

DETERMINING OBJECT ATTITUDE FROM EXTENDED GAUSSIAN IMAGES

James J. Little*

Department of Computer Science
The University of British Columbia
Vancouver, B.C., Canada V6T 1W5

Abstract

The Extended Gaussian Image (EGI) of an object records the variation of surface area with surface orientation, uniquely representing convex objects. The inversion problem for polyhedra (from an EGI to a description in terms of vertices and faces) has been solved, by an iterative algorithm [Little, 1983]. The algorithm depends upon the mixed volume, a geometric construction linking the areas and positions of the faces of an object. A robust method for determining object attitude from the EGI is developed here, using this construction. Experiments show the method's insensitivity to small attitude differences.

1 Introduction

Orientation maps can be generated by binocular stereo [Baker and Binford, 1981] [Grimson, 1981], photometric stereo [Woodham, 1980], shape from shading [Horn, 1975], or by differentiation of laser range images [Brou, 1984]. By translating the surface normals of an object to a common point, a representation of the distribution of surface orientation is formed, called the Extended Gaussian Image (EGI) [Horn, 1984]. Figure 1 shows the Extended Gaussian Image of a polyhedron and the corresponding object.

Horn and Ikeuchi [1984] demonstrated the feasibility of using EGIs for attitude determination, comparing the EGI of a prototype with a sensed EGI. The reconstruction method using the mixed volume suggests a new method for attitude determination. This method is practical and more robust than direct comparison of EGIs. A detailed discussion of this method is presented in [Little, 1985].

* This research was supported in part by a UBC University Graduate Fellowship and by the Natural Science and Engineering Research Council of Canada grant A3390.

2 Object Models

Objects will have surfaces tiled with planar facets. The intersection of half-spaces forms a convex polyhedron. A bounded convex polyhedron is a *polytope*. The direction of the outward facing normal ω on a face of the polytope is the surface *orientation*, which can be identified with a point on the unit sphere U . The set of orientations of the faces of a polytope, Ω , is referenced by indices $1 \dots n$.

A plane J *supports* a convex body C if C lies on one side of the plane and C has at least one point in common with J . For any orientation ω there is a unique support plane J for C with equation $\langle \omega, x \rangle = c$. A support function of orientation can be defined:

$$h(\omega) = c, \text{ where } \langle \omega, x \rangle = c \text{ is a supporting plane}$$

A polytope in this setting can be represented by $H = (h_1, h_2 \dots h_n) = (h(\omega_1), h(\omega_2) \dots h(\omega_n))$, the values of the support function at Ω .

Two polytopes P and Q in R^3 are *homothetic* if

$$P = \{x \mid x = \lambda * y + t, y \in Q, \lambda \in R, \lambda > 0, t \in R^3\}$$

Homotheticity is invariance under translation and scaling.

The areas and orientations of the faces describe a function $A(\omega)$; if the polytope has a face with orientation ω , its value is the area of the face, otherwise, zero. This function is represented as a vector A of the values of $A(\omega_i)$.




Figure 1: Extended Gaussian Image and Object

The Extended Gaussian Image of a polytope P can be represented as a set of vectors N : $\{n_i | n_i = \omega_i A_i, A_i = \mathcal{A}(\omega_i)\}$. Minkowski[1897] showed that when N sums to 0, N uniquely represents a polytope, up to translation.

1 Recovering Shape from an EGI

The reconstruction method derives from Minkowski's theorem[1897], arising from the study of mixtures(convex sums) of convex bodies. The *mixed volume* of polytopes P and Q , $V(P, Q)$, is the product of the support function of P and the area function of Q :

$$V(P, Q) = 1/3 \langle H_P, A_Q \rangle$$

As a consequence of the theory of mixtures of polytopes:

$$V(P, Q)^3 \geq V(Q)^2 V(P) \quad (1)$$

where $V(P)$ is the volume of P . Equality holds only when P and Q are homothetic.

Minkowski's theorem states that, among all polytopes P of fixed volume, the P which minimizes $V(P, Q)$ is homothetic to Q . The reconstruction method[Little,1983] minimizes the mixed volume to recover object shape.

2 Determining Attitude

Determining the attitude of a known object is equivalent to finding a rotation R , such that $R(EGI_{prototype}) = EGI_{sensed}$. Attitude in R^3 can be identified with a rotation $R(\theta, n)$, where θ is the angle by which an object is rotated about an axis n in R^3 .

To access the area function it is necessary to quantize orientation, tessellating U . The projection of the icosahedron onto U generates a tessellation. The number of facets in a tessellation U can be increased by subdividing the triangular faces of the icosahedron into smaller triangles. At frequency i , each face is mapped into i^2 triangles.

Attitude is quantized by the rotation group Γ_{60} which brings the vertices of the icosahedron into correspondence. The difference between two attitudes is the angle of the rotation taking one into the other. The minimum difference among the 60 attitudes in Γ_{60} is 72° . For a detailed discussion of methods for representing and quantizing attitude in R^3 and quantizing orientation on U , see Brou [1984].

2.1 Comparing Area Functions

To determine object attitude from the EGI, Horn and Ikeuchi [Horn and Ikeuchi,1984] match the prototype EGI and the sensed EGI at a discrete set of attitudes. At each sample attitude the matching measure between two area distributions is:

$$\sum_{i=1}^m (A_{R(i)} - A_i)^2 \quad (2)$$

where A_i is the sensed EGI area function at direction ω_i ; and $A_{R(i)}$ is the prototype area function evaluated at the direction to which ω_i is transformed by R . This amounts to maximizing the correlation of A and A . The attitude which minimizes the measure in Equation 2 is selected.

Area matching can fail when the object attitude is offset from a sample attitude[Brou, 1984]. A polytope with m faces fills at most m cells on U . As the resolution of orientation increases, the number of empty cells increases. If the attitude offset is near the cell resolution, the faces of the prototype and the object may not lie in the same cells, even at the correct attitude (the attitude in Γ_{60} nearest the actual attitude). The minimum matching error increases and the minimum may not indicate the correct attitude.

2.2 Mixed Volumes

Reconstruction builds a polytope whose area function fits the EGI; attitude determination seeks an attitude which rotates the sensed EGI into correspondence with the prototype EGI. Both determine a fit between an EGI and a model. Rotating a polytope preserves volume, so among all P' which are rotated versions of P , that P' which has the same attitude as P minimizes the mixed volume. To determine object attitude, minimize the mixed volume:

$$1/3 \sum_{i=1}^m (H_{R(i)} * A_i)$$

H_i is the prototype support function and A_i the sensed area function.

2.3 Support and Area Functions

For polytopes, the area function $\mathcal{A}(\omega)$ is discontinuous; in contrast, the support function $\mathcal{H}(\omega)$ is continuous on U . The comparative smoothness of the support function can be seen by examining its relation to the area function, most easily for polygons. The area function of a polygon, for example, consists of a finite set of non-zero points in the

interval $0 - 2\pi$. Its autocorrelation is 0 almost-everywhere. Not surprisingly, when the attitude of the sensed object is slightly different from the prototype (or any attitude in the sample set), the value of the area matching is maximum. Discretizing A by sampling U is helpful, as Brou remarks, introducing a smoothing effect, widening each of the pulses to the resolution. With discretization, small changes in attitude do not affect the area function, until the size of the attitude difference exceeds the resolution on U , when the errors recur.

The support function for a polygon can be written in terms of the area function A :

$$H(\omega) = \cos(\omega) \int_0^\omega -A(\psi) \sin(\psi) d\psi + \sin(\omega) \int_0^\omega A(\psi) \cos(\psi) d\psi$$

when the point on the polygon where $\omega = 0$ coincides with the origin. The support function is an integral transform of A and is much smoother. Because A varies rapidly, shifting A by a small rotation changes its correlation with the prototype A significantly. The support function varies less and the error rate is correspondingly lower. For a polygon, the support function is piecewise sinusoidal, as is the mixed volume.

3 Experiments

Initial tests employ the complete EGI, the total spherical distribution, to demonstrate feasibility. To examine the effects of attitudes different from the test set, at a specified resolution, the attitude difference between the object and the prototype was varied. The results can be evaluated under two criteria: the magnitude of the rotation between the selected attitude and the correct attitude, and, where the method ranks the correct attitude (it should have rank 0 – lower is better).

The EGI of the polytope shown in Figure 1 was rotated about the coordinate axes by angles of 5, 10, 15 and 20 degrees and its attitude determined by both methods. At frequency 2, the area matching method fails at angles as small as 15.0° ; both fail sometimes when the angle is 20.0° . When the mixed volume method fails, its failure is less severe in two respects: the attitude selected is closer to the correct attitude and the rank of the correct attitude is lower. Area matching fails in 4 out of 12 cases and MV in only 1 out of 12.

The results for the 20° rotation are shown in another form in Figure 2. The graph shows, on the vertical axis,




Figure 2: Mixed volume (solid) vs. area (dotted)

the values of area matching, in a dotted line, and mixed volume, in a solid line; the horizontal axis shows the attitudes (from Γ_{60}) at which the matchings were evaluated. The horizontal axis is separated by tick marks into regions which are 0, 72, 120, 144 and 180° from the prototype attitude. Point A is the minimum found by MV, correctly identifying the object attitude. Point B, the value of the area matching method for the correct attitude, is rather high; there are 21 attitudes with lower values. The area function achieves a minimum (point C) at an attitude in the 72° set.

At frequency 3, the MV method errs in 2 out of 12 cases (both at 20° offset), while the area matching method fails for 6 out of 12 cases (some as little as 10° offset). In general, the area matching method performs worse at a finer resolution. Finally, at frequency 5, the number of errors increases, for area matching, from 6 out of 12 to 8/12, while the mixed volume method fails twice. Three other objects, of varying eccentricity, were tested. The MV method succeeded for all offsets, while area matching failed 8 times at frequency 3 and 14 times at frequency 5.

In all cases, MV ranks the correct attitude in the top 5, suggesting that more sophisticated procedures can be applied to these to determine correct attitude. Area matching performs much more poorly at ranking the correct attitude. The better behaviour of MV at lower resolutions also suggests that MV would be useful in a coarse-to-fine search strategy.

3.1 Visible Subsets of EGIs

In practice, the sensed portion of a surface corresponds to a single visible hemisphere of the EGI. With only partial information, experiments with several prototypes contrasted the MV technique with the area method. In no case does the MV method fail where the area method succeeds; in all cases, the ranking of the MV method is lower than the




Figure 3: Non-convex polygon and its reconstruction

area method. There were 3 times as many errors by the area matching method.

Using MV for attitude determination depends on the fact that volume $V(P)$ remains constant (see Equation 1). Ignoring this effect incurred no penalty in these experiments, because the contribution of invisible faces to the mixed volume is negligible. A reconstruction method makes possible construction of suitable invisible faces.

4 Non-Convexity and Attitude

To extend the use of EGIs for non-convex objects, consider using the support function of the convex object reconstructed from the area function of a non-convex object. Dashed lines in Figure 3 delimit the convex polygon corresponding to the area function of this non-convex polygon. If, hypothetically, the entire area function of the sensed object were available, then recovering the attitude of the non-convex polygon would be indistinguishable from recovering attitude of the convex polygon. Using complete EGIs and ignoring the effects of self-occlusion, this method proved effective for a set of sample polygons, including the polygon in the Figure 3. This holds the promise that mixed volume method can be applied to attitude determination for non-convex objects, especially where the effects of self-occlusion are small.

5 Conclusions

An EGI supplies the sensed area function for an object. The object model permits calculation of the support function. The mixed volume combines these two functions; it is minimized when the attitude of the model matches that of the sensed object. Because of the smoothness of the support function, the mixed volume technique is robust for determining the attitude of convex objects. The experiments described here demonstrate its effectiveness and its

insensitivity to small changes in attitude which affect other methods. This method can also be applied to non-convex objects when the effects of self-occlusion are small.

References

- [1] H.H. Baker and T.O. Binford, "Depth From Edge and Intensity Based Stereo", *Proceedings of the Seventh IJCAI*, pp. 631-636, Vancouver, 1981
- [2] P. Brou, "Using the Gaussian Image to Find the Orientation of Objects", *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 89-125, Winter, 1984
- [3] W.E.L. Grimson, "From Images to Surfaces: A Computational Study of the Human Early Visual System", MIT Press, Cambridge, Mass, 1981
- [4] B.K.P. Horn, "Obtaining Shape From Shading Information", *The Psychology of Computer Vision*, ed. P.H. Winston, pp. 115-155, McGraw-Hill, New York, 1975
- [5] B.K.P. Horn, "Extended Gaussian Images", *Proceedings of the IEEE*, pp. 1671-1686, December, 1984
- [6] B.K.P. Horn and K.I. Ikeuchi, "The Mechanical Manipulation of Randomly Oriented Parts", *Scientific American*, August 1984
- [7] J.J. Little, "An Iterative Method for Reconstructing Convex Polyhedra from Extended Gaussian Image", *Proceedings of AAAI-83*, pp. 247-250, 1983
- [8] J.J. Little, "Recovering Shape and Determining Attitude from Extended Gaussian Images", Ph.D. Thesis, University of British Columbia, Vancouver, April 1985
- [9] Herman Minkowski, "Allgemeine Lehrsätze über die konvexe Polyeder", *Nachr. Ges. Wiss. Göttingen*, pp. 198-219, 1897
- [10] R.J. Woodham, "Photometric Method for Determining Surface Orientation from Multiple Images", *Optical Engineering*, vol. 19, pp. 139-144, 1980