

Extended Gaussian Images

BERTHOLD K. P. HORN

This is a primer on extended Gaussian images. Extended Gaussian images are useful for representing the shapes of surfaces. They can be computed easily from:

1. *needle maps obtained using photometric stereo; or*
2. *depth maps generated by ranging devices or binocular stereo. Importantly, they can also be determined simply from geometric models of the objects. Extended Gaussian images can be of use in at least two of the tasks facing a machine vision system:*

1. *recognition, and*
 2. *determining the attitude in space of an object.*
- Here, the extended Gaussian image is defined and some of its properties discussed. An elaboration for nonconvex objects is presented and several examples are shown.*

I. INTRODUCTION

In order to recognize an object and to determine its attitude in space, it is necessary to have a way of representing the shape of its surface. Giving the distance to the surface along parallel rays on a regularly spaced grid provides one way of doing this. This simple representation is called a depth map. A range finder produces surface descriptions in this form, as does automated binocular stereo [1]. Unfortunately, depth maps do not transform in a simple way when the object rotates. (For one thing, interpolation must be used to get a new depth map on a regularly spaced grid.)

Alternatively, surface orientation might be given for points on the surface on some regular sampling grid. This grid may conveniently correspond to the picture cells in an image. Such a simple representation is called a needle map (Fig. 1) [2]. Photometric stereo is a method for recovering surface orientation using multiple images taken with different lighting conditions [3]–[8]. It produces surface descriptions in this form. A needle map also is not directly helpful when it comes to comparing surfaces of objects that may be rotated relative to one another. (Both depth maps and needle maps depend on the *position* of the object as well as its attitude.)

The extended Gaussian image, on the other hand, does make it easy to deal with the varying attitude of an object in space as we shall see [2], [9]–[13]. For one thing, it is insensitive to the position of the object. Some information appears to be discarded in the formation of the extended Gaussian image. Curiously, in the case of convex objects,

Manuscript received December 31, 1983; revised July 18, 1984. This paper describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Office of Naval Research (ONR) under ONR Contract N00014-77C-0389 and in part by the Advanced Research Projects Agency of the Department of Defense under ONR Research Contract N00014-80-C-0505.

The author is with the Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.




Fig. 1. A needle map shows unit surface normals at points on the surface on a regular grid. Normals which point towards the viewer will be seen as dots, while tilted surface patches give rise to normals which are shown as lines pointing in the direction of steepest descent.

the representation is nevertheless unique. That is, no two convex objects have the same extended Gaussian image.

This representation of the shape of the surface of an object allows one to match information obtained from image or range sensors with that contained in computer models of the objects and has proven most useful in work on automatic bin picking [14]–[16]. A recent report describes a system that picks one object out of a jumbled pile of similar objects using this approach [17]. We start our discussion with objects having planar faces. Later we consider smoothly curved objects. Methods for computing discrete approximations of extended Gaussian images, called orientation histograms, are presented too. Orientation histograms can be computed from experimental data or mathematical descriptions of the objects. Sections marked with an asterisk may be omitted on first reading or if your interest in the mathematical details is limited.

II. DISCRETE CASE: CONVEX POLYHEDRA

Minkowski showed in 1897 that a convex polyhedron is fully specified (up to translation) by the area and orientation of its faces [18]–[20]. We can represent area and orientation of the faces conveniently by point masses on a sphere. Imagine moving the unit surface normal of each face so that its tail is at the center of a unit sphere. The

head of the unit normal then lies on the surface of the unit sphere. This sphere is called the Gaussian sphere and each point on it corresponds to a particular surface orientation. The extended Gaussian image of the polyhedron is obtained by placing a mass at each point equal to the surface area of the corresponding face (Fig. 2).




Fig. 2. The extended Gaussian image of polyhedron can be thought of as a collection of point masses on the Gaussian sphere. Each mass is proportional to the area of the corresponding face. Point masses on the visible hemisphere are shown as solid marks, the others as open marks. The center of mass (shown as the symbol \otimes) must be at the center of the sphere if the polyhedron is a closed subject.

It seems at first, as if some information is lost in this mapping, since the *position* of the surface normals is discarded. Viewed another way, no note is made of the shape of the faces or their adjacency relationships. It can nevertheless be shown that (up to translation) the extended Gaussian image uniquely defines a convex polyhedron [9]. An iterative algorithm has recently been invented for recovering a convex polyhedron from its extended Gaussian image [21].

A. Properties of the Extended Gaussian Image

The extended Gaussian image is not affected by translation of the object. Rotation of the object induces an equal rotation of the extended Gaussian image, since the unit surface normals rotate with the object.

Mass distributions which lie entirely within one hemisphere, that is, are zero in the complementary hemisphere, do not correspond to closed objects. As we shall see, the center of mass of an extended Gaussian image has to lie at the origin. This is clearly not possible if a whole hemisphere is empty. Also, a mass distribution which is nonzero only on a great circle of the sphere corresponds to the limit of a sequence of cylindrical objects of increasing length and decreasing diameter (Fig. 3). We will exclude such pathological cases and confine our attention to closed, bounded objects [9], [20].

Some properties of the extended Gaussian image are important: First, the total mass of the extended Gaussian image is obviously just equal to the total surface area of the polyhedron. If the polyhedron is closed, it will have the same projected area when viewed from any pair of opposite directions. This allows us to compute the location of the center of mass of the extended Gaussian image.

Imagine viewing a convex polyhedron from a great distance. Let the direction from the object towards the viewer




Fig. 3. A mass distribution confined to a great circle corresponds to the limit of a sequence of cylindrical objects of increasing length and decreasing diameter. Such pathological mass distributions can be avoided if we confine our attention to bounded objects.

be given by the unit vector \hat{v} . A face, with unit normal \hat{s}_i , will be visible only if $\hat{s}_i \cdot \hat{v} \geq 0$. Suppose that the surface area of this face is O_i . Due to foreshortening it will appear only as large as would a face of area

$$(\hat{s}_i \cdot \hat{v}) O_i$$

normal to \hat{v} (Fig. 4). The total apparent area of the visible surface is

$$A(\hat{v}) = \sum_{\{\hat{s}_i \cdot \hat{v} > 0\}} (\hat{s}_i \cdot \hat{v}) O_i$$

when viewed from the direction \hat{v} . The total apparent area of the visible surface when viewed from the opposite




Fig. 4. A surface element appears smaller because of foreshortening. The apparent area is the true area times the cosine of the angle between the surface normal and the vector pointing towards the viewer.

direction is

$$A(-\hat{v}) = \sum_{\{\hat{s}_i \cdot \hat{v} < 0\}} (\hat{s}_i \cdot \hat{v}) O_i.$$

This should be the same, that is, $A(\hat{v}) = A(-\hat{v})$. Consequently,

$$\sum_{\text{all } i} (\hat{s}_i \cdot \hat{v}) O_i = \left[\sum_{\text{all } i} \hat{s}_i O_i \right] \cdot \hat{v} = 0$$

where the sum now is over all faces of the object. This holds true for all view vectors, \hat{v} , so we must have

$$\sum_{\text{all } i} \hat{s}_i O_i = 0.$$

That is, the center of mass of the extended Gaussian image is at the origin.

An equivalent representation, called a spike model, is a collection of vectors each of which is parallel to one of the surface normals and of length equal to the area of the corresponding face. The result regarding the center of mass is equivalent to the statement that these vectors must form a closed chain when placed end to end (Fig. 5).




Fig. 5. Vectors parallel to the normals of the faces of a polyhedron, and of length equal to the areas of the corresponding faces, form a closed chain when placed end to end.

B. Reconstruction of a Tetrahedron^(*)

Faces that share a common edge are said to be adjacent. The masses on the Gaussian sphere corresponding to two adjacent faces need not be each others closest neighbors. Recovering a polyhedron from its extended Gaussian image is not easy in the general case because it is hard to determine which faces are adjacent [21]. Finding the actual offsets of each of the faces from the center of mass of the polyhedron is not as hard.

The structure of a tetrahedron, however, is very simple: Every face is adjacent to the other three. The *shape* of the tetrahedron is completely determined by the surface normals of the four faces, only the *size* of the tetrahedron remaining to be determined. In other words: There is only one degree of freedom left. Another way to look at it is to note that the four faces must have areas that place the center of mass of the extended Gaussian image at the origin, as we have just seen. This condition places three constraints on the four parameters.

Let the given unit surface normals be \hat{a} , \hat{b} , \hat{c} , and \hat{d} , and the areas of the corresponding faces, A , B , C , and D (Fig. 6). We have to determine the distances, a , b , c , and d , of




Fig. 6. A tetrahedron with vertices A , B , C , and D . We are to find the distances of the faces from the center of mass, given the areas and surface normals of the faces.

these faces from the center of mass of the tetrahedron. From these distances we can, if desired, compute the positions of the vertices A , B , C , and D , simply by intersecting three of the planes at a time. The notation here is that the face opposite vertex A has area A and unit surface normal \hat{a} , and so on.

The perpendicular distance of the center of area of a triangle from one of the sides is one third the perpendicular distance of the vertex opposite that side. Similarly, in a

tetrahedron, the distance from the center of mass to a particular face is equal to one quarter of the distance of the vertex opposite that face. We start by finding a formula for the distance of the face with area D , say, from the opposite vertex d . The desired distance, d , will be just a quarter of the result obtained in this fashion. The remaining three distances, a , b , and c , can then be computed using formulas obtained by cyclical permutation of the variables.

The position of the reconstructed tetrahedron is arbitrary, since the extended Gaussian image is insensitive to translation. To make the result unique, we might place the center of mass at the origin. To reduce the size of the expressions to be manipulated here, however, it is convenient to move the tetrahedron so that one vertex, D , say, is at the origin. The distances of the faces from the center of mass are obviously not affected by this.

Suppose for now that we know the locations of the vertices A , B , and C relative to D . We can then compute the directions of the six edges of the tetrahedron, by taking all of the distinct pairwise differences of the four vertex positions. Four surface normals can then be found by taking cross-products of these edge-direction vectors. We actually need only four of the edge vectors forming a closed circuit to do this. The results can then be normalized to obtain unit surface normals

$$\hat{a} = -\frac{B \times C}{|B \times C|} \quad \hat{b} = -\frac{C \times A}{|C \times A|} \quad \hat{c} = -\frac{A \times B}{|A \times B|}$$

and

$$\hat{d} = \frac{(A - C) \times (B - A)}{|(A - C) \times (B - A)|} = \frac{A \times B + B \times C + C \times A}{|A \times B + B \times C + C \times A|}$$

Now the perpendicular distance of the plane with area D from the origin can be found by taking the dot-product of any of the three vertices, A , B , and C with the unit normal \hat{d} . Thus

$$4d = \hat{d} \cdot A = \hat{d} \cdot B = \hat{d} \cdot C = \frac{[ABC]}{|A \times B + B \times C + C \times A|}$$

The area of the facet opposite the origin is also easy to compute

$$D = \frac{1}{2} |(A - C) \times (B - A)| = \frac{1}{2} |A \times B + B \times C + C \times A|$$

Our task is to express the offset d in terms of the area D and the given unit surface normals. The two formulas above do not allow us to do that directly, because we do not know what the value of $[ABC]$ is. This quantity, by the way, is six times the volume, V , of the tetrahedron, or

$$V = \frac{1}{3} (4d) D = \frac{1}{6} [ABC]$$

We proceed by considering the four distinct triple products of the four unit surface normals. First of all

$$[\hat{a} \hat{b} \hat{c}] = -\frac{[ABC]^2}{|A \times B| |B \times C| |C \times A|}$$

since $[(x \times y)(y \times z)(z \times x)] = [xyz]^2$. Then, by similar reasoning,

$$[\hat{a} \hat{b} \hat{d}] = \frac{[ABC]^2}{|B \times C| |C \times A| |A \times B + B \times C + C \times A|}$$

since $[xy(x+y+z)] = [xyz]$. Formulas for $[b\hat{c}\hat{d}]$ and $[c\hat{a}\hat{d}]$ can be found by cyclical permutation of the variables.

Multiplying the three formulas found this way together we get

$$\begin{aligned} & [a\hat{b}\hat{d}][b\hat{c}\hat{d}][c\hat{a}\hat{d}] \\ &= \frac{[ABC]^6}{(A \times B)^2(B \times C)^2(C \times A)^2 |A \times B + B \times C + C \times A|^3} \end{aligned}$$

and so

$$\begin{aligned} \frac{[a\hat{b}\hat{d}][b\hat{c}\hat{d}][c\hat{a}\hat{d}]}{[a\hat{b}\hat{c}]^2} &= \frac{[ABC]^2}{|A \times B + B \times C + C \times A|} \\ &= (4d)^2(2D). \end{aligned}$$

So that finally

$$4d = \frac{\sqrt{(2D)[a\hat{b}\hat{d}][b\hat{c}\hat{d}][c\hat{a}\hat{d}]}}{-[a\hat{b}\hat{c}]}$$

The other distances, a , b , and c , can be computed using similar formulas obtained by cyclical permutation of the variables.

III. CONTINUOUS CASE: SMOOTHLY CURVED OBJECTS

The ideas presented in the previous section can be extended to apply to smoothly curved surfaces.

A. Gaussian Image

One can associate a point on the Gaussian sphere with a given point on a surface by finding the point on the sphere which has the same surface normal (Fig. 7) [20], [22], [25]. Thus it is possible to map information associated with




Fig. 7. The Gaussian image of an object is obtained by associating with each point on its surface the point on the Gaussian sphere which has the same surface orientation. The mapping is invertible if the object has positive Gaussian curvature everywhere.

points on the surface onto points on the Gaussian sphere. In the case of a convex object with positive Gaussian curvature everywhere, no two points have the same surface normal. The mapping from the object to the Gaussian sphere in this case is invertible: Corresponding to each point on the Gaussian sphere there is a unique point on the surface. (If the convex surface has patches with zero Gaussian curvature, curves or even areas on it may correspond to a single point on the Gaussian sphere.)

One useful property of the Gaussian image is that it rotates with the object. Consider two parallel surface normals, one on the object and the other on the Gaussian sphere. The two normals will remain parallel if the object

and the Gaussian sphere are rotated in the same fashion. A rotation of the object thus corresponds to an equal rotation of the Gaussian sphere.

B. Gaussian Curvature

Consider a small patch δO on the object. Each point in this patch corresponds to a particular point on the Gaussian sphere. The patch δO on the object maps into a patch, δS say, on the Gaussian sphere (Fig. 8). If the surface is strongly




Fig. 8. A patch on the object maps into a patch on the Gaussian sphere. The Gaussian curvature is the limit of the ratio of the area of the patch on the Gaussian sphere to the area of the patch on the subject as these become smaller and smaller.

curved, the normals of points in the patch will point into a wide fan of directions. The corresponding points on the Gaussian sphere will be spread out. Conversely, if the surface is planar, the surface normals are parallel and map into a single point.

These considerations suggest a suitable definition of curvature. The Gaussian curvature is defined to be equal to the limit of the ratio of the two areas as they tend to zero. That is,

$$K = \lim_{\delta O \rightarrow 0} \frac{\delta S}{\delta O} = \frac{dS}{dO}$$

From this differential relationship we can obtain two useful integrals. Consider first integrating K over a finite patch O on the object:

$$\int_O \int K dO = \int_S \int dS = S$$

where S is the area of the corresponding patch on the Gaussian sphere. The expression on the left is called the integral curvature. This relationship allows one to deal with surfaces which have discontinuities in surface normal.

Now consider instead integrating $1/K$ over a patch S on the Gaussian sphere

$$\int_S \int 1/K dS = \int_O \int dO = O$$

where O is the area of the corresponding patch on the object. This relationship suggests the use of the *inverse* of the Gaussian curvature in the definition of the extended Gaussian image of a smoothly curved object, as we shall see. It also shows, by the way, that the integral of $1/K$ over the whole Gaussian sphere equals the total area of the object.

C. Alternate Definition of Gaussian Curvature

Consider a plane which includes the surface normal at some point on a smooth surface. The surface cuts this plane




Fig. 9. Normal sections of the surface are made with planes which include the surface normal. The planes corresponding to the largest and smallest values of curvature are referred to as the principal planes. The Gaussian curvature is equal to the product of the largest and smallest values of curvature.

along a curve called a normal section (Fig. 9) [19], [22]–[25]. Let the curvature of the normal section be denoted by κ_N . Consider the one-parameter family of planes containing the surface normal. Suppose that θ is the angle between a particular plane and a given reference plane. Then κ_N varies with θ in a periodic fashion. In fact, if we measure θ from the plane that gives maximum curvature, then it can be shown that

$$\kappa_N(\theta) = \kappa_1 \cos^2 \theta + \kappa_2 \sin^2 \theta$$

where κ_1 is the maximum and κ_2 is the minimum curvature. These two values of κ are called the principal curvatures. The corresponding planes are called the principal planes. The two principal planes are orthogonal, provided that the principal curvatures are distinct (Fig. 9).

It turns out that

$$K = \kappa_1 \kappa_2$$

is equal to the Gaussian curvature introduced earlier. This is clearly zero for a plane. It is equal to $1/R^2$ for a spherical surface of radius R , since the curvature of any normal section is $1/R$.

A ruled surface is one which can be generated by sweeping a line through space. A hyperboloid provides one example of such a surface. Developable surfaces are special cases of ruled surfaces [22], [23], [25]. Cylindrical and conical surfaces are examples of developable surfaces (Fig. 10). On a developable surface at least one of the two principal




Fig. 10. A conical surface is an example of a developable surface. On it the Gaussian curvature is everywhere zero, because (at least) one of the principal curvatures is zero.

curvatures is zero at all points. Consequently, the Gaussian curvature is zero everywhere too.

D. The Extended Gaussian Image

We can define a mapping which associates the inverse of the Gaussian curvature at a point on the surface of the object with the corresponding point on the Gaussian sphere. Let u and v be parameters used to identify points on the original surface. Similarly, let ξ and η be parameters used to identify points on the Gaussian sphere. (These could be longitude and latitude, for example.) Then we define the extended Gaussian image as

$$G(\xi, \eta) = \frac{1}{K(u, v)}$$

where (ξ, η) is the point on the Gaussian sphere which has the same normal as the point (u, v) on the original surface. It can be shown that this mapping is unique (up to translation) for convex objects. That is, there is only one convex object corresponding to a particular extended Gaussian image [9], [19], [26]. The proof is unfortunately nonconstructive and no direct method for recovering the object is known.

E. Properties of the Extended Gaussian Image

The center of mass of the extended Gaussian image of a smoothly curved object is at the origin. We show this in a way similar to that used earlier for extended Gaussian images of polyhedral objects. Consider viewing a convex object from a great distance. Let the direction from the object towards the viewer be given by the unit vector \hat{v} . A surface patch with unit normal \hat{s} will be visible only if $\hat{s} \cdot \hat{v} \geq 0$. Suppose its surface area is δO (Fig. 4). Due to foreshortening it will appear only as large as would a patch of area

$$(\hat{s} \cdot \hat{v}) \delta O$$

normal to \hat{v} . Let $H(\hat{v})$ be the unit hemisphere for which $\hat{s} \cdot \hat{v} \geq 0$. Then the apparent area of the visible surface is

$$A(\hat{v}) = \int_{H(\hat{v})} G(\hat{s}) (\hat{s} \cdot \hat{v}) dS$$

when viewed from the direction \hat{v} . The apparent area of the visible surface when viewed from the opposite direction is

$$A(-\hat{v}) = \int_{H(-\hat{v})} G(\hat{s}) (\hat{s} \cdot -\hat{v}) dS.$$

This should be the same, that is, $A(\hat{v}) = A(-\hat{v})$. Consequently,

$$\int_S G(\hat{s}) (\hat{s} \cdot \hat{v}) dS = \left(\int_S G(\hat{s}) \hat{s} dS \right) \cdot \hat{v} = 0$$

where the integral now is over the whole sphere S . This holds true for all view vectors, \hat{v} , so we must have

$$\int_S G(\hat{s}) \hat{s} dS = 0.$$

That is, the center of mass of the extended Gaussian image is at the origin. (This, by the way, is not a very helpful constraint in practice, since one usually only sees one side of the object.)

Another property of the extended Gaussian image is also easily demonstrated. The total mass of the extended Gauss-

ian image equals the total surface area of the object. If one wishes to deal with objects of the same shape but differing size one may normalize the extended Gaussian image by dividing by the total mass.

We can think of the extended Gaussian image in terms of mass density on the Gaussian sphere. It is possible then to deal in a consistent way with places on the surface where the Gaussian curvature is zero, using the integral of $1/K$ shown earlier. A planar region, for example, corresponds to a point mass. This in turn corresponds to an impulse function on the Gaussian sphere with magnitude proportional to the area of the planar region.

A mass distribution has inertia about an axis passing through its center of mass that depends on the direction of the axis. This inertia takes on three stationary values, for three particular orthogonal directions, called the principal axes of the object. It is tempting to imagine that one can find the attitude of an object by lining up the principal axes of inertia of the observed extended Gaussian image and the one computed from the geometric model [9]. This would be rather straightforward, requiring only the calculation of the eigenvectors of a three by three inertia matrix. In practice, one typically has information only about the visible hemisphere and thus cannot compute the required first and second moments over the whole sphere.

F. Objects that are not Convex

Three things happen when the surface is nonconvex:

1. The Gaussian curvature for some points will be negative.
2. More than one point on the object will contribute to a given point on the Gaussian sphere.
3. Parts of the object may be obscured by other parts.

We chose to extend the definition of the extended Gaussian image in this case to be the sum of the *absolute* values of the inverses of the Gaussian curvature at all points having the same surface orientation

$$G(\xi, \eta) = \sum_i \frac{1}{|K(u_i, v_i)|}$$

This definition is motivated by the method used to compute the extended Gaussian image in the discrete case, as we will see later.

The above extension makes sense if there are a finite, or at most a countable, number of points on the surface with the same orientation. At times, however, all points on a curve or even an area on the surface have parallel surface normals. In this case we may use

$$G(\hat{n}) = \int_S \delta(\hat{n} - \hat{s}) \int_O \delta(\hat{o} - \hat{s}) dO dS$$

where \hat{n} is a unit vector on the Gaussian sphere, while \hat{s} is a unit vector on the surface of the object. The integration is over the whole surface of the object O and δ is the unit impulse function defined on a sphere.

We can be more specific, if we let $r(u, v)$ be a vector giving the point on the surface corresponding to the parameters u and v , then

$$G(\xi, \eta) \cos \eta = \iint \delta(\xi - \theta(u, v), \eta - \phi(u, v)) \cdot |r_u \times r_v| du dv$$

where $\theta(u, v)$ and $\phi(u, v)$ are the latitude and longitude of

the point on the Gaussian sphere which has the same orientation as the surface does at the point (u, v) . A planar region of area A will thus contribute an impulse of weight A to the extended Gaussian image, while a cylindrical region will give rise to an impulse wall along a great circle at right angles to the axis of the cylinder. The integral of the impulse wall will be equal to the area of the cylindrical region.

Usually we think of the extended Gaussian image as a fixed entity associated with an object. In the case of non-convex objects we might want to alter the definition to include only those parts of the surface visible from a particular direction. This would make the (modified) Gaussian image dependent on the viewpoint. We avoid this potential complication here.

G. Examples of Extended Gaussian Images^(*)

The extended Gaussian image of a sphere of radius R is

$$G(\xi, \eta) = R^2$$

as discussed already.

Perhaps slightly more interesting is the case of an ellipsoid with semi-axes a , b , and c lined up with the




Fig. 11. Ellipsoid with contours obtained by cutting the surface with three orthogonal planes passing through pairs of points where the Gaussian curvature has stationary values.

coordinate axes (Fig. 11). An equation for its surface can be written

$$\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 + \left(\frac{z}{c}\right)^2 = 1.$$

More useful for our purposes here is a parametric form

$$\begin{aligned} x &= a \cos \theta \cos \phi \\ y &= b \sin \theta \cos \phi \\ z &= c \sin \phi. \end{aligned}$$

A normal at the point

$$\mathbf{r} = (a \cos \theta \cos \phi, b \sin \theta \cos \phi, c \sin \phi)^T$$

on the surface is given by

$$\mathbf{n} = (bc \cos \theta \cos \phi, ca \sin \theta \cos \phi, ab \sin \phi)^T$$

as will be shown later. The Gaussian curvature turns out to be equal to

$$\begin{aligned} K &= \left[\frac{abc}{(bc \cos \theta \cos \phi)^2 + (ca \sin \theta \cos \phi)^2 + (ab \sin \phi)^2} \right]^2 \\ &= \left[\frac{abc}{n^2} \right]^2 \end{aligned}$$

where $n^2 = \mathbf{n} \cdot \mathbf{n}$.

If we let ξ be the longitude and η be the latitude on the Gaussian sphere, then a unit normal at the point (ξ, η) on the sphere is given by (Fig. 12)

$$\hat{n} = (\cos \xi \cos \eta, \sin \xi \cos \eta, \sin \eta)^T.$$




Fig. 12. Latitude and longitude can be used to identify points on the Gaussian sphere. Each point on the Gaussian sphere corresponds to a unique surface orientation.

Now $n = n\hat{n}$. Identifying terms in the two expressions for surface normals at corresponding points on the ellipsoid and the Gaussian sphere we get

$$\begin{aligned} bc \cos \theta \cos \phi &= n \cos \xi \cos \eta \\ ca \sin \theta \cos \phi &= n \sin \xi \cos \eta \\ ab \sin \phi &= n \sin \eta \end{aligned}$$

so that

$$n^2 [(a \cos \xi \cos \eta)^2 + (b \cos \xi \sin \eta)^2 + (c \sin \eta)^2] = (abc)^2$$

and finally, substituting for n^2 in the equation for K , we get

$$G(\xi, \eta) = \frac{1}{K} = \left[\frac{abc}{(a \cos \xi \cos \eta)^2 + (b \sin \xi \cos \eta)^2 + (c \sin \eta)^2} \right]^2.$$

The extended Gaussian image, in this case, varies smoothly and has the stationary values

$$\left(\frac{bc}{a}\right)^2 \left(\frac{ca}{b}\right)^2 \text{ and } \left(\frac{ab}{c}\right)^2$$

at the points where r is equal to $(\pm 1, 0, 0)^T$, $(0, \pm 1, 0)^T$, and $(0, 0, \pm 1)^T$, respectively. These results can be easily checked by sectioning the ellipsoid using the xy , yz , and zx planes. The Gaussian curvature, in this case, equals the product of the curvatures of the resulting ellipses. One then uses the fact that the maximum and minimum curvatures of an ellipse with semi-axes a and b are a/b^2 and b/a^2 .

Later we will derive the extended Gaussian image of a torus, an object that is not convex.

IV. DISCRETE APPROXIMATION: NEEDLE MAPS

Consider the surface broken up into small patches of equal area. Let there be ρ patches per unit area. Erect a surface normal on each patch. Consider the polyhedral object formed by the intersection of the tangent planes perpendicular to these surface normals. It approximates the original surface. The smaller the patches, the better the approximation.

The extended Gaussian image of the original (smoothly curved) convex object is approximated by impulses corresponding to the small patches. The magnitude of each impulse is about $1/\rho$, corresponding to the area of the patch it rests on (Fig. 13). Strongly curved areas will distribute their impulses over a large region on the Gaussian sphere, while areas which are nearly planar will have them concentrated in a small region. In fact, the number of impulses per unit area on the Gaussian sphere approaches ρ times the absolute value of the Gaussian curvature as we make ρ larger and larger. This can be shown using the integral of $1/K$ given earlier.




Fig. 13. Mapping of discrete patches on an object onto the Gaussian sphere. The patches in this case correspond to a regular tessellation of the image plane. Since the patches lie on a conical surface they contribute to the extended Gaussian image along a small circle.

The tessellation of the surface can be based on an arbitrary division into triangular patches as long as the magnitude of each impulse on the Gaussian sphere is made proportional to the area of the corresponding patch on the surface. Alternatively, one can divide the surface up according to the division of the image into picture cells. In this case one has to take into account that the area occupied in the image by a given patch is affected by foreshortening. The actual surface area is proportional to $1/(\hat{s} \cdot \hat{v})$, where \hat{s}_i is the normal of the patch, while \hat{v} is the vector pointing towards the viewer (Fig. 4).

Measurements of surface orientation from images will not be perfect, since they are affected by the noise in brightness measurements. Similarly, surface orientations obtained from range data will be somewhat inaccurate. Consequently, the impulses on the Gaussian sphere will be displaced a little from their true positions. The expected density on the Gaussian sphere will nevertheless tend to be equal to the inverse of the Gaussian curvature. One cannot, however, expect the impulses corresponding to a planar surface to be coincident. Instead, they will tend to form a small cluster. To be precise, the effect of noise is to smear out the information on the sphere. The extended Gaussian image is convolved with a smoothing function of width proportional to the magnitude of the noise.

A. Using Object Models

Extended Gaussian images also have to be computed for surfaces of prototypical object models. In this case it is best to find a convenient way to parameterize the surface and break it up into many small patches. Suppose the surface is

given in terms of two parameters u and v as $\mathbf{r}(u, v)$. Then at the point (u, v) we see that \mathbf{r}_u and \mathbf{r}_v are two tangents (Fig. 14). The cross product of these tangents is normal to the surface. The unit normal

$$\hat{\mathbf{n}} = \frac{\mathbf{r}_u \times \mathbf{r}_v}{|\mathbf{r}_u \times \mathbf{r}_v|}$$




Fig. 14. A surface normal can be computed by taking the cross product of two tangent vectors. The tangent vectors can be obtained by differentiation of the parametric form of the equation of the surface.

allows us to determine to which point on the Gaussian sphere this patch corresponds. Suppose that we divided the range of u into segments of size δu and the range of v into segments of size δv . Then the area of the patch

$$\delta A = |\mathbf{r}_u \times \mathbf{r}_v| \delta u \delta v$$

can be used to determine what contribution it makes to the corresponding place on the Gaussian sphere. Note that we do not have to explicitly compute the Gaussian curvature or take second partial derivatives.

V. TESSELLATION OF THE GAUSSIAN SPHERE: ORIENTATION HISTOGRAMS

It is useful to divide the sphere up into cells in order to represent the information on the Gaussian sphere in a computer. Ideally the cells should satisfy the following criteria:

1. all cells should have the same area;
2. all cells should have the same shape;
3. the cells should have regular shapes that are compact;
4. the division should be fine enough to provide good angular resolution;
5. for some rotations, the cells should be brought into coincidence with themselves.

Cells which are compact combine information only from surface patches which have nearly the same orientation. Elongated cells of the same area combine information from surface patches which have more widely differing orientations. The area of a regular polygon with n sides inscribed in a circle of radius r is

$$\pi r^2 \left[\frac{\sin(2\pi/n)}{(2\pi/n)} \right]$$

So the area of a hexagon inscribed in a circle is $(3\sqrt{3}/2)r^2$, twice that of a triangle inscribed in the same circle. Tessellations with near-triangular cells will thus combine information from orientations which are $\sqrt{2}$ times as far from the average as do tessellations using near-hexagonal cells.

If cells occur in a regular pattern, the relationship of a cell to its neighbors will be the same for all cells. Such arrange-

ments are to be preferred. Unfortunately, it is not possible to simultaneously satisfy the criteria listed above.

A simple tessellation consists of a division into latitude bands, each of which is then further divided along longitudinal strips (Fig. 15). The cells could be made more nearly equal in area by having fewer at higher latitudes, or by making the latitude bands wider there, or both. One advantage of this scheme is that it makes it easy to compute to which cell a particular surface normal should be assigned. Still, this arrangement does not come close to satisfying the criteria stated above. In particular, the cells




Fig. 15. The Gaussian sphere can be divided into cells along meridians and lines of longitude. The resulting cells do not have the same areas, however, and only align with each other for certain rotations about the axis through the poles.

are brought into alignment only for a few rotations about the axis of the globe. Rotations about any other axis cannot bring the cells into alignment.

A. Tessellations Based on Regular Polyhedra

Better tessellations may be found by projecting regular polyhedra onto the unit sphere after bringing their center to the center of the sphere [27]. Regular polyhedra are uniform and have faces which are all of one kind of regular polygon. (They are also called the Platonic solids) [19], [20], [28]–[32]. The vertices of a regular polyhedron are congruent. A division obtained by projecting a regular polyhedron has the desirable property that the resulting cells all have the same shape and area. Also, all cells have the same geometric relationship to their neighbors. In the case of the dodecahedron, the cells are even fairly well rounded. The dodecahedron, however, has only twelve cells (Fig. 16(a)). Even the icosahedron, with twenty triangular cells, provides too coarse a sampling of orientations (Fig. 16(b)). Furthermore, its cells are not well rounded. Unfortunately, there




Fig. 16. Tessellation of the Gaussian sphere using (a) the regular dodecahedron and (b) the regular icosahedron.

are only five regular solids (tetrahedron, hexahedron, octahedron, dodecahedron, and icosahedron).

One can go a little further by considering semi-regular polyhedra. A semi-regular polyhedron has regular polygons as faces, but the faces are not all of the same kind [19], [20], [28]–[32]. (They are also called the Archimedean polyhedra.) As for regular polyhedra, the vertices are congruent. There can be either two or three different types of faces and these have different areas. An illustration of a tessellation using a semi-regular polyhedron is provided by a soccer ball (Fig. 17(a)). It is based on the truncated icosahedron, a semi-regular polyhedron which has 12 pentagonal faces and 20 hexagonal faces.




Fig. 17. Tessellation of the Gaussian sphere using (a) the truncated icosahedron and (b) the pentakis dodecahedron.

Unfortunately there are only 13 semi-regular polyhedra. (The five truncated regular polyhedra, cuboctahedron, icosidodecahedron, snub cuboctahedron, snub icosidodecahedron, truncated cuboctahedron, rhombicuboctahedron, truncated icosidodecahedron, and the rhombicosidodecahedron.) Overall, these objects do not provide us with fine enough tessellations. The snub icosidodecahedron has the largest number of faces, but each of its 80 triangles is much smaller than each of its 12 pentagons.

The edges of a semi-regular polyhedron are all the same length. One consequence of this is that the different types of faces have different areas. The area of a regular polygon of n sides and edge-length e equals

$$\frac{ne^2}{4 \tan(\pi/n)}$$

so it is very roughly proportional to n^2 . This is a problem generally with semi-regular polyhedra. It is sometimes possible to derive a new polyhedron which has the same adjacency relationships between faces as a given semi-regular polyhedron but also has faces of equal area. The shapes of some of these faces then are no longer regular, however.

If we desire a finer subdivision still, we can consider splitting each face of a given tessellation further into triangular facets. If, for example, we split each pentagonal face of a dodecahedron into five equal triangles we obtain a pentakis dodecahedron with 60 faces (Fig. 17(b)). This happens to be the dual of the truncated icosahedron, discussed above. If we apply this method instead to the truncated icosahedron we construct an object with 180 faces. This object, as well as the pentakis dodecahedron, form suitable bases for further subdivision, as we shall show later.

To see how fine a division we might need, let us calculate the angular spread of surface normals which map into a particular cell. If there are n equal cells, then each one will

have area

$$A = (4\pi)/n$$

since the total area of the unit sphere is 4π . (This area equals the solid angle of the cone formed by the cell when connected to the center of the sphere.) The shape which minimizes the angular spread for given surface area is the circular disc. The area of a circular disc on the unit sphere is

$$A = 2\pi(1 - \cos \theta)$$

where θ is the half-angle of the cone formed by the disc when connected to the center of the sphere. If θ is small, the area can be approximated by

$$A \approx \pi\theta^2.$$

Thus if there are many cells and if they could be made circular, the angular spread would be

$$\theta \approx 2/\sqrt{n}.$$

The best we can hope for, however, are near-hexagonal cells. The area of a hexagon inscribed in a circle of radius r is $(3\sqrt{3}/2)r^2$, as already mentioned. The area of the circle, πr^2 , is about 20 percent more. So a hexagonal shape has a spread which is

$$\sqrt{\frac{2\pi}{3\sqrt{3}}} = 1.0996 \dots$$

as large as that of a circular shape of equal area. A lower bound on the angular spread for a tessellation with n cells then is

$$\theta = \sqrt{\frac{4\pi}{3\sqrt{3}n}} = \frac{2.1993 \dots}{\sqrt{n}}.$$

For $n = 60$, for example, the spread is at greater than 16.2° . One should also remember that the spread for triangular cells is even more, namely $\sqrt{2}$ times that for hexagonal cells.

B. Geodesic Domes

To proceed further, we can divide the triangular cells into four smaller triangles according to the well known geodesic dome constructions [27], [30], [33]. We attain high resolution by relenting on several of the criteria given above (Fig. 18). Specifically, the cells of a geodesic tessellation do not all have the same area and shape. The cells are also not compact, being shaped like (irregular) triangles. The *duals* of geodesic domes are better in this respect, since they have facets that are mostly (irregular) hexagons, with a




Fig. 18. Tessellation of the Gaussian sphere using a frequency-four geodesic tessellation based on the icosahedron. (There are $16 \times 20 = 320$ faces.)

dozen (regular) pentagons thrown in. Tessellations of arbitrary fineness can be constructed in this fashion. The pentakis dodecahedron is a good starting point for a geodesic division, as is the object constructed earlier from the truncated icosahedron by dividing the faces into triangles.

Each of the edges of the triangular cells of the original polyhedron are divided into f sections, where f is called the frequency of the geodesic division. The result is that each face is divided into f^2 (irregular) triangles. Tessellations where the frequency is a power of two are particularly well suited to the method suggested here, as we see next.

One has to be able to efficiently compute to which cell a particular surface normal belongs. In the case of the tessellations derived from regular polyhedra, one first computes the dot-product of the given unit vector and the vector to the center of each cell. (These reference vectors correspond to the vertices of the *dual* of the original regular polyhedron.) This gives one the cosine of the angle between the two. The closest reference vector is the one which gives the *largest* dot-product. The given vector is then assigned to the cell corresponding to that reference vector.

In the case of a geodesic dome, it is possible to proceed hierarchically, particularly if the frequency is a power of two. The geodesic dome is based on some regular polyhedron. The appropriate facet of this polyhedron is found as above. Next, one determines into which of the triangles of the first division of this facet the given unit normal falls. This can be done by considering which dot-product has the *second* largest value. No new dot-products need to be computed. The process is then repeated with the four triangles into which this facet is divided, and so on. In practice, lookup table methods can be used, which, while not exact, are very quick.

Let the area occupied by one of the cells on the Gaussian sphere be ω (in the case of the icosahedron $\omega = 4\pi/20$). The expected number of surface normals mapped into a cell equals

$$\rho\omega|\bar{C}|$$

for a convex object, where \bar{C} is the average of $G(\xi, \eta)$ over the cell.

It is clear that the extended Gaussian image can be computed locally. One simply counts the number of surface normals that belong in each cell. The expression for the Gaussian curvature, on the other hand, includes first and second partial derivatives of the surface function. In practice, estimates of derivatives are unreliable, because of noise. It is important therefore that the extended Gaussian image can be computed *without* estimating the derivatives.

The values in the cells can be thought of as an orientation histogram. It has recently been brought to my attention that this is analogous to a scheme used for histogramming directions of dendrites on neurons [34].

The result can be displayed graphically using normal vectors on each of the cells to represent the weight of the accumulated surface normals. A frequency-two subdivision of the pentakis dodecahedron provides 240 cells, enough for most practical purposes (Fig. 19). The angular spread in this case is about 11.5° . An alternative way to present the extended Gaussian image graphically is by means of a grey-level image where brightness in each cell is proportional to the count. The surface of the Gaussian sphere may be projected stereographically instead of orthographically in order to preserve the shapes of the cells. Their areas will be scaled unequally however.




Fig. 19. Orientation histogram collected on a geodesic dome derived from the pentakis dodecahedron. (There are $12 \times 5 \times 4 = 240$ faces.) This is a discrete approximation of the extended Gaussian image. The length of the vector attached to the center of a cell is proportional to the number of surface normals on the surface of the original object which have orientations falling within the range of directions spanned by that cell.

A further refinement of the orientation histogram has us store the sum of the vectors, scaled according to the area of the corresponding patch, rather than just the sum of the areas of the patches. This requires three times as much memory space, but provides more accuracy. In fact, in the case of polyhedra, this representation is exact.

VI. SOLIDS OF REVOLUTION

In the case of the surface of a solid of revolution, the Gaussian curvature is rather easy to determine. The solid of revolution can be produced by rotating a (planar) generating curve about an axis (Fig. 20). Let the generating curve be specified by the perpendicular distance from the axis, $r(s)$, given as a function of arc length s along the curve. Let θ be the angle of rotation around the axis. Now consider




Fig. 20. A solid of revolution can be generated by rotating a curve around an axis. The curve can be specified by giving the distance from the axis as a function of the arc length along the curve.

the Gaussian sphere positioned so that its axis is aligned with the axis of the solid of revolution. Let ξ be the longitude and η be the latitude on the Gaussian sphere.

We can let ξ correspond to θ . That is, a point on the object produced when the generating curve has rotated through an angle θ has a surface normal that lies on the Gaussian sphere at a point with longitude $\xi = \theta$.

A. Gaussian Curvature of Solid of Revolution

Consider a small patch on the Gaussian sphere lying between ξ and $\xi + \delta\xi$ in longitude and between η and $\eta + \delta\eta$ in latitude. Its area is

$$\cos \eta \delta\xi \delta\eta.$$

We need only determine the area of the corresponding patch on the object. It is

$$r \delta\theta \delta s$$

where δs is the change in arc distance along the generating curve corresponding to the change $\delta\eta$ in surface orientation. The Gaussian curvature is the limit of the ratio of the two areas as they tend to zero. That is,

$$K = \lim_{\substack{\delta\eta \rightarrow 0 \\ \delta\xi \rightarrow 0}} \frac{\cos \eta \delta\xi \delta\eta}{r \delta\theta \delta s} = \lim_{\delta\eta \rightarrow 0} \frac{\cos \eta \delta\eta}{r \delta s} = \frac{\cos \eta}{r} \frac{d\eta}{ds}$$

since $\delta\xi = \delta\theta$. The curvature of the generating curve, κ_G , is just the rate of change of direction with arc length along it [22]–[24]. So

$$\frac{d\eta}{ds} = \kappa_G$$

and hence

$$K = \frac{\kappa_G \cos \eta}{r}.$$

It is easy to see that (Fig. 21), $\sin \eta = -r_s$, where r_s is the




Fig. 21. The figure shows the relationships between the infinitesimal increments in arc length along the curve, distance from the axis of rotation, and distance along the axis of rotation.

partial derivative of r with respect to s . Differentiating with respect to s we get

$$\cos \eta \frac{d\eta}{ds} = \frac{d}{ds}(-r_s) = -r_{ss}$$

and so we obtain the simple formula

$$K = -\frac{r_{ss}}{r}.$$

In the case of a sphere of radius R , for example, we have $r = R \cos(s/R)$ for $-(\pi/2)R < s < +(\pi/2)R$. Thus $r_{ss} = -(r/R^2)$ and $K = 1/R^2$.

For some purposes it is more useful to express the radius

r as a function of the distance along the axis, rather than as a function of arc length along the curve. Let the distance along the axis be denoted by z . It is easy to see that (Fig. 21) $\tan \eta = -r_z$, and so, differentiating with respect to s ,

$$\sec^2 \eta \frac{d\eta}{ds} = \frac{d}{ds}(-r_z) = -r_{zz} \frac{dz}{ds}$$

where, from the figure, we see that $\cos \eta = z_s$, so that

$$\kappa_G \cos \eta = \frac{d\eta}{ds} \cos \eta = -r_{zz} \cos^3 \eta$$

and finally

$$K = \frac{r_{zz}}{-r(1 + r_z^2)^2}$$

since

$$\sec^2 \eta = 1 + r_z^2.$$

B. Alternate Derivation of Gaussian Curvature of a Solid of Revolution^(*)

We first need to review Meusnier's theorem [22]–[24]. Consider a normal section of a surface at a particular point. It is obtained by cutting the surface with one of the planes including the local normal (Fig. 22). Suppose that the curva-




Fig. 22. The curvature of the curve obtained by cutting a surface using an inclined plane is greater than that obtained by cutting it using a plane which includes the surface normal. Meusnier's theorem tells us that the ratio of the two curvatures is equal to the cosine of the angle between the two planes.

ture of the curve in which the surface cuts this plane is κ_N . Now imagine tilting the plane away from the normal by an angle η (using the local tangent as an axis to rotate about). The new plane will cut the surface in a curve with higher curvature. In fact, it can be shown that the new curve has curvature

$$\kappa_N / \cos \eta.$$

It is easy to see this in the case of a sphere, since a plane including the center cuts the sphere in a great circle, while an inclined plane cuts it in a small circle of radius proportional to the cosine of the angle of inclination.

Now, let us return to the surface of revolution. It is not hard to show that one of the principal curvatures at a point on the surface will correspond to a cut through the surface by a plane which includes the axis of revolution. The curve obtained in this way is just the generating curve of the solid

of revolution. So one of the two principal curvatures is equal to the curvature κ_G of the generating curve at the corresponding point.

Now consider a plane perpendicular to the axis of revolution through the same surface point (Fig. 23). It cuts the surface in a circle. The curvature in this plane equals $(1/r)$,




Fig. 23. If the solid of revolution is cut by a plane perpendicular to the axis of rotation, a circle is obtained. The curvature in this plane is just the inverse of the distance of the surface from the axis. The curvature of the corresponding normal section can be obtained using Meusnier's theorem.

where r is the radius of the solid of revolution at that point. This horizontal plane, however, is *not* a normal section. Suppose that the normal makes an angle η relative to this plane. (The local tangent plane also makes an angle η relative to the axis of revolution.) Now construct the plane including the local normal which intersects the horizontal plane in a line perpendicular to the axis. This plane will be inclined η relative to the one we have just studied. It also produces the second principal normal section sought after. By Meusnier's theorem we see that the curvature of the curve found in this normal section is $\kappa_N = (1/r) \cos \eta$. Finally, the Gaussian curvature is found by multiplication to be

$$K = \frac{\kappa_G \cos \eta}{r}$$

In the case of a sphere of radius R , for example, we have $r = R \cos \eta$ and $\kappa_G = 1/R$, so that $K = 1/R^2$, as expected.

To make this result more usable, erect a coordinate system with the z -axis aligned with the axis of revolution. The generating curve is given as $r(z)$. Let the first and second derivatives of r with respect to z be denoted by r_z and r_{zz} , respectively. It is easy to see that (Fig. 21), $\tan \eta = r_z$, so that

$$\cos \eta = \frac{1}{\sqrt{1 + r_z^2}}$$

Furthermore

$$\kappa_G = -\frac{r_{zz}}{(1 + r_z^2)^{3/2}}$$

so that finally

$$K = -\frac{r_{zz}}{r(1 + r_z^2)^2}$$

In order to use this result in deriving extended Gaussian images it is necessary to identify points on the surface with

points on the Gaussian sphere. Suppose that we introduce a polar angle θ such that

$$x = r \cos \theta \quad \text{and} \quad y = r \sin \theta.$$

Then a unit normal to the surface is given by

$$\frac{(\cos \theta, \sin \theta, -r_z)^T}{\sqrt{1 + r_z^2}}$$

Equating this to the unit normal on the Gaussian sphere

$$(\cos \xi \cos \eta, \sin \xi \cos \eta, \sin \eta)^T$$

we get

$$\xi = \theta \quad \text{and} \quad \tan \eta = -r_z.$$

C. Extended Gaussian Image of a Torus

As an illustration we will now determine the extended Gaussian image of a torus. Let the torus have major axis R and minor axis ρ (Fig. 24). A point on the surface can be




Fig. 24. A torus obtained by spinning a circle around an axis. The resulting object is not convex. Its extended Gaussian image can be computed nevertheless.

identified by θ and s , where θ is the angle around the axis of the torus; while s is the arc length along the surface measured from the plane of symmetry. Then

$$r = R + \rho \cos(s/\rho)$$

and

$$r_{ss} = -(1/\rho) \cos(s/\rho)$$

so that

$$K = -\frac{r_{ss}}{r} = \frac{1}{\rho} \frac{\cos(s/\rho)}{R + \rho \cos(s/\rho)}$$

Two points, P and P' (Fig. 24), separated by π in θ , have the same surface orientation on the torus. The surface normal at one of these places points away from the axis of rotation, while it points towards the axis at the other place. Accordingly, two points on the object,

$$(\theta, s) = (\xi, \rho \eta) \quad \text{and} \quad (\theta, s) = (\xi + \pi, \rho(\pi - \eta))$$

correspond to the point (ξ, η) on the Gaussian sphere. The curvatures at these two points have opposite signs

$$K_+ = +\frac{1}{\rho} \frac{\cos \eta}{R + \rho \cos \eta} \quad \text{and} \quad K_- = -\frac{1}{\rho} \frac{\cos \eta}{R - \rho \cos \eta}$$

The torus is not a convex object, so more than one point on its surface contributes to a given point on the extended Gaussian image. If we add the *absolute* values of the

inverses of the curvature we get

$$G(\eta, \xi) = \frac{1}{K_+} - \frac{1}{K_-} = 2R\rho \sec \eta.$$

If we had added the inverses algebraically instead we would have obtained

$$\frac{1}{K_+} + \frac{1}{K_-} = 2\rho^2.$$

which is twice the result for a sphere of radius ρ .

The same results could have been found using

$$K = \frac{\kappa_G \cos \eta}{r}$$

since $\kappa_G = -1/\rho$ and $r = R \pm \rho \cos \eta$, so that

$$K = \frac{\pm \cos \eta}{\rho(R \pm \rho \cos \eta)}.$$

The extended Gaussian image of a torus has singularities at the poles. These correspond to the two rings on which the torus would rest if it were dropped onto a plane. All of the points on one of these rings have the same surface orientation.

We can think of the Gaussian sphere as covered by two sheets of a Riemann surface, one corresponding to the inner half of the torus, closer to its axis of symmetry, the other corresponding to the outer half. The two sheets are connected to one another at the poles, branch points corresponding to the two rings mentioned above. There the Gaussian curvature changes sign.

We may also note at this point that all tori with the same surface area, $(4\pi^2\rho R)$, have the same extended Gaussian image.

D. The Unique Convex Object with $G(\xi, \eta) = 2 \sec \eta^{(*)}$

While all tori with surface area $4\pi^2$ have the same extended Gaussian image

$$G(\xi, \eta) = 2 \sec \eta$$

there is only one convex object which has that extended Gaussian image. It is a solid of revolution since $G(\xi, \eta)$ is independent of ξ . So we have on the one hand

$$K = 1/2 \cos \eta$$

and on the other hand

$$K = \frac{\kappa_G \cos \eta}{r}$$

so that

$$\kappa_G = r/2.$$

The equation states that the curvature of the generating curves varies linearly with the distance from the axis of rotation. This deceptively simple equation represents a nonlinear second-order differential equation for r in terms of z since

$$\kappa_G = -\frac{r_{zz}}{(1+r_z^2)^{3/2}}$$

so that

$$r_{zz} = -\frac{r}{2}(1+r_z^2)^{3/2}.$$

Now

$$\frac{d}{dz} \frac{1}{\sqrt{1+r_z^2}} = -\frac{r_z r_{zz}}{(1+r_z^2)^{3/2}}$$

and

$$\frac{d}{dz} \frac{r^2}{4} = \frac{r r_z}{2}$$

so that

$$\frac{d}{dz} \frac{1}{\sqrt{1+r_z^2}} = \frac{d}{dz} \frac{r^2}{4}$$

or

$$\frac{1}{\sqrt{1+r_z^2}} = \frac{r^2 + c^2}{4}$$

where c^2 is a constant of integration. We now have reduced the problem to a nonlinear first-order differential equation for r in terms of z . If the object is to be convex and smooth at its poles, we expect $r_z \rightarrow \infty$ as $r \rightarrow 0$. Thus $c = 0$. Next note that the term on the left equals $\cos \eta$. So we also have

$$\cos \eta = \frac{r^2}{4}$$

or, using an earlier expression for κ_G ,

$$\kappa_G = -\sqrt{\cos \eta}.$$

This is an implicit equation for the curve of least energy [35]! The curve of least energy is the curve which minimizes the integral of the square of the curvature κ_G . It can be solved for z in terms of r to yield

$$z = \sqrt{2} [2E(\cos^{-1}(r/2), 1/\sqrt{2}) - F(\cos^{-1}(r/2), 1/\sqrt{2})]$$

where E and F are incomplete elliptic integrals. If we let s be the arc length along the curve we can also write the solution in Whewell form

$$s = \sqrt{2} F(\cos^{-1} \sqrt{\cos \eta}, 1/\sqrt{2})$$

or Césaro form

$$s = \sqrt{2} F(\cos^{-1}(-\kappa_G), 1/\sqrt{2}).$$

The length of the curve from the pole to the equator is

$$\sqrt{2} K(1/\sqrt{2}) = \sqrt{2} K(\sin(\pi/4)) = \frac{\Gamma(1/4)^2}{2\sqrt{2}\pi}$$

where K is the complete elliptic integral of the first kind, and Γ is the gamma function [24]. The height from equator to the pole is

$$W = \frac{(2\pi)^{3/2}}{\Gamma(1/4)^2}$$

while the maximum radius is

$$H = 2.$$

The minimum radius of curvature equals one, so that a circle tangent at the outermost point is also tangent at the origin [35]. This circle, when rotated about the vertical axis, produces a torus with the same extended Gaussian image (Fig. 25). Both objects have total surface area $4\pi^2$.




Fig. 25. The unique convex object with the same extended Gaussian image as a torus has an interesting shape. It is a solid of revolution whose generating curve is the curve of least energy. This is the shape which a uniform bar constrained to pass through two points in space with given orientation will adopt.

VII. GAUSSIAN CURVATURE IN THE GENERAL CASE

When the object is not a solid of revolution we need to work a little harder to obtain the Gaussian curvature. Let $x = x(u, v)$, $y = y(u, v)$, and $z = z(u, v)$ be parametric equations for points on a given surface. Let $\mathbf{r} = (x, y, z)^T$ be a vector to a point on the surface. Then

$$\mathbf{r}_u = \frac{\partial \mathbf{r}}{\partial u} \quad \text{and} \quad \mathbf{r}_v = \frac{\partial \mathbf{r}}{\partial v}$$

are two tangents to the surface, as already noted earlier. The cross-product of these two vectors

$$\mathbf{n} = \mathbf{r}_u \times \mathbf{r}_v$$

will be perpendicular to the local tangent plane (Fig. 14). The length of this normal vector squared equals

$$n^2 = \mathbf{n} \cdot \mathbf{n} = (\mathbf{r}_u \cdot \mathbf{r}_u)(\mathbf{r}_v \cdot \mathbf{r}_v) - (\mathbf{r}_u \cdot \mathbf{r}_v)^2$$

since $(\mathbf{a} \times \mathbf{b}) \cdot (\mathbf{c} \times \mathbf{d}) = (\mathbf{a} \cdot \mathbf{c})(\mathbf{b} \cdot \mathbf{d}) - (\mathbf{a} \cdot \mathbf{d})(\mathbf{b} \cdot \mathbf{c})$. A unit vector $\hat{\mathbf{n}} = \mathbf{n}/n$ can be computed using this result.

A. Gaussian Curvature from Variation in Normals(*)

The Gaussian curvature is the limit of the ratio of the area of a patch on the Gaussian sphere to the area of the corresponding patch on the surface, as the area shrinks to zero. Consider an infinitesimal triangle formed by the three points on the surface corresponding to (u, v) , $(u + \delta u, v)$, and $(u, v + \delta v)$. The lengths of two sides of this triangle are

$$|\mathbf{r}_u| \delta u \quad \text{and} \quad |\mathbf{r}_v| \delta v$$

while the sine of the angle between these sides equals

$$\frac{|\mathbf{r}_u \times \mathbf{r}_v|}{|\mathbf{r}_u| |\mathbf{r}_v|}$$

so that an outward normal with size equal to the area of the triangle is given by

$$\frac{1}{2} (\mathbf{r}_u \times \mathbf{r}_v) \delta u \delta v = \frac{1}{2} \mathbf{n} \delta u \delta v.$$

To determine the area of the corresponding triangular patch on the Gaussian sphere we need to find the unit surface normals at the three points. The unit surface normals will be

$$\hat{\mathbf{n}}, \quad \hat{\mathbf{n}} + \hat{\mathbf{n}}_u \delta u, \quad \text{and} \quad \hat{\mathbf{n}} + \hat{\mathbf{n}}_v \delta v$$

if we ignore terms of higher order in δu and δv . Here $\hat{\mathbf{n}}_u$ and $\hat{\mathbf{n}}_v$ are the partial derivatives of $\hat{\mathbf{n}}$ with respect to u and v . Note that $\hat{\mathbf{n}}_u$ and $\hat{\mathbf{n}}_v$ are perpendicular to $\hat{\mathbf{n}}$. The area of the patch on the Gaussian sphere equals the magnitude of

$$\frac{1}{2} (\hat{\mathbf{n}}_u \times \hat{\mathbf{n}}_v) \delta u \delta v$$

by reasoning similar to that used in determining the area of the original patch on the given surface. We need to find $\hat{\mathbf{n}}_u$ and $\hat{\mathbf{n}}_v$ to compute this area. Now

$$\hat{\mathbf{n}}_u = \frac{\partial \mathbf{n}}{\partial u} \frac{1}{n} = \frac{n \mathbf{n}_u - \mathbf{n} n_u}{n^2}.$$

From $n^2 = \mathbf{n} \cdot \mathbf{n}$ we get

$$n \mathbf{n}_u = \mathbf{n} \cdot \mathbf{n}_u$$

so that

$$\hat{\mathbf{n}}_u = \frac{(\mathbf{n} \cdot \mathbf{n}) \mathbf{n}_u - (\mathbf{n} \cdot \mathbf{n}_u) \mathbf{n}}{n^3} = \frac{(\mathbf{n} \times \mathbf{n}_u) \times \mathbf{n}}{n^3}$$

and

$$\hat{\mathbf{n}}_v = \frac{(\mathbf{n} \cdot \mathbf{n}) \mathbf{n}_v - (\mathbf{n} \cdot \mathbf{n}_v) \mathbf{n}}{n^3} = \frac{(\mathbf{n} \times \mathbf{n}_v) \times \mathbf{n}}{n^3}$$

since $(\mathbf{a} \times \mathbf{b}) \times \mathbf{c} = (\mathbf{a} \cdot \mathbf{c}) \mathbf{b} - (\mathbf{b} \cdot \mathbf{c}) \mathbf{a}$. Then

$$\begin{aligned} \hat{\mathbf{n}}_u \times \hat{\mathbf{n}}_v &= \frac{n^2}{n^6} [(\mathbf{n} \cdot \mathbf{n})(\mathbf{n}_u \times \mathbf{n}_v) + (\mathbf{n} \cdot \mathbf{n}_u)(\mathbf{n}_v \times \mathbf{n}) \\ &\quad + (\mathbf{n} \cdot \mathbf{n}_v)(\mathbf{n} \times \mathbf{n}_u)] \end{aligned}$$

or

$$\hat{\mathbf{n}}_u \times \hat{\mathbf{n}}_v = \frac{1}{n^4} [n \mathbf{n}_u \mathbf{n}_v] \mathbf{n}$$

since

$$\begin{aligned} [abc] \mathbf{p} &= (\mathbf{a} \cdot \mathbf{p})(\mathbf{b} \times \mathbf{c}) + (\mathbf{b} \cdot \mathbf{p})(\mathbf{c} \times \mathbf{a}) \\ &\quad + (\mathbf{c} \cdot \mathbf{p})(\mathbf{a} \times \mathbf{b}) \end{aligned}$$

where $[abc] = (\mathbf{a} \times \mathbf{b}) \cdot \mathbf{c}$.

This shows that the patch on the Gaussian sphere has the same orientation as the patch on the surface, as it should.

An outward pointing normal of size equal to the area is given by

$$\frac{1}{2} \frac{1}{n^3} [\mathbf{nn}_u \mathbf{n}_v] \mathbf{n} \delta u \delta v.$$

The ratio of the two areas, the Gaussian curvature, is then just

$$K = \frac{[\mathbf{nn}_u \mathbf{n}_v]}{n^4}.$$

Now

$$\mathbf{n} = \mathbf{r}_u \times \mathbf{r}_v$$

so that

$$\mathbf{n}_u = \mathbf{r}_{uu} \times \mathbf{r}_v + \mathbf{r}_u \times \mathbf{r}_{uv}$$

and

$$\mathbf{n}_v = \mathbf{r}_{uv} \times \mathbf{r}_v + \mathbf{r}_u \times \mathbf{r}_{vv}.$$

Now using

$$(\mathbf{a} \times \mathbf{b}) \times (\mathbf{c} \times \mathbf{d}) = [\mathbf{abd}] \mathbf{c} - [\mathbf{abc}] \mathbf{d}$$

or

$$(\mathbf{a} \times \mathbf{b}) \times (\mathbf{c} \times \mathbf{d}) = [\mathbf{acd}] \mathbf{b} - [\mathbf{bcd}] \mathbf{a}$$

we get

$$\begin{aligned} \mathbf{n}_u \times \mathbf{n}_v = & -[\mathbf{r}_{uu} \mathbf{r}_v \mathbf{r}_{vv}] \mathbf{r}_v + [\mathbf{r}_{uu} \mathbf{r}_v \mathbf{r}_{vv}] \mathbf{r}_u - [\mathbf{r}_{uu} \mathbf{r}_v \mathbf{r}_u] \mathbf{r}_{vv} \\ & + [\mathbf{r}_u \mathbf{r}_{uv} \mathbf{r}_v] \mathbf{r}_{uv} + [\mathbf{r}_u \mathbf{r}_{uv} \mathbf{r}_{vv}] \mathbf{r}_u \end{aligned}$$

so that

$$[\mathbf{nn}_u \mathbf{n}_v] = \mathbf{n} \cdot (\mathbf{n}_u \times \mathbf{n}_v) = [\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uu}] [\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{vv}] - [\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uv}]^2$$

and finally

$$K = \frac{[\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uu}] [\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{vv}] - [\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uv}]^2}{|\mathbf{r}_u \times \mathbf{r}_v|^4}.$$

This result can be used to derive the expression for curvature of a solid of revolution in a more rigorous fashion.

B. Fundamental Forms of a Surface^(*)

Let, as before,

$$\hat{\mathbf{n}} = \frac{\mathbf{r}_u \times \mathbf{r}_v}{|\mathbf{r}_u \times \mathbf{r}_v|}$$

be the unit surface normal vector. The first fundamental form of a surface gives the square of the element of distance as [24]

$$ds^2 = |d\mathbf{r}|^2 = E(u, v) du^2 + 2F(u, v) du dv + G(u, v) dv^2.$$

The second fundamental form of a surface gives the normal curvature using the equation [24]

$$-d\mathbf{r} \cdot d\hat{\mathbf{n}} = L(u, v) du^2 + 2M(u, v) du dv + N(u, v) dv^2.$$

The coefficients can be expressed in terms of derivatives of \mathbf{r} as follows:

$$E = \mathbf{r}_u \cdot \mathbf{r}_u$$

$$F = \mathbf{r}_u \cdot \mathbf{r}_v$$

and

$$G = \mathbf{r}_v \cdot \mathbf{r}_v$$

and

$$L = \mathbf{r}_u \cdot \hat{\mathbf{n}}_u$$

$$M = \mathbf{r}_u \cdot \hat{\mathbf{n}}_v = \mathbf{r}_v \cdot \hat{\mathbf{n}}_u$$

and

$$N = \mathbf{r}_v \cdot \hat{\mathbf{n}}_v$$

or

$$L = \frac{[\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uu}]}{\sqrt{EG - F^2}}$$

$$M = \frac{[\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{uv}]}{\sqrt{EG - F^2}}$$

and

$$N = \frac{[\mathbf{r}_u \mathbf{r}_v \mathbf{r}_{vv}]}{\sqrt{EG - F^2}}$$

so that [24]

$$K = \frac{LN - M^2}{EG - F^2}.$$

Finally, if the surface is given as $z(x, y)$, the above reduces to the familiar

$$K = \frac{z_{xx}z_{yy} - z_{xy}^2}{(1 + z_x^2 + z_y^2)^2}.$$

C. Application of the General Formula to the Ellipsoid^(*)

In the case of the ellipsoid we have, as discussed before

$$\mathbf{r} = (a \cos \theta \cos \phi, b \sin \theta \cos \phi, c \sin \phi)^T$$

$$\mathbf{r}_\theta = (-a \sin \theta \cos \phi, b \cos \theta \cos \phi, 0)^T$$

$$\mathbf{r}_\phi = (-a \cos \theta \sin \phi, -b \sin \theta \sin \phi, c \cos \phi)^T$$

and

$$\mathbf{r}_{\theta\theta} = (-a \cos \theta \cos \phi, -b \sin \theta \cos \phi, 0)^T$$

$$\mathbf{r}_{\theta\phi} = (a \sin \theta \sin \phi, -b \cos \theta \sin \phi, 0)^T$$

$$\mathbf{r}_{\phi\phi} = (-a \cos \theta \cos \phi, -b \sin \theta \cos \phi, -c \sin \phi)^T.$$

A surface normal can be found by taking cross-products

$$\mathbf{n} = \mathbf{r}_\theta \times \mathbf{r}_\phi = (bc \cos \theta \cos \phi, ca \sin \theta \cos \phi, ab \sin \phi)^T \cos \phi$$

and the coefficients of the first fundamental form are

$$E = \mathbf{r}_\theta \cdot \mathbf{r}_\theta = (a^2 \sin^2 \theta + b^2 \cos^2 \theta) \cos^2 \phi$$

$$F = \mathbf{r}_\theta \cdot \mathbf{r}_\phi = (a^2 + b^2) \sin \theta \cos \theta \sin \phi \cos \phi$$

$$G = \mathbf{r}_\phi \cdot \mathbf{r}_\phi = (a^2 \cos^2 \theta + b^2 \sin^2 \theta) \sin^2 \phi + c^2 \cos^2 \phi.$$

Hence

$$\begin{aligned} EG - F^2 = & [(bc \cos \theta \cos \phi)^2 \\ & + (ca \sin \theta \cos \phi)^2 + (ab \sin \phi)^2] \cos^2 \phi. \end{aligned}$$

To compute the coefficients of the second fundamental form we need

$$[\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\theta\theta}] = \mathbf{n} \cdot \mathbf{r}_{\theta\theta} = -abc \cos^3 \phi$$

$$[\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\theta\phi}] = \mathbf{n} \cdot \mathbf{r}_{\theta\phi} = 0$$

$$[\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\phi\phi}] = \mathbf{n} \cdot \mathbf{r}_{\phi\phi} = -abcc \cos \phi$$

and so

$$[\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\theta\theta}] [\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\phi\phi}] - [\mathbf{r}_\theta \mathbf{r}_\phi \mathbf{r}_{\theta\phi}]^2 = (abc \cos^2 \phi)^2.$$

Finally then

$$LN - M^2 = \frac{(abc \cos^2 \phi)^2}{EG - F^2}$$

and

$$K = \frac{LN - M^2}{EG - F^2} = \left[\frac{abc}{(bc \cos \theta \cos \phi)^2 + (ca \sin \theta \cos \phi)^2 + (ab \sin \phi)^2} \right]^2$$

This result was used earlier in the discussion of the extended Gaussian image of the ellipsoid.

VIII. SUMMARY AND CONCLUSIONS

We have defined the extended Gaussian image, discussed its properties, and given examples. Methods for determining the extended Gaussian images of polyhedra, solids of revolution, and smoothly curved objects in general were shown. The orientation histogram, a discrete approximation of the extended Gaussian image, was described along with a variety of ways of tessellating the sphere. Machine vision methods for obtaining the surface orientation information required to build an orientation histogram are discussed elsewhere [1], [3]–[8]. Extended Gaussian images based on object models can be matched with those derived from experimental data. The application of extended Gaussian images to object recognition and, more importantly, to finding the attitude in space of an object, are discussed in a recent article [17].

ACKNOWLEDGMENT

The author wishes to thank E. Grimson and T. Lozano-Perez who made a number of very helpful suggestions after reading a draft of this paper.

REFERENCES

- [1] W. E. L. Grimson, *From Images to Surfaces*. Cambridge, MA: M. I. T. Press, 1981.
- [2] B. K. P. Horn, "Sequins and quills—Representations for surface topography," M.I.T. A.I. Lab. Memo 536, May 1979.
- [3] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from a single view," in *Image Understanding Systems and Industrial Applications (Proc. S.P.I.E. 22nd Annu. Tech. Symp.)*, vol. 155, pp. 136–143, Aug. 1978.
- [4] B. K. P. Horn, R. J. Woodham, and W. M. Silver, "Determining shape and reflectance using multiple images," M.I.T. A.I. Lab. Memo 490, Aug. 1978.
- [5] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Opt. Eng.*, vol. 19, no. 1, pp. 139–144, Jan./Feb. 1980.
- [6] W. M. Silver, "Determining shape and reflectance using multiple images," S. M. thesis, Dep. Elec. Eng. Comput. Sci., M.I.T., June 1980.
- [7] K. Ikeuchi, "Determining surface orientation of specular surfaces using the photometric stereo method," *IEEE Trans.*

- Pattern Anal. Mach. Intell.*, vol. PAMI-3, no. 6, pp. 661–669, Nov. 1981.
- [8] E. N. Coleman and R. Jain, "Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry," *Comput. Graph. Image Process.*, vol. 18, no. 4, pp. 309–328, 1982.
- [9] D. A. Smith, "Using enhanced spherical images," M.I.T. A.I. Lab. Memo No. 530, May 1979.
- [10] R. Bajcsy, "Three-dimensional scene analysis," in *Proc. 5th Int. Pattern Recognition Conf.* (Miami, FL, Dec. 1980), pp. 1064–1074.
- [11] K. Ikeuchi, "Recognition of objects using the extended Gaussian image," in *Proc. IJCAI-81* (Vancouver, B. C., Canada, Aug. 1981), pp. 595–600.
- [12] C. Dane and R. Bajcsy, "Three-dimensional segmentation using the Gaussian image and spatial information," in *Proc. IEEE Comput. Soc. Conf. on Pattern Recognition and Image Processing* (Dallas, TX, Aug. 1981), pp. 54–56.
- [13] D. H. Ballard and D. Sabbah, "detecting object orientation from surface normals," in *Proc. Int. Pattern Recognition Conf.* (München, Germany, Dec. 1981), pp. 63–67.
- [14] P. Brou, "Finding objects in depth maps," Ph.D. dissertation, M.I.T. Dep. Elec. Eng. Comput. Sci., Sept. 1983.
- [15] K. Ikeuchi, "Determining the attitude of an object from a needle map using the extended Gaussian image," M.I.T. A.I. Lab. Memo No. 714, Apr. 1983.
- [16] K. Ikeuchi, B. K. P. Horn, S. Nagata, T. Callahan and O. Feingold, "Picking an object from a pile of objects," M.I.T. A.I. Lab. Memo No. 726, May 1983.
- [17] B. K. P. Horn and K. Ikeuchi, "Mechanically manipulating randomly oriented parts," *Scientific American*, vol. 251, no. 2, pp. 100–113, Aug. 1984.
- [18] H. Minkowski, "Allgemeine Lehrsätze über die konvexen Polyeder," *Nachrichten von der Königlichen Gesellschaft der Wissenschaften, mathematisch-physikalische Klasse, Göttingen*, pp. 198–219, 1897. (Also in H. Minkowski, *Gesammelte Abhandlungen*. Leipzig, Germany: 1911).
- [19] A. V. Pogorelov, *Differential Geometry*. Noordhoff-Groningen: The Netherlands, 1956.
- [20] L. A. Lyusternik, *Convex Figures and Polyhedra*. New York: Dover, 1963.
- [21] J. J. Little, "An iterative method for reconstructing convex polyhedra from extended Gaussian images," in *Proc. Nat. Conf. on Artificial Intelligence* (Washington, D C, Aug. 1983), pp. 247–254.
- [22] D. Hilbert and S. Cohn-Vossen, *Geometry and the Imagination*. New York: Chelsea, 1952.
- [23] B. O'Neill, *Elementary Differential Geometry*. New York: Academic Press, 1966.
- [24] G. A. Korn and T. M. Korn, *Mathematical Handbook—For Scientists and Engineers*. New York: McGraw-Hill, 1968.
- [25] M. P. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice Hall, 1976.
- [26] A. Alexandroff, "Existence and uniqueness of a convex surface with a given integral curvature," *C. R. (Dokl.) de l'Académie des Sciences de l'URSS*, vol. 35, no. 5, pp. 131–134, 1942.
- [27] M. J. Wenninger, *Spherical Models*. Cambridge, England: Cambridge Univ. Press, 1979.
- [28] L. Fejes Toth, *Regular Figures*. New York: Pergamon, 1964.
- [29] M. J. Wenninger, *Polyhedron Models*. Cambridge, England: Cambridge Univ. Press, 1971.
- [30] H. S. M. Coxeter, *Regular Polytopes*. New York: Dover, 1973.
- [31] H. Kenner, *Geodesic Math—And How To Use It*. Los Angeles, CA: Univ. of California Press, 1976.
- [32] P. Pearce and S. Pearce, *Polyhedra Primer*. New York: Van Nostrand Reinhold, 1978.
- [33] A. Pugh, *Polyhedra—A Visual Approach*. Los Angeles, CA: Univ. of California Press, 1976.
- [34] C. M. Brown, "Representing the orientation of dendritic fields with geodesic tessellations," Internal Rep. TR-13, Computer Sci. Dep., Univ. of Rochester, Rochester, NY, 1977.
- [35] B. K. P. Horn, "The curve of least energy," *ACM Trans. Math. Software*, vol. 9, no. 4, pp. 441–460, Dec. 1983.