Shape Reconstruction From Brightness Functions

Richard J. Gardner[a] and Peyman Milanfar[b]

[a]Mathematics Department
Western Washington University, Bellingham, WA 98225
[b]Electrical Engineering Department
University of California, Santa Cruz, CA 95064

(Invited paper appearing in the Proceedings of SPIE Conference 4474, Advanced Signal Processing Algorithms, Architectures, and Implementations XI, San Diego, CA, July 29 - August 3, 2001)

ABSTRACT

In this paper we address the problem of reconstructing the shape of a convex object from measurements of the *areas* of its shadows in several directions. This type of very weak measurement is sometime referred to as the *brightness* function of the object, and may be observed in an imaging scenario by recording the total number of pixels where the object's image appears. Related measurements, collected as a function of viewing angle, are also referred to as "lightcurves" in the astrophysics community, and are employed in estimating the shape of atmosphereless rotating bodies (e.g. asteroids).

We address the problem of shape reconstruction from brightness functions by constructing a least-squares optimization framework for approximating the underlying shapes with polygons in two dimensions, or polyhedra in three dimensions, from *noisy*, and possibly *sparse* measurements of the brightness values.

Keywords: Inverse problem, shape, polygon, polyhedron, brightness, curvature, lightcurve.

1. INTRODUCTION

The problem addressed here is that of reconstruction of a planar or 3-dimensional shape from noisy, and possibly sparse, measurements of its brightness function. In the planar case, the brightness function gives the lengths of the orthogonal projections (i.e., shadows) of the shape onto lines; in three dimensions, it gives the areas of the projections of the shape onto planes*. The shapes considered are bounded by a simple closed curve or surface. An approximation to the shape is effected by reconstructing a polygon or polyhedron that best matches the measured brightness data in a least-squares sense.

It is worth noting at the outset that data consisting of brightness measurements is quite weak. For example, such data contains no information about the location of the shape, since any translate of it will have the same brightness function. Also, it is not possible to detect holes or dents in a planar shape; indeed the shape's brightness function will agree with that of its convex hull. For this reason, we shall restrict attention to convex bodies, that is, compact convex sets with nonempty interiors.

One example of this problem in an imaging scenario is the case of a severely ill-resolved object in the far field of a camera. More specifically, imagine an object with lambertian surface that is so poorly resolved as to have its entire image fall on a single pixel at any given time. In this case, the available measurements simply consist of the time history of the tracked pixel intensity. In particular, if the object is rotating, then this time history is proportional at any given time to orthogonal projection of the object in the direction in which it is exposed to the camera. The measurement of this time history for an atmosphereless body is effectively what astronomers refer to as the "lightcurve" of the object. The question then is how to approximate the shape of the object from such weak data.

Direct correspondence to milanfar@ee.ucsc.edu. This work was supported in part by National Science Foundation Grants DMS-9802388 and CCR-9984246.

*For ease of presentation and due to limited space, the discussion of experimental results will be given only for the 2-dimensional case.

The problem of reconstruction from brightness functions is quite an old one in the mathematical area of convex geometry. It is also an important unsolved problem in *geometric tomography* (see[3]), the area of mathematics concerning the retrieval of information about a geometric object from data concerning its sections or projections. (Here again, and throughout the paper, the term *projection* will be used in the usual mathematical sense of orthogonal projection or shadow, and not, as in the computer tomography literature, for what we would prefer to call an X-ray.) It has long been known that the brightness function is even too weak to distinguish between convex bodies with centroid at the origin. If a convex body is not symmetric, then its reflection in the origin is a different body with the same brightness function. The lack of uniqueness is actually much more serious. The existence of nonspherical convex bodies of constant brightness (i.e., whose brightness functions have the same value in every direction) shows that even the unit ball is not determined among all convex bodies by its brightness function! We shall have more to say about this nonuniqueness in Section 3.

In the planar case, the brightness function is just a phase-shifted version of the more familiar width function that measures the distance between calipers holding the convex body (see Section 8 for more details); the reconstruction problem from either type of measurement describes one and the same problem. (In the 3-dimensional case, however, there is no such simple relationship between the two functions.) The width function is in turn closely related to another well-known function, the support function, that measures the distance from a fixed point inside the convex body to its tangent lines (or tangent planes, as appropriate). The remarks concerning nonuniqueness above apply equally to the width function. (The first examples of the more familiar nonspherical convex bodies of constant width were discovered by Euler!) However, it is known that any convex body is determined by its support function, so support function measurements constitute much stronger data. In fact, reconstruction from (noisy) support measurements has already been studied fairly extensively; see, for example,[5] ,[9] and.[10]

2. APPROACH TO SHAPE PARAMETERIZATION

There exist many ways to parameterize a shape, of which the Cartesian coordinate representation is only one. A particularly convenient and useful representation is the so-called *extended circular image* (in the 3-dimensional case called the *extended Gaussian image* in the computer vision literature). In this representation, the shape is encoded in terms of its curvature as a function of the normal direction to the boundary. We seek to reconstruct a convex body by first reconstructing its extended Gaussian image from the raw brightness data, and then obtaining from this a more direct representation of the object such as its support function or Cartesian representation.

For planar bodies with twice continuously differentiable boundaries, the radius of curvature $f(u)$ at the point on the boundary with outer unit normal vector u is well defined. In higher dimensions, similar smoothness assumptions guarantee the existence for all u of the Gauss curvature of the body at the point on the boundary with outer normal u, and we may consistently define $f(u)$ to be the reciprocal of this Gauss curvature. Then there is a relationship between $f(u)$ (which is just the extended circular or Gaussian image mentioned above, and also called the *curvature function* in convex geometry) and the brightness function $b(v)$, where the unit vector v denotes the "viewing direction." This is described by a well-known integral transformation, which we call the *geometric cosine transform.* Specifically, we have

$$b(v) = \frac{1}{2} \int |u^T v| f(u) du, \tag{1}$$

where du, here and throughout the paper, denotes integration over the unit sphere with respect to its natural measure.

The formula (1) requires some further explanation. Note that the term $|u^T v|$ is the absolute value of the cosine of the angle between the unit normal u to the boundary and the viewing direction vector v, hence the term geometric *cosine* transform. (Since the term "cosine transform" in signal and image processing has a different meaning, we use the adjective "geometric" to avoid confusion.) In general, the function f essentially defines the surface area (or arclength) function for the body. In fact, by replacing f by a suitable measure called the *surface area measure* of the body, (1) can be extended to any convex body, without any smoothness assumptions. If, for example, the body is a polygon, the function f must be replaced by the measure in the unit circle that consists of a finite sum of point masses, each placed at the outer unit normal to an edge and with weight equal to the length of that edge. The resulting formula ((29) below) is a sum whose geometric meaning is transparent, and was in fact known to Cauchy. See[3] for the general formula.

The inverse problem of interest is to solve a two-step problem:

- Solve the integral equation (1) to obtain values $f(u)$ from the "measurements" $b(v)$.
- From the computed values $f(u)$, find an explicit representation of the planar or 3-dimensional convex body.

It is interesting that the first step has also been considered by Kiderlen[6] in determining the directional distribution of fiber processes.

3. UNIQUENESS ISSUES

We noted in the introduction that if K is a convex body, its reflection $-K$ in the origin has the same brightness function. This geometrically obvious fact follows from (1), since with $w = -u$ we have

$$\begin{aligned} b_{-K}(v) &= \frac{1}{2}\int |u^T v| f_{-K}(u) du \\ &= \frac{1}{2}\int |u^T v| f_K(-u) du \\ &= \frac{1}{2}\int |w^T v| f_K(w) dw = b_K(v), \end{aligned} \tag{2}$$

where b_K, f_K and b_{-K}, f_{-K} are the brightness functions and extended Gaussian images of K and $-K$, respectively. To progress beyond this simple observation, the following theorem is helpful, and indeed plays a crucial role in our problem.

PROPOSITION 1. (Minkowski's existence theorem.) *A function f on the unit sphere is the extended Gaussian image of some convex body if and only if the support of f is not contained in a hyperplane and*

$$\int u f(u) du = 0. \tag{3}$$

See,[3] and for a proof see.[14] We stress that technically f should in general be a measure rather than a function, specifically, a nonnegative finite Borel measure in the unit sphere. The hypothesis in Proposition 1 that the support of f is not contained in a hyperplane through the origin merely ensures that the corresponding convex body is not degenerate (i.e., "flat").

By Proposition 1, for any $0 \le t \le 1$ there is a convex body K_t whose extended Gaussian image f_{K_t} is given by

$$f_{K_t}(u) = (1-t) f_K(u) + t f_{-K}(u). \tag{4}$$

The existence of K_t follows directly from the fact that since both f_K and f_{-K} satisfy the hypotheses of Proposition 1, f_{K_t} does also. Now K_t has the same brightness function as K, since by (1) and (2),

$$\begin{aligned} b_{K_t}(v) &= \frac{1}{2}\int |u^T v| f_{K_t}(u) du \\ &= \frac{1-t}{2}\int |u^T v| f_K(u) du + \frac{t}{2}\int |u^T v| f_{-K}(u) du \\ &= (1-t) b_K(v) + t b_{-K}(v) = b_K(v). \end{aligned}$$

An inequality called the *Kneser-Süss inequality* (see[14]) implies that if $-K$ is not a translate of K, then the volumes of the convex bodies K_t are strictly increasing for $0 \le t \le 1/2$ (and strictly decreasing for $1/2 \le t \le 1$), and hence there will be infinitely many different (noncongruent) convex bodies whose brightness function equals that of K. Any body such that K is a translate of $-K$ is called *centrally symmetric.*

To summarize, there is a serious lack of uniqueness in the inverse problem if the unknown convex body is not centrally symmetric with center at the origin, that is, *origin symmetric.* Moreover, most origin-symmetric convex bodies are also not determined, among all convex bodies, by their brightness function. We noted in the introduction that even the unit ball is not so determined (see[3] for a picture of a nonspherical convex body of constant brightness, discovered by Blaschke). It can actually be shown that parallelotopes (in two and three dimensions, these are the parallelograms and parallelepipeds) with their centers at the origin are the only convex bodies that are determined in this sense; see[3] and the references given there.

What is true is that any origin-symmetric convex body is determined among all *origin-symmetric* convex bodies by its brightness function. This is a consequence of *Aleksandrov's projection theorem* (see[3]). More generally, one can say that given any convex body K, there is a unique origin-symmetric convex body with the same brightness function. It turns out that this origin-symmetric body, often called the *Blaschke body* and denoted by ∇K, is just $K_{1/2}$, the member of the family defined above corresponding to $t = 1/2$. By,[3] the Blaschke body ∇K has another special property: It has the maximum volume of all convex bodies with the same brightness function as K.

Thus uniqueness is restored in the inverse problem if it is known a priori that the shape is an origin-symmetric convex body.

If we assume that the centroid of the convex body is fixed, at the origin, say, then all the nonuniqueness in the two-step procedure outlined at the end of Section 2 resides in the first step. This is because *Aleksandrov's uniqueness theorem*[3] implies that any two convex bodies with the same extended Gaussian image must be identical, up to translation. This also means that the convex body whose existence is guaranteed by Proposition 1 is unique, up to translation.

4. POLYGONS AND POLYHEDRA

Henceforth we shall use the generic name *polytope* to mean a polygon or polyhedron as appropriate, and *volume* to mean length, area, or volume according to the dimension of the object.

The basic approach taken here to the shape reconstruction problem is to approximate the underlying shape by a polytope, which we can reconstruct using the measured brightness values. If we assume that the inverse problem will be solved for a polytope, we can take $f(u)$ to be a measure consisting of a finite sum of point masses, or, equivalently, to have the form of a distribution described in terms of Dirac delta functions. Specifically, for a convex polytope with N facets, we have

$$f(u) = \sum_{k=1}^{N} a_k \, \delta(u - u_k), \tag{5}$$

where a_k denotes the volume of the kth facet and u_k denotes the outer unit normal to the kth facet. (The term *facet* applies to an edge in two dimensions and a face in three dimensions.) Substituting the above expression into (1), we get

$$b(v) = \frac{1}{2} \sum_{k=1}^{N} a_k \, |u_k^T v|. \tag{6}$$

We close this section with a result that will be important for one of our algorithms. Suppose that V is a finite set of unit vectors. The hyperplanes through the origin orthogonal to the vectors in V divide the space into a finite set of polyhedral cones, which intersect the unit sphere in a finite set of regions. Let $U = U(V)$ be the finite set of vertices of these regions, and call these points the *nodes* corresponding to V. The following result was proved by Campi, Colesanti, and Gronchi.[2]

PROPOSITION 2. *Let K be a convex body and let V be a finite set of unit vectors. Among all convex bodies that have the same brightness function as K at the directions in V, there is a unique convex body that has maximal volume. Moreover, this body is an origin-symmetric convex polytope with each of its facets orthogonal to a node in U.*

5. THE ESTIMATION PROBLEM

With equation (1) in place, we are in a position to describe the computational inverse problem we are going to solve. We begin by measuring the function $b(v)$ at M locations v_m in the unit sphere, where the measurements may be corrupted by random error. By (1), we can describe the data by

$$\tilde{b}(v_m) = \frac{1}{2} \int |u^T v| f(u) du \; + \; n(v_m), \tag{7}$$

where $n(v_m)$ denotes the random measurement error, $m = 1, \ldots, M$. If it is known that the object is a polytope whose extended Gaussian image is given by (5), then we can replace (7) by

$$\tilde{b}(v_m) = \frac{1}{2} \sum_{k=1}^{N} a_k \, |u_k^T v_m| \; + \; n(v_m), \tag{8}$$

for $m = 1, \ldots, M$. The problem can now be viewed as an estimation problem where we seek best estimates of the unknown parameters a_k and u_k from the measured data $\tilde{b}(v_m)$.

We shall consider two versions of this problem. The first is when we know that the convex body to be reconstructed is a polytope with a prescribed number N of facets, and the second is when we know nothing a priori about the convex body.

For the first version of the problem, we shall assume that $M \geq N$. Collecting the data samples into vector form, we can write

$$\tilde{b} = C(U)\, a \, + \, n, \tag{9}$$

where

$$\begin{aligned}
\tilde{b} &= \left[\tilde{b}(v_1), \ldots, \tilde{b}(v_M)\right]^T, & (10)\\
a &= \left[a_1, \ldots, a_N\right]^T, & (11)\\
U &= \left[u_1, \ldots, u_N\right]^T, & (12)\\
n &= \left[n(v_1), \ldots, n(v_M)\right]^T, \text{ and} & (13)\\
C_{m,k} &= |u_k^T v_m|/2. & (14)
\end{aligned}$$

With this notation, we seek the least squares estimate of the unknowns as

$$(\widehat{a}, \widehat{U}) = \arg \min_{(a,U)} \; \|\tilde{b} \, - \, C(U)a\|^2. \tag{15}$$

As an estimation or optimization problem, (15) is nonlinear in the unknown parameters u_k. Without the absolute value in (8), the problem would be quite similar to those encountered in the antenna array processing community,[7] and well-known methods employing generalized eigenvalue techniques could be used to solve first for the u_k, and subsequently for the a_k. As it is, however, the nonlinearity must be dealt with directly.

For the second version of our problem, where nothing is known a priori about the convex body K to be reconstructed, we attempt to reconstruct the origin-symmetric polytope, whose existence is guaranteed by Proposition 2, with the same brightness function as K in the directions v_m, $m = 1, \ldots, M$. In this version of the problem, therefore, we seek

$$\widehat{a} = \arg \min_{a} \; \|\tilde{b} \, - \, C(U)a\|^2, \tag{16}$$

where $U = \{u_k : k = 1, \ldots, N\}$ is now the *known* set of nodes (defined before Proposition 2) corresponding to the vectors v_m, $m = 1, \ldots, M$. Notice that the nonlinear part of the least squares problem is no longer necessary.

6. GEOMETRIC CONSTRAINTS

Since we wish to output a convex polytope, we need to impose some geometric constraints on the minimization problems (15) and (16). While such constraints tend to make the solution of the optimization problems more complex, they are quite useful in limiting the effect of noise, essentially *regularizing* the problem.

Recall that a is a vector containing the volumes a_k of the facets of a polytope with outer normals u_k, $k = 1, \ldots, N$. Since $a_k \geq 0$, we obtain the linear inequality constraint

$$a \geq 0. \tag{17}$$

A second linear constraint is obtained by observing that by Proposition 1, convex polytopes must satisfy (3). In view of (5), this becomes the vector equation

$$a^T U = \sum_{k=1}^{N} a_k u_k = 0. \tag{18}$$

Thus in the n-dimensional case we have altogether $N + n$ real linear constraints.

If a priori information is available that the unknown convex polytope is origin symmetric, then N is even and to ensure an origin-symmetric output we can replace (18) by the stronger combination of the two constraints

$$a_{N/2+k} = a_k \tag{19}$$

and

$$u_{N/2+k} = -u_k \tag{20}$$

for $k = 1, \ldots, N/2$.

7. THE RECONSTRUCTION ALGORITHMS

Algorithm 1 (number of facets prescribed)

If a polytope P with a prescribed number N facets is to be reconstructed, we may as well also assume that its centroid is at the origin. Our reconstruction algorithm takes as input a natural number N, unit vectors v_m, $m = 1, \ldots, M$, and noisy brightness measurements $\tilde{b}_m$, $m = 1, \ldots, M$ of P at these directions. It then proceeds as follows:

Step 1. Choose an initial guess, a convex polytope P_0 with N facets and centroid at the origin.

Step 2. Solve the constrained optimization problem defined by (15), subject to constraints (17) and (18) (or (17), (19), and (20) if an origin-symmetric output is desired). This yields an optimal $\widehat{a}$ and $\widehat{U}$ specifying facet areas and outer normals of an optimal convex polytope $\widehat{P}$ with N facets and centroid at the origin; in short, the extended Gaussian image of $\widehat{P}$.

Step 3. Reconstruct $\widehat{P}$ from its extended Gaussian image.

Algorithm 2 (number of facets not prescribed)

If nothing is known a priori about the convex body to be reconstructed, our reconstruction algorithm takes as input unit vectors v_m, $m = 1, \ldots, M$ and noisy brightness measurements $\tilde{b}_m$, $m = 1, \ldots, M$ of the body at these directions. It then proceeds as follows:

Step 0. Calculate the nodes u_k, $k = 1, \ldots, N$.

Step 1. Choose an initial guess, a convex polytope P_0 with centroid at the origin and with N facets whose outer unit normal vectors are the nodes.

Step 2. Solve the constrained optimization problem defined by (16), subject to constraints (17) and (18) (or (17), (19), and (20) if an origin-symmetric output is desired). This yields an optimal $\widehat{a}$ specifying facet areas of an optimal convex polytope $\widehat{P}$ with N facets and centroid at the origin. The extended Gaussian image of $\widehat{P}$ is $(\widehat{a}_k, u_k)$, $k = 1, \ldots, N$.

Step 3. Reconstruct $\widehat{P}$ from its extended Gaussian image.

In Step 1, one possible systematic method of selecting the initial guess can be based on *Cauchy's surface area formula*, which states that

$$S = \frac{1}{\kappa_{d-1}} \int b(v) dv, \tag{21}$$

where S denotes the surface area of a d-dimensional convex body and κ_d is the the volume of the d-dimensional unit ball; see, for example.[3] Together with the *isoperimetric inequality* (see[3])

$$V \leq \kappa_d \left(\frac{S}{d\kappa_d} \right)^{d/(d-1)}, \tag{22}$$

where V is the volume of the body, this gives an upper bound for the volume in terms of the brightness function. The initial guess P_0 can be chosen to satisfy this volume estimate or the surface area estimate (21).

Steps 2 and 3 will be discussed for the planar and 3-dimensional cases in the following sections.

8. RECONSTRUCTION OF PLANAR SHAPES

We begin by supplying further information concerning the brightness, width, and support functions of a planar convex body. Suppose the body is bounded by a curve $\mathcal{C}$ parameterized as $(x(\alpha),\, y(\alpha))$, by the angle α which the normal to the curve makes with the positive x-axis. Since the data we consider are unaffected by translating the shape, we may for convenience choose the origin to be interior to $\mathcal{C}$. The *support function* $h(\alpha)$ of $\mathcal{C}$ (or the body bounded by it) is then the perpendicular distance from the origin to the tangent to $\mathcal{C}$ at $(x(\alpha),\, y(\alpha))$. The support function is a periodic function with period 2π. Under suitable smoothness assumptions, we can write the parameterization of the curve in terms of $h(\alpha)$ and its first derivative as

$$x(\alpha) = h(\alpha)\cos\alpha - h'(\alpha)\sin\alpha \tag{23}$$

$$y(\alpha) = h(\alpha)\sin\alpha + h'(\alpha)\cos\alpha. \tag{24}$$

(See[3] or.[1]) The *width function* $w(\alpha)$ of $\mathcal{C}$ (or the body bounded by it) is the distance between the two tangent lines to $\mathcal{C}$ that are perpendicular to α. This can be written as

$$w(\alpha) = h(\alpha) + h(\alpha + \pi), \tag{25}$$

which describes a periodic function with period π.

Since the brightness function $b(\alpha)$ is the length of the "shadow" of the body in the direction of the normal at α, it is just a phase-shifted version of the width function:

$$b(\alpha) = w(\alpha + \pi/2), \tag{26}$$

which again defines a periodic function with period π. It is known that any convex body is uniquely determined by its support function (see,[1] ,[14] or[16]).

In the planar case, the expression (1) for the brightness function becomes

$$b(\alpha) = \frac{1}{2}\int_0^{2\pi} |\cos(\alpha-\theta)| f(\theta)d\theta, \tag{27}$$

where f is the extended circular image of the convex body. For a convex polygon the latter takes the form

$$f(\alpha) = \sum_{k=1}^{N} a_k\, \delta(\alpha - \theta_k), \tag{28}$$

where θ_k denotes the angle of the normal to the kth edge. The brightness function of the same polygon is then

$$b(\alpha) = \frac{1}{2}\sum_{k=1}^{N} a_k\, |\cos(\alpha - \theta_k)|. \tag{29}$$

For example, the square with vertices at $(\pm 1, \pm 1)$ has $a_k = 2$ and $\theta_k = (k-1)\pi/2$ for $k = 1, 2, 3, 4$, so that

$$b(\alpha) = 2\left(|\cos\alpha| + |\sin\alpha|\right),$$

which is easy to verify directly.

In the formulas above, $0 \le \alpha < 2\pi$, but from the point of view of taking measurements, we can assume that $0 \le \alpha < \pi$, since the brightness function is even.

We now discuss the reconstruction algorithms outlined in Section 7 as they are presently implemented.

We can write each input unit vector $v_m = (\cos\alpha_m, \sin\alpha_m)^T$, where $0 \le \alpha_m < \pi$.

In Step 0 of Algorithm 2, the nodes are simply the $2M$ unit vectors in directions $\alpha_m \pm \pi/2$, $m = 1, \ldots, M$; in this case we set $N = 2M$ and list the nodes as u_k, $k = 1, \ldots, N$.

The initial guess in Step 1 does not make use of the surface area and volume estimates given in Section 7; apart from having N edges and centroid at the origin, the initial convex polygon is chosen more or less arbitrarily. (It would also be possible to refine this initial guess by looking at the largest and smallest measured values of the brightness function to estimate the eccentricity and direction of elongation of the shape.)

In Step 2, we write each unit vector $u_k = (\cos\theta_k, \sin\theta_k)^T$, where $0 \le \theta_k < 2\pi$, $k = 1, \ldots, N$, and replace (12) by

$$\Theta = [\theta_1, \ldots, \theta_N]^T. \tag{30}$$

With corresponding changes in notation, the basic constrained optimization problem in Algorithm 1 becomes

$$\begin{aligned} (\widehat{a}, \widehat{\Theta}) \quad &= \quad \arg\min_{(a,\Theta)} \|\tilde{b} - C(\Theta)a\|^2, \\ \text{such that} \quad & \quad a \ge 0, \\ & \quad a^T[\cos\theta_1, \ldots, \cos\theta_N]^T = 0, \\ \text{and} \quad & \quad a^T[\sin\theta_1, \ldots, \sin\theta_N]^T = 0. \end{aligned} \tag{31}$$

The obvious modification is made if Algorithm 2 is to be used. We use the function `fmincon` from MATLAB's optimization toolbox to solve this problem. This function employs a variety of nonlinear optimization techniques such as line search or Newton's method.

In Step 3 we reconstruct the output polygon $\widehat{P}$ from its extended circular image, a list of lengths $\widehat{a}_k$ and outer normal angles $\widehat{\theta}_k$ of its edges, ordered so that these angles increase with k. This is easily effected by finding vertices inductively, as follows. Let w_0 be the origin and $w_k = w_{k-1} + a_1(\cos(\theta_k + \pi/2), \sin(\theta_k + \pi/2))$ for $k = 1, \ldots, N$. If it is desired that the centroid of $\widehat{P}$ is at the origin, it is easy to compute the appropriate translation.

9. RECONSTRUCTION OF THREE-DIMENSIONAL SHAPES

Write the input unit vectors in spherical polar coordinates as

$$v_m = (\cos\alpha_m \sin\beta_m, \sin\alpha_m \sin\beta_m, \cos\beta_m)^T,$$

where $0 \le \alpha_m < 2\pi$ and $0 \le \beta_m < \pi/2$, $m = 1, \ldots, M$.

In Step 0 of Algorithm 2, the nodes are the $M(M-1)$ unit vectors at the intersections of the M great circles in the unit sphere orthogonal to v_m, $m = 1, \ldots, M$. In this case we set $N = M(M-1)$ and list the nodes as u_k, $k = 1, \ldots, N$.

For Step 1 of Algorithm 1, suppose first that N is odd. We generate a random set of unit vectors $u_k^{(0)}$, $k = 1, \ldots, N$ and random real numbers $a_k^{(0)}$, $k = 1, \ldots, N-3$, between some preassigned $\varepsilon > 0$ and 1. These vectors and numbers are substituted for the corresponding u_k's and a_k's in the constraint (18), which represents three scalar equations, and these are solved for the remaining three variables to yield $a_{N-2}^{(0)}$, $a_{N-1}^{(0)}$, and $a_N^{(0)}$. If all these three numbers are between ε and 1, the solution is accepted; otherwise, the process is repeated. The resulting vectors and numbers represent the outer unit normals and facet areas of an initial guess convex polyhedron P_0 with N facets. If N is even, Step 1 is easier; we can either generate randomly or specify a set of unit vectors $u_k^{(0)}$, $k = 1, \ldots, N/2$, let $u_{N/2+k}^{(0)} = -u_k^{(0)}$ for $k = 1, \ldots, N/2$, and let $a_k^{(0)} = 1/N$, $k = 1, \ldots, N$. For Step 1 of Algorithm 2, we can simply define $a_k^{(0)} = 1/N$, $k = 1, \ldots, N$, since N is even, and the known vectors u_k satisfy (20).

In Step 2, we write

$$u_k = (\cos\theta_k \sin\varphi_k, \sin\theta_k \sin\varphi_k, \cos\varphi_k)^T,$$

where $0 \le \theta_k < 2\pi$ and $0 \le \varphi_k \le \pi$, $k = 1, \ldots, N$. Then the data (8) becomes

$$b(\alpha_m, \beta_m) = \frac{1}{2}\sum_{k=1}^{N} a_k |\sin\beta_m \sin\varphi_k \cos(\alpha_m - \theta_k) + \cos\beta_m \cos\varphi_k| + n(\alpha_m, \beta_m), \tag{32}$$

for $m = 1, \ldots, M$. The least-squares problem is adapted accordingly, with constraint (17) as before and constraint (18) becoming

$$\sum_{k=1}^{N} a_k \cos\theta_k \sin\varphi_k = 0,$$

$$\sum_{k=1}^{N} a_k \sin\theta_k \sin\varphi_k = 0,$$

and

$$\sum_{k=1}^{N} a_k \cos\varphi_k = 0.$$

In the symmetric case, constraints (19) and (20) are used instead.

The main extra difficulty encountered in the 3-dimensional situation lies in Step 3, namely, to reconstruct the output convex polyhedron $\widehat{P}$ from its extended Gaussian image, the areas $\widehat{a}_k$ and outer normals $\widehat{u}_k$ of its faces. Luckily, this is a problem of interest in its own right, and as such has been addressed by several authors.

In principle, the problem appears to have been solved first by Little.[12] Little's algorithm was modified and refined by Lemordant, Tao, and Zouaki[11] and by Kaasalainen, Lamberg, Lumme, and Boswell (see[8] and the references given there). The first group of authors were motivated by computer vision and the second group by astrophysics. A quite different solution was offered by Sumbatyan and Troyan[15] for the purposes of reconstruction of a cavity from ultrasound; the algorithm presented is based on solving nonlinear P.D.E.'s and is not of interest to us here. None of these three groups appear to have been aware of the work of the other two. More recent work includes that of Xu and Suk,[17] who embellish Little's algorithm in order to handle nonconvex polygons and polyhedra.

We briefly describe Little's algorithm, which is based on Minkowski's original proof of Proposition 1 (for convex polytopes). The algorithm takes as input pairs (a_k, u_k), $k = 1, \ldots, N$, and attempts to reconstruct a convex polyhedron P with centroid at the origin whose faces have areas a_k and outer normals u_k. By Proposition 1, it is only necessary that the vectors u_k span 3-dimensional space and that

$$\sum_{k=1}^{n} a_k u_k = 0. \tag{33}$$

If $l = [l_1, \ldots, l_N]^T$ is a vector of nonnegative real numbers, let $V(l)$ denote the volume of the polyhedron $P(l)$ whose faces have outer normals u_k and are contained in planes at distances l_k from the origin, $k = 1, \ldots, N$. (We have $l_k = h(u_k)$, where h is the support function of $P(l)$.) Minkowski's proof shows that the vector of distances from the origin to the planes containing the faces of P is the solution of the optimization problem

$$\text{minimize} \qquad a^T l = \sum_{k=1}^{N} a_k l_k, \tag{34}$$

$$\text{such that} \qquad V(l) = 1 \tag{35}$$

$$\text{and} \qquad l \geq 0. \tag{36}$$

The algorithm begins with $l = [1, \ldots, 1]^T$ and a reduced gradient method is used in the above optimization problem to improve on this initial guess. If l is the current estimate, the convex polyhedron $P(l)$ is only known by its $\mathcal{H}$-*representation*, that is, the equations of the planes containing its faces. The $\mathcal{V}$-*representation* of $P(l)$, a list of its vertices, can then be computed by a clever use of polar duality. After this, the centroid of $P(l)$ is computed and $P(l)$ is translated so that its centroid is at the origin. In view of the constraint (35), the volume $V(l)$ is also computed, and $P(l)$ is scaled by a factor $V(l)^{-1/3}$ so that its volume becomes one. The objective function $a^T l$ is then evaluated and the process stops if the improvement is less than a prescribed value. Finally, the output polyhedron must be reconstructed from the optimal vector $\widehat{l}$.

We employ the modified version of Little's algorithm due to Lemordant, Tao, and Zouaki.[11] The crucial improvement uses duality in programming theory to replace the above optimization problem by an equivalent one:

$$\text{maximize} \qquad V(l)^{1/3}, \tag{37}$$

$$\text{such that} \qquad a^T l = 1 \tag{38}$$

$$\text{and} \qquad l \geq 0. \tag{39}$$

The exponent 1/3 has been inserted (it does not alter the solution, of course) because using the well-known Brunn-Minkowski inequality from convex geometry (see, for example,[3] ,[14] or[16]), one can show that $V(l)^{1/3}$ is a concave function. Thus the new optimization problem involves a concave objective function and linear constraints. The basic modified algorithm solves this problem by the reduced gradient or Newton's method to produce an optimal vector $\widehat{l}$. As in Little's algorithm, this involves the computation of $V(l)$ at each stage, but not the computation of the centroid of $P(l)$. Also, the volume computation is done by Laguerre's algorithm (see[11]), which requires only the $\mathcal{H}$-representation and not the $\mathcal{V}$-representation. In fact it is only necessary to calculate a $\mathcal{V}$-representation once, for the reconstruction of the output polyhedron from $\widehat{l}$.

To summarize, the algorithm from[11] has three basic ingredients: (1) maximizing a concave function subject to linear constraints, (2) calculating the volume of a convex polyhedron from its $\mathcal{H}$-representation, and (3) reconstructing a convex polyhedron from its $\mathcal{H}$-representation. The latter can be broken into two parts: (3a) converting the $\mathcal{H}$-representation of a convex polyhedron to its $\mathcal{V}$-representation and (3b) computing the convex hull of a set of points in three dimensions.

10. NUMERICAL EXPERIMENTS

In this section we present some numerical examples of the performance of the 2-dimensional algorithm described in Section 8. We stress that although in this case measurements of an origin-symmetric body are equivalent to support measurements, our algorithms are equally applicable in the 3-dimensional situation.

10.1. Example 1

We display the reconstruction of the affinely regular 14-gon shown in the lower part of Figure 1. We simulated 100 measurements of $b(\alpha)$ at equally-spaced values in the range $[0, \pi]$. These values were then corrupted with Gaussian white noise of variance 0.2 before being fed to the algorithm.

The input affinely regular 14-gon is shown in the botton picture, along with the initial guess, and the output polygon produced by Algorithm 1 (with the symmetry constraints). The reconstruction and true polygons are nearly indistinguishable. The upper plot of Figure 1 shows the brightness functions for these polygons as well as the noisy measured brightness function. Also recorded is the Hausdorff error, that is, the Hausdorff distance (see, for example,[3]) between the input and output polygons.

10.2. Example 2

In this example, we reconstruct an affinely regular octagon from noisy measurements of its brightness function. We input 20 measurements of the brightness function, considerably fewer than in the previous example, corrupted with Gaussian white noise of a greater variance, 0.5. Again, Algorithm 1 with symmetry constraints was used, with the result is depicted in Figure 2. The output polygon matches the input polygon quite closely, but it is a hexagon; this illustrates the possibility of degenerate edges in the output.

11. CONCLUSIONS AND PROGRAM OF FUTURE WORK

It would appear that the Algorithm 1 does a reasonable job of reconstructing planar convex bodies from quite noisy brightness data. Algorithm 2 has yet to be implemented, but this is straightforward and will be done very soon.

Even in the planar case, several issues remain.

1. It may be possible to find conditions under which the algorithms are provably convergent.

2. In the origin-symmetric case, the results should be compared with earlier work on reconstruction from noisy support line measurements.

3. The least squares problem (15) is of a special type known as *separable*. In view of this, there is the possibility of improving performance by a variant of the Newton method, where each of the unknowns a and U is treated separately in each iteration. See for example, the article by Golub and Pereyra.[4]

4. We could study the role of convexity in our problem. What can be done when it is removed?

5. In Algorithm 1, we can study the "underdetermined" problem when $M < N$.

Turning to the 3-dimensional case, we have begun improving our implementation of the reconstruction from the extended Gaussian image, described in the previous section, so that it compares the extended Gaussian image of the output polyhedron with that of the input. This facility should be available shortly. The 3-dimensional versions

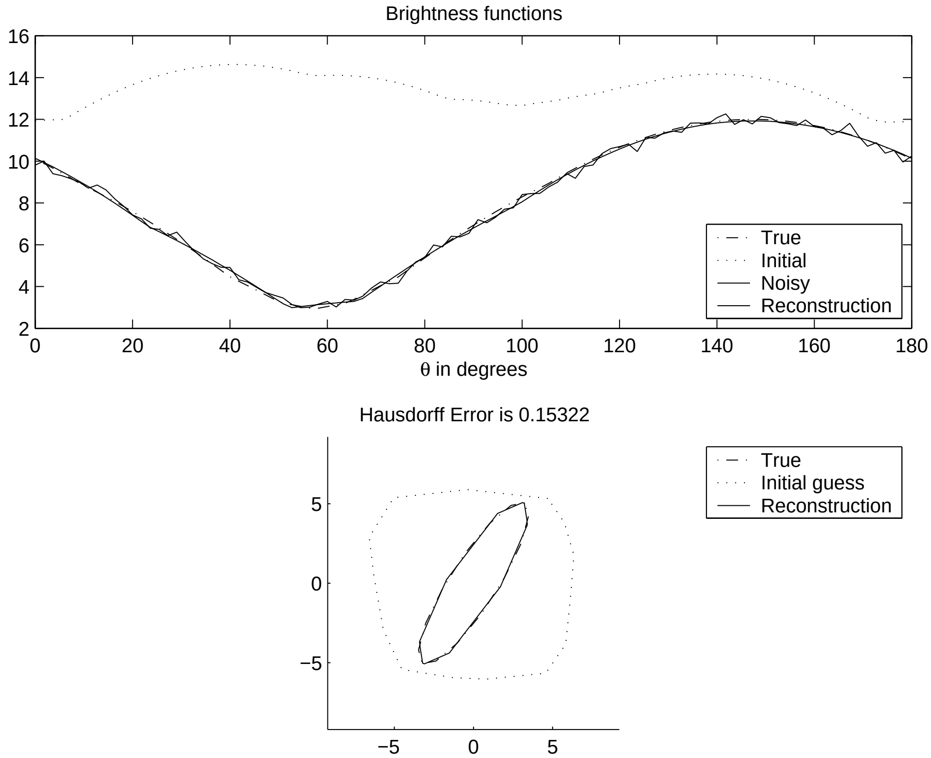

Figure 1. Reconstruction of the 14-gon of Example 1

of Algorithms 1 and 2 need to be implemented and hooked up. We anticipate doing this in the very near future; all the ingredients are in place and described in detail above. It is important to note here that in Step 3 we shall be dealing with the reconstruction of a convex polyhedron from *noisy* values of the Gaussian image, a feature that seem to have been ignored or only lightly touched upon by previous authors. Thus we hope to establish, perhaps for the first time, a practical methodology for the treatment of a problem that is bound to be plagued by numerical and measurement inaccuracies in practice. Also, it is worth mentioning that the problem of shape reconstruction from local curvature measurements is of interest in the robotics community where recently developed tactile sensors (e.g.[13]) enable a robotic hand to “grasp” or “touch” an object and make local measurements of shape.

REFERENCES

1. T. Bonnesen and W. Fenchel. *Theory of Convex Bodies.* BCS Associates, Moscow, Idaho, U.S.A., 1987. German original: Springer, Berlin, 1934.
2. S. Campi, A. Colesanti, and P. Gronchi. Convex bodies with extremal volumes having prescribed brightness in finitely many directions. *Geom. Dedicata*, 57:121–133, 1995.
3. R. J. Gardner. *Geometric Tomography.* Cambridge University Press, New York, 1995.
4. G. H. Golub and V. Pereyra. The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM J. Numer. Anal.*, 10:413–432, 1973.
5. W. C. Karl. *Reconstructing Objects from Projections.* PhD dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1991.
6. M. Kiderlen. Non-parametric estimation of the directional distribution of stationary line and fibre processes. *Adv. Appl. Prob.*, 33:6–24, 2001.
7. H. Krim and M. Viberg. Two decades of array signal processing. *IEEE Signal Processing Magazine*, 13:67–94, 1996.
8. L. Lamberg. On the Minkowski problem and the lightcurve operator. *Ann. Acad. Sci. Fennicae,* Math. Diss. vol. 87, 1993.

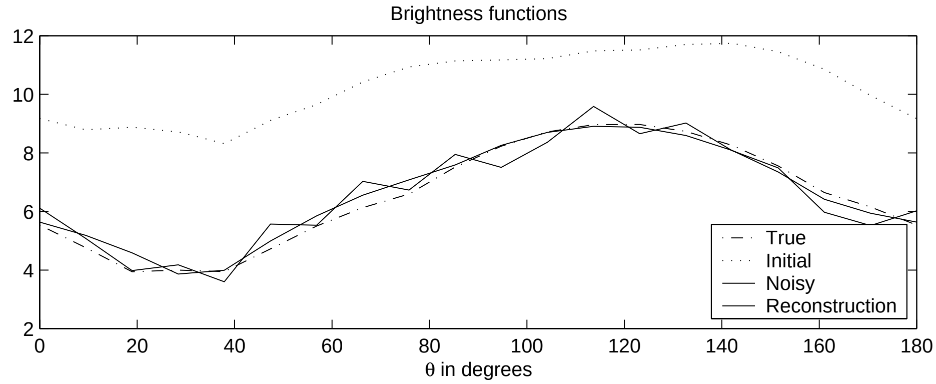

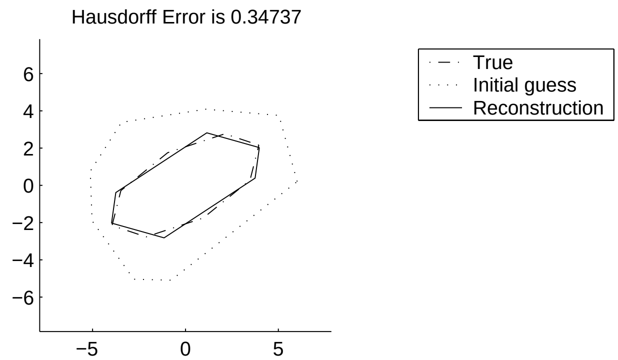

Figure 2. Reconstruction of the octagon of Example 2

9. A. S. Lele. *Convex Set Reconstruction from Support Line Measurements and Its Application to Laser Radar Data.* Msc dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1991.
10. 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. America A*, 9:1693–1714, 1992.
11. J. Lemordant, P. D. Tao, and H. Zouaki. Modélisation et optimisation numérique pour la reconstruction d'un polyèdre à partir de son image gaussienne généralisée. *RAIRO Modél. Math. Anal. Numér.*, 27:349–374, 1993.
12. 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*, pages 247–250, 1983.
13. E. Nicolson and R. Fearing. The reliability of curvature estimates from linear elastic tactile sensors. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Nagoya, Japan, May 1995.
14. R. Schneider. *Convex Bodies: The Brunn-Minkowski Theory.* Cambridge University Press, Cambridge, 1993.
15. M. A. Sumbatyan and E. A. Troyan. Reconstruction of the shape of a convex defect from a scattered wave field in the ray approximation. *J. Appl. Maths. Mechs.*, 56:464–468, 1992.
16. R. Webster. *Convexity.* Oxford University Press, Oxford, 1994.
17. Zhijing Xu and Minsoo Suk. Representation and reconstruction of polygons and polyhedra using hierarchical extended gaussian images. *Ann. Math. and Artificial Intelligence*, 13:377–399, 1995.