The Support Function, Curvature Functions and 3-D Attitude Determination

Ying Li Robert J. Woodham

Technical Report 92-32
November 1992

Laboratory for Computational Intelligence
Department of Computer Science
The University of British Columbia
Vancouver, B. C. V6T 1Z2
Canada

Abstract

Attitude determination finds the rotation between the coordinate system of a known object and that of a sensed portion of its surface. Orientation-based representations record 3-D surface properties as a function of position on the unit sphere. They are useful in attitude determination because they rotate in the same way as the object rotates. Three such representations are defined using, respectively, the support function and the first and second curvature functions. The curvature representations are unique for smooth, strictly convex objects. The support function representation is unique for any convex object.

The essential mathematical basis for these representations is provided. The paper extends previous results on convex polyhedra to the domain of smooth, strictly convex surfaces. Combinations of the support function of a known object with curvature measurements from a visible surface transform attitude determination into an optimization problem for which standard numerical solutions exist.

Dense measurements of surface curvature are required. Surface data can be obtained from laser range finding or from shape-from-shading methods, including photometric stereo. A proof-of-concept system has been implemented and experiments conducted on a real object using surface orientation and surface curvature data obtained directly from photometric stereo.

1 Introduction

Shape representations are required to support both recognition and localization tasks. Recognition identifies the object. Localization determines the three translational and the three rotational degrees of freedom of the object in space. Localization is required for many robot vision tasks including directing a robot arm to grasp an object, navigation and camera calibration. Attitude determination solves for the three rotational degrees of freedom between the coordinate system of a known object and that of a viewer. Thus, attitude determination is a sub-problem of localization.

One approach useful both for recognition and for attitude determination is to record 3-D surface properties as a function of position on the unit sphere. These representations are termed *orientation-based* because one associates each point on the sphere with the unit vector from the center of the sphere to that point. Orientation-based representations are a compact description of 3-D object shape. Orientation-based representations have the desirable property that the object and the representation rotate together. This makes them ideal candidates for the task of attitude determination.

An orientation-based representation must define the mapping between surface points and points on the sphere. The standard way is to use the *Gauss map*. The Gauss map takes each surface point to the point on the sphere corresponding to the normal to the tangent plane at that point. The Gauss map is unique for smooth (i.e., C^2), strictly convex objects[1]. Representations used in computer vision include: the Extended Gaussian Image (EGI), defined as the reciprocal of the Gaussian curvature [1], and the support function, defined as distance from an origin to the tangent plane [2]. For polyhedra, the EGI specifies the area of each face as a function of face orientation. The EGI has been used for both recognition and attitude determination of polyhedra [3, 4, 5]. The support function appears explicitly in one of the methods described [5].

Mathematics defines other representations of 3-D shape based on the Gauss map. For example, the first and second curvature functions are defined, respectively, as the sum of the principal radii of curvature and the product of the principal radii of curvature[2]. These curvature functions possess desirable mathematical properties when combined with the support function and are the curvature measures used here[3].

Several other local curvature measures can easily be defined as orientation-based representations including the Gaussian and the mean curvature, popularized by Besl and Jain [6, 7], and Koenderink's curvedness and shape index [8]. It has proven difficult to extend representations based on the Gauss map beyond the convex case since, in general, the Gauss map is many-to-one. Approaches have been described to decompose non-convex surfaces into regions for which the Gauss map is unique and to augment the information recorded to handle the many-to-one nature of the mapping [9, 10]. Alternatively, new orientation-based

[1] A point set is strictly convex if, given any two points in the set, the open line segment between the points lies in the interior of the set. A convex C^2 surface is strictly convex if and only if its Gaussian curvature is everywhere positive.

[2] Thus, for smooth objects, the EGI and the second curvature function are equivalent.

[3] For polyhedra, mathematics defines first and second area functions. These two area functions are analogous to the two curvature functions defined for smooth strictly convex surfaces. In particular, for polyhedra, the EGI is equivalent to the second area function.

representations can be defined by choosing a different way to establish the mapping between surface point and point on the sphere. Attitude determination has been demonstrated for a more general class of "starshaped" objects using the radial function and the dilation map [11].

The novel contribution here is to demonstrate 3-D shape matching for C^2 strictly convex surfaces. This demonstration consists of theoretical foundations, algorithm development and experimental proof-of-concept using real objects and surface data obtained from an existing photometric stereo system.

Section 2 formalizes the necessary mathematical results concerning the support function, curvature functions and inequalities between mixed volumes. The mixed volumes used combine the support function of one object with a curvature function of another. Section 3 defines attitude determination in terms of mixed volumes and shows how the problem can be transformed into an optimization problem for which standard numerical solutions exist. Section 4 describes the implementation and presents some experimental results. To be effective, dense estimates of surface shape are required. Dense surface data can be obtained from laser range finding or from shape-from-shading methods, including photometric stereo. Photometric stereo is particularly well-suited to the task since it provides robust local estimates of both surface orientation and surface curvature, without the need to explicitly determine depth [12]. The experiments test numerical solutions for three cases: 1) attitude determination when both the model and sensed surface are given in known analytic form; 2) attitude determination when the sensed surface, then the model and then both are discretized versions of a known analytic form; and 3) attitude determination for a real object with surface orientation and curvature data obtained directly from photometric stereo and model data given in known analytic form. Cases 1 and 2 represent simulation studies that were essential to software development, error analysis and tests of robustness. They are not reported in detail here. Instead, Section 4 describes case 3 results on a real object. Section 5 summarizes the findings.

2 The Support Function and Curvature Functions

It is convenient to define the support function first for arbitrary points, $v \in R^3$, and then to specialize the definition to points, u, on the (unit) sphere.

Definition 2.1 Let $C \subset R^3$ be a nonempty bounded set. The *support function* $H(C;v)$ of C is the real-valued function defined by

$$H(C;v) = \sup\{\langle x, v\rangle \mid x \in C,\ v \in R^3\}\ .$$

The support function of a set, C, is positively homogeneous of degree one. That is, $H(C;\lambda v) = \lambda H(C;v)$, for $\lambda > 0$. Thus, representing the support function over the unit sphere, $\|v\| = 1$, is sufficient to determine the function over the whole space, R^3. Let S^2 and S^1 denote, respectively, the unit sphere in R^3 and the unit circle in R^2. The support function of a compact convex set $C \subset R^3$ maps the point $u \in S^2$, treated as a vector, to the signed distance between the origin and the tangent plane of C with outer normal u. Figure 1 shows the 2-D example of a convex polygon and its support function defined for

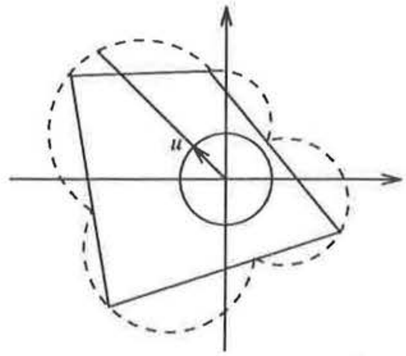

Figure 1: The support function of a convex polygon.

vectors $u \in S^1$. The value of the support function is the distance between the origin and the dashed arc along the direction determined by u.

An ellipsoid is an example of a smooth, strictly convex object and is the prototype shape considered here. Let $E_{a,b,c}$ be the ellipsoid

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1 \ . \tag{1}$$

The support function of $E_{a,b,c}$ is

$$H(E_{a,b,c}; v) = \sqrt{a^2 {v_1}^2 + b^2 {v_2}^2 + c^2 {v_3}^2} \ , \tag{2}$$

for $v = (v_1, v_2, v_3) \in R^3$.

Definition 2.2 Let λ_1 and λ_2 be non-negative real numbers and let C_1 and C_2 be convex objects. Then, the convex object, C, where

$$C = \lambda_1 C_1 + \lambda_2 C_2 = \{\lambda_1 x_1 + \lambda_2 x_2 \mid x_i \in C_i \, , \, i = 1, 2\} \tag{3}$$

is called the *linear combination* or *mixture* of C_1 and C_2.

The support function of C is the same linear combination of the support functions of C_1 and C_2. That is, $H(C; v) = \lambda_1 H(C_1; v) + \lambda_2 H(C_2; v)$. Consider the single point set $\{a\}$ and suppose $\lambda_1 = \lambda_2 = 1$. Then, $H(C + \{a\}; v) = H(C; v) + H(\{a\}; v) = H(C; v) + \langle a, v \rangle$. This illustrates the behavior of the support function under translation by a vector a. Equivalently, this illustrates how the support function depends on the choice of coordinate system origin. The difference is the term $\langle a, v \rangle$, where a is the (vector) difference between any two choices of origin.

The support function of any set is convex. Conversely, any convex function that is positively homogeneous of degree one is the support function of a convex body. If C_1, C_2 are nonempty compact convex sets in R^3 with $H(C_1; v) = H(C_2; v)$ for every non-zero $v \in R^3$, then $C_1 = C_2$. The support function uniquely characterizes convex sets. In fact, any compact convex set C can be represented by its support function as

$$C = \{x \mid \langle x, v \rangle \leq H(C; v) \, , v \in R^3 \, , v \neq 0\} \ .$$

A non-convex set and its convex hull have the same support function. This demonstrates that the support function is unique only for convex sets.

Let C be a smooth, strictly convex 3-D object. Let $x(C;u)$ be the unique point on the surface of C with outer normal u, $u \in S^2$. Let $r_1(C;u)$ and $r_2(C;u)$ be the two principal radii of curvature of C at $x(C;u)$. Then, the two curvature functions, $F_1(C;u)$ and $F_2(C;u)$, can be defined as follows:

$$\begin{aligned} F_1(C;u) &= r_1(C;u) + r_2(C;u) , \\ F_2(C;u) &= r_1(C;u) r_2(C;u) . \end{aligned}$$

The second curvature is the reciprocal of the Gaussian curvature and the first curvature is equal to twice the mean curvature divided by Gaussian curvature. The curvature functions of $E_{a,b,c}$ can be shown to be

$$F_1(E_{a,b,c};u) = \frac{a^2b^2({u_1}^2+{u_2}^2)+b^2c^2({u_2}^2+{u_3}^2)+a^2c^2({u_1}^2+{u_3}^2)}{(a^2{u_1}^2+b^2{u_2}^2+c^2{u_3}^2)^{3/2}} , \qquad (4)$$

$$F_2(E_{a,b,c};u) = \frac{a^2b^2c^2({u_1}^2+{u_2}^2+{u_3}^2)}{(a^2{u_1}^2+b^2{u_2}^2+c^2{u_3}^2)^2} , \qquad (5)$$

for points $u = (u_1, u_2, u_3) \in S^2$.

The important property that the support function and the curvature functions share with other orientation-based representations is that they rotate in the same way as the object rotates. Let R denote an arbitrary rotation. Then,

$$\begin{aligned} H(R(C);u) &= H(C;R^{-1}(u)) , && (6) \\ F_i(R(C);u) &= F_i(C;R^{-1}(u)) ,\ i=1,2 , && (7) \end{aligned}$$

for points $u = (u_1, u_2, u_3) \in S^2$.

The notion of mixed volume plays an important role in the studies of convex bodies in higher dimensional spaces. The following is the definition by Busemann [13] (page 43) restricted to R^3.

Definition 2.3 (Busemann [13] page 43.) Let C_i, $i = 1, 2, \ldots, n$, be n 3-D convex bodies. Denote the volume of C by $V(C)$. Let $\lambda_i \geq 0$, $i = 1, 2, \ldots, n$, be numbers and let C be the mixture

$$C = \sum_{i=1}^{n} \lambda_i C_i .$$

Then the volume of C can be expressed as

$$V(C) = \sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{n} V_{ijk} \lambda_i \lambda_j \lambda_k$$

which is a polynomial of degree three in the variables λ_i. The coefficients V_{ijk} are uniquely determined by requiring that they be symmetric in their subscripts. It follows that the coefficient V_{ijk} depends only on the bodies C_i, C_j and C_k, not on any of the remaining n bodies. V_{ijk} is called the *mixed volume* of C_i, C_j and C_k and is written as $V(C_i, C_j, C_k)$.

Let C_1 and C_2 be two smooth, strictly convex 3-D objects. Let B^3 be the unit ball[4] in R^3. Two mixed volumes, $V(C_1, C_2, B^3)$ and $V(C_1, C_2, C_2)$ are of particular interest because

$$V(C_1, C_2, B^3) = \frac{1}{6} \int_{S^2} H(C_1; \omega) F_1(C_2; \omega) d\omega , \tag{8}$$

$$V(C_1, C_2, C_2) = \frac{1}{3} \int_{S^2} H(C_1; \omega) F_2(C_2; \omega) d\omega . \tag{9}$$

These two mixed volumes combine the support function of C_1 and the curvature functions of C_2. It should be noted that equations (8) and (9) do not depend on the choice of origin used to define the support function. As will be seen, these two mixed volumes serve as similarity measures for 3-D attitude determination.

Definition 2.4 Two sets P and Q are *homothetic* if and only if $P = \{a\} + \lambda Q$ for some $a \in R^3$ and $\lambda > 0$.

Objects that are homothetic differ by at most a translation and a scaling. The following two theorems are fundamental:

Theorem 2.1 (Busemann [13] page 49.) Let C_1, C_2 and C_3 be 3-D convex bodies. Then,

$$V^2(C_1, C_2, C_3) \geq V(C_1, C_1, C_3)\, V(C_2, C_2, C_3) ,$$

with equality if and only if C_1 and C_2 are homothetic.

Theorem 2.2 (Minkowski's Inequality. Busemann [13] page 48.) Let C_1 and C_2 be 3-D convex bodies. Then,

$$V^3(C_1, C_2, C_2) \geq V(C_1)\, V^2(C_2) .$$

with equality if and only if C_1 and C_2 are homothetic.

3 Solutions to the Attitude Determination Problem

Definition 3.1 *Attitude determination* is the problem of finding a rotation, R, such that $R(C_1)$ and C_2 are homothetic, where C_1 is a known 3-D model and C_2 is an instance of C_1 under an unknown rotation, translation and scaling.

Throughout, assume that C_1 is a given object model defined in a standard coordinate system and that C_2 is a measured instance of C_1 subject to unknown rotation, translation and scaling. By Theorem 2.1,

$$V(R(C_1), C_2, B^3) \geq \sqrt{V(R(C_1), R(C_1), B^3)\; V(C_2, C_2, B^3)} \tag{10}$$

with equality if and only if $R(C_1)$ and C_2 are homothetic. Further, it is known that $V(C, C, B^3)$ is equal to 1/3 the surface area of a convex body C. Surface area is invariant under rotation. Therefore, $V(R(C_1), R(C_1), B^3) = V(C_1, C_1, B^3)$ and the minimum value of

[4]The unit ball, B^3, is the unit sphere, S^2, plus all its interior points.

(10) is known independent of R. Accordingly, $V(R(C_1), C_2, B^3)$ achieves this minimum if and only if $R(C_1)$ and C_2 are homothetic. Similarly, by Theorem 2.2,

$$V(R(C_1), C_2, C_2) \geq \sqrt[3]{V(R(C_1))\, V^2(C_2)} \tag{11}$$

with equality if and only if $R(C_1)$ and C_2 are homothetic. Volume is invariant under rotation. Therefore, $V(R(C_1)) = V(C_1)$ and the minimum value of (11) is known independent of R. $V(R(C_1), C_2, C_2)$ achieves this minimum if and only if $R(C_1)$ and C_2 are homothetic. Now, define functions of R as follows:

$$\varphi(R) \triangleq V(R(C_1), C_2, B^3)\,, \tag{12}$$

$$\psi(R) \triangleq V(R(C_1), C_2, C_2)\,. \tag{13}$$

The functions $\varphi(R)$ and $\psi(R)$ attain their known minima if and only if $R(C_1)$ and C_2 are homothetic. Therefore, by Definition 3.1, the 3-D attitude determination problem can be solved if and only if the minima of $\varphi(R)$ or $\psi(R)$ can be found. Either of these minima is an equivalent solution to the 3-D attitude determination problem. The global minimum of $\varphi(R)$ is $\sqrt{V(C_1, C_1, B^3)\, V(C_2, C_2, B^3)}$ and that of $\psi(R)$ is $\sqrt[3]{V(C_1)\, V^2(C_2)}$, both of which are known independent of R. By the if-and-only-if conditions of Theorems 2.1 and 2.2, these global minima are unique, modulo any rotational symmetries that C_1 possesses.

A rotation, R, can be represented as a triple (ϕ, θ, Ω) interpreted to mean a counter-clockwise rotation by angle Ω around unit vector $(\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$. When R is represented in this way, $\varphi(R)$ and $\psi(R)$ are functions of the three variables $\phi, \theta, \Omega \in R^3$ and are written as $\varphi(\phi, \theta, \Omega)$ and $\psi(\phi, \theta, \Omega)$. Thus, the problem of 3-D attitude determination is transformed into two equivalent optimization problems:

$$\text{minimize:} \quad \varphi(\phi, \theta, \Omega)\,, \quad (\phi, \theta, \Omega) \in R^3\,, \tag{14}$$

$$\text{minimize:} \quad \psi(\phi, \theta, \Omega)\,, \quad (\phi, \theta, \Omega) \in R^3\,. \tag{15}$$

Since the objective functions are periodic and bounded, solutions to both of these optimization problems necessarily exist.

It is important to note that the required mixed volumes are well defined as long as C_1 and C_2 are 3-D convex bodies. Thus, optimization problems (14) and (15), derived from Theorems 2.1 and 2.2, apply to polyhedra too. If C_1 and C_2 are smooth and strictly convex, then the objective functions $\varphi(R)$ and $\psi(R)$ can be written explicitly, according to Equations (8) and (9), as

$$\begin{aligned} \varphi(R) &= \frac{1}{6}\int_{S^2} H(R(C_1);\omega)\, F_1(C_2;\omega)\, d\omega\,, \\ \psi(R) &= \frac{1}{3}\int_{S^2} H(R(C_1);\omega)\, F_2(C_2;\omega)\, d\omega\,. \end{aligned}$$

In practice, sensed surface data typically is obtained from a single viewpoint. Thus, the points at which the curvature functions of C^2 are known span only a hemisphere. To proceed, it is necessary to "complete" the visible surface, thus converting it into a convex body. Suppose the viewpoint is in the positive z direction. Further, suppose that the

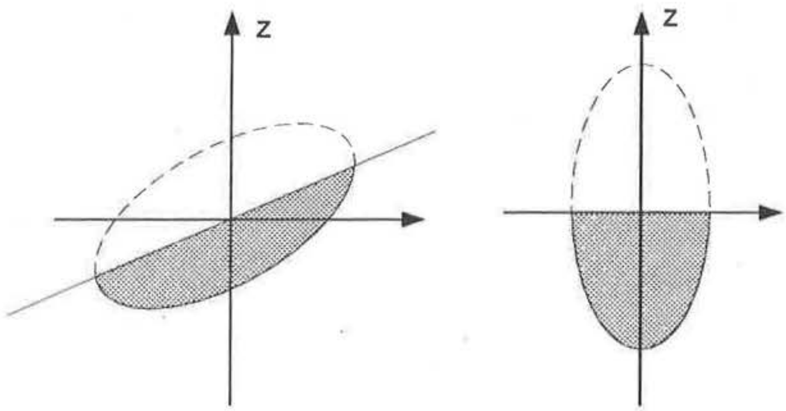

(a) $R(C_1)'$ (b) C_2'

Figure 2: Only the data on a hemisphere is used in optimization when the object is viewed from a single viewpoint.

occluding boundaries of C_2 and $R(C_1)$ each lie in a plane[5]. Assume coordinate systems are assigned so that the plane for C_2 is $z = 0$ and so that the plane for $R(C_1)$ contains the origin. Let $R(C_1)'$ be the convex body bounded by points of $R(C_1)$ that are visible in the positive z direction and by the plane containing the occluding boundary. Let C_2' be the convex body bounded by points of C_2 that are visible in the positive z direction and by the plane $z = 0$. Figure 2 provides a 2-D example where C_1 and C_2 are ellipses and where $R(C_1)'$ and C_2' are the shaded regions shown. Of course, Theorems 2.1 and 2.2 still apply to $R(C_1)'$ and C_2'. Therefore,

$$\frac{V(R(C_1)', C_2', B^3)}{\sqrt{V(R(C_1)', R(C_1)', B^3)}} \geq \sqrt{V(C_2', C_2', B^3)}\,, \tag{16}$$

and

$$\frac{V(R(C_1)', C_2', C_2')}{\sqrt[3]{V(R(C_1)')}} \geq \sqrt[3]{V^2(C_2')}\,, \tag{17}$$

with equality if and only if $R(C_1)'$ and C_2' are homothetic. The minimum values of (16) and (17) still are known independent of R. But, neither $R(C_1)'$ nor C_2' is smooth and strictly convex so that the mixed volumes in (16) and (17) are more difficult to derive. One can split the analysis into the two hemispheres S^{2-} and S^{2+}, where S^{2-} and S^{2+} denote, respectively, the hemispheres corresponding to $z < 0$ and $z > 0$. Details are provided in [15]. The planar regions of $R(C_1)'$ and C_2' introduce new area and mixed area terms into the mixed volumes that, while slowly varying, do depend on R. Ignoring these terms effects the accuracy of the mixed volumes $V(R(C_1)', C_2', B^3)$, $V(R(C_1)', R(C_1)', B^3)$ and $V(R(C_1)', C_2', C_2')$. The

[5] Marr [14] provides if and only if conditions for an occluding contour to be planar, independent of viewpoint. Here, this is equivalent to assuming the surface is quadratic.

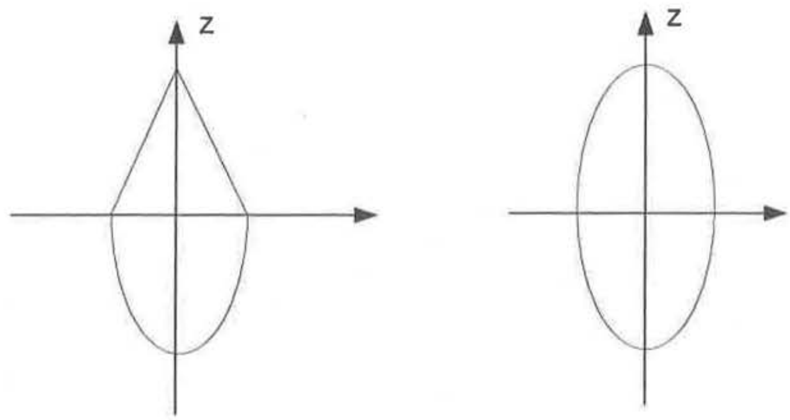

Figure 3: Two 2-D convex objects that do match when viewed in the positive z direction but that do not match over the whole unit circle.

---

effect, however, is small. When the analysis is confined to the hemisphere S^{2-}, the objective functions corresponding to (16) and (17) become:

$$\overline{\varphi}(R) = \frac{\frac{1}{6}\int_{S^{2-}} H(R(C_1);\omega)\, F_1(C_2;\omega)\, d\omega}{[\frac{1}{6}\int_{S^{2-}} H(R(C_1);\omega)\, F_1(R(C_1);\omega)\, d\omega]^{\frac{1}{2}}}, \tag{18}$$

$$\overline{\psi}(R) = \frac{\frac{1}{3}\int_{S^{2-}} H(R(C_1);\omega)\, F_2(C_2;\omega)\, d\omega}{[\frac{1}{3}\int_{S^{2-}} H(R(C_1);\omega)\, F_2(R(C_1);\omega)\, d\omega]^{\frac{1}{3}}}. \tag{19}$$

Minimizing $\overline{\varphi}(R)$ and $\overline{\psi}(R)$ only approximates the minimizing solutions to (16) and (17). But, experiments on synthetic data suggest that the approximation is both accurate and robust. Even if perfect, minimizing $\overline{\varphi}(R)$ and $\overline{\psi}(R)$ does not solve the attitude determination problem, as defined in Definition 3.1, since part of the object is never seen and therefore may not be matched correctly. It does solve the attitude determination problem correctly to the extent possible, given the data available. Figure 3 depicts two 2-D convex objects that match when viewed in the positive z direction but that do not match over the whole unit circle.

4 Experiments

Experiments have been conducted on a real object to demonstrate the feasibility of the approach. The object is the ellipsoid $E_{3,5,9}$. Using Equations (2), (4) and (5), the support function and the curvature functions are

$$H(E_{3,5,9};u) = \sqrt{9{u_1}^2 + 25{u_2}^2 + 81{u_3}^2}, \tag{20}$$

$$F_1(E_{3,5,9};u) = \frac{954{u_1}^2 + 2250{u_2}^2 + 2754{u_3}^2}{(9{u_1}^2 + 25{u_2}^2 + 81{u_3}^2)^{3/2}}, \tag{21}$$

$$F_2(E_{3,5,9}; u) \;=\; \frac{18225({u_1}^2 + {u_2}^2 + {u_3}^2)}{(9{u_1}^2 + 25{u_2}^2 + 81{u_3}^2)^2}\,, \tag{22}$$

for points $u = (u_1, u_2, u_3) \in S^2$.

A polyvinylchlorid ellipsoid was custom machined by an automated, numerically controlled milling machine using numerical data derived analytically for $E_{3,5,9}$. A sphere of comparable size was machined from the same material to serve as a calibration object for an existing photometric stereo system.

Three light source photometric stereo uses three images of an object taken with the identical imaging geometry but under different conditions of illumination. Reflectance is measured empirically. A sphere is a useful calibration object because the full visible hemisphere of surface orientations is present and each orientation is readily determined by simple geometric analysis of the object silhouette. Calibration produces a lookup table mapping measured intensity triples, (E_1, E_2, E_3), to known surface orientations, represented by the gradient, (p, q). Once calibrated, photometric stereo determines, by table lookup, the gradient, (p, q), for each intensity triple, (E_1, E_2, E_3), measured from the surface of any other object made of the same material as the calibration object. Gradient estimation is robust because the three independent intensity measurements overdetermine the two unknown surface orientation parameters.

In principle, dense curvature information can be obtained by differentiating the gradient, (p, q). Photometric stereo provides a better way to obtain local curvature estimates. The lookup table constructed for photometric stereo implicitly represents three image irradiance equations, $E_i(x, y) \;=\; R_i(p, q)$, $i = 1, 2, 3$, where (x, y) are the image coordinates and each $R_i(p, q)$, $i = 1, 2, 3$, is a *reflectance map*. As part of calibration, the reflectance maps, $R_i(p, q)$, are made explicit because the partial derivatives of $R_i(p, q)$ with respect to p and q are used for curvature estimation. Thus, for each intensity triple, (E_1, E_2, E_3), calibration determines the gradient, (p, q), and the six partial derivatives of $R_i(p, q)$, $i \;=\; 1, 2, 3$. Photometric stereo then combines the partial derivatives of $E_i(x, y)$ with respect to x and y with the partial derivatives of $R_i(p, q)$ to estimate the surface Hessian matrix. Curvature estimation is robust because the six independent intensity gradient measurements overdetermine the three unknown surface curvature parameters. Combining the Hessian and the gradient determines the principal curvatures, k_1 and k_2. Details of the photometric stereo system are described in [12]. In the experiments described here, the gradient, (p, q), obtained from photometric stereo determines the mapping from sensed surface point to the point $u \in S^{2-}$ and the principal curvatures, k_1 and k_2, determine values for the first and second curvature functions, $F_1(C; u)$ and $F_2(C; u)$.

Since only a single viewpoint is used, the objective functions to optimize are the approximations, $\overline{\varphi}(R)$ and $\overline{\psi}(R)$, defined in Equations (18) and (19). Two experiments are described using the images shown in Figures 4 and 5 respectively. The first example is solved using $\overline{\varphi}(R)$ and the second using $\overline{\psi}(R)$.

Let C_1 be the model, $E_{3,5,9}$, in its standard attitude and let C_2 be the sensed object with $C_2 = \lambda\, R_0(C_1) + \{a\}$ where R_0 is a fixed but unknown rotation, $\lambda > 0$ is an unknown scale factor and $a \in R^3$ is an unknown translation. It is known that optimization is independent of λ. Therefore, without loss of generality, let $\lambda = 1$. The functions $F_i(C_2; u) = F_i(R_0(C_1); u)$, $i = 1, 2$, are estimated at all surface points visible to the camera. The rotation, R_0, is esti-

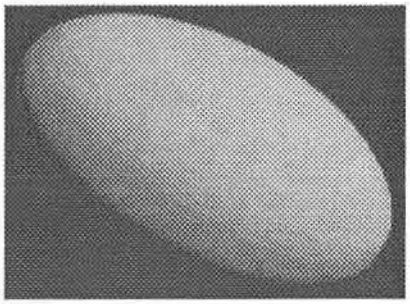

(a) first light source

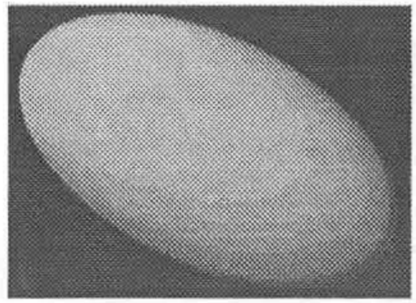

(b) second light source

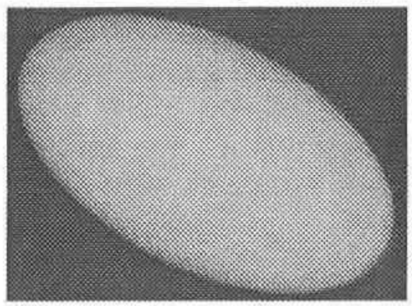

(c) third light source

Figure 4: Images of $E_{3,5,9}$ used with the objective function $\overline{\varphi}(R)$, the combination of the support function and the first curvature function.

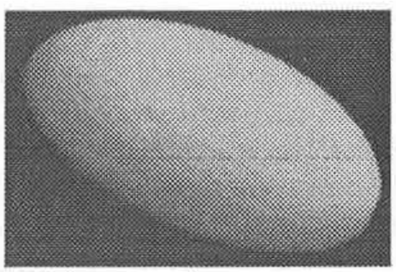

(a) first light source

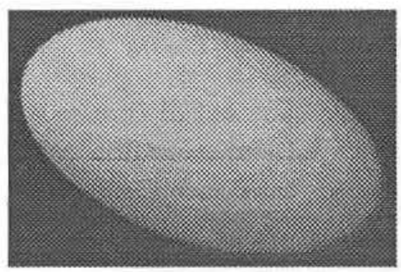

(b) second light source

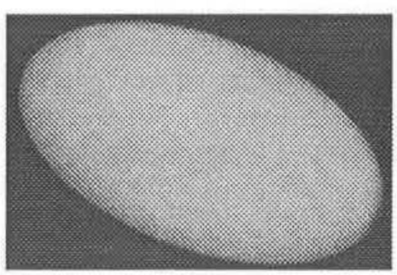

(c) third light source

Figure 5: Images of $E_{3,5,9}$ used with the objective function $\overline{\psi}(R)$, the combination of the support function and the second curvature function.

mated by minimizing $\overline{\varphi}(R)$ and $\overline{\psi}(R)$. For comparison to ground truth, R_0 also is determined *a priori* as part of the experimental setup.

The surface integrals, given by Equations (18) and (19), are transformed into the volume integrals:

$$\overline{\varphi}(R) = \frac{\frac{1}{6}\int_{\frac{\pi}{2}}^{\pi}\int_0^{2\pi} H(R(C_1);x(\phi,\theta))F_1(R_0(C_1);x(\phi,\theta))d\theta d\phi}{[\frac{1}{6}\int_{\frac{\pi}{2}}^{\pi}\int_0^{2\pi} H(R(C_1);x(\phi,\theta))F_1(R(C_1);x(\phi,\theta))d\theta d\phi]^{\frac{1}{2}}},$$

$$\overline{\psi}(R) = \frac{\frac{1}{3}\int_{\frac{\pi}{2}}^{\pi}\int_0^{2\pi} H(R(C_1);x(\phi,\theta))F_2(R_0(C_1);x(\phi,\theta))d\theta d\phi}{[\frac{1}{3}\int_{\frac{\pi}{2}}^{\pi}\int_0^{2\pi} H(R(C_1);x(\phi,\theta))F_2(R(C_1);x(\phi,\theta))d\theta d\phi]^{\frac{1}{3}}},$$

where $x(\phi,\theta) = (\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$. The numerical integration routine QB01AD from Harwell [16] is used to evaluate the 2-D volume integrals. The interpolation routine of Renka [17, 18] is used to interpolate the irregularly spaced values of $F_i(R_0(C_1);u)$, $i = 1,2$, obtained from photometric stereo, as required. As a result $F_i(R_0(C_1);u)$, $i = 1,2$, become C^1 functions over the hemisphere S^{2-}.

The nonlinear programming subroutine NLPQL [19] is used to find the minima of $\overline{\varphi}(R)$ and $\overline{\psi}(R)$. NLPQL can be used to solve optimization problems with constraints, optimization problems with simple bounds, or unconstrained optimization problems. It requires an estimate of the gradient of the objective function. Here, the gradients of the objective functions, $\overline{\varphi}(R)$ and $\overline{\psi}(R)$, are estimated by simple forward differencing. Convergence is achieved either when the Kuhn-Tucker conditions (see [20] page 51) are satisfied to within a specified accuracy or when the objective functions are not improved significantly given that the constraints are satisfied to within the specified accuracy. NLPQL does not guarantee that the minima found are global. But, in this application, the true minima are known. Therefore, the validity of the minima found by the subroutine is established.

In the experiments, the optimization process was executed 256 times, each time corresponding to a different initial guess for object rotation. A very large bound was given to NLPQL to effectively make a constrained optimization into an unconstrained optimization. All the initial guesses converged to points with the same minimum values of $\overline{\varphi}(R)$ or $\overline{\psi}(R)$ and with the same object attitudes. Thus, for the object tested, $E_{3,5,9}$, the method is robust with respect to the initial guess.

The position of the sensed object is established manually for each experiment. The true rotation of the object with respect to its standard attitude also is estimated manually. The estimated rotation is used as a rough measure of accuracy to evaluate the rotation found by the optimization process. A way to visualize the optimization result is to superimpose the rotated model onto the image of the object. As examples, the results of the two optimizations each with initial guess $(0.1, 0.2, 0.3)$ are shown in Figure 6 and Figure 7, respectively. In the figures, the black and white shows the silhouette of the object and the wire frame shape in gray is the rotated model projected onto the image plane.

Since the three axes of the ellipsoid, $E_{3,5,9}$, all are different, it has few symmetries. Therefore, it is possible to evaluate the optimization results by comparing rotation matrices.

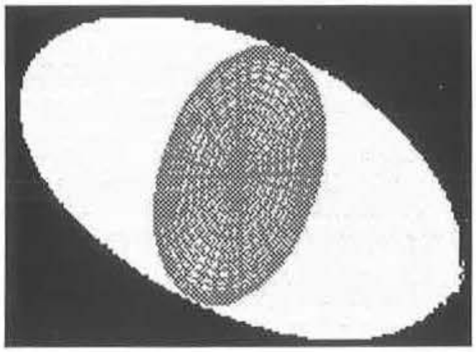

(a) initial guess.

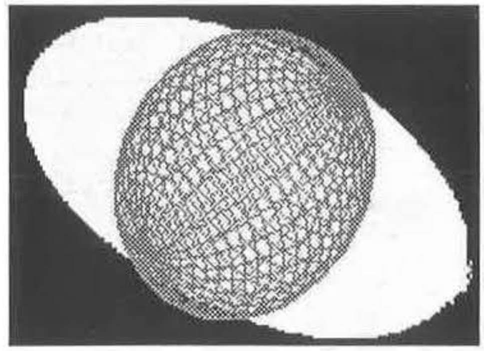

(b) 3rd iteration.

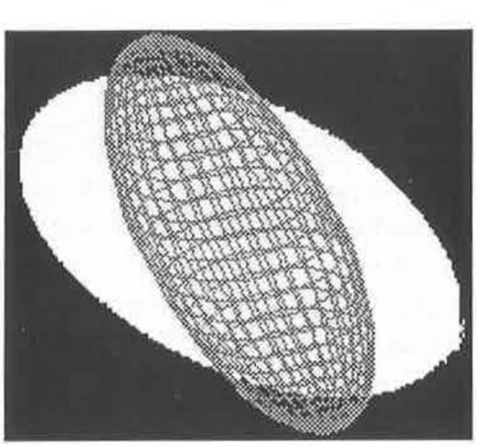

(c) 4th iteration.

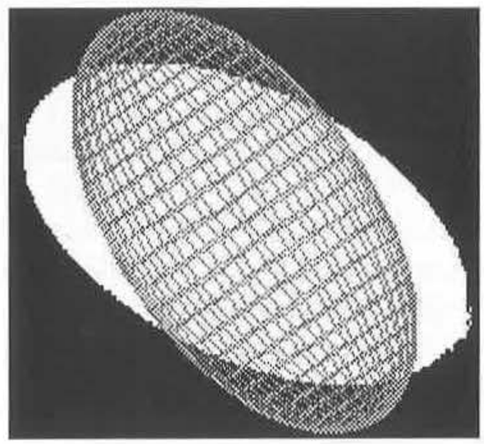

(d) 6th iteration.

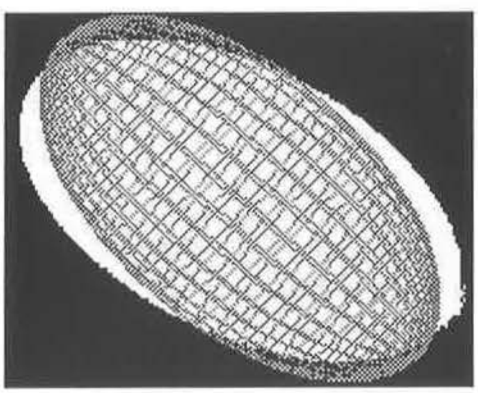

(e) 9th iteration.

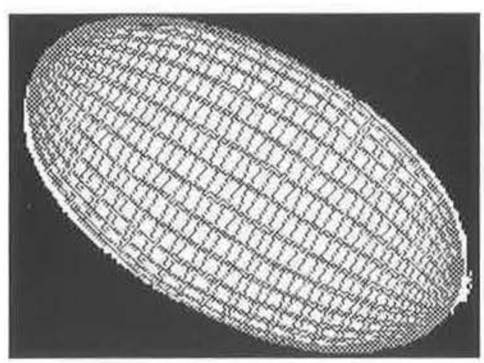

(f) final result.

Figure 6: Results of real data experiments using the support function and the first curvature function.

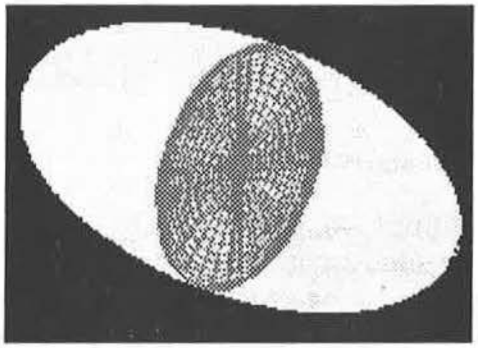

(a) initial guess.

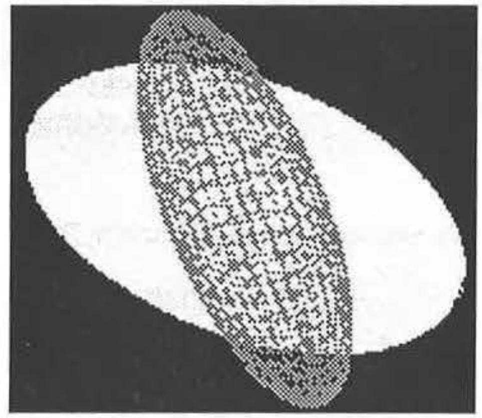

(b) 1st iteration.

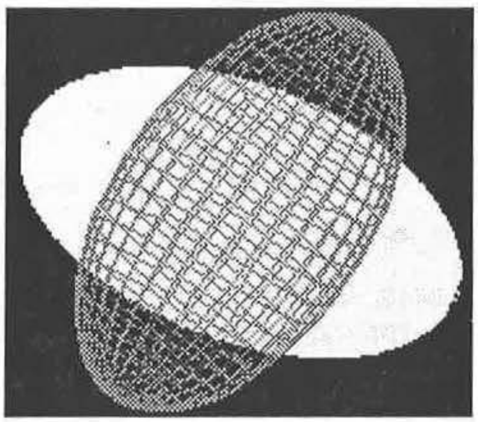

(c) 2nd iteration.

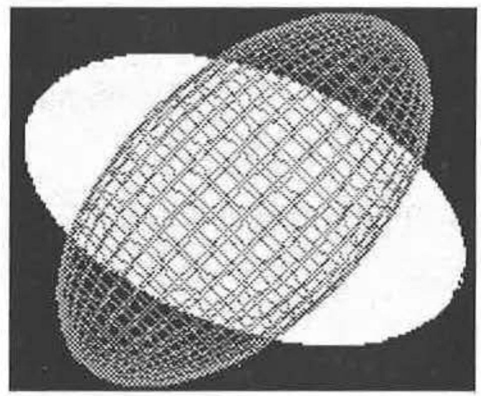

(d) 3rd iteration.

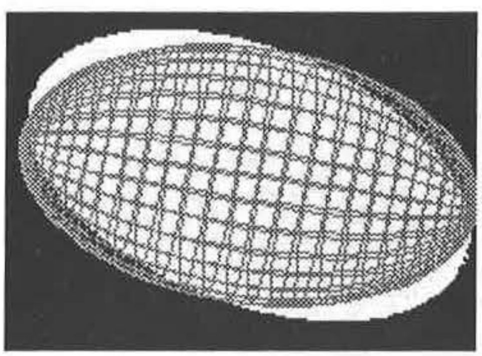

(e) 4th iteration.

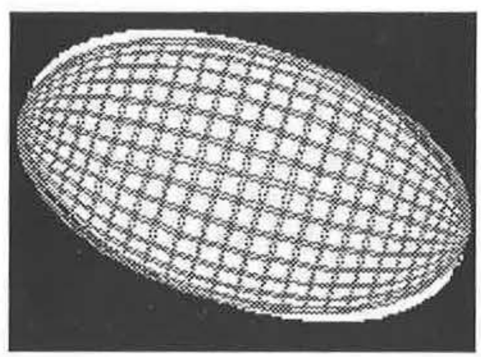

(f) final result.

Figure 7: Results of real data experiments using the support function and the second curvature function.

The estimated *a priori* rotation matrix for the ellipsoid imaged in Figure 4 is

$$\begin{bmatrix} 0 & -0.4627155767 & -0.8865067936 \\ 0 & 0.8865067936 & -0.4627155767 \\ 1 & 0 & 0 \end{bmatrix}.$$

The rotation estimated by optimizing $\overline{\varphi}(R)$ is the matrix

$$\begin{bmatrix} 0.0591327494 & -0.5181631311 & -0.8532351889 \\ 0.0264112871 & 0.8552437408 & -0.5175524972 \\ 0.9979006773 & 0.0080692626 & 0.0642582690 \end{bmatrix}$$

or this matrix multiplied on the right by a matrix corresponding to a reflection about one or more axis. These matrices all define the same attitude since the ellipsoid is symmetric about each coordinate axis in its standard attitude. The estimated *a priori* rotation matrix for the ellipsoid imaged in Figure 5 is

$$\begin{bmatrix} 0 & -0.3701448041 & -0.9289740707 \\ 0 & 0.9289740707 & -0.3701448041 \\ 1 & 0 & 0 \end{bmatrix}.$$

The rotation estimated by optimizing $\overline{\psi}(R)$ is the matrix

$$\begin{bmatrix} -0.0371509607 & -0.2698579447 & -0.9621831926 \\ 0.0206643797 & 0.9624345419 & -0.2707263137 \\ 0.9990959863 & -0.0299406615 & -0.0301789208 \end{bmatrix}$$

or, as before, this matrix multiplied on the right by a matrix corresponding to a reflection about one or more axis.

5 Conclusions

The desirable property that all orientation-based representations share is that the object and the representation rotate together. This makes an orientation-based representation well-suited to the task of attitude determination. Photometric stereo provides dense, local estimates of both surface orientation and surface curvature. Imaging from an unknown viewpoint determines a visible hemisphere of the representation. For smooth, strictly convex objects, matching the visible hemisphere to the full spherical model can be formulated as a single, uniform optimization process. Within this framework, depth need not be represented explicitly.

The method transforms the 3-D attitude determination problem for smooth, strictly convex objects into an optimization problem for which standard numerical solutions exist. An important additional property of the optimization is the fact that the value of the extremum is known *a priori*. Thus, one can always assess the validity of the solution found by the optimization.

The algorithm demonstrated here explicitly assumes that the object surface is smooth and strictly convex, thus ensuring that the Gauss map is one-to-one. The theorems on mixed volume apply to all convex bodies. A (technically) different treatment is required to develop algorithms for convex objects with developable or planar surfaces. It should be noted, however, that the support function need not treat smooth objects and polytopes differently. The support function is defined for all points on the unit sphere regardless whether the object is smooth, a polytope or a combination. This is one of the reasons conjectured for its success in polyhedral shape matching [5].

Good matching results have been obtained. Precise determination of accuracy and robustness requires more quantitative work. Accuracy assessment must take into account sensor calibration, *a priori* determination of the "correct" attitude of the presented object and uncertainty in the "shape-from" method used to acquire the raw orientation and curvature data. Based on the experimental work performed to date, the overall accuracy of the method is consistent with what one can expect, given these other factors. It would be helpful to agree upon a metric for rotation space to quantify differences between the correct and the estimated object attitude. The method is robust because it is a true 3-D method that employs dense surface data, not just data from 2-D contours or other sparse sets of features. Indeed, for the test object, $E_{3,5,9}$, 2-D contours alone do not determine 3-D attitude.

An obvious question is, "Which works best, the first or the second curvature function?". Experiments using purely synthesized data suggest that attitude determination using the first curvature function is slightly more accurate. This probably is related to nothing more than the observation that, all else equal, there is less uncertainty in the sum of two numbers than there is in their product. This observation is merely anecdotal since, of course, one also needs to take into account properties of the particular numerical optimization routines used. Theory suggests a possible advantage to using the second curvature function when analysis is confined to the visible hemisphere S^{2-}. In this case, $\overline{\psi}(R)$ may better approximate the ratio of the true mixed volumes required than does $\overline{\varphi}(R)$. However, experiments to date on both real and synthetic data do not distinguish between $\overline{\varphi}(R)$ and $\overline{\psi}(R)$ either in terms of accuracy or robustness.

In the implementation described, optimization proceeds using a large number of initial guesses. The correct attitude is found even when there is no *a priori* knowledge of object attitude. At the same time, optimization benefits from a good initial guess. This suggests that the approach also is well-suited to motion tracking and navigation tasks where the solution at one time can be used as the initial guess for the solution at the next sample time.

Acknowledgment

Rod Barman assisted in the preparation and mounting of the test objects. Stewart Kingdon provided essential software for image acquisition. Jim Little contributed to many discussions on attitude determination. Robert Renka provided the spherical interpolation routine and converted it to double precision for our use. Tom Nicol suggested using the optimization routine NLPQL and also helped to convert it for local use. Ian Cavers provided the integration routine QB01AD from Harwell. The authors also thank Branislav Klco of Studio Apropos!, Vancouver, BC, for machining the test objects. Support for the work described was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), by the Canadian Institute for Advanced Research (CIAR) and by the Institute for Robotics and

Intelligent Systems (IRIS), a Canadian network of centres of excellence.

References

[1] B. Horn, "Extended Gaussian images," *Proc. IEEE*, vol. 72, pp. 1671–1686, 1984.

[2] V. S. Nalwa, "Representing oriented piecewise C^2 surfaces," *International Journal of Computer Vision*, no. 2, pp. 131–154, 1989.

[3] P. Brou, "Using the Gaussian image to find the orientation of objects," *International Journal of Robotics Research*, vol. 3, no. 4, pp. 89–125, 1984.

[4] B. K. P. Horn and K. Ikeuchi, "The mechanical manipulation of randomly oriented parts," *Scientific American*, vol. 251, no. 2, pp. 100–111, August, 1984.

[5] J. J. Little, "Determining object attitude from extended Gaussian images," in *Proc. 9th Int. Joint Conf. on Artificial Intelligence*, (Los Angeles, CA), pp. 960–963, 1985.

[6] P. J. Besl and R. C. Jain, "Three-dimensional object recognition," *ACM Computing Surveys*, vol. 17, pp. 75–145, 1985.

[7] P. J. Besl and R. C. Jain, "Invariant surface characteristics for 3D object recognition in range images," *Computer Vision Graphics and Image Processing*, vol. 33, pp. 33–80, 1986.

[8] J. J. Koenderink, *Solid Shape*. Cambridge, MA: MIT Press, 1990.

[9] P. Liang and J. S. Todhunter, "Representation and recognition of surface shapes in range images: A differential geometry approach," *Computer Vision Graphics and Image Processing*, vol. 52, pp. 78–109, 1990.

[10] S. B. Kang and K. Ikeuchi, "Determining 3-D object pose using the Complex Extended Gaussian Image," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1991*, pp. 580–585, 1991.

[11] Y. Li and R. J. Woodham, "Starshaped sets, the radial function and 3-D attitude determination," TR-92-27, UBC Dept. of Computer Science, Vancouver, BC, 1992.

[12] R. J. Woodham, "Surface curvature from photometric stereo," in *Physics-Based Vision: Principles and Practice (Vol III: Shape Recovery)* (L. Wolff, S. Shafer, and G. Healey, eds.), pp. 121–155, Boston, MA: Jones and Bartlett Publishers, Inc., 1992.

[13] H. Busemann, *Convex Surfaces*, vol. 6 of *Interscience Tracts in Pure and Applied Mathematics*. New York: Interscience Publishers Inc., 1958.

[14] D. Marr, "Analysis of occluding contour," *Proc. R. Soc. Lond. B*, vol. 197, pp. 441–475, 1977.

[15] Y. Li, "Orientation-based representations of shape and attitude determination," PhD thesis, UBC Dept. of Computer Science, Vancouver, BC, 1992.

[16] Advanced Computing Department, AEA Industrial Technology, Harwell Laboratory, Oxfordshire, England, *Harwell Subroutine Library Specifications*, September 1990.

[17] R. J. Renka, "Interpolation of data on the surface of a sphere," *ACM Transactions on Mathematical Software*, vol. 10, pp. 417–436, 1984.

[18] R. J. Renka, "Algorithm 623 interpolation on the surface of a sphere," *ACM Transactions on Mathematical Software*, vol. 10, pp. 437–439, 1984.

[19] K. Schittkowski, "NLPQL: A FORTRAN subroutine solving constrained nonlinear programming problems," *Annals of Operations Research*, vol. 5, pp. 485–500, 1985/6.

[20] R. Fletcher, *Practical Methods of Optimization*, vol. 2. Chichester: John Wiley & Sons, 1981.