

A Study On Applications And Techniques Of Surface Re-Construction

¹Subhanga Kishore Das, ²Dr. Sabyasachi Pattnaik
HOD, Department of Comp. Sc. & Engg, BOSE, Cuttack-753007,
Reader, Dept. Of I & CT, F.M University, Balasore, Orissa-756 019
subha_1971@hotmail.com, spattnaik40@yahoo.co.in

Abstract- This paper describes a general method for automatic reconstruction of accurate, concise, piecewise smooth surfaces from unorganized 3D points. Instances of surface reconstruction arise in numerous scientific and engineering applications, including reverse-engineering, the automatic generation of CAD models from physical objects etc.

Previous surface reconstruction methods have typically required additional knowledge, such as structure in the data, known surface genus, or orientation information. In contrast, the method outlined in this paper requires only the 3D coordinates of the data points. From the data, the method is able to automatically infer the topological type of the surface, its geometry, and the presence and location of features such as boundaries, creases, and corners. The surface reconstruction method has three major phases:

Initial surface estimation, Mesh optimization, and piecewise smooth surface optimization.

In this paper emphasis has been given on the initial surface estimation.

1.0 INTRODUCTION

Computer-aided geometric design and computer-aided manufacturing systems are used in numerous industries to design and create physical objects from digital models. However, the reverse problem, that of inferring a digital description from an existing physical object, has received much less attention. We refer to this problem as reverse-engineering or, more specifically, 3D scanning. There are various properties of a 3D object that one may be interested in recovering, including its shape, its color, and its material properties. This paper addresses the problem of recovering 3D shape, also called *surface reconstruction*

The goal of surface reconstruction can be stated as follows: Given a set of sample points X assumed to lie on or near an unknown surface U , create a surface model S approximating U (see Figure 1.1).




Figure 1.1: Example of surface reconstruction.

As we shall see shortly, previous reconstruction methods have usually been crafted to exploit characteristics of specific problem instances. In contrast, this paper examines the surface reconstruction problem in a general form that makes few assumptions about the sample X and the unknown surface U . In the *general surface reconstruction problem* we consider, the points X may be noisy, and no structure or other information is assumed within them. The surface U (assumed to be a manifold¹) may have arbitrary topological type², including boundaries, and may contain sharp features such as the creases and corners present in the surface of Figure 1.1a. Since the points X may be a noisy sampling, we do not attempt to interpolate them, but will instead find an approximating surface.

A surface reconstruction procedure cannot guarantee recovering U exactly, since it is only given information about U through a finite set of sample points. The reconstructed surface S should have the same topological type as U , and be everywhere close to U . In this thesis we will evaluate the reconstruction method by considering examples where the true underlying surface U is known and can be compared visually and quantitatively to the reconstruction.

1.1 Surface Re-Construction applications

Here some general surface reconstruction problems have been considered by a number of application areas in science and engineering, including 3D scanning, surface reconstruction from contours, and surface sketching. The

next three sections discuss each of these application areas in more detail.

1.1.1 3D scanning

One of the most important applications of surface reconstruction is *3D scanning*—the measurement and modeling of shape and other visual properties. Our vision is to put shape on an equal footing with these other media. We would like to acquire, represent, analyze, and recreate 3D shapes with ease.

There are numerous methods for acquiring shape information. For instance, in computer vision, registration of landmarks in multiple views is used to infer object shape. In a different technique called shape from shading, the intensity of light reflected from the object's surface provides knowledge of surface orientation, and with further processing, the global shape of the object.

In the manufacturing industries, mechanical touch probes mounted on coordinate measuring machines are used to record points on surfaces such as car bodies and airplane wings. The resulting measurements are very accurate, but the technique is extremely slow and limited to materials that can withstand mechanical contact. Cheaper, less accurate hand-held 3D digitizing probes determine position using magnetic fields.

Recently, mechanical probes are being replaced by laser range scanners. Laser range scanners illuminate the object with a laser beam, and measure distance using either triangulation, interference or time-of-flight principles (for an extensive survey of range imaging sensors, see Besl [4]). Laser range scanners are promising because they can provide dense, accurate range data at high bandwidths.

Range scanning systems typically produce *range images*—rectangular grids of distances from the sensor to the object being scanned. If the sensor and object are fixed, only objects that are “point viewable” can be fully digitized. More sophisticated systems, such as those produced by Cyberware Laboratory, Inc., are capable of digitizing cylindrical objects by rotating either the sensor or the object. To adequately scan objects of more complicated topological type, such as the object depicted in Figure 1.1a (a surface of genus 3), multiple range images must be generated. Although the resulting data contains structure within each range image, merging the data to reconstruct a useful surface representation is a non-trivial task.

3D scanning applications The development of fast, inexpensive 3D scanning systems opens up a vast range of applications such as **reverse engineering, industrial design, analysis and simulation, populating virtual environments, 3D faxing etc.**

To fully realize the potential of 3D scanning, it is essential to develop general, automatic, efficient, and robust surface

reconstruction algorithms for converting the data points that 3D scanners produce into useful models.

1.1.2 Contour data

Another application area involves the reconstruction of surfaces from contours. In many medical studies it is common to slice biological specimens into thin layers with a microtome. The outlines of the structures of interest are then digitized to create a stack of contours. In manufacturing, similar stacks of contours are also produced by cross-section CAT scans of mechanical parts. The *surfaces from contours* problem attempts to recover the three dimensional structures from the stacks of parallel two-dimensional contours. Although the problem has received a good deal of attention [3,9], there remain severe limitations with current methods. Perhaps foremost among these is the difficulty of automatically dealing with branching structures. While algorithms addressing the general surface reconstruction problem may not be as successful as methods specialized for contour data, they need not consider such special cases.

In a related problem, ultrasound sensing is used to study the shape of the heart [8]. Contour images of the heart are obtained after insertion of a probe into the esophagus. Unlike the microtome data, the ultrasound contours are not parallel. Moreover, the probe is able to generate many sets of contours from different directions and from different positions. Algorithms for solving the surfaces from contours problem cannot be easily applied to this type of data.

1.1.3 Surface sketching

A number of researchers, including Schneider [9] and Eisenman, have investigated the creation of curves in R^2 by tracing the path of a stylus or mouse as the user sketches the desired shape. Sachs describe a system, called 3-Draw, that permits the creation of free-form curves in R^3 by recording the motion of a stylus fitted with a Polhemus sensor. This can be extended to the design of free-form surfaces by ignoring the order in which positions are recorded, allowing the user to move the stylus arbitrarily back and forth over the surface. The problem is then to construct a surface representation faithful to the unordered collection of points.

1.2 Surface reconstruction Algorithms

1.2.1 Algorithms assuming fixed topological type

A common restriction of surface reconstruction methods is that they assume that the topological type of the surface is known a priori.

Parametric reconstruction methods Parametric methods represent the reconstructed surface as an embedding $f: \hat{C} \rightarrow R^3$ of a 2-dimensional parameter domain \hat{C} . Since the domain \hat{C} and the surface $f(\hat{C})$ are homeomorphic, parametric reconstruction methods inherently require knowledge of the topological type of the surface. Moreover,

to converge correctly, they also require an initial embedding $f_0(\wedge)$ that is “sufficiently close” to U . Equivalently, they assume a “good” initial parameterization of the points X in \wedge . This presents a problem since such an initial parameterization may be difficult to construct.

Function reconstruction Terms like “surface fitting” appear in reference to two distinct classes of problems: surface reconstruction and function reconstruction. The goal of surface reconstruction was stated earlier. The goal of function reconstruction may be stated as follows: Given a surface D , a set $X_i \in D$, and a set $Y_i \in R$, determine a function

$$f: D \rightarrow R, \text{ such that } f(x_i) \sim y_i.$$

The domain surface D is most commonly a plane, in which case the problem is a standard one considered in approximation theory. The case where D is a sphere has also been treated extensively. Foley [2] defines radial basis functions centered on points scattered over a sphere. Schudy and Ballard [1, 2] use spherical harmonics to fit a surface as a function over a spherical domain. Sclaroff and Pentland [5] describe a hybrid implicit/parametric surface fitting method that involves fitting a function over a deformed superquadric. Some recent work under the title *surfaces on surfaces* addresses the case when D is a general curved surface such as the skin of an airplane [2]. Function reconstruction methods can be used for surface reconstruction in simple, special cases, where the surface to be reconstructed is, roughly speaking, the graph of a function over a *known* surface D . It is important to recognize just how limited these special cases are — for example, not every surface homeomorphic to a sphere is the graph of a function over the sphere. The point is that function reconstruction must not be misconstrued to solve the general surface reconstruction problem.

1.2.2 Algorithms exploiting structure information

Many surface reconstruction algorithms exploit structure in the data. For instance, algorithms solving the surfaces from contours problem make heavy use of the fact that the data points are organized into contours, and that the contours lie in parallel planes. Similarly, algorithms to reconstruct surfaces from multiple range images typically exploit the adjacency relationship of the data within each range image. These approaches have the drawback that they must deal with special cases using ad hoc techniques. It is therefore difficult to apply them to similar but different problems. For instance, methods solving the surfaces from contours problem cannot be used when presented with several sets of intersecting contours.

1.2.3 Algorithms exploiting orientation information

Knowledge of the orientation of the surface at each data point is extremely valuable in surface reconstruction. In fact, automatically determining such orientation is one of

the main challenges in our method. Many previous reconstruction methods assume that such orientation information is supplied with the data. When the data points X are obtained from volumetric data, the gradient of this data can provide orientation information that helps guide the reconstruction.

1.2.4 Algorithms for triangulating noise-free data

Some recent computational geometry methods come close to addressing the general surface reconstruction problem. They find meshes of arbitrary topological type that interpolate sets of unorganized points. Since they interpolate the data, their main limitation is that they require the data to be noise-free.

1.2.5 Implicit surface fitting algorithms

Several methods fit algebraic implicit surfaces (zero sets of polynomial functions) to sets of points. However, the intent of these methods is not to reconstruct surfaces but to either recognize objects or infer their orientations in a scene. These fitting methods cannot be used directly for surface reconstruction because the topological type of algebraic surfaces is highly unpredictable; in most cases, fitting an algebraic surface to a set of points results in numerous surface sheets that happen to pass near the data but only connect up far away (e.g. Figure 1.2). One approach to controlling the topological type of an implicit surface is to triangulate space and define a *piecewise* algebraic function of low degree over the resulting simplices.




Figure 1.2: Problem with fitting an algebraic surface.

1.3 Overview of the surface reconstruction method

Surface reconstruction algorithms have typically been designed to exploit additional knowledge in specific problem instances. In contrast, our approach is to pose a unifying general problem. This approach has both theoretical and practical merit. On the theoretical side, abstracting to a general problem often sheds light on the truly critical aspects of the problem. On the practical side, a single algorithm that solves the general problem can be used to solve any specific problem instance.

We have developed a method for automatically reconstructing an accurate, concise piecewise smooth

surface S from a set of points X , where X is an unorganized, noisy sample of an unknown surface U ; the unknown surface U can have arbitrary topological type (including boundaries), and may contain tangent plane discontinuities such as creases and corners; no other information, such as structure in the data or orientation information, is provided. A major difficulty in this general surface reconstruction problem is that the topological type of U is not known a priori and must be inferred from the points. To tackle this difficulty, we have partitioned the reconstruction problem: we first robustly determine the topological type of the surface, and only then concern ourselves with the accuracy and conciseness of the model. Our reconstruction method consists of three successive phases, as illustrated in Figures 1.3 and 1.4.

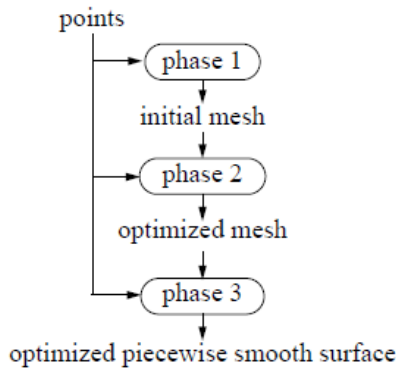


Figure 1.3: The three phases of our surface reconstruction method.

Phase 1: initial surface estimation: From an unorganized set of points (Figure 1.4a), phase 1 constructs an initial dense mesh (Figure 1.4b). The goal of this phase is to determine the topological type of the surface, and to produce an initial estimate of its geometry.

Phase 2: mesh optimization: Starting with the dense mesh created in phase 1, phase 2 reduces the number of faces and improves the fit to the data points (Figure 1.4c). We cast this problem as optimization of an energy function that explicitly models the trade-off between the competing goals of accuracy and conciseness. The free variables in the optimization are the number of vertices in the mesh, their connectivity, and their positions.

Phase 3: piecewise smooth surface optimization: In phase 3, the surface representation is changed from a piecewise linear one (meshes) to a piecewise smooth one. We introduce of a new piecewise smooth representation based on subdivision. These surfaces are ideal for surface reconstruction, as they are simple to implement, can model sharp features concisely, and can be fit using an extension of the phase 2 optimization algorithm.

(a) Unorganized points X (b) Result of phase 1: initial dense mesh (c) Result of phase 2: optimized mesh (d) Result of

phase 3: piecewise smooth surface Starting with the optimized mesh produced in phase 2, phase 3 fits an accurate, concise piecewise smooth subdivision surface (Figure 1.4d), again by optimizing an energy function that trades off accuracy and conciseness. In addition to varying the geometry and size of the surface representation, phase 3 also optimizes over the number and locations of sharp features. The automatic detection and recovery of sharp features in the surface is an essential part of phase 3. Phase 2 could in principle be eliminated, but has proven useful for two reasons:

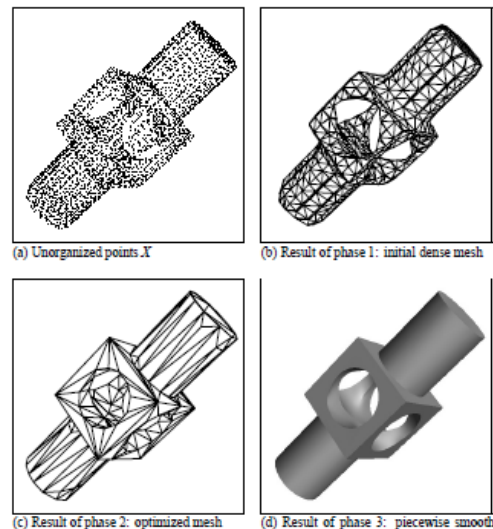


Figure 1.4: Example of the three phases of the surface reconstruction method.

first, it is computationally more efficient to optimize over a piecewise linear surface in the early stages of optimization, and second, initial estimates of sharp features are much more robust when obtained from the phase 2 mesh.

1.4 Conclusion

This paper has addressed the problem of reconstructing “surfaces”—orientable 2-dimensional manifolds embedded in R^3 . curves (1-dimensional manifolds). Future research should explore the reconstruction of more general manifolds, such as non-orientable manifolds and higher dimensional manifolds, as well as non-manifold sets.

This reconstruction method may be generalized to allow reconstruction of non-orientable manifolds. In phase 1, although a non-orientable manifold cannot be defined as the zero set of a globally defined signed distance function, it is possible to use such a description locally. Instead of globally orienting the tangent planes as we do now, it may be possible to determine their relative orientations on a cube by cube basis.

That is, when generating the contour within a cube, the tangent planes contributing to the function values at the cube's vertices can be oriented relative to each other by considering only a local neighborhood of the Riemannian Graph. In phases 2 and 3, the current implementation requires the surfaces to be orientable only because of the current half-edge data structure used to represent meshes . Using a different data structure would remove this restriction.

REFERENCES

- [1] E. L. Allgower and P. H. Schmidt. *An algorithm for piecewise linear approximation of an implicitly defined manifold*. *SIAM Journal of Numerical Analysis*, 22:322–346, April 1985.
- [2] R. E. Barnhill, K. Opitz, and H. Pottman. *Fat surfaces: a trivariate approach to triangle-based interpolation on surfaces*. *CAGD*, 9(5):365–378, November 1992.
- [3] J. L. Bentley. *Multidimensional divide and conquer*. *Comm. ACM*, 23(4):214–229, 1980.
- [4] P. J. Besl. *Active, optical range imaging sensors*. *Machine Vision and Applications*, 1(2):127–152, 1988.
- [5] P. J. Besl and H. D. McKay. *A method for registration of 3-D shapes*. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.
- [6] Jules Bloomenthal. *Polygonization of implicit surfaces*. *Computer Aided Geometric Design*, 5(4):341–355, November 1988.
- [7] Ruud M. Bolle and Baba C. Vemuri. *On three-dimensional surface reconstruction methods*. *IEEE Trans. Pat. Anal. Mach. Intell.*, 13(1):1–13, January 1991.
- [8] Y. Breseler, J. A. Fessler, and A. Macovski. *A Bayesian approach to reconstruction from incomplete projections of a multiple object 3D domain*. *IEEE Trans. Pat. Anal. Mach. Intell.*, 11(8):840–858, August 1989.
- [9] James F. Brinkley. *Knowledge-driven ultrasonic three-dimensional organ modeling*. *IEEE Trans. Pat. Anal. Mach. Intell.*, 7(4):431–441, July 1985.