \dots, n(\alpha_M)]^T$.

As in Section 3.1.2, the approach is that illustrated in Fig. 1; the noisy diameter function measurements are used to obtain the EGI of an approximating convex polygon, and then the vertices of this polygon can be calculated from (11). Using Cauchy’s projection formula (6), we can rewrite (22) as

$$\mathbf{y} = \mathbf{C}(\Theta)\mathbf{a} + \mathbf{n}, \quad (23)$$

where \mathbf{y} and \mathbf{n} are as in (22), $\mathbf{a} = [a_1, \dots, a_N]^T$, $\Theta = [\theta_1, \dots, \theta_N]^T$, and the $M \times N$ matrix $\mathbf{C}(\Theta)$ has entries $\mathbf{C}(\Theta)_{m,k} = |\cos(\alpha_m - \theta_k)|/2$.

Since the unknown body K is origin symmetric, we seek an approximating convex polygon that is also origin symmetric, and hence has an even number N of sides. Assuming N is prescribed, we may solve the constrained *non-linear* least squares problem

$$(\hat{\mathbf{a}}, \hat{\Theta}) = \arg \min_{(\mathbf{a}, \Theta)} \|\mathbf{y} - \mathbf{C}(\Theta)\mathbf{a}\|^2, \quad (24)$$

subject to the non-negativity constraint (4), and the constraints

$$a_k = a_{N/2+k} \quad \text{for } k = 1, \dots, N/2, \quad (25)$$

and

$$\theta_k = \theta_{N/2+k} + \pi \quad \text{for } k = 1, \dots, N/2, \quad (26)$$

for origin symmetry. (Note that (25) and (26) imply (5).) This algorithm is the maximum likelihood estimator (MLE) of the unknown parameters \mathbf{a} and Θ . However, the non-linear nature of (24) renders this algorithm computationally very expensive. Without adaptation, it is therefore of little practical use. One possible adaptation is prompted by the observation that the least squares problem (24) is of a special type known as separable. It might then be feasible to use the variable projection method of Golub and Pereyra [7] where each of the unknowns \mathbf{a} and Θ is treated separately in each iteration.

Fortunately, there are other viable methods we can employ. The outer normal angles can be prescribed, for example by setting the angles in Θ to be $\alpha_m \pm \pi/2$, $m = 1, \dots, M$, so that $N = 2M$. The resulting *linear* estimation algorithms were proposed by Gardner and Milanfar [5, 6] (where versions are also given that allow more generally the reconstruction of an n -dimensional convex body from noisy measurements of its brightness function; see [4] and [6] for an extensive analysis). These algorithms construct an approximating origin-symmetric convex polygon whose edges number up to $2M$. (Note that some of the optimal edge lengths may be zero.) They are suitable for situations in which the unknown body K is not itself a polygon, since the number of edges of the output will generally increase with the number of measurements. If, however, the number of edges is to be prescribed, various pruning methods can be considered. These are discussed in the next section.

3.3. Pruning

The linear diameter function algorithm proposed in [5, 6] constructs a polygon whose edges number up to twice the number of measurements. Similarly, the support function algorithm (15) can result in a polygon with as many edges as the number of measurements. Hence, an extra operation may be needed to prune the edges of this polygon down to or below a user-defined number. The challenge is to achieve this while causing minimal change in the overall shape. We have studied three approaches, namely, thresholding, clustering (suggested to us by Gil Kalai), and decimation.

In *thresholding*, we discard edges with lengths smaller than a preset value and similarly discard vertices if the difference between the normal directions of its adjacent sides is less than another preset value. This method can be used as a post-processing step af-

ter use of the linear diameter function algorithm, since the origin symmetry then allows edges and vertices to be discarded in antipodal pairs. Thus the constraints (25) and (26), and therefore (5) remain valid, guaranteeing a convex polygon output. Without adaptation, it is not suitable for the support function algorithms, where the following two methods are preferable. Note that even for the diameter function algorithm, extra work would be needed to find the thresholds that result in an output with exactly a given number of edges.

Clustering is a method whereby the vertices of the polygon are divided into a fixed number of disjoint groups (clusters), using hierarchical or k -means clustering algorithms as in [1] and [11], for example. A representative vertex is chosen in each cluster, and these representative vertices are the vertices of the output convex polygon.

Decimation is used for polygonal mesh simplification by the graphics community [17, 29]. Here we iteratively delete the smallest edge of the polygon by merging it with its adjacent edges. The merger affects the EGI vectors of adjacent edges in a way depending on the geometry (contrary to thresholding, where the length of the EGI vectors remains unchanged). More details about decimation are provided in [20].

Any of these three operations can be regarded as an optional extra step in Fig. 1. The graphs in Fig. 6 show the effect of clustering and decimation in pruning the edges in reconstructions of a regular polygon with $N = 16$ edges from noisy measurements of its diameter function from $M = 45$ equally spaced directions. The clustering operation was performed using hierarchical dendograms and single-linkage cost function (see [1] for more explanation). The post-processing operations were used to prune the number of edges from (up to) 90 to exactly 10, and for reasons mentioned above thresholding is not convenient for this




Figure 6. Comparison of clustering and decimation methods.

purpose. The graphs plot the average Hausdorff distance δ_H between the input and output polygons versus the noise power σ in Monte Carlo simulation of 1000 runs.

Figure 6 demonstrates that the pruning operations perform well and that decimation seems to perform better than clustering.

4. Statistical Analysis

4.1. The (Unconstrained) Cramér-Rao Lower Bound

The *Cramér-Rao lower bound* (CRLB) provides the theoretical lower bound on the variance of any unbiased estimator and hence can be used to study the fundamental limits of parametric estimation problems; see, for example, [12, pp. 27–35]. Since no unbiased estimator can have lower variance than the CRLB, it provides a convenient benchmark for comparison in the performance of a specific estimator.

Consider a measurement vector \mathbf{y} given by

$$\mathbf{y} = f(\Psi) + \mathbf{n}, \quad (27)$$

where $\Psi = [\psi_1, \dots, \psi_P]^T$ is the vector of parameters to be estimated and \mathbf{n} is the noise. Let $\hat{\Psi}$ be the vector of estimated parameters and denote by $\mathbf{J}(\Psi)$ the $P \times P$ Fisher information matrix (FIM), whose entries are given by the following expression involving the joint probability density function (p.d.f.) $p(\mathbf{x}; \Psi)$ of the observed data:

$$\mathbf{J}(i, j) = \mathbf{J}(\Psi)_{i,j} = -E \left(\frac{\partial^2 \ln p(\mathbf{x}; \Psi)}{\partial \psi_i \partial \psi_j} \right), \quad (28)$$

for $i, j = 1, \dots, P$, where $E(\cdot)$ denotes the expected value.

The CRLB (see [12, pp. 30–44]) states that the covariance matrix $\text{Cov}(\hat{\Psi})$ satisfies

$$\text{Cov}(\hat{\Psi}) \geq \mathbf{J}^{-1}(\Psi). \quad (29)$$

We first apply this bound to the problem of reconstructing from noisy support function measurements. Note that the definition (1) of the support function uses the supremum function and so is not as convenient for the statistical analysis as the formula (13) based on the EGI parametrization. Therefore the support function measurements are modelled by (19), with the assumptions behind (21) stated in Section 3.1.2. Recall that the noise is Gaussian with variance σ^2 . The measurement

vector is $\mathbf{y} = [h_K(\alpha_1), \dots, h_K(\alpha_M)]^T$ and the parameter vector is $\Psi = [a_1, \dots, a_N, \theta_1, \dots, \theta_N]^T$, so that $P = 2N$. It follows from (18) that the joint p.d.f. for the support function data is

$$p_h(\mathbf{y}; \Psi) = \prod_{m=1}^M \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[\frac{-1}{2\sigma^2} \left(h_K(\alpha_m) - \sum_{k=1}^N d_{t_m k} a_k \sin(\alpha_m - \theta_k) \right)^2 \right]. \quad (30)$$

The FIM will be a $2N \times 2N$ matrix consisting of four $N \times N$ blocks $\mathbf{J}_1, \mathbf{J}_2, \mathbf{J}_3$, and \mathbf{J}_4 . Specifically, for $i, j = 1, \dots, N$, we have

$$\begin{aligned} \mathbf{J}_1(i, j) &= \mathbf{J}(i, j) = -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial a_i \partial a_j} \right) \\ &= \frac{1}{\sigma^2} \sum_{m=1}^M d_{t_m i} d_{t_m j} s_{m,i} s_{m,j}, \end{aligned} \quad (31)$$

$$\begin{aligned} \mathbf{J}_2(i, j) &= \mathbf{J}(i, j + N) = -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial a_i \partial \theta_j} \right) \\ &= -\frac{a_j}{\sigma^2} \sum_{m=1}^M d_{t_m i} d_{t_m j} s_{m,i} c_{m,j}, \end{aligned} \quad (32)$$

$$\begin{aligned} \mathbf{J}_3(i, j) &= \mathbf{J}(i + N, j) = -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial \theta_i \partial a_j} \right) \\ &= -\frac{a_i}{\sigma^2} \sum_{m=1}^M d_{t_m i} d_{t_m j} s_{m,j} c_{m,i}, \end{aligned} \quad (33)$$

and

$$\begin{aligned} \mathbf{J}_4(i, j) &= \mathbf{J}(i + N, j + N) = -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial \theta_i \partial \theta_j} \right) \\ &= \frac{a_i a_j}{\sigma^2} \sum_{m=1}^M d_{t_m i} d_{t_m j} c_{m,i} c_{m,j}, \end{aligned} \quad (34)$$

where $c_{m,k} = \cos(\alpha_m - \theta_k)$ and $s_{m,k} = \sin(\alpha_m - \theta_k)$. For details, see [20]. Thus \mathbf{J} is a function of the EGI vector parameters and the set of measurement directions.

For diameter functions, we need to calculate the CRLB for only half the number of parameters, due to the origin symmetry of the reconstruction; the other parameters are then defined by (25) and (26). Thus the measurement vector is $\mathbf{y} = [b_K(\alpha_1), \dots, b_K(\alpha_M)]^T$ and the parameter vector is

$\Psi = [a_1, \dots, a_{N/2}, \theta_1, \dots, \theta_{N/2}]^T$, so that $P = N$. Also, (22) can be written (compare (23)) in the form

$$\mathbf{y} = 2\mathbf{C}(\Theta)\mathbf{a} + \mathbf{n}, \quad (35)$$

where now $\mathbf{a} = [a_1, \dots, a_{N/2}]$ and $\Theta = [\theta_1, \dots, \theta_{N/2}]$. By (35), the joint p.d.f. for the diameter function data is

$$p_b(\mathbf{y}; \Psi) = \prod_{m=1}^M \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[\frac{-1}{2\sigma^2} \left(b_K(\alpha_m) - \sum_{k=1}^{N/2} a_k |\cos(\alpha_m - \theta_k)| \right)^2 \right] \quad (36)$$

and the FIM is an $N \times N$ matrix consisting of four $N/2 \times N/2$ blocks we will again label \mathbf{J}_1 , \mathbf{J}_2 , \mathbf{J}_3 , and \mathbf{J}_4 . Specifically, for $i, j = 1, \dots, N/2$, we have

$$\begin{aligned} \mathbf{J}_1(i, j) = \mathbf{J}(i, j) &= -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial a_i \partial a_j} \right) \\ &= \frac{1}{\sigma^2} \sum_{m=1}^M |c_{m,i}| |c_{m,j}|, \end{aligned} \quad (37)$$

$$\begin{aligned} \mathbf{J}_2(i, j) = \mathbf{J}(i, j + N/2) &= -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial a_i \partial \theta_j} \right) \\ &= \frac{a_j}{\sigma^2} \sum_{m=1}^M |c_{m,i}| s_{m,j} \operatorname{sgn}(c_{m,j}), \end{aligned} \quad (38)$$

$$\begin{aligned} \mathbf{J}_3(i, j) = \mathbf{J}(i + N/2, j) &= -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial \theta_i \partial a_j} \right) \\ &= \frac{a_i}{\sigma^2} \sum_{m=1}^M |c_{m,j}| s_{m,i} \operatorname{sgn}(c_{m,i}), \end{aligned} \quad (39)$$

and

$$\begin{aligned} \mathbf{J}_4(i, j) = \mathbf{J}(i + N/2, j + N/2) \\ &= -E \left(\frac{\partial^2 \ln p(\mathbf{y}; \Psi)}{\partial \theta_i \partial \theta_j} \right) \\ &= \frac{a_i a_j}{\sigma^2} \sum_{m=1}^M s_{m,i} s_{m,j} \operatorname{sgn}(c_{m,i}) \operatorname{sgn}(c_{m,j}), \end{aligned} \quad (40)$$

where $c_{m,k} = \cos(\alpha_m - \theta_k)$, $s_{m,k} = \sin(\alpha_m - \theta_k)$, and $\operatorname{sgn}(f)$ returns +1 or -1 according to whether f takes positive or negative values, respectively. For details, see [20].

For diameter function measurements, the lower bound on the variance of the estimated shape parameters for any input polygon can be obtained from

(29) and the inverse of the FIM from (37)–(40). However, for support function measurements, the FIM from (31)–(34) is singular with rank $(2N - 2)$, yielding an infinite lower bound. The reason for this behavior is that the equality constraint (5) has not been taken into account. For diameter function measurements, only the inequality constraint (4) comes into play, in view of the formulation in (35), where (25) and (26) guarantee that (5) holds. Such inequality constraints induce open subsets in the parameter space without any isolated boundaries and so do not affect the CRLB; see [8] for more details. On the other hand, equality constraints can reduce the dimension of the parameter space and cannot be ignored in the calculations. The next section addresses this issue.

4.2. The Constrained Cramér-Rao Lower Bound

In this section we use the results of Stoica and Ng [26] for a *constrained Cramér-Rao lower bound* (CCRLB) analysis with a singular FIM.

Let $\mathbf{g}_i(\Psi) = 0$, $i = 1, \dots, Q$, where $\Psi \in \Re^P$ and $Q < P$, be equality constraints defined by continuously differentiable functions, and let $\mathbf{g}(\Psi) = [\mathbf{g}_1(\Psi), \dots, \mathbf{g}_Q(\Psi)]^T$ be the corresponding constraint vector. The gradient matrix of \mathbf{g} is a $Q \times P$ matrix defined by

$$\mathbf{G}(\Psi) = \frac{\partial \mathbf{g}(\Psi)}{\partial \Psi^T},$$

where \mathbf{G} should have full rank. Let \mathbf{U} be a $P \times (P - Q)$ matrix with orthogonal columns in the null space of $\mathbf{G}(\Psi)$. Then $\mathbf{G}(\Psi)\mathbf{U} = 0$. The analysis in [26] shows that, under the given constraints, the covariance matrix of the estimated parameters satisfies

$$\operatorname{Cov}(\hat{\Psi}) \geq \mathbf{U}(\mathbf{U}^T \mathbf{J} \mathbf{U})^{-1} \mathbf{U}^T. \quad (41)$$

As seen from (41), the lower bound is obtained by projecting the singular FIM onto the constraint subspace.

We are working with two equality constraints given by (5), and the gradient matrix of these constraints is the $2 \times 2N$ matrix

$$\mathbf{G}(\Psi) = \begin{bmatrix} \cos \theta_1 & \cdots & \cos \theta_N & -a_1 \sin \theta_1 & \cdots & -a_N \sin \theta_N \\ \sin \theta_1 & \cdots & \sin \theta_N & a_1 \cos \theta_1 & \cdots & a_N \cos \theta_N \end{bmatrix}. \quad (42)$$

After carrying out the analysis discussed in this section, we will be able to obtain the CCRLB for the shape parameters. However, our final aim is to reconstruct the 2-D shape from support-type function data and not simply to find its EGI. Thus what we really need is to find bounds with respect to the 2-D shape and not just the parameters. Rather than the quality of estimates of Ψ , we need a global quality measure of the entire reconstructed boundary. This can be achieved by developing asymptotic confidence regions around the true shape, a topic introduced in the next section.

4.3. Confidence Regions

4.3.1. Background. Asymptotic global confidence regions are used to analyze and visualize the performance of 2-D parametric shape estimators. Assuming a maximum likelihood estimator (MLE) operating in the asymptotic regime, the CRLB for the shape parameters can be used to define a confidence region around the true boundary of the shape, as in [27]. Note that the MLE is asymptotically normal, unbiased, and asymptotically attains the CRLB; see [12, pp. 164–167]. In our analysis, we follow [27] (but using the CCRLB instead of the CRLB for support function measurements).

Consider a 2-D shape parameterization of the form

$$\mathbf{s}(t; \Psi) = [s_x(t; \Psi), s_y(t; \Psi)]^T,$$

where s_x and s_y give the x and y coordinates of a point \mathbf{s} on the object's boundary, indexed by $t \in [0, T]$, and where Ψ is a P -dimensional parameter vector. At each point along the boundary (i.e., for all $t \in [0, T]$) we determine a local confidence region $U_\beta(t) \subset \mathbb{R}^2$ centered at the true point $\mathbf{s}(t) = \mathbf{s}(t; \Psi)$ (see Fig. 7). The size of the local confidence region depends on β , which in turn depends on the chosen *local confidence level* $\eta \in [0, 1]$; if $\hat{\mathbf{s}}(t)$ is estimated using an MLE operating in the asymptotic regime, then

$$\Pr\{\hat{\mathbf{s}}(t) \in U_\beta(t)\} = \eta.$$

The first step towards generating the local confidence region for a point $\mathbf{s}(t; \Psi)$ is to calculate the $P \times P$ covariance matrix $\mathbf{C}_\Psi = \text{Cov}(\Psi)$. From this we calculate the 2×2 matrix

$$\mathbf{C}_s(t) = \nabla_\Psi \mathbf{s}(t; \Psi) \mathbf{C}_\Psi [\nabla_\Psi \mathbf{s}(t; \Psi)]^T, \quad (43)$$




Figure 7. Local and global confidence regions.

where the $2 \times P$ matrix $\nabla_\Psi \mathbf{s}(t; \Psi)$ is the gradient of $\mathbf{s}(t; \Psi)$ with respect to Ψ . The matrix $\mathbf{C}_s(t)$ gives the individual variances and the covariance between the x and y coordinates of $\mathbf{s}(t)$ when estimated using the MLE operating in the asymptotic regime. Next, a local confidence level $\eta \in [0, 1]$ is selected. The *local confidence region* $U_\beta(t)$ for the point $\mathbf{s}(t)$ is

$$U_\beta(t) = \{\mathbf{x} \in \mathbb{R}^2 : (\mathbf{x} - \mathbf{s}(t))^T \mathbf{C}_s(t)^{-1} (\mathbf{x} - \mathbf{s}(t)) \leq \beta^2\}, \quad (44)$$

where $\beta > 0$ is calculated by assuming that the left-hand side of the inequality in (44) is a Chi-square random variable of degree 2 such that the probability that it is less than or equal to β^2 is η . For each t , $U_\beta(t)$ is an ellipse centered at $\mathbf{s}(t)$, also referred to as a *local confidence ellipse*, whose size depends on η . An asymptotic *global confidence region* U_β given by

$$U_\beta = \bigcup_{t \in [0, T]} U_\beta(t) \quad (45)$$

can now be obtained by moving $U_\beta(t)$ along the boundary. This global confidence region (see Fig. 7) defines an uncertainty band around the entire boundary of the true shape, in which the estimated shape can lie.

4.3.2. Local Confidence Regions. Once the EGI of the approximating polygon has been obtained from the support-type function measurements, we may use (11) to obtain

$$\mathbf{s}(t; \Psi) = \left[\begin{array}{c} \sum_{k=1}^N d_{tk} a_k \cos(\theta_k + \pi/2), \\ \sum_{k=1}^N d_{tk} a_k \sin(\theta_k + \pi/2) \end{array} \right]^T, \quad (46)$$

where the discrete values $t = 1, \dots, N$ of t each index a vertex of the polygon. From this, we find the $2 \times 2N$

gradient matrix $\nabla_{\Psi} \mathbf{s}(t; \Psi)$ to be

$$\begin{bmatrix} d_{t1} \cos(\theta_1 + \pi/2) & \cdots & d_{tN} \cos(\theta_N + \pi/2) & -d_{t1} a_1 \sin(\theta_1 + \pi/2) & \cdots & -d_{tN} a_N \sin(\theta_N + \pi/2) \\ d_{t1} \sin(\theta_1 + \pi/2) & \cdots & d_{tN} \sin(\theta_N + \pi/2) & d_{t1} a_1 \cos(\theta_1 + \pi/2) & \cdots & d_{tN} a_N \cos(\theta_N + \pi/2) \end{bmatrix}. \quad (47)$$

Using (43), (44), and (47), the asymptotic local confidence ellipses $U_{\beta}(t)$ can be calculated for $t = 1, \dots, N$.

In the case of support function measurements, the matrix \mathbf{C}_{Ψ} in (43) is estimated from the CCRLB given by the right-hand side of (41). For diameter function measurements, we can instead use the CRLB given by the right-hand side of (29), but before doing so the $N \times N$ FIM obtained from (37)–(40) must be inverted and replicated to form a $2N \times 2N$ matrix (see [20] for more details).

Figure 8 shows the local confidence ellipses for a polygon with $N = 7$ sides whose support function is measured at $M = 180$ equally spaced angles. (Here and below, the entire interval, $[0, 2\pi]$ in this case, is divided into M bins of equal size and measurement vectors are positioned at the center of each bin.) The noise is Gaussian with power $\sigma = 0.2$ and the user-defined local confidence level is $\eta = 0.73$. From (12) we know that each measurement at an angle α_m lying between adjacent outer normal angles θ_t and θ_{t+1} will depend only on the position of the vertex of the polygon between these two angles. Therefore the greater the size of the interval $[\theta_t, \theta_{t+1}]$, the more information we have about the corresponding vertex of the polygon and the smaller the local confidence ellipse will be. This is illustrated by Fig. 8.

Figure 9 depicts the local confidence ellipses for an origin-symmetric, regular polygon with $N = 12$ sides whose diameter function is measured from $M = 36$ equally spaced viewing angles in the range $[0, \pi]$. The




Figure 8. Local confidence regions for shape from support function.

noise power is $\sigma = 0.1$ and the local confidence level is $\eta = 0.73$. As expected from the symmetry of the

polygon and measurement angles, the local confidence ellipses are all congruent. (Note that had the vertices of the output polygon been specified by (7) instead of (11), the size of the local confidence ellipses would have increased with the index t ; see [20, 21].)

Decreasing the number of measurements provides less information about the shape and leads to worse shape estimates. This is illustrated by Fig. 10, obtained using the same shape and parameters as in Fig. 8 except that now $M = 72$ instead of 180. We observe that the confidence ellipses in Fig. 10 are larger in size compared to Fig. 8, illustrating the increased uncertainty. An opposite effect is observed by increasing the number of measurements.

In Fig. 11, the underlying polygon and parameters are as in Fig. 9, but the $M = 1000$ measurement angles are randomly drawn from the von Mises distribution $M(\pi/3, 5)$ (see [19]) with mean $\pi/3$ and concentra-




Figure 9. Local confidence regions for shape from diameter function.




Figure 10. Effect of the number of support function measurements.




Figure 11. Measurements angles sampled from a von Mises distribution.




Figure 12. Local confidence regions (equally spaced measurement directions).

tion parameter $\kappa = 5$. The high value of κ leads to a very heavy concentration of measurement angles near $\pi/3$. A diameter function measurement provides the greatest information about the vertices of the polygon contained in the support lines parallel to the measurement direction. This explains the very different sizes of the local confidence ellipses in Fig. 11, which are largest at the vertices that lie in support lines parallel to the relatively few measurement directions orthogonal to $\pi/3$.

Figure 12 is obtained using an affinely regular polygon with $N = 8$ sides whose support function is measured at $M = 72$ equally spaced angles, with noise power $\sigma = 0.1$. The local confidence ellipses vary in size, the largest occurring at vertices 1 and 5 where the “curvature” is smallest. This is compatible with earlier remarks. If the user has prior information about the eccentricity of the input polygon, the distribution of the measurement angles can be changed accordingly to compensate; more measurements should be taken in directions corresponding to low curvature. This is demonstrated in Fig. 13, where a total of 72 measurements are made in 12 equally spaced directions each in the intervals $[0, \pi/6)$, $[\pi/6, 5\pi/6)$ and $[5\pi/6, \pi)$,




Figure 13. Local confidence regions with a better set of viewing directions.

and the 36 antipodal directions. The local confidence ellipses have become more nearly equal in size.

4.3.3. Global Confidence Regions. In the previous section, local confidence regions were constructed only at the vertices of the input polygon. This restriction, a consequence of (46) being valid only for discrete values of t , means that the formula (45) for global confidence regions is not immediately accessible.

To deal with this difficulty, we define

$$\mathbf{s}(t; \Psi) = (\lceil t \rceil - t)\mathbf{s}(\lfloor t \rfloor; \Psi) + (t - \lfloor t \rfloor)\mathbf{s}(\lceil t \rceil; \Psi), \quad (48)$$

for all $1 \leq t \leq N + 1$, where, as usual, $\lfloor t \rfloor$ and $\lceil t \rceil$ denote the greatest integer less than or equal to t and the least integer greater than t , respectively, and where $t = 1$ and $t = N + 1$ both index the first vertex of the polygon. From (46) and (48) we obtain

$$\mathbf{s}(t; \Psi) = \begin{bmatrix} \sum_{k=1}^N d'_{tk} a_k \cos(\theta_k + \pi/2), \\ \sum_{k=1}^N d'_{tk} a_k \sin(\theta_k + \pi/2) \end{bmatrix}^T, \quad (49)$$

for $1 \leq t \leq N + 1$, where

$$d'_{tk} = (\lceil t \rceil - t)d_{\lfloor t \rfloor k} + (t - \lfloor t \rfloor)d_{\lceil t \rceil k}.$$

Using (49), we can obtain the local confidence region $U_\beta(t)$ for any point on the boundary of the polygon. A global confidence region U_β is then defined by (45).

Figure 14 shows the global confidence region obtained in the manner just described for the same polygon and parameters as in Fig. 8. Local confidence ellipses for the polygon were generated for t varying




Figure 14. Global confidence region for shape from support function.




Figure 15. Global confidence region for shape from diameter function.

from 1 to 9 in steps of 0.005, and a sample of these ellipses is also depicted. Similarly, Fig. 15 shows the global confidence region for the polygon and parameters used in Fig. 9, with t sampled at steps of 0.005, along with the local confidence regions at the vertices.

5. Performance and Experimental Results

The optimal algorithm is the MLE operating in the asymptotic regime, since it is normal, unbiased, and attains the CRLB. Though the algorithms we have considered above are not necessarily unbiased, their performance can be measured effectively by comparing them with the MLE operating in the asymptotic regime, using the confidence regions developed earlier for this algorithm. The performance analysis can be carried out with the aid of both local and global confidence regions. For local confidence regions we check if each vertex of the estimated polygon is contained in the corresponding local confidence ellipse; see [20] for an analysis.

Here we present a performance analysis for the algorithms from Section 3, using global confidence regions. Let \hat{s} be the estimated shape boundary resulting from a given algorithm and define the corresponding error

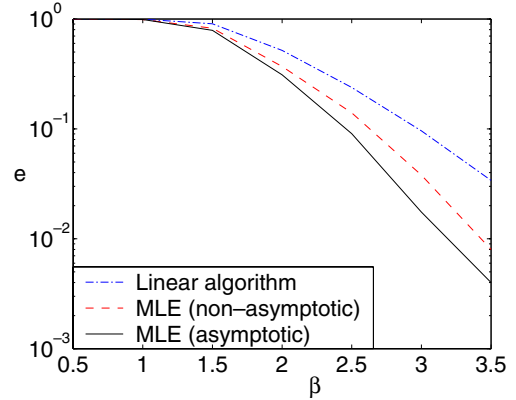


Figure 16. Performance evaluation of diameter function algorithm.

probability $e = e(\beta)$ by

$$e = 1 - \Pr\{\hat{s} \in U_\beta\}. \quad (50)$$

Thus e is the probability that the estimated shape boundary \hat{s} does *not* lie completely inside the global confidence region U_β . By Monte Carlo simulation, we can compare performances of algorithms by plotting for each the error probability for a range of values of β . We use a log scale on the vertical axis for better illustration.

Figure 16 provides a performance evaluation of the diameter function algorithms, based on 1000 runs of each algorithm for a regular polygon with $N = 10$ sides, $M = 45$ equally spaced measurement angles, and noise power $\sigma = 0.04$. The graphs give the relative performance of the linear diameter function algorithm given in [5], the non-linear algorithm based on (24) (MLE (non-asymptotic)), and the MLE (asymptotic). The latter is produced as follows. Since the MLE is asymptotically normal, unbiased, and attains the CRLB, we first draw an EGI sample from the normal distribution $N(\Psi, \mathbf{C}_\Psi)$ where the matrix \mathbf{C}_Ψ is as described in Section 4.3.2 for diameter function measurements. This sample will always satisfy (5) because of the constraints imposed on \mathbf{C}_Ψ . However, some of the edge lengths may be negative, in which case the sample is discarded and a new sample is drawn. From the EGI sample we obtain the estimated boundary \hat{s} by (46), and the value of e in (50) is estimated from 1000 independent trials of this type.

Of course, in using the linear diameter function algorithm instead of the non-linear one, the user pays a price in terms of the quality of the estimates. This is visible in Fig. 16. However, this figure also indicates that the linear diameter function algorithm does a rea-




Figure 17. Underlying polygon used for the analysis in Fig. 18.




Figure 19. Shape from support function measurements.




Figure 18. Performance evaluation of support function algorithms.




Figure 20. True, noisy, and estimated support function measurements.

sonably good job when compared to the optimal MLE estimator.

Figure 18 provides a similar evaluation of the support function algorithms, based on 1000 runs of each algorithm for a polygon with $N = 5$ sides, $M = 72$ equally spaced measurement angles, and noise power $\sigma = 0.1$. The underlying polygon has barycenter at $\mathbf{0}$ as shown in Fig. 17. The graphs in Fig. 18 give the relative performance of the algorithms of Prince and Willsky (PW), Lele, Kulkarni, and Willsky (LKW), the non-linear algorithm based on (21) (MLE (non-asymptotic)), and the MLE (asymptotic). However, we stress that the comparison is not really a fair one, since on the one hand the PW algorithm produces outputs with outer normals in the measurement directions, and on the other, the LKW algorithm assumes, as prior knowledge, *unknown* outer normal angles. In view of this it is hardly surprising that the PW algorithm performs much worse, and the LKW algorithm much better, than the MLE algorithms.

Figures 19 and 21 illustrate true (solid line) and estimated (dotted line) polygons resulting from the linear support function algorithm based on (15) and the

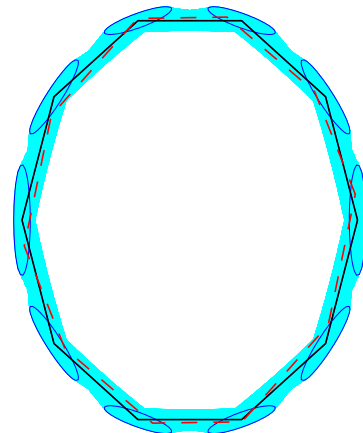


Figure 21. Shape from diameter function measurements.

linear diameter function algorithm from [5]. In each case the noise power is $\sigma = 0.1$ and local and global confidence regions corresponding to local confidence level $\eta = 0.9$ are shown for comparison. For Fig. 19 the true polygon has $N = 8$ sides with support function measured at $M = 72$ equally spaced angles. The true polygon in Fig. 21 is an affinely regular poly-




Figure 22. True, noisy, and estimated diameter function measurements.

gon with $N = 10$ sides whose diameter function is measured at $M = 36$ equally spaced angles; the extra post-processing step of decimation described in Section 3.3 was used to prune the number of edges down to 10. Figs. 20 and 22 depict the true, noisy, and estimated measurements in each case.

6. Conclusion and Future Work

In this paper we discussed non-linear and linear algorithms for estimating a planar convex shape from finitely many noisy measurements of its support or diameter function. We carried out a systematic statistical analysis of these algorithms via computation of the Cramér-Rao lower bound on the estimated parameters. Using this bound, we found local and global confidence regions for the underlying shape, providing a valuable visual demonstration of the effect of experimental parameters on the quality of estimated shape. Finally, we presented a performance analysis of the linear algorithm for shape from diameter functions that indicates this algorithm operates reasonably well.

Acknowledgment

We thank Markus Kiderlen for some very helpful remarks.

References

1. R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*, Wiley-Interscience: New York, 1973.
2. N.I. Fisher, P. Hall, B. Turlach, and G.S. Watson, "On the estimation of a convex set from noisy data on its support function," *J. Amer. Statist. Assoc.*, Vol. 92, No. 437, pp. 84–91, 1997.
3. R.J. Gardner, *Geometric Tomography*, Cambridge University Press: New York, 1995.
4. R.J. Gardner, M. Kiderlen, and P. Milanfar, "Convergence of algorithms for reconstructing convex bodies and directional measures," *Ann. Statist.*, to appear.
5. R.J. Gardner and P. Milanfar, "Shape reconstruction from brightness functions," in *SPIE Conference on Advanced Signal Processing Algorithms, Architectures, and Implementations X*, Vol. 4474, August 2001, San Diego, CA.
6. R.J. Gardner and P. Milanfar, "Reconstruction of convex bodies from brightness functions," *Discrete Comput. Geom.*, Vol. 29, pp. 279–303, 2003.
7. G.H. Golub and V. Pereyra, "The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate," *SIAM J. Numer. Anal.*, Vol. 10, pp. 413–432, 1973.
8. J.D. Gorman and A.O. Hero, "Lower bounds for parametric estimation with constraints," *IEEE Trans. Inform. Theory*, Vol. 26, No. 6, pp. 1285–1301, 1990.
9. P. Hall and B. Turlach, "On the estimation of a convex set with corners," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 21, No. 3, pp. 225–234, 1999.
10. B.K.P. Horn, *Robot Vision*, The MIT Press: Cambridge, MA, 1986.
11. V. Hosel and S. Walcher, "Clustering techniques: A brief survey," Technical Report, Institute of Biomathematics and Biometry, 2001.
12. S. Kay, *Fundamentals of Statistical Signal Processing*, Prentice-Hall: Englewood Cliffs, NJ, 1993.
13. A.S. Lele, S.R. Kulkarni, and A.S. Willsky, "Convex polygon estimation from support line measurements and applications to target reconstruction from laser radar data," *J. Optical Soc. Amer. A*, Vol. 9, pp. 1693–1714, 1992.
14. S.Y.R. Li, "Reconstruction of polygons from projections," *Inform. Process. Lett.*, Vol. 28, pp. 235–240, 1988.
15. M. Lindenbaum and A. Bruckstein, "Blind Approximation of Planar Convex Sets," *IEEE Trans. Robotics and Automation*, Vol. 10, No. 4, pp. 517–529, 1994.
16. J.J. Little, "An iterative method for reconstructing convex polyhedra from extended Gaussian images," in *Proc. AAAI National Conf. Artificial Intelligence*, Washington, D.C., 1983, pp. 247–250.
17. S.K. Lodha and R. Franke, "Scattered data techniques for surfaces," in *Proc. Dagstuhl Conf. on Scientific Visualization*, IEEE Computer Society Press, 1999, pp. 182–222.
18. E. Mammen, J.S. Marron, B.A. Turlach, and M.P. Wand, "A general projection framework for constrained smoothing," *Statist. Sci.*, Vol. 16, No. 3, pp. 232–248, 2001.
19. K. V. Mardia and P. E. Jupp, *Directional Statistics*, John Wiley and Sons: Chichester, 2000.
20. A. Poonawala, "Reconstructing shapes from support and brightness functions," Masters Thesis, Computer Engineering Department, University of California Santa Cruz, 2004.
21. A. Poonawala, P. Milanfar, and R.J. Gardner, "A statistical analysis of shape reconstruction from areas of shadows," in *36th Asilomar Conference on Signal Systems and Computers*, Vol. 1, pp. 916–920, November 2002.
22. J.L. Prince and A.S. Willsky, "Estimating convex sets from noisy support line measurements," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 12, No. 4, pp. 377–389, 1990.
23. A.S. Rao and K.Y. Goldberg, "Shape from diameter: Recognizing polygonal parts with a parallel-jaw gripper," *Intl. J. of Robotics Research*, Vol. 13, No. 1, pp. 1637, 1994.

24. T.J. Richardson, "Approximation of planar convex sets from hyperplane probes," *Discrete Comput. Geom.*, Vol. 18, pp. 151–177, 1997.
25. R. Schneider, *Convex Bodies: The Brunn-Minkowski Theory*, Cambridge University Press: Cambridge, 1993.
26. P. Stoica and B.C. Ng, "On the Cramér-Rao bound under parametric constraints," *IEEE Signal Processing Letters*, Vol. 5, No. 7, pp. 177–179, 1998.
27. J. C. Ye, Y. Bresler, and P. Moulin, "Asymptotic global confidence regions in parametric shape estimation problems," *IEEE Trans. Inform. Theory*, Vol. 46, No. 5, pp. 1881–1895, 2000.
28. J.C. Ye, Y. Bresler, and P. Moulin, "Cramér-Rao bounds for parametric shape estimation in inverse problems," *IEEE Trans. on Image Processing*, Vol. 12, No. 1, pp. 71–83, 2003.
29. P. Zhingeng, Z. Kun, and S. Jiaoying, "A new mesh simplification algorithm based on triangle collapses," *J. Comput. Sci. Tech.*, Vol. 16, No. 1, pp. 57–63, 2001.



Amyn Poonawala received the B.E. degree from the University of Mumbai, India, in 2001, and the M.S. degree from the University of California, Santa Cruz (UCSC), in 2004, both in computer engineering. He is currently pursuing the Ph.D. degree in computer engineering at UCSC. His technical interests include statistical signal and image processing and inverse problems in microlithography.



Peyman Milanfar received the B.S. degree in electrical engineering/mathematics from the University of California, Berkeley, in

1988, and the S.M., E.E., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, in 1990, 1992, and 1993, respectively. Until 1999, he was a Senior Research Engineer at SRI International, Menlo Park, CA. He is currently Associate Professor of Electrical Engineering at the University of California, Santa Cruz. He was a Consulting Assistant Professor of computer science at Stanford University from 1998–2000, and a visiting Associate Professor there in 2002. His technical interests are in statistical signal and image processing, and inverse problems. He won a National Science Foundation CAREER award in 2000, was associate editor for the *IEEE Signal Processing Letters* from 1998 to 2001, and is a Senior member of the IEEE.



Richard Gardner holds B.Sc. and Ph.D. degrees in mathematics from University College London and was awarded a D.Sc. degree from the University of London in 1988 for contributions to measure theory and convex geometry. He has held positions at universities and research institutions in several countries and has been Professor of Mathematics at Western Washington University since 1991. He founded geometric tomography, an area of geometric inverse problems involving data concerning sections by and projections on lines or planes, and published a book on the subject in 1995.