

Polyhedra of Minimal Area as 3D Object Models

Joseph O'Rourke

Electrical Engineering Department
Johns Hopkins University
Baltimore, MD 21218

ABSTRACT

Polyhedra of minimal surface area are suggested as natural candidates for object models. The problem of computing such a polyhedron from a set of 3D points that are its vertices is explored. An approximate algorithm is suggested, and results of a preliminary implementation are discussed.

1. INTRODUCTION

Ullman's "Structure from Motion Theorem" [16] and similar schemes [1] [8] [9] enable the three dimensional (3D) coordinates of corresponded points on the surface of a rigid object to be extracted from a sequence of images. After the 3D coordinates of a set of points has been found, a natural next step is to form a 3D object model for the set of points [9] [3]. The object model could then be used for prediction and verification in later frames of the image sequence, or for recognition purposes, or for shape analysis.

Let us ignore the other information that might be extracted from the images, and ask the following question: What is the most reasonable 3D object model for a set of 3D points? Posed in this stark form, there is surely no best answer for all applications. First, there are a large variety of representations for solid objects [14], and second, there are a variety of definitions of "most reasonable" for each representation.

This paper explores one possible object model: *polyhedra*. It is proposed that a polyhedron of minimal surface area is the most "natural" polyhedral model for a given set of points. This concept is explored theoretically, and an algorithm for computing an approximation to this polyhedron is described. Preliminary results of an implementation of the algorithm are presented.

2. POLYHEDRA OF MINIMAL SURFACE AREA

We will say that a simple polyhedron^{*} P is a *minimal surface area polyhedron* for a set of 3D points S iff (a) the set of vertices of P is identical to S , and (b) no other simple polyhedron whose vertex set is precisely S has a smaller surface area.

One could also define polyhedra of maximal volume, but these seem to be less "natural" than polyhedra of minimal surface area. A well-known fact from elementary physics may partially explain the naturalness of minimal area surfaces: the tension in certain surfaces (e.g., soap bubble films) is proportional to their area, so that they tend to form the surface of minimal area wherever they have free movement.

The idea that maximizing the volume or minimizing the surface area may result in natural polyhedral shapes is not new. Keppel [7] interpolated triangular faces between parallel slices by maximizing the sandwiched volume, and Fuchs *et al* [4] attacked the

same problem by minimizing an arbitrary function of the tiles, for example, the surface area. The fact that only a single cylindrical band of faces have to be computed makes the interpolation problem far more tractable than the general case under consideration here.

A few basic properties of minimal area polyhedra are claimed here without proof:

- (1) A minimal area polyhedron *exists* for every set of 4 or more 3D points that are not all coplanar.
- (2) For a given set of 3D points, a minimal area polyhedron is not necessarily *unique*.
- (3) The *convex hull* is a minimal surface area polyhedron for the set of hull vertices, since the convex hull is the *only* simple polyhedron for its vertices.
- (4) The computation of the 2D analog of the minimal area polyhedron, the minimal perimeter polygon for a set of 2D points, is *NP-hard*; I conjecture that the 3D problem is *NP-hard* as well.

The reason that the 2D problem is *NP-hard* is that the minimal perimeter polygon is identical to the Euclidean Traveling Salesperson Path, and the computation of this has been shown to be *NP-hard* [5] [12].

If the computation of the minimal area polyhedron is in fact a fundamentally intractable problem, heuristics and sub-optimal algorithms are inevitable. A number of Traveling Salesperson Problem (TSP) heuristics are known [2] [15], but it is not clear how these may be directly applied to the 3D problem. However, their sense may be mimicked, and in the next section, an algorithm is proposed for computing an approximation to the polyhedron of minimal surface area via a type a "greedy" algorithm.

3. ALGORITHM: SHRINKING FROM THE HULL

That the convex hull is a minimal area polyhedron for its vertices ((3) above) suggests that it be used as the starting point of an algorithm that systematically modifies a given polyhedron to include new points. The assumption is that most objects encountered in the real world have large convex sections that coincide with the hull, and so an algorithm which starts with the convex hull will be fast and accurate on these objects.

The two dimensional convex hull is sometimes explained as the shape that a rubber band would take were it to encompass a set of pegs placed at the point locations. The physical analogy can be extended to non-convex three dimensional shapes: there is a type of electronic plastic shielding that shrinks to conform to whatever shape it encloses when heated.[†] This is roughly the behavior intended for the algorithm below:

^{*} The notion of a "simple polyhedron" may be made precise by a combinatorial topology definition, e.g., [6]. We intend the definition to require that the surface have genus 0, i.e., there are no holes.

[†] This analogy was suggested to me by Bernd Neumann.

```

Compute the convex hull;
while (there are internal points) do
  begin
    (1) Choose a point internal to the hull;
    (2) Modify the current polyhedron to include it;
    (2) Make local adjustments to the polyhedron
        in the vicinity of the modified region.
  end
end

```

The internal point chosen in step (1) is the one that is the "best" in the following sense. Each internal point is associated with the closest face that it orthogonally projects into, or just the closest face if it projects into none. An area stretch factor is computed for each point as the ratio of the area that would result if the closest face were modified as in step (2) to the area of the closest face. The "best" internal point is the one with the smallest stretch factor. The modification in step (2) consists of removing the face of the polyhedron closest to the chosen point, and replacing it with a number of faces that reach in towards the point and include it as a vertex (see Figure 1).




Figure 1. Shrinking in one polyhedron face to include an internal point.

This modification essentially puts a dent into the polyhedron, and it is quite possible that it will disturb the minimality of the surface area in the neighborhood of the dent. Step (3) of the algorithm is intended to improve the polyhedron in the vicinity of the alteration via a simple local transformation, called here a *flip* transformation (see Figure 2).




Figure 2. The *flip* transformation deletes the front faces ABC and ACD, and replaces them with the back faces ABD and BCD.

The flip transformation should only be performed if it "improves" the surface. Currently the improvement is measured by a heuristic mix of three terms: (a) the area reduction, (b) the perimeter reduction, and (c) the reduction in the sharpness of the face angles. If it is performed, then the neighboring faces should be considered for the flip operation also, and so on, recursively. The cascading is guaranteed to halt eventually, since only inward flips are permitted.

Note that the algorithm is essentially a "best-first" or "nearest neighbor" algorithm, but with local improvements after each greedy

step.




Figure 3. 53 points: 26 forming the cylinder end caps, and 27 distributed on a spherical dent.

A preliminary version of the algorithm has been implemented, and three samples of its performance are shown in Figures 3, 4, and 5. In each of the examples, the convex hull of the input point set is a 26 point cylinder determined by two 12 point rings plus a point in the center of each cap. The examples differ in the number and placement of the points internal to the hull. In Figure 3, 27 internal points form a spherical dent in the side of the cylinder. The 32 internal points in Figure 4 are distributed in 3 rings about the cylinder axis, the middle one having a smaller radius than its symmetrically placed neighbors. Figure 5's 50 internal points lie on the surfaces of two intersecting hemispheres. In all three examples, the algorithm has produced a surface close to the "natural" one, although some minor quirks are evident in the surface shown in Figure 5.




Figure 4. 58 points: 26 forming the cylinder end caps, and 32 arranged in three rings.

4. DISCUSSION

Relative to naive search, which requires exponential time, the shrinking algorithm is quite fast. The three dimensional convex hull of a set of n points can be computed in time $O(n \log n)$ by the divide-and-conquer method of Preparata and Hong [13].* The remainder of the algorithm has a cost of $O(fn)$ in the current implementation, where f is the number of the n points internal to the hull and f is the average number of flip transformations per step.† This is certainly no worse than $O(n^3)$, and may be $O(n^2)$.

The main theoretical drawback of the algorithm is that, in contrast to some of the TSP heuristics [2] [15], it is not guaranteed to achieve a surface area within a fixed percentage of the minimal area. The algorithm can be led astray by certain inputs; this is typical behavior of algorithms that seek a global optimum via heuristic-guided local search.

Finally, it is rarely the case that a 3D object model must be constructed solely from a set of 3D points: usually, other information is available. For example, the location of certain sharp edges may be known, or a silhouette may be available. The point of this note is to suggest that surface area minimization might be used as an aid in the formation of 3D object models when the other available constraints are insufficient for completely determining the 3D shape.




Figure 5. 76 points: 26 forming the cylinder end caps, and 50 distributed on intersecting hemispherical surfaces.

ACKNOWLEDGEMENTS

The work reported in this communication was begun at H.-H. Nagel's laboratory at the University of Hamburg. I thank Professor Nagel and his colleagues for their hospitality and for providing a stimulating research environment.

REFERENCES

- [1] Bonde, T. and Nagel, H.-H., "Deriving a 3-D Description of a Moving Rigid Object from Monocular TV-Frame Sequences," Workshop on Computer Analysis of Time-Varying Imagery, University of Pennsylvania, Philadelphia, 1979, pp. 44-45.
- [2] Christofides, N., "Worst-case analysis of a new heuristic for the Traveling Salesman Problem," Symposium on Algorithms and Complexity, Department of Computer Science, Carnegie-Mellon University, Apr. 1976.
- [3] Dreschler, L. and Nagel, H.-H., "Using 'Affinity' for Extracting Images of Moving Objects from TV-Frame Sequences," Universität Hamburg, Fachbereich Informatik, Bericht Nr. 44, Feb. 1978.
- [4] Fuchs, H., Z. Kedem, M., and Uselton, S. P., "Optimal Surface Reconstruction from Planar Contours," *CACM*, Vol. 20, No. 10, Oct. 1977, pp. 693-702.
- [5] Garey, M. R., Graham, R. L., and Johnson, D. S., "Some NP-Complete Geometric Problems," May 1976, pp. 10-22. *Proceedings of the Eighth Annual ACM Symposium on Theory of Computing*.
- [6] Giblin, P. J., *Graphs, Surfaces, and Homology*, Chapman and Hall, 1977, especially pp. 51-57.
- [7] Keppel, E., "Approximating Complex Surfaces by Triangulation of Contour Lines," *IBM Journal of Research and Development*, Vol. 19, No. 1, Jun. 1975, pp. 2-11.
- [8] Meiri, A. Z., "On Monocular Perception of 3-D Moving Objects," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-2, No. 6, Nov. 1980, pp.582-583.
- [9] Nagel, H.-H., "From Digital Picture Processing to Image Analysis," International Conference on Image Analysis and Processing, Pavia, Italy, Oct. 1980.
- [10] Neumann, B., "Motion Analysis of Image Sequences for Object Grouping and Three Dimensional Reconstruction," *Proceedings of the International Conference on Pattern Recognition*, Miami, Florida, Dec. 1980, pp. 1262-1265.
- [11] O'Rourke, J., "Polyhedral Object Models from 3D Points," Universität Hamburg, Fachbereich Informatik, Technical Report IfI-HH-M-77/80, Jul. 1980.
- [12] Papadimitriou, C. H. and Steiglitz, K., "Some Complexity Results for the Traveling Salesman Problem," Eighth Annual ACM SIGACT Symposium, May 1976, pp. 1-9.
- [13] Preparata, F. P. and Hong, S. J., "Convex Hulls of Finite Sets of Points in Two and Three Dimensions," *CACM*, Vol. 20, No. 2, Feb. 1977, pp. 87-93.
- [14] Requicha, A. A. G., "Representations for Rigid Solids: Theory, Method, and Systems," *ACM Computing Surveys*, Vol. 12, No. 4, Dec. 1980.
- [15] Rosenkrantz, D. J., Stearns, R. E., and Lewis, P. M., "Approximate Algorithms for the Traveling Salesperson Problem," *Proceedings of the 15th SWAT Conference*, 1974, pp. 33-42.
- [16] Ullman, S., *The Interpretation of Visual Motion*, Cambridge, Mass., MIT Press.

* This has been implemented in Pascal, and is remarkably fast, requiring only 17 seconds of CPU time (on a DEC-KL10) to compute the hull of 256 points. See [11] for details.

† For the examples shown in Figures 3-5, $f \approx 3$.