

An Iterative Method for Reconstructing Convex Polyhedra from Extended Gaussian Images

James J. Little*

Department of Computer Science
University of British Columbia

ABSTRACT

In computing a scene description from an image, a useful intermediate representation of a scene object is given by the orientation and area of the constituent surface facets, termed the Extended Gaussian Image (EGI) of the object. The EGI of a convex object uniquely represents that object. We are concerned with the computational task of reconstructing the shape of scene objects from their Extended Gaussian Images, where the objects are restricted to convex polyhedra. We present an iterative method for reconstructing convex polyhedra from their Extended Gaussian Images.

I INTRODUCTION

The representation of an object by the orientation of its surface arises in many computer vision problems. Specifically, needle maps [Horn,1982], the " $2\frac{1}{2}$ -D sketch" [Marr,1976], and intrinsic images [Barrow and Tenenbaum,1978] all represent the orientation of an object at the points of the image. Orientation is specified as a vector pointing in the direction of the surface normal. Orientation maps can form the output of stereo processing from several images [Grimson,1981, Baker and Binford, 1981], photometric stereo [Woodham,1980], or any of the so-called "shape from" methods, such as shape from shading [Horn,1975, Ikeuchi and Horn,1981], shape from contour [Marr,1977], shape from texture [Kender,1979, Witkin,1981], and shape from edge interpretation [Mackworth,1973, Kanade,1981, Sugihara,1982]. By translating the surface normals of an object to a common point of application, a representation of the distribution of surface orientation is formed, called the Extended Gaussian Image (EGI).

Ikeuchi [Ikeuchi,1981] discussed the use of the EGI for recognizing objects in an industrial environment. For each unknown object, the EGI of its visible hemisphere is formed by a propagation of constraints method [Ikeuchi and Horn,1981]. The EGI is then compared to the EGIs of objects stored in a library. The best-matched prototype identifies the object.

Since it can be shown that the EGI does not uniquely identify a concave object, the EGI representation applies only to convex objects. In this discussion we will consider only convex polyhedra, which are formed by the intersection of a finite number of half-spaces. A bounded convex polyhedron will be termed a polytope. The EGI of a polytope P can be interpreted as a set of vectors, one for each face in the polytope.

The length of each vector is the area of the corresponding face in the polytope. Minkowski [Minkowski,1897] showed that the EGI of a convex object uniquely specifies the object up to a translation. Further, he proved that any set of vectors whose sum is zero represents the EGI of a convex object.

A natural question then arises: can one describe an algorithm for reconstructing a convex polytope from its EGI? Minkowski's proof of existence and uniqueness is not strictly constructive; it only provides an indirect route to the solution. Ikeuchi proposed an algorithm for generating the polytope corresponding to a given EGI with n faces, as follows. The solution is found by determining $L=(l_1, l_2, \dots, l_n)$, the n -vector of distances of the faces of the polytope from the origin. The vector L , together with the orientations of the faces, defines the locations in three-space of the half-spaces forming the polytope $P(L)$. The areas of the faces of $P(L)$, its volume and its centre of gravity can be computed. In the following discussion, we will consider that any polytope will be translated so that its centre of gravity coincides with the origin.

In Ikeuchi's algorithm for solving the reconstruction problem, the process is subdivided into n distinct cases; in the i^{th} case, face i is the farthest from the origin. When face i is chosen as maximum, l_i is set to 1.0; all other l_j vary between 0.0 and 1.0. The $n-1$ dimensional space of distances is quantized (at spacing d). Each of the d^{n-1} locations in this space specifies the locations of the n faces in three space. The polytope can be constructed, and the areas of its faces determined. These areas are scaled and compared with the objective.

No analysis of the accuracy of the algorithm is supplied. Ikeuchi's method minimizes the sum of the square differences between the calculated areas of the polytope and the given areas in the EGI. It is not clear that the polytope which results from this minimization (after normalizing) will have the same structure as the desired polytope. In addition, the method is very expensive. To double the resolution of the algorithm, one must increase the number of evaluation points exponentially.

II CONSTRUCTIVE METHODS

To find a constructive solution, first consider the two-dimensional case. The EGI of a polygon is a system of vectors emanating from the origin. If the system sums to zero, then it represents a convex polygon. Figure 1 shows a two-dimensional EGI; the reconstructed polygon is rotated by $-\frac{\pi}{2}$.

* This research was supported in part by a UBC University Graduate Fellowship and an NSERC Postgraduate Scholarship.




Figure 1 The EGI of a Convex Polygon and Its Reconstruction

To construct the polygon, given the system of vectors, one proceeds as follows:

Assume the vectors $\{ \mathbf{v}_i \}$ are given from 1 to n in anti-clockwise order. Take \mathbf{v}_1 , rotate it by $\frac{\pi}{2}$ and place its tail at some point in the plane. For the remaining vectors, in order, rotate \mathbf{v}_i by $\frac{\pi}{2}$ and place its tail at the head of \mathbf{v}_{i-1} . Because the system sums to zero, the head of \mathbf{v}_n will close with the tail of \mathbf{v}_1 . By definition, the length of each vector is the length of the corresponding edge in the polygon, and its orientation is normal to that of the edge. Hence each edge in the reconstructed polygon will be the correct length and at the proper orientation.

The two-dimensional method does not directly extend to higher dimensions. In two dimensions, the adjacencies among the facial elements, the edges, is clear from the EGI. In three dimensions the adjacency relationships are not given by the EGI and must form part of the solution. In that case, how can the solution be formulated?

A result of [Tutte,1962] states that the number of different adjacency relations for polytopes with n triangular faces is asymptotically exponential in n . The number of general polytopes (with faces having any number of sides) is larger. Hence any method which examines all possible adjacency relations will take exponential time.

III MINKOWSKI'S PROOF

Minkowski's proof provides clues for finding a reconstruction method. The original proof considers polytopes in any dimension d ; we will describe the proof in 3-space for clarity. For a polytope P in R^3 , the following set of vectors is formed: $U(P) = \{ \mathbf{u}_i \mid 1 \leq i \leq n \}$ where each \mathbf{u}_i is a non-zero vector emanating from the origin parallel to the outward normal of face i of P . The length of each \mathbf{u}_i is the area of face i , A_i . This set of vectors corresponds to the EGI given above. A set of vectors U is equilibrated if and only if they sum to zero and no two vectors are positively proportional, i.e., no two are linear multiples of a common unit vector. An equilibrated set of vectors U is fully equilibrated if and only if it spans R^3 . Minkowski's polytope reconstruction theorem shows that

- 1) if P is a polytope in R^3 not contained in any plane then the $U(P)$ is fully equilibrated and
- 2) if U is a fully equilibrated system of vectors, then there exists a polytope P unique within a translation such that U is the EGI of P .

This description is taken from [Grunbaum,1967,p.332].

Let \mathbf{L} be the n -vector of distances from the origin of the faces of the polytope $P(\mathbf{L})$. In the proof of condition (2), Minkowski shows that \mathbf{L} minimizes

$$f(\mathbf{L}) = \sum_{i=1}^n A_i l_i \quad (1)$$

where A_i is the area of face i given by the EGI and l_i is the distance of face i from the origin, subject to the constraint that the volume of $P(\mathbf{L})$, $V(\mathbf{L})$, is greater than or equal to one. By the Brunn-Minkowski theorem [Grunbaum,1967], the subset of R^n given by $\{ \mathbf{L} \mid V(\mathbf{L}) \geq 1 \}$ is convex. Convexity of the constraint set implies that the minimum of the objective function $f(\mathbf{L})$, since it is linear, will lie on the boundary of the convex set, where $V(\mathbf{L})=1$, and that a local minimum of $f(\mathbf{L})$ is the global minimum. Reconstructing a polytope from its EGI can be accomplished by solving a suitably formulated constrained minimization problem.

IV THE ITERATIVE METHOD


In an iterative solution to a constrained optimization problem, a sequence of points is generated which satisfy the constraints, i.e. are feasible, and which converges to the optimum [Gill et al., 1981]. To devise a minimization method, we, first, give a procedure for constructing $P(\mathbf{L})$, and, second, show that, from any $P(\mathbf{L})$, we can find an \mathbf{L}' such that $V(\mathbf{L}')$ is 1. This permits the generation of a convergent sequence of feasible points by starting from an initial point, taking a step toward the minimum, restoring feasibility, and repeating.

A. Constructing $P(\mathbf{L})$

To construct a polytope $P(\mathbf{L})$, we form the intersection of the n half-spaces specified by the vector \mathbf{L} . Brown [1978] describes a method for transforming the problem of intersecting n half-spaces into a convex hull problem. Brown uses the dual transform, described in the vision literature by [Huffman,1971, Mackworth,1973, Draper,1981]. The dual transform takes a plane with equation

$$A x + B y + C z + 1 = 0 \quad (2)$$

into the point (A,B,C) in R^3 (see figure 2). The planes of P do not pass through the origin so equation (2) is defined for all faces. The n planes forming P correspond to n points in R^3 , for which the algorithm of Preparata and Hong [1978] determines the convex hull in $O(n \log n)$ time. Any face of the convex hull of the dual points corresponds to a vertex of P . Any two points incident on an edge in dual (P) correspond to a pair of faces of P which share an edge. In sum, the adjacency information in the dual provides the adjacency information for P . Hence we can construct the vertices and edges of P . The centroid of P must coincide with the origin so its centre of gravity must be computed; each l_i is augmented by the scalar product of the centre of gravity, a point in R^3 , and the normal vector of face i .



1 = ACB 2 = AEDC 3 = DEF 4 = ABFE 5 = BCDF

Figure 2 A polytope and its dual

B. Restoring Feasibility

Once $P(\mathbf{L})$ has been constructed, it is straightforward to determine a corresponding point \mathbf{L}' which is feasible. The volume $V(\mathbf{L})$ of a 3-d polytope $P(\mathbf{L})$ is a homogeneous polynomial in \mathbf{L} of degree 3. The formula for the gradient of $V(\mathbf{L})$ can be derived from this polynomial. The gradient is used in computing the minimizing step. From a given $P(\mathbf{L})$ we can compute the volume $V(\mathbf{L})$ and scale \mathbf{L} by $V(\mathbf{L})^{\frac{1}{3}}$, yielding a polytope $P(\mathbf{L}')$ with unit volume.

C. Determining a Minimizing Step

Constrained optimization is a well-studied problem, so many methods are available for determining the step direction and magnitude [Gill et al., 1981]. The reduced gradient method is a simple method which was chosen for implementation. By taking a step in R^n in the hyperplane perpendicular to $\mathbf{G}(\mathbf{L})$, we will remain close to the constraint surface $V(\mathbf{L})=1$. The step is in the direction which minimizes $f(\mathbf{L})$, that is, in the direction of the projection of the vector \mathbf{A} , the n-vector of areas of the faces given by the EGI, onto the hyperplane perpendicular to $\mathbf{G}(\mathbf{L})$. This step is a multiple of:

$$\langle \mathbf{A}, \mathbf{G}(\mathbf{L}) \rangle \mathbf{G}(\mathbf{L}) - \mathbf{A}, \text{ where } \langle x, y \rangle \text{ is the inner product} \quad (3)$$

D. The Method

The iterative method for reconstructing a convex polyhedron from its EGI combines the procedures described above. The procedure is formulated as follows:

- 1) Set \mathbf{L} to $(1,1,\dots,1)$.
- 2) Construct $P(\mathbf{L})$:
 - 1) Transform the n planes given by \mathbf{L} into M , a set of n points in R^3 , using the dual transform.
 - 2) Compute the convex hull of M , call it $CH(M)$.
 - 3) Determine the adjacency relations of $P(\mathbf{L})$ from $CH(M)$. Calculate the locations of the vertices of $P(\mathbf{L})$.
- 3) Compute the centroid of $P(\mathbf{L})$. Translate the centroid of $P(\mathbf{L})$ to the origin. Compute $V(\mathbf{L})$ and the gradient of V , $\mathbf{G}(\mathbf{L})$. Scale \mathbf{L} by $V(\mathbf{L})^{\frac{1}{3}}$ to make its volume unity.
- 4) Evaluate $f(\mathbf{L})$; if the decrease in f is less than a pre-specified value, terminate. Otherwise, compute a step using equation (3), update \mathbf{L} , and repeat, starting at step 2.

V PERFORMANCE

An example polytope has been reconstructed from its EGI (figure 3). The polytope to which the EGI corresponds is shown in figure 4.




Figure 3 Stereo View of the EGI of a Distorted Octahedron




Figure 4 Stereo View of the Original Polytope

The faces of the polytope are parallel to those of a regular octahedron, while the distances of the faces from the origin have been altered. The polytope constructed initially is shown (in stereo) in figure 5.




Figure 5 Stereo View of Initial Polytope

The initial polytope is an octahedron, in which each face is adjacent to three others. In the course of the minimization, intermediate polytopes exhibit changing adjacency structures. The adjacency structure at an early stage becomes identical to that of the target polytope. The final reconstructed polytope is shown in figure 6; the value of \mathbf{L} for this polytope is :

$$(0.336, 0.699, 0.519, 1.137, 1.222, 0.517, 0.460, 0.443)$$

and its adjacency structure is:

FACE : ADJACENT TO FACES

1	:	2	3	4	8	5	6
2	:	1	6	3			
3	:	1	2	6	7	8	4
4	:	1	3	8			
5	:	1	8	6			
6	:	1	5	8	7	3	2
7	:	3	6	8			
8	:	1	4	3	7	6	5




Figure 6 Stereo View of the Reconstructed Polytope

The reconstructed polytope has the same adjacency structure as the original polytope. An advantage of this minimization formulation is its indifference to the adjacency relations in the polytope. A correct adjacency structure is guaranteed by Minkowski's original argument.

The iterative reconstruction method terminated on the fourteenth step, when the value of the objective function $f(\mathbf{L})$ had decreased by less than 0.002% on successive steps. The distances of the planes vary on average less than 0.9% from the original; the maximum difference is 4.2%.

The requirements of the reconstruction procedure can be factored into two components: the number of iterations required to find an acceptable solution and the number of operations per iteration. Each iteration requires $O(n \lg n)$ operations to compute the convex hull of the n dual points. In addition, $O(n)$ operations are necessary to evaluate the volume. Each iteration thus requires $O(n \lg n)$ computations. The number of iterations depends on the constrained minimization method used. The convergence rate of an iterative method is said to be linear if the error at step i , ϵ_i , satisfies the following formula:

$$\lim_{i \rightarrow \infty} \frac{|\epsilon_{i+1}|}{|\epsilon_i|^r} = \gamma \quad (4)$$

where $\gamma < 1$ and $r=1$. A reduced gradient method [Gill et al., 1981] was implemented; its convergence rate is linear. When the exponent r in equation (4) is 2, the convergence rate is said to be quadratic. To achieve quadratic convergence, the Hessian matrix of $V(\mathbf{L})$ or an approximation to the Hessian must be used, which requires $O(n^2)$ operations. Thus reducing the number of steps by improving the convergence rate requires expending more resources per step.

VI SUMMARY

The iterative method presented answers an open question on the inversion of the Extended Gaussian Image representation. The formulation of polyhedron reconstruction as a minimization problem has permitted the use of existing powerful minimization techniques.

ACKNOWLEDGMENTS

Thanks go to William Firey for helpful suggestions, to David Kirkpatrick and Alan Mackworth for discussions, and to Robert Woodham for his valuable advice.

REFERENCES

- H.H. Baker and T.O. Binford, "Depth From Edge and Intensity Based Stereo," *Proc. Seventh International Joint Conference on Artificial Intelligence*, pp. 631-636 (1981).
- H.G. Barrow and J.M. Tenenbaum, "Recovering Intrinsic Scene Characteristics From Images," pp. 3-26 in *Computer Vision Systems*, ed. E.M. Riseman, Academic Press, New York (1978).
- K. Q. Brown, "Fast Intersection of Half Spaces," CMU Technical Report CMU-CS-78-129 (1978).
- S.W. Draper, "The use of gradient and dual space in line-drawing interpretation," *Artificial Intelligence* **17** pp. 461-508 (1981).
- P.E. Gill, W. Murray, and M.H. Wright, *Practical Optimization*, Academic Press, New York, New York (1981).
- W.E.L. Grimson, *From Images to Surfaces: A Computational Study of the Human Early Visual System*, MIT Press, Cambridge, Mass (1981).
- Branko Grunbaum, *Convex Polytopes*, John Wiley and Sons, Ltd. , London and New York (1967).
- B.K.P. Horn, "Obtaining Shape From Shading Information," pp. 115-155 in *The Psychology of Computer Vision*, ed. P.H. Winston, McGraw-Hill, New York (1975).
- B.K.P. Horn, "Sequins and Quills - a representation for surface topography," in *Representation of 3-dimensional Objects*, ed. R. Bajcsy, Springer-Verlag, Berlin and New York (1982).
- D.A. Huffman, "A duality concept for the analysis of polyhedral scenes," in *Machine Intelligence*, ed. B. Meltzer and D. Michie, Edinburgh Univ. Press, Edinburgh, U.K. (1971).
- K.I. Ikeuchi, "Recognition of 3-D Objects Using the Extended Gaussian Image," *Proceedings of the Seventh IJCAI*, pp. 595-600 (1981).
- K.I. Ikeuchi, and B.K.P. Horn, "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence* **17**(1981).
- T. Kanade, "Recovery of the Three Dimensional Shape of an Object from a Single View," *Artificial Intelligence* **17** pp. 409-461 (1981).
- J.R. Kender, "Shape From Texture : an Aggregation Transform That Maps a Class of Textures Into Surface Orientation," *Proceeding of the Sixth International Joint Conference on Artificial Intelligence*, pp. 475-480 (1979).
- A.K. Mackworth, "Interpreting Pictures of Polyhedral Scenes," *Artificial Intelligence* **4**(2) pp. 121-137 (1973).
- D. Marr, "Early Processing of Visual Information," *Phil. Trans. Royal Society of London* **275B**(942) pp. 483-524 (1976).
- David Marr, "Analysis of Occluding Contour," *Proc. Royal Soc. London* **B**(197) pp. 441-475 (1977).
- Herman Minkowski, "Allgemeine Lehrsätze über die konvexe Polyeder," pp. 198-219 in *Nachr. Ges. Wiss. Göttingen*, (1897).
- F.P. Preparata and S.J. Hong, "Convex Hulls of Finite Sets of Points in Two and Three Dimensions," *CACM* **20** pp. 87-93 (1977).
- K. Sugihara, "Mathematical Structures of Line Drawings of Polyhedrons - Toward Man-Machine Communication by Means of Line Drawings," *Pattern Analysis and Machine Intelligence* **4** pp. 458-468 (1982).
- W.T. Tutte, "A Census of Planar Triangulations," *Canadian Journal of Math.* **14** pp. 21-38 (1962).
- A.P. Witkin, "Recovering Surface Shape and Orientation from Texture," *Artificial Intelligence* **17** pp. 17-47 (1981).
- R.J. Woodham, "Photometric Method for Determining Surface Orientation from Multiple Images," *Optical Engineering* **19** pp. 139-144 (1980).