

Generating Models of Solid Objects by Matching 3D Surface Segments

Michael Potmesil

Bell Laboratories
Holmdel, New Jersey 07733

Abstract

This paper summarizes a method for constructing surface models of arbitrarily-shaped, solid objects by spatially matching 3D surface segments describing the objects. A surface segment of a solid object is a partial description of the object's surface, usually limited to the surface area that can be reconstructed from two corresponding images of the object in an arbitrary orientation. Each surface segment is defined in its 3D coordinate system. The matching algorithm uses heuristic search to align overlapping surface segments of an object into a common 3D coordinate system. It searches for a rigid 3D transformation which minimizes shape differences between a pair of surface segments. Once this transformation is determined, the surface segments are merged into a new segment. This match-and-merge process is iteratively repeated for pairs of available surface segments until a complete model of the object is assembled.

1.0 Introduction

A recurring problem in computer vision and related fields has been automatic generation of computer models of the surfaces of arbitrarily-shaped, physical, three-dimensional objects. This paper contains a brief summary of a process for constructing such models [8]. The process of modeling a 3D object is divided here into two major steps:

- (1) acquisition and hierarchical structuring of 3D surface segments of the object, followed by
- (2) matching and merging of the surface segments into a complete model of the object.

The surface models are composed of networks of curved, topologically rectangular, parametric patches. The *bicubic parametric patch* [3] is used as the basic *surface element*. Adjacent patches with positional and derivative continuities across their boundaries form *sheets of composite patches*. Given enough of these composite patches, an arbitrarily-shaped surface can be represented to an arbitrary precision. The 3D surface data, required to define these patches, are obtained by photogrammetric methods from multiple images of surfaces illuminated by a pattern of orthogonal lines defined on a unit-square grid. The reconstructed surface patches are then converted into a *quadtree* hierarchical structure by recursively merging four adjacent patches into a new coarser patch at the next higher level of the tree.

The surface-acquisition method generates only a description of a *surface segment* - a surface area that is parametrized by a projected pattern of lines, and also visible in at least two images of the object. Several such segments must be obtained to cover the complete surface area of a typical object. They are obtained by changing the object's orientation and repeating the surface acquisition process. Each surface segment is, therefore, defined in its own 3D object space. The surface segments are obtained, however, in such a way that they partially overlap to allow their eventual alignment by matching the common surface areas.

In the novel step in this modeling process, the individual surface segments of an object are transformed into a common 3D space by a matching procedure and then their overlapping sections are eliminated by a merging procedure. A heuristic-search algorithm is used to register spatially two partially overlapping surface segments into a common 3D space by minimizing distances of shape-features on the surface segments. This match-and-merge process is iteratively repeated for all 3D surface segments until a complete model of the object, in a single 3D object space, is generated.

2.0 Generation of 3D Surface Segments

2.1 Acquisition of 3D Surface Data

Surface information is obtained by photogrammetric and image-processing techniques from multiple images of the measured object illuminated by a controlled-light source [4,7,8]. To model an existing surface with 3D parametric patches, the surface is parametrized with a grid of 2D orthogonal isoparametric lines. These lines trace boundary curves of the patches being computed. The parametrized surface is then recorded in at least two 2D images. The projected lines are identified in these images, and then reconstructed back into 3D using camera geometry of the images. A calibration stand, in which the measured object is placed, permits determination of the camera geometry directly from the image data.

2.2 Hierarchical Structuring of Surface Data

The surface patches provided by the reconstruction method are converted into a hierarchical quadtree representation. The bottom level of the tree contains the original patches;

the higher levels, generated by patch merging, contain coarser and coarser approximations to the original surface. Each node of a quadtree consists of a patch, an invisible bounding volume - either a sphere or a parallelepiped - cast around the patch, and a normal vector which is the average of the normal-vector field [2] of the surface elements within the node. The normal vector of a quadtree node gives the orientation of the surface within it and is used in the matching process. This quadtree format of bivariate surface elements allows, in general, logarithmic rather than linear searching and sorting time of the individual elements.

A single ray-casting procedure is used to access this geometrical surface representation. It computes all intersections of a ray - a line - and the surface description. At each intersection it also evaluates the surface-normal vector, the surface-normal curvature, and other surface information. This procedure is used in all matching, merging, and other algorithms. Since the type of geometrical surface representation is confined to this procedure, other surface representations, such as planar or quadric, can easily be added [9], and all other procedures are independent of the actual surface representation.

3.0 Matching and Merging of Surface Segments

3.1 Matching Algorithm

Matching of 3D surfaces refers, in the context of this work, to finding a spatial registration of two 3D surface descriptions that maximizes their shape similarities [8]. The goal of the matching algorithm is to compute a rigid 3D transformation which spatially registers (aligns) two surface segments defined in two different object coordinate systems. This transformation occurs when the orientation and the shape differences between the two surfaces are minimized (i.e., the shape similarity is maximized). If a measurement of the shape similarities is acceptable then the surfaces are said to match, and the transformation aligns them. In order to find such a match, the surface segments must completely or partially overlap, i.e., there must be a set of surface points common to both of them. The orientation and shape difference are evaluated at a number of *evaluation points* on the first segment from information provided by the ray-casting procedure. For a ray in the direction of the surface-normal vector at an evaluation point of the first segment this procedure computes:

- (1) *positional difference*, the 3D Euclidean distance from the feature point to the nearest point on the second segment;
- (2) *orientation difference*, the angular difference of the surface-normal vectors of the two points; and
- (3) *curvature difference*, the magnitude difference of the surface-normal curvatures of the two points.

The evaluation points are selected at each level of the quadtree representation, either at patch control points, or at surface points with maximum curvature. A heuristic-search algorithm [6] with an evaluation function controls

the generation of new transformations. From the distances of the evaluation points on the first segment to the intersected points on the second segment the algorithm evaluates the current match and generates new transformations for the successor nodes in the search process. The initial estimates of the match are generated by aligning the normal vectors of the top quadtree nodes. New nodes in the search are generated by aligning the evaluation points with the intersected points, and finally by varying the transformation parameters.

3.2 Merging Algorithm

The basic merging algorithm [8] reparametrizes two or more surface segments by mapping a new parametric grid on them. This approach is similar to that used in Section 2 to acquire 3D surface data by physically projecting an actual grid on the object being measured. Here the projection is done to mathematical representations of the surface segments. The projected grid is usually defined in a plane, and orthographic projection is used. Alternatively, the projected grid could be defined on a spherical or cylindrical surface surrounding the surface segments being merged.

A similar approach can be used to convert closed surface segments describing a solid object into a *volumetric* representation. Rays projected from parametric grid are intersected with the surfaces of the object. The intersections and the sampling distances of the rays define a stack of parallel parallelepipeds, which can further be converted into the *octree* representation [9].

4.0 An Example of Model Generation

This example illustrates the entire modeling process of obtaining a 3D description of a solid object. The object is a blob-shaped balsa model of a car (apparently meant for the late 1980's). The object was positioned in six different object systems in which the top, bottom, left, right, front, and back sides, respectively, of the object could be photographed. In each object system, the object was consecutively illuminated by three parameter systems. While illuminated by a parameter system, two images of the object were recorded. The computer model was therefore generated from 36 images of the object.

Four views of the original object are shown in Figure 4.1. A proof print of the 36 images of the illuminated object is given in Figure 4.2. The square marks in these images are the camera calibration marks from which the camera geometry was determined. Each row contains six images of the object in one object system; each pair of images within a row contains the object illuminated by the same parameter system. There were 18 3D surface segments [Figure 4.3] reconstructed from these 36 images. Since every three segments were defined in the same object system, they were directly merged into one segment. In Figure 4.4 each drawing in the left column shows three surface segments in the same object system before merger, and the corresponding drawing in the right column shows




Figure 4.1 Four views of the modeled object.




Figure 4.2 Proof print of 36 images of the modeled object.

the merged surface segment. At last, the six surface segments shown in Figure 4.4 were matched and aligned into the same object system. This matching process consisted of five iterations; the coordinate system of the first surface segment remained stationary while the other five segments, one at a time, were matched with the existing surface description [Figure 4.5]. Four arbitrary views of the matched segments are shown in this figure together with the axes of the object coordinate system.

5.0 Conclusions

A method for generation of 3D models of solid objects was presented. It is based on a 3D matching process which







Figure 4.3 18 3D surface segments reconstructed from 36 images in Figure 4.2.




top side



bottom side




left side




right side

Figure 4.4 18 surface segments of Figure 4.3 merged into six new segments.



front side



back side

Figure 4.4 (CONTINUED)





Figure 4.5 Surface segments of Figure 4.4 matched into a single model.

spatially registers partial surface descriptions. The model representation is versatile enough to be used in application tasks in scene analysis, robotics, CAD/CAM, computer graphics, and other areas. Two immediate applications are (a) conversion to octree representation for fast object analysis and display [5], and (b) generation of characteristic views for object recognition [1].

The matching algorithm has been applied here only to the problem of assembling models of objects from overlapping surface segments. It can be generalized and expanded to perform also these tasks:

- (1) *object recognition* - a surface description will be matched with a set of models to determine whether the surface is a part of one of the models,
- (2) *object segmentation* - a model will be matched with surface or volume-shape primitives; when a match is obtained, the matched region is segmented and the process is repeated in order to decompose a complicated object into a structure of simple shapes, and
- (3) *matching of non-overlapping segments* - fragments of a solid object, such as broken pottery pieces, will be matched and assembled into the original object.

6.0 Acknowledgments

The work presented in this paper was performed by the author as a part of his doctoral research in the Image Processing Laboratory, Rensselaer Polytechnic Institute, Troy, New York. It was supported in part by the National Science Foundation, Automation, Bioengineering and Sensing Program, under Grant ENG-7904821, and in part by the Office of Naval Research under Contract N00014-82-K-0301, Professor Herbert Freeman, principal investigator. This support is gratefully acknowledged.

7.0 References

- [1] Chakravarty, I., "The Use of Characteristic Views as a Basis for Recognition of Three-Dimensional Objects," IPL-TR-034, Ph.D. Thesis, Image Processing Laboratory, RPI, Troy, NY, October 1982
- [2] do Carmo, M. P., *Differential Geometry of Curves and Surfaces*, Prentice Hall, New Jersey, 1976
- [3] Faux, I. D., and Pratt, M. J., *Computational Geometry for Design and Manufacture*, John Wiley and Sons, New York, 1979
- [4] Freeman, H., and Potmesil, M., "Curved Surface Representation Utilizing Data Extracted from Multiple Photographic Images," *Proceedings of a Workshop on the Representation of Three-Dimensional Objects*, University of Pennsylvania, May 1979, H1-H26
- [5] Meaghar, J. R., "Geometric Modeling Using Octree Encoding," *Computer Graphics and Image Processing*, 19, 1982, 129-147
- [6] Nilsson, N. J., *Problem Solving Methods in Artificial Intelligence*, McGraw Hill, New York, 1970
- [7] Potmesil, M., "Generation of 3D Surface Descriptions from Images of Pattern-Illuminated Objects," *Proceedings of IEEE Computer Society Conference on Pattern Recognition and Image Processing*, Chicago, Illinois, August 1979, 553-559
- [8] Potmesil, M., "Generating Three-Dimensional Surface Models of Solid Objects from Multiple Projections," IPL-TR-033, Ph.D. Thesis, Image Processing Laboratory, RPI, Troy, NY, October 1982
- [9] Roth, S. D., "Ray Casting for Modeling Solids," *Computer Graphics and Image Processing*, 18, (1), January 1982, 109-144