Starshaped Sets, the Radial Function and 3-D Attitude Determination

Ying Li   Robert J. Woodham

Technical Report 92-27
October 1992

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

Abstract

Attitude characterizes the three rotational degrees of freedom between the coordinate system of a known object and that of a viewer. Orientation-based representations record 3-D surface properties as a function of position on the unit sphere. The domain of the representation is the entire sphere. Imaging from a single viewpoint typically determines a hemisphere of the representation. Matching the visible region to the full spherical model for a known object estimates 3-D attitude.

The radial function is used to define a new orientation-based representation of shape. The radial function is well-defined for a class of sets called *starshaped* in mathematics. A starshaped set contains at least one interior point from which all boundary points are visible. The radial function records the distance from the origin of the coordinate system to each boundary point. The novel contribution of this paper is to extend previous mathematical results on the matching problem for convex objects to starshaped objects. These results then allow one to transform the attitude determination problem for starshaped sets into an optimization problem for which standard numerical solutions exist. Numerical optimization determines the 3-D rotation that brings a sensed surface into correspondence with a known model.

The required surface data can be obtained, for example, from laser range finding or from shape-from-shading. A proof-of-concept system has been implemented and experiments conducted on real objects using surface data derived from photometric stereo.

Keywords: shape, orientation-based representation, attitude determination, starshaped set, convexity, radial function, range data, photometric stereo

1 Introduction

Dense representations of surface shape support recognition, localization and inspection tasks. Recognition determines object identity. Localization specifies the three translational and the three rotational degrees of freedom of the object in space, as is necessary, for example, to direct a robot arm to grasp the object. Inspection verifies the suitability of the object for a particular task and typically requires careful attention to surface detail. Dense representations of surface shape can be obtained from laser ranging, shape-from-shading or photometric stereo.

Attitude characterizes the three rotational degrees of freedom between the coordinate system of a known object and that of a viewer. One approach that has proven useful for both recognition and attitude determination is to represent 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 differ in the way the mapping between surface points and points on the sphere is established. One way is via 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. Figure 1a illustrates the Gauss map. Mathematics defines many representations based on the Gauss map. Several have been used in computer vision including: 1) the Extended Gaussian Image (EGI) defined as the reciprocal of the Gaussian curvature [1]; 2) the support function defined as distance from an origin to the tangent plane [2]; and 3) the first and second curvature functions defined, respectively, as the sum of the principal radii of curvature and the product of the principal radii of curvature [3][1]. The EGI has been used for both recognition and attitude determination [4, 5, 6, 3]. The support function appears explicitly in two of the methods used [6, 3]. Other local curvature representations that can easily be fit to this orientation-based framework include the Gaussian and the mean curvature, popularized by Besl and Jain [7, 8], and Koenderink's curvedness and shape index [9]. The Gauss map is unique for convex objects. 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 [10, 11].

One can go beyond convexity by choosing a different way to establish the mapping between surface point and point on the sphere. Let the origin be in the interior of the object. Without loss of generality, assume that the size of the object is large enough so that the sphere fits entirely within the object. The mapping between surface point and the sphere is obtained as the intersection of the ray to the point with the sphere. Call this map the *dilation map*. Figure 1b illustrates the dilation map. Unfortunately, in general, the dilation map is not unique either. But, for a suitable choice of origin, it is unique when the object is what in mathematics is called a starshaped set [12, 13]. A set, S, is starshaped if and only if there exists at least one point in S that "sees" every other point in S. Two points in S

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

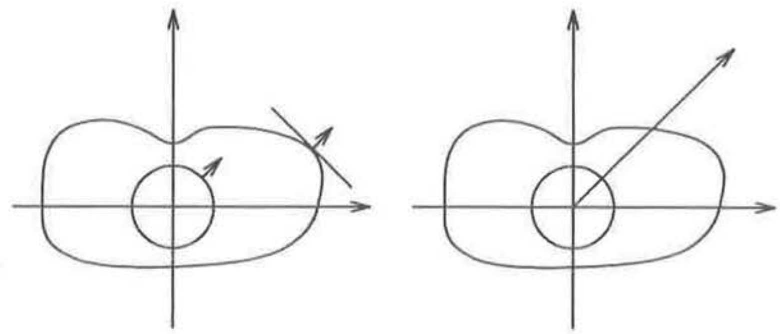

(a) Gauss map (b) dilation map

Figure 1: Maps between surface points and points on the sphere, illustrated in 2-D.

"see" each other if and only if the line joining the two points lies entirely within S.

The *radial function* is the distance from the origin to the point. The radial function for 2-D sets has been used to recognize objects based on their 2-D bounding contours [14]. The radial function also has been used in a medical application to fit the shape of a heart to a 3-D spherical harmonic model [15]. For a suitable choice of origin, the radial function defines a new orientation-based representation since the radial function can be considered a function of points on the sphere under the dilation map. This is the representation used here for attitude determination. The novel contribution is to extend previous mathematical results on the matching problem for convex objects to starshaped objects. These results then allow one to transform the attitude determination problem for starshaped sets into an optimization problem for which standard numerical solutions exist. Numerical optimization determines the 3-D rotation that brings a sensed surface into correspondence with a known model.

Section 2 formalizes the necessary mathematical results concerning the radial function and starshaped sets. Section 3 defines the attitude determination problem and shows how it can be transformed into an optimization problem for which standard numerical solutions exist. Section 4 describes the implementation and presents some experimental results. 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 with sensed data obtained, via photometric stereo, from real objects 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 the case 3 results on real objects. Section 5 summarizes the results.

2 The Radial Function

It is convenient to define the radial function first for arbitrary non-zero points $x \in R^3$ and then to specialize the definition to points on the (unit) sphere. Lutwak [16] defines the radial function for convex bodies. Here Lutwak's original definition is extended, without change, to starshaped sets.

Definition 2.1 The *radial function* of a starshaped set, S, in R^3 is defined as

$$\rho(S;x) \triangleq \sup\{\lambda > 0 |\ \lambda x \in S\}\ ,\ \text{for } x \in R^3,\ x \neq 0 \tag{1}$$

The radial function of a starshaped set does depend on the choice of origin in the coordinate system. If the origin is not interior to the set, S, then the radial function is not defined for every point $x \in R^3$. The radial function is well-defined for compact starshaped sets when the origin is inside the kernel of the starshaped set[2]. When S is a compact starshaped set in R^3 with origin, O, in the interior of its kernel, the radial function is

$$\rho(S;x) = \|\xi_x\|/\|x\|\ ,\ \text{for } x \in R^3,\ x \neq 0 \tag{2}$$

where ξ_x is the (unique) point of intersection of the ray $\overrightarrow{Ox}$ with the boundary of S.

If the radial function of S is $\rho(S;x)$, then that of λS is $\lambda\rho(S;x)$, $\lambda > 0$. Also, the radial function is positively homogeneous of degree minus one with respect to x. That is, $\rho(S;\lambda x) = \frac{1}{\lambda}\rho(S;x)$, $\lambda > 0$. Thus, representing the radial function over the unit sphere is sufficient to determine the function over the whole space, R^3. Figure 2 shows a 2-D starshaped set and its radial function defined for points, (u,v), on the unit circle. The analytic expressions between the dashed lines define the radial function for points on the unit circle, $u^2 + v^2 = 1$, in the corresponding region.

For a given choice of origin, the radial function is unique for compact starshaped sets. That is, if S_1, S_2 are nonempty compact starshaped sets in R^3 such that $\rho(S_1;x) = \rho(S_2;x)$ for every $x \in R^3$, then $S_1 = S_2$. Thus, a compact starshaped set S can be characterized by its radial function as $\{x \mid \rho(S;x) \geq 1\}$. For non-starshaped sets, uniqueness is no longer guaranteed. Figure 3(a) shows a 2-D non-starshaped set that has the same radial function as the 2-D starshaped set shown in Figure 3(b).

The *distance function* defined for convex bodies by Minkowski (Bonnesen-Fenchel [17] page 23) is the reciprocal of the radial function of the convex body. Valentine [13] extended the definition of distance function to starshaped sets and proved that a starshaped set is convex if and only if its distance function is convex (Valentine [13], page 32). The radial function and its reciprocal, the distance function, both are shape representations well-suited to starshaped sets.

The important property that the radial function shares with other orientation-based representations is that it rotates in the same way as the object rotates. That is,

$$\rho(R(S);x) = \rho(S;R^{-1}(x))\ , \tag{3}$$

[2]The kernel of a starshaped set, S, is the set of all points $x \in S$ such that every point of S can be seen by x. The kernel is a convex set.

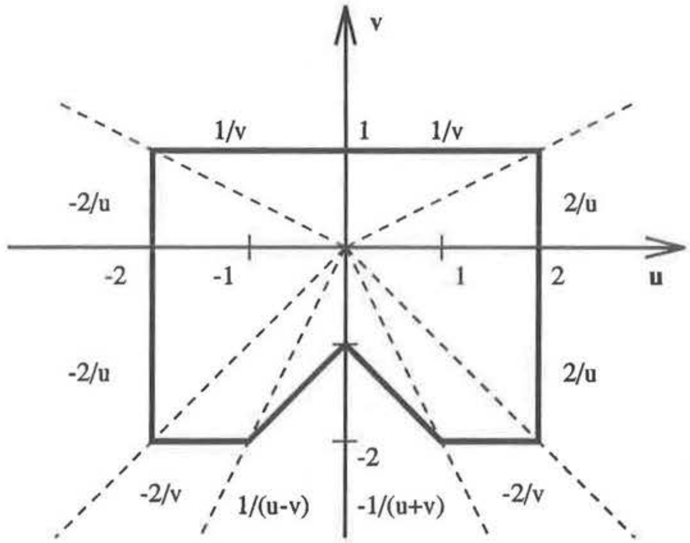

Figure 2: A 2-D starshaped set and its radial function, defined for points, (u, v), on the unit circle.

---

where R is a rotation. This is the property that makes it possible to use the radial function to solve the 3-D attitude determination problem.

A quantity that is associated with compact starshaped sets is the dual mixed volume. It combines radial functions and serves as a similarity measurement in attitude determination. The following definition and theorem again are obtained by extending the results of Lutwak [16] from convex objects to starshaped sets. Lutwak's definition of the dual mixed volume can be applied to starshaped sets, without change.

Definition 2.2 Let S^2 denote the unit sphere in R^3. Let S_1, S_2, S_3 be compact starshaped sets in R^3 with the origin in the interior of their kernels. The *dual mixed volume* of S_1, S_2, S_3 is defined as

$$\tilde{V}(S_1, S_2, S_3) \triangleq \frac{1}{3} \int_{S^2} \rho(S_1; x)\rho(S_2; x)\rho(S_3; x)dx \ .$$

Similarly, Lutwak's inequality for the dual mixed volume follows, without change.

Theorem 2.1 Let S_1, S_2, S_3 be compact starshaped sets in R^3 with the origin in the interior of their kernels.

$$\tilde{V}^3(S_1, S_2, S_3) \le \tilde{V}(S_1, S_1, S_1)\tilde{V}(S_2, S_2, S_2)\tilde{V}(S_3, S_3, S_3) \ ,$$

with equality if and only if S_1, S_2, S_3 are all dilations of each other (with the origin as the center of dilation).

Lutwak's proof for convex bodies uses Hölder's inequality for integrals (see Hardy [18] page 140). The essential requirement is that $\rho(S_i; x), i = 1, 2, 3$, be non-negative, not identically zero and measurable on S^2. The radial function, $\rho(S; x)$, of a compact starshaped

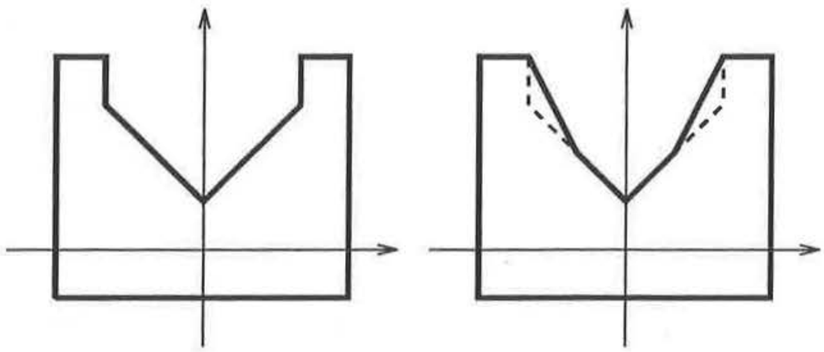

(a) non-starshaped (b) starshaped

Figure 3: A non-starshaped set that has the same radial function as a starshaped set.

set, S, with the origin in the interior of its kernel, satisfies these conditions. Thus Lutwak's proof applies to starshaped sets too.

3 Solutions to the Attitude Determination Problem

Definition 3.1 The attitude determination problem is defined as finding a rotation R such that $R(S) = S'$, where S is a model, S' is a measured object that is obtained by an unknown rotation from S.

Let S_1 and S_2 be two compact starshaped sets in R^3 with the origin in the interiors of their kernels. Then Definition 2.2 and Theorem 2.1 give

$$\tilde{V}(S_1, S_2, S_2) \equiv \frac{1}{3}\int_{S^2} \rho(S_1; x)\rho^2(S_2; x)dx \leq \sqrt[3]{\tilde{V}(S_1, S_1, S_1)\tilde{V}^2(S_2, S_2, S_2)}\ ,$$

with equality if and only if S_1, S_2 are dilations of each other (with the origin as the center of dilation). This means that among all starshaped sets S_1 of the same dual mixed volume $\tilde{V}(S_1, S_1, S_1)$, those that are dilations of S_2 yield the maximum of $\tilde{V}(S_1, S_2, S_2)$.

Define $\chi(R)$, a function of rotation R, by

$$\chi(R) \triangleq \tilde{V}(R(S_1), S_2, S_2) = \frac{1}{3}\int_{S^2} \rho(R(S_1); x)\rho^2(S_2; x)dx\ . \tag{4}$$

Then

$$\chi(R) \leq \sqrt[3]{\tilde{V}(R(S_1), R(S_1), R(S_1))\tilde{V}^2(S_2, S_2, S_2)}\ .$$

Note that $\tilde{V}(R(S_1), R(S_1), R(S_1))$ does not depend on R, that is,

$$\tilde{V}(R(S_1), R(S_1), R(S_1)) = \tilde{V}(S_1, S_1, S_1) ,$$

due to the rotation property (3) of the radial function. Then by Theorem 2.1 $\chi(R)$, as a function of R, reaches maximum if and only if $R(S_1) = \lambda S_2$, for some $\lambda > 0$. By Definition 3.1, the problem of attitude determination can be solved if and only if a maximal point of $\chi(R)$ can be found. The maximal point is a solution to the attitude determination problem. The global maximum of $\chi(R)$ is $\sqrt[3]{\tilde{V}(S_1, S_1, S_1)\tilde{V}^2(S_2, S_2, S_2)}$. By the if-and-only-if condition in Theorem 2.1, the global maximum is unique, modulo any rotational symmetries that S_1 possesses.

Rotation can be represented as the triple (ϕ, θ, Ω) where this is taken to mean counterclockwise rotation by angle Ω around unit vector $(\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$. When R is represented this way, the function $\chi(R)$ becomes $\chi(\phi, \theta, \Omega)$, a function of three variables that has domain R^3. The problem of attitude determination then becomes the following optimization problem:

$$\text{maximize: } \chi(\phi, \theta, \Omega) , \quad (\phi, \theta, \Omega) \in R^3 .$$

Since the objective function of this optimization problem is periodic and bounded, a solution to the optimization problem necessarily exists.

This approach described so far assumes that the radial function of both objects is known over the whole unit sphere. When $\rho(S_2; x)$ is known only at points clustered in one region of the unit sphere, the objective function, $\chi(R)$, defined in Equation (4), may not be appropriate. For example, when $\rho(S_2; x)$ is obtained from sensed data from a single viewpoint, the points at which the radial function is known will be at most a hemisphere. An altered objective function is defined to solve the attitude determination problem when the radial function is known only on a portion of the sphere.

Let V denote the smallest union of spherical polytopes[3] in S^2 that contains all the points x on S^2 where $\rho(S_2; x)$ is available. Again, by Hölder's inequality,

$$\int_V \rho(R(S_1); x)\rho^2(S_2; x)dx \leq \left[\int_V \rho^3(R(S_1); x)dx\right]^{\frac{1}{3}} \left[\int_V \rho^3(S_2; x)dx\right]^{\frac{2}{3}} , \tag{5}$$

with equality if and only if $\rho(R(S_1); x)$ and $\rho(S_2; x)$ are proportional to each other over V. Simply substituting V for S^2 in Equation (4), however, is not sufficient since the right side of (5) now depends on R (when V is not equal to S^2). Define a new objective function as

$$\overline{\chi}(R) = \frac{1}{3} \cdot \frac{\int_V \rho(R(S_1); x)\rho^2(S_2; x)dx}{[\int_V \rho^3(R(S_1); x)dx]^{\frac{1}{3}}} . \tag{6}$$

[3]A spherical polytope is the intersection of a finite number of closed hemispheres which is not empty and contains no pairs of antipodal points of S^2.

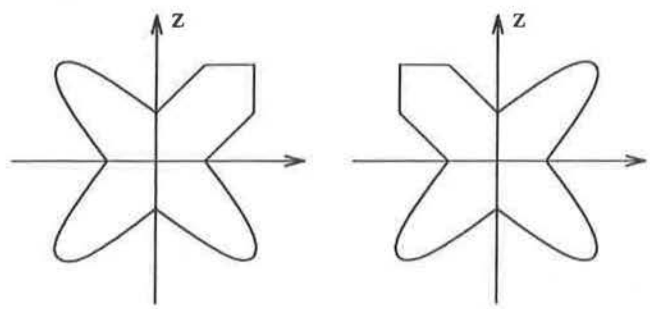

Figure 4: Two 2-D starshaped sets that match when viewed in the positive z direction but that do not match over the whole unit circle.

Hölder's inequality implies that $\overline{\chi}(R)$ achieves a maximum if and only if $R(S_1)$ and S_2 are dilations of each other over V. The maximal value is $\frac{1}{3}\left[\int_V \rho^3(S_2;x)dx\right]^{\frac{2}{3}}$. Thus, the part of the object where the radial function is sensed is matched to the model by maximizing $\overline{\chi}(R)$.

Strictly speaking, this does not solve the attitude determination problem, as defined in Definition 3.1, since part of the object is not 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 4 depicts two 2-D starshaped figures that match when viewed in the positive z direction but that do not match over the whole unit circle.

4 Experiments

Experiments were conducted on two real starshaped objects to demonstrate the feasibility of the approach. The objects are nicknamed the "peanut" and the "pillow", respectively. The radial functions of the two objects, defined for points on the unit sphere with standard spherical coordinates (ϕ, θ), are the following spherical harmonics:

$$\text{peanut} \;:\; \rho(\phi,\theta) = 1 + 3\cos^2(\phi)\,, \tag{7}$$
$$\text{pillow} \;:\; \rho(\phi,\theta) = 4 + 3\sin(2\theta)\sin^2(\phi)\,. \tag{8}$$

The parametric equations of the object surface, given in terms of the radial function, $\rho(\phi,\theta)$, and parameters ϕ and θ are

$$\begin{aligned} x &= \rho(\phi,\theta)\sin\phi\cos\theta\,, \\ y &= \rho(\phi,\theta)\sin\phi\sin\theta\,, \\ z &= \rho(\phi,\theta)\cos\phi\,. \end{aligned}$$

The peanut is a solid of revolution. A side view sketch is shown in Figure 5. Three side views of the pillow are shown in Figure 6. The two objects were custom fabricated from numerical data sampled from their radial functions. The objects are made of polyvinylchlorid and were machined by an automated, numerically controlled milling machine. A third object, a

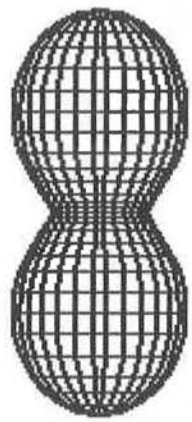

Figure 5: Experimental shape: the peanut. It is a solid of revolution.

---

sphere, was machined from the same material to serve as a calibration object for photometric stereo.

Photometric stereo was used to obtain surface data [19]. Three images of each object are taken under three different lighting conditions with the same imaging geometry. The images are shown in Figure 7 and Figure 8 respectively for the peanut and the pillow. Photometric stereo uses reflectance data obtained from the calibration sphere to determine surface gradient information at each visible point. The relative height of each surface point was obtained by reconstructing depth from gradient using the method of Harris [20]. Finally, the radial function was computed as the distance from a fixed point inside the kernel of the object to each visible surface point. The radial function does depend on the choice of this fixed point. By convention, the origin of the object coordinate system is taken to be the center of gravity of the object, whenever possible.

Since surface data are acquired from a single viewpoint, the radial function is not known over the entire sphere. Thus the objective function $\overline{\chi}(R)$, defined in Equation (6), is used instead of $\chi(R)$, defined in Equation (4).

Let S_1 be either the peanut or the pillow in a standard attitude, $S_2 = R_0(S_1)$, where R_0 is a fixed but unknown rotation. One can think of S_1 as the model and S_2 as the object. Let S^{2-} denote the set of all points x on S^2 such that the intersection of the ray $\overrightarrow{Ox}$ and the object surface is a point visible to the camera. Then the objective function, $\overline{\chi}(R)$, defined by Equation (6), is

$$\overline{\chi}(R) = \frac{1}{3} \cdot \frac{\int_{S^{2-}} \rho(R(S_1);x)\rho^2(R_0(S_1);x)dx}{[\int_{S^{2-}} \rho^3(R(S_1);x)dx]^{\frac{1}{3}}} . \tag{9}$$

The function $\rho(R(S_1);x) = \rho(S_1;R^{-1}(x))$ is known for both the peanut and the pillow from Equations (7) and (8). The function $\rho(R_0(S_1);x)$ is estimated for all surface points visible to the camera. The rotation, R_0, is estimated by maximizing $\overline{\chi}(R)$. (For comparison to ground

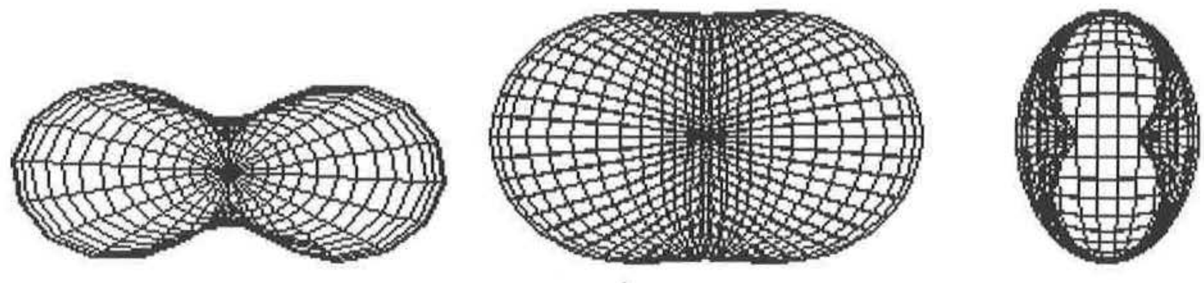

(a) top view (b) front view (c) side view

Figure 6: Experimental shape: the pillow.

truth, R_0 also is determined *a priori* as part of the experimental setup).

By the rotation property (3) of the radial function,

$$\overline{\chi}(R) = \frac{1}{3} \cdot \frac{\int_{S^{2-}} \rho(S_1; R^{-1}(x))\rho^2(R_0(S_1); x)dx}{[\int_{S^{2-}} \rho^3(S_1; R^{-1}(x))dx]^{\frac{1}{3}}} . \tag{10}$$

Suppose the viewpoint is in the positive z direction and suppose S^{2-} is the hemisphere corresponding to $z < 0$. Then, the surface integral (10) can be transformed into a volume integral:

$$\overline{\chi}(R) = \frac{1}{3} \cdot \frac{\int_{\frac{\pi}{2}}^{\pi} \int_0^{2\pi} \rho(S_1; R^{-1}(x(\phi,\theta)))\rho^2(R_0(S_1); x(\phi,\theta))\sin\phi d\theta d\phi}{\left[\int_{\frac{\pi}{2}}^{\pi} \int_0^{2\pi} \rho^3(S_1; R^{-1}(x(\phi,\theta)))\sin\phi d\theta d\phi\right]^{\frac{1}{3}}} , \tag{11}$$

where $x(\phi, \theta) = (\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$.

The numerical integration routine QB01AD from Harwell [21] is used to calculate the 2-D integral (11). The interpolation routine of Renka [22, 23] is used to interpolate values of $\rho(R_0(S_1); x)$ over S^{2-}, as required. The result determines a C^1 function from an irregular distribution of data samples on the sphere.

A regular tessellation of the sphere is used when a radial function is discretized from a known analytic form. The tessellation used is based on a geodesic dome built from an icosahedron with a desired frequency [24]. Figure 9 shows the 8-frequency geodesic dome. It has 642 vertices, 1920 arcs, and 1280 facets.

The nonlinear programming subroutine NLPQL [25] is used to find the maxima of $\overline{\chi}(R)$. 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 gradient of the objective function, $\overline{\chi}(R)$, is estimated by simple forward differencing. Convergence is achieved either when the Kuhn-Tucker conditions (see [26] page 51) are satisfied to within a specified accuracy or when

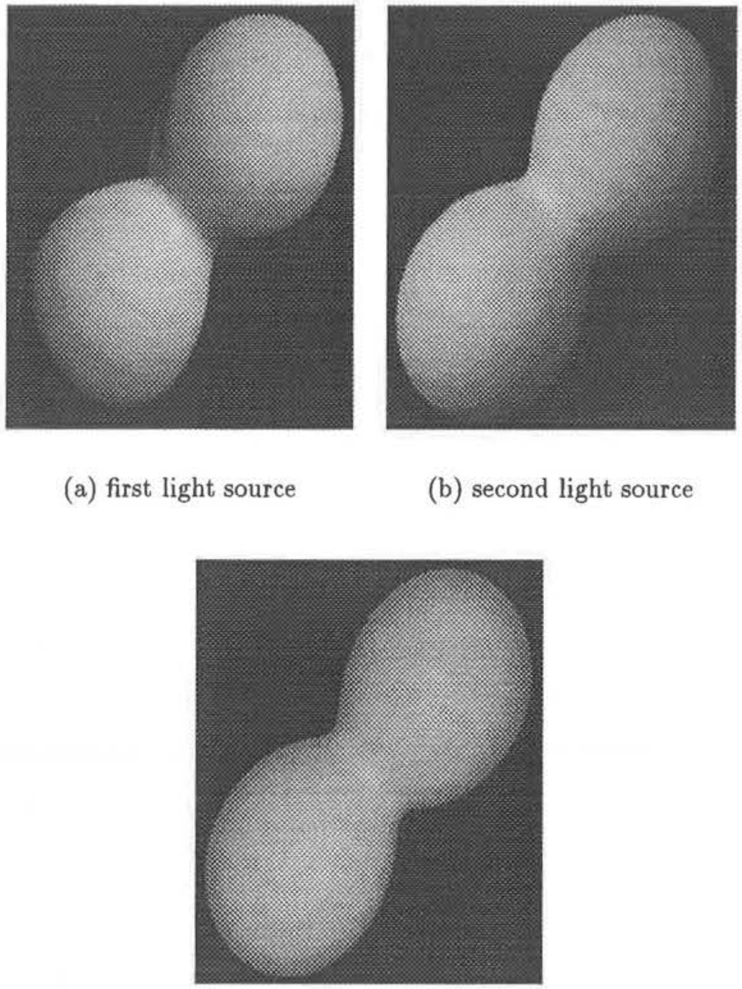

(a) first light source (b) second light source

(c) third light source

Figure 7: Images of peanut under three light sources.

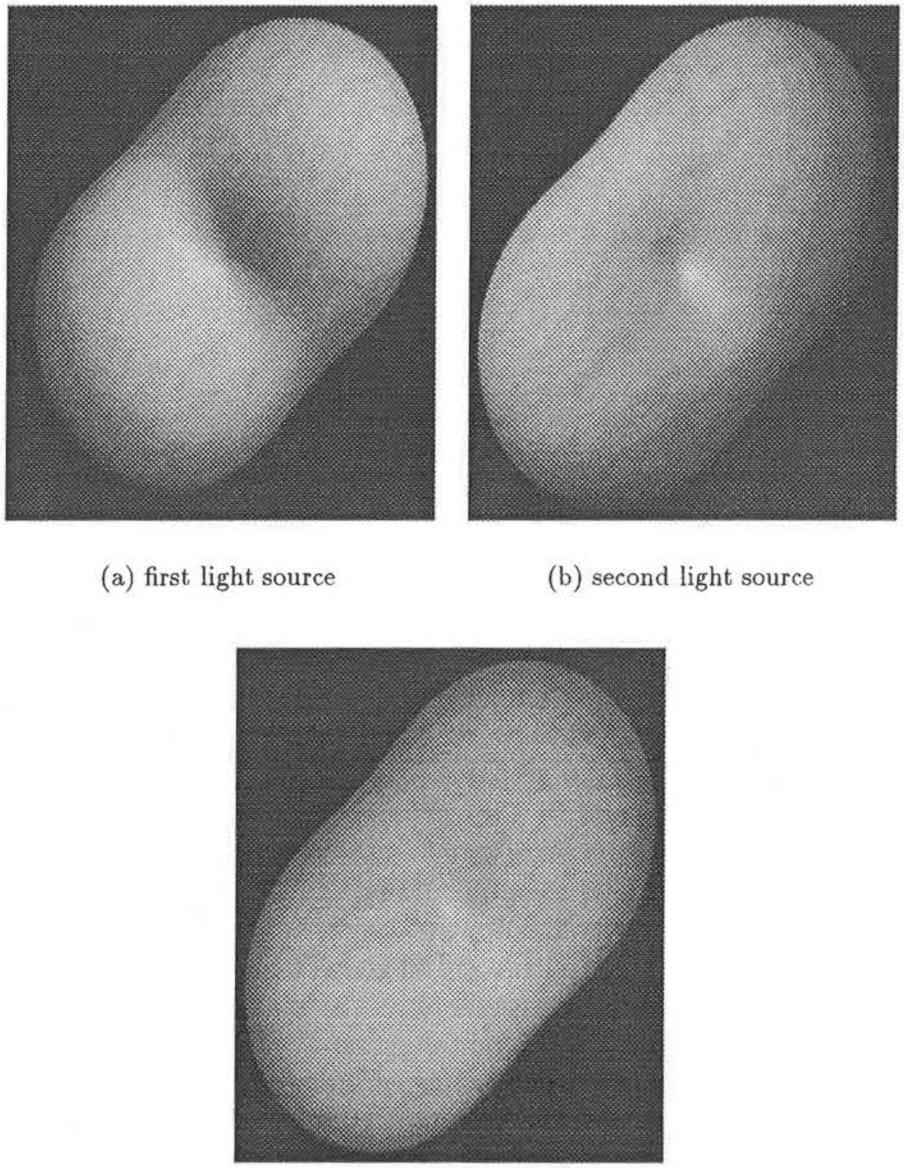

(a) first light source (b) second light source

(c) third light source

Figure 8: Images of pillow under three light sources.

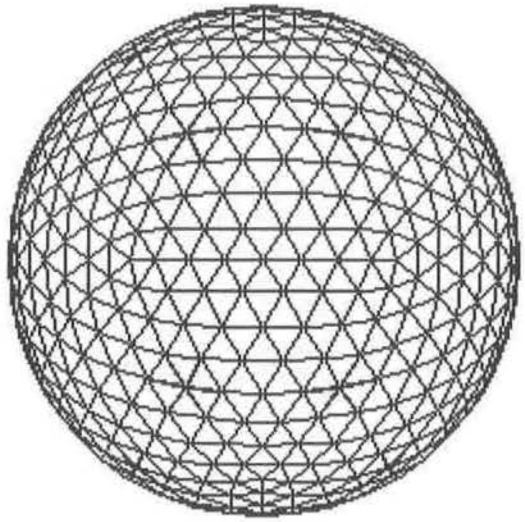

Figure 9: An 8-frequency geodesic dome.

the objective function is not improved significantly given that the constraints are satisfied to within the specified accuracy. NLPQL does not guarantee that the maximum found is global. But, in this application, the global maximum of $\overline{\chi}(R)$ is known. Therefore, it is known whether the maximum found by the subroutine is the global maximum.

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. For both the peanut and the pillow, all the initial guesses converged to points with the same maxiumum value of $\chi(R)$ and with the same object attitude. Thus, for the objects tested, 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.

When an object is highly symmetric, like the peanut, it is difficult to evaluate the optimization results by comparing rotation matrices because different matrices correspond to the identical object attitude. The optimization process for the peanut with initial guess $(0.1, 0.2, 0.3)$ in radians is shown in Figure 10. In the figures, the black and white shows the silhouette of the object and the wire frame shape in gray is the rotated model.

When an object has few symmetries, like the pillow, it is possible to evaluate the optimization results by comparing rotation matrices. The estimated *a priori* rotation matrix for the pillow imaged in Figure 8 is

$$\begin{bmatrix} -0.5074229658 & -0.5061199538 & 0.6973983984 \\ 0.5418144635 & 0.4419270649 & 0.7149388481 \\ -0.6700440447 & 0.7406369305 & 0.0499791693 \end{bmatrix}. \tag{12}$$

The rotation matrices estimated by optimization are either

$$\begin{bmatrix} -0.5419370962 & -0.4890334995 & 0.6834840307 \\ 0.5313071435 & 0.4307655958 & 0.7294886712 \\ -0.6511658031 & 0.7584769202 & 0.0263791359 \end{bmatrix}, \tag{13}$$

or

$$\begin{bmatrix} -0.4890250728 & -0.5419388271 & -0.6834886874 \\ 0.4308033845 & 0.5312805744 & -0.7294857062 \\ 0.7584608902 & -0.6511860404 & -0.0263404405 \end{bmatrix}. \tag{14}$$

Matrix (14) is equal to matrix (13) multiplied by a rotation of 180 degree around direction $(1, 1, 0)$. They define the same attitude since the pillow as imaged in Figure 8 is symmetric about the line determined by the origin and the direction $(1, 1, 0)$. The optimization process for the pillow with initial guess $(0.1, 0.2, 0.3)$ in radians is shown in Figure 11.

5 Conclusions

Orientation-based representations are a compact description of 3-D object shape. A 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. Dense surface data measured from an unknown viewpoint determines a visible hemisphere of the representation. Matching the visible hemisphere to the full spherical model can be formulated as a single, uniform optimization process. In particular, one does not need multiple viewpoint dependent representations of a modeled object.

A new orientation-based representation has been introduced based on the radial function and the dilation map. A method that uses this representation to determine 3-D attitude has been theoretically justified and empirically demonstrated for the class of objects known as (compact) starshaped sets. Starshaped sets are described in the mathematical literature and have been used before in computational geometry to study problems of visibility. Starshaped objects appear here as a useful generalization of convexity that extends, in a principled way, previous work on shape matching.

The method transforms the attitude determination problem for starshaped sets 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 treatment given here for orientation-based representations implicitly assumes that the object surface is smooth (i.e., C^2). Curvature-based representations under the Gauss

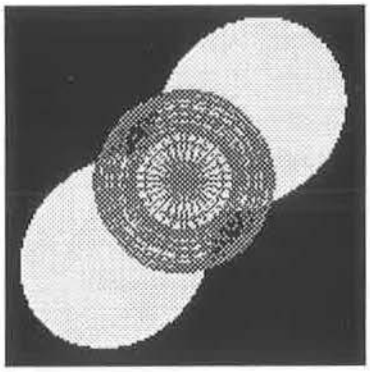

(a) initial guess.

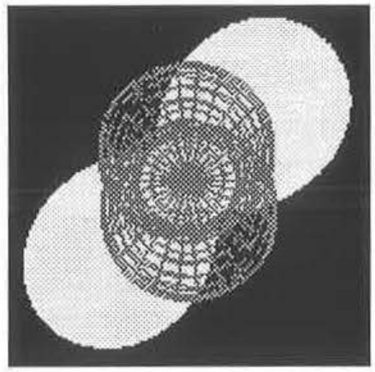

(b) 3rd iteration.

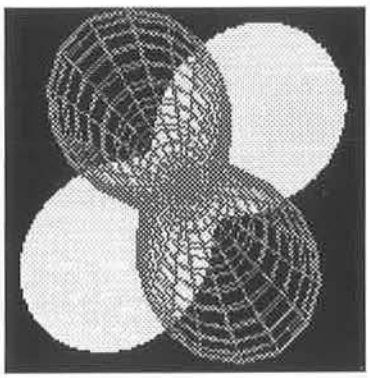

(c) 5th iteration.

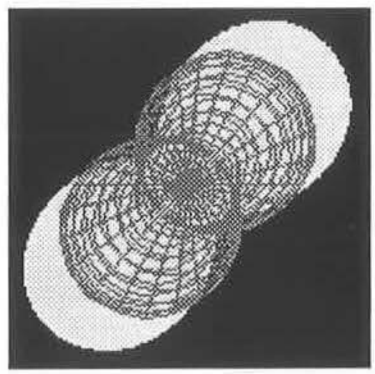

(d) 7th iteration.

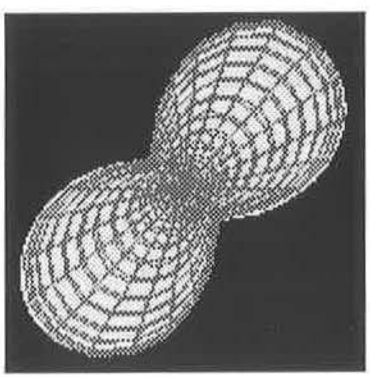

(e) 8th iteration.

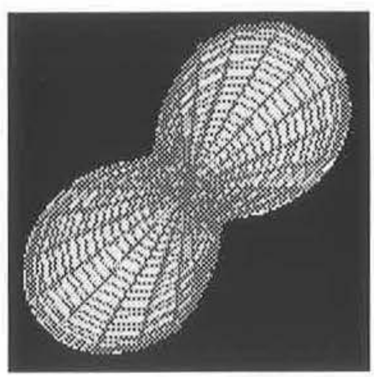

(f) final result.

Figure 10: Results of real data 3-D attitude determination for the peanut using its radial function.

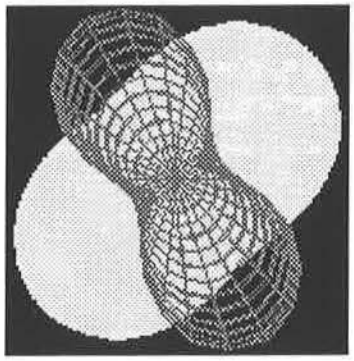

(a) initial guess.

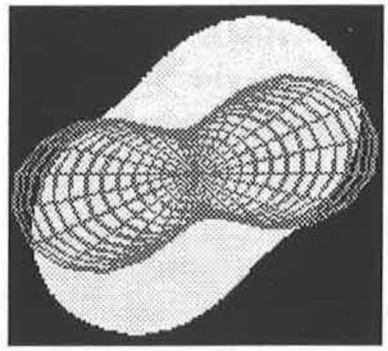

(b) 3rd iteration.

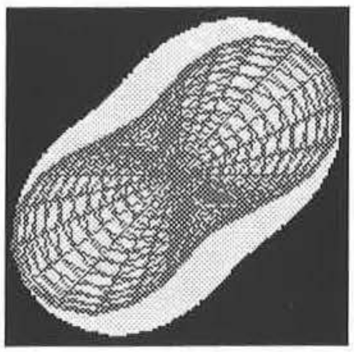

(c) 6th iteration.

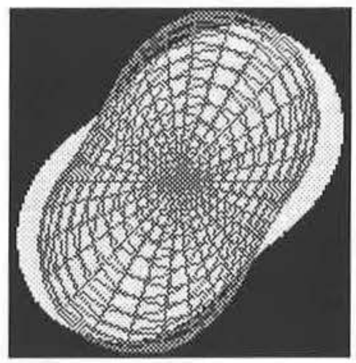

(d) 7th iteration.

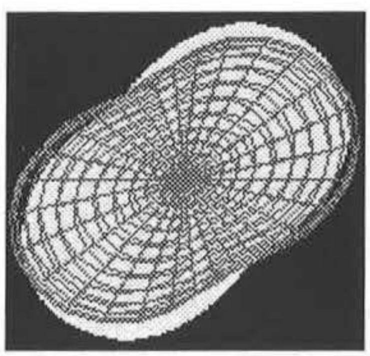

(e) 9th iteration.

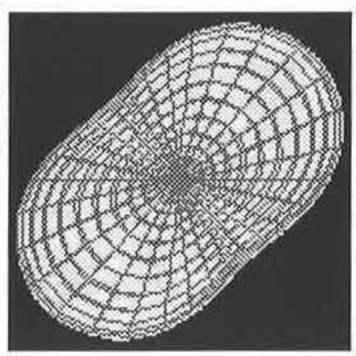

(f) final result.

Figure 11: Results of real data 3-D attitude determination for the pillow using its radial function.

map require a (technically) different treatment for the case of polyhedra. For example, the second area function, which is the analog of the EGI for polyhedra, is defined only for a discrete set of orientations, one for each face of the polyhedron. With the radial function, one need not treat smooth objects and polytopes differently. The radial function is defined for all points on the unit sphere regardless whether the object is smooth, a polytope or a combination[4].

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 surface data. Based on the experimental work performed to date, the overall accuracy of the method is consistent with the best 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. The radial function does, of course, depend on the choice of coordinate system origin. For most object shapes, the radial function will not change significantly for (slight) changes in the location of the origin. Intuitively, this suggests that the method is stable with respect to choice of origin, provided the origin is within the kernel of the starshaped set. Experiments on synthetic data support this intuition. The method maintains accuracy and robustness with respect to the choice of coordinate system origin.

The approach is intended for recognition, localization and inspection tasks using dense surface data that can be obtained from laser ranging, shape-from-shading or photometric stereo. The work here used data obtained from photometric stereo. It would be useful to experiment with other sources of dense surface data. 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 time t can be used as the initial guess at time $t + \Delta t$.

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

[4]This also is true for the support function and is one of the reasons conjectured for its success in polyhedral shape matching [6].

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] Y. Li, "Attitude determination by the support function and the curvature functions," in *Vision Interface 92*, (Vancouver, BC), pp. 209–214, 1992.

[4] 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.

[5] 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.

[6] 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.

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

[8] 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.

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

[10] 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.

[11] 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.

[12] M. Breen, "A characterization theorem for bounded starshaped sets in the plane," *Proceedings of the American Mathematical Society*, vol. 94, pp. 693–698, 1985.

[13] F. A. Valentine, *Convex Sets*. McGraw-Hill Series in Higher Mathematics, New York: McGraw-Hill Book Company, 1964.

[14] N. S. Friedland and A. Rosenfeld, "Compact object recognition using energy-function-based optimization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 770–777, July 1992.

[15] B. B. Schudy and D. H. Ballard, "A computer model for extracting moving heart surfaces from four-dimensional cardiac ultrasound data," in *Proceedings of Sixth Conference on Computer Applications in Radiology and Computer/Aided Analysis of Radiological Images*, pp. 366–376, 1979.

[16] E. Lutwak, "Dual mixed volumes," *Pacific Journal of Mathematics*, vol. 58, pp. 531–538, 1975.

[17] T. Bonnesen and W. Fenchel, *Theory of Convex Bodies.* Moscow, Idaho, USA: BCS Associates, 1987. Translated from German (Theorie der konvexen Körper) and edited by L. Boron, C. Christenson and B. Smith.

[18] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities.* Cambridge: Cambridge University Press, 1964.

[19] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Optical Engineering*, vol. 19, pp. 139–144, 1980.

[20] J. G. Harris, "The coupled depth/slope approach to surface reconstruction," Technical Report 908, Artificial Intelligence Laboratory, MIT, 1986.

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

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

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

[24] J. D. Clinton, "Advanced structural geometry studies, part I - polyhedral subdivision concepts for structural applications," CR 1734, NASA, 1971.

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

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