Reconstruction of Surfaces From Scattered Points

Benjamin F. Gregorski, Bernd Hamann, and Kenneth I. Joy

Center for Image Processing and Integrated Computing (CIPIC), Department of Computer Science, University of California, Davis, CA 95616-8562, USA
{gregorsk,hamann,joy}@cs.ucdavis.edu

1 Introduction

Digitization devices and scanners generate very large point sets representing complicated geometric models. Data sets typically result from multiple scans, frequently even multiple viewpoints. We discuss a method that constructs a B-Spline surface from scattered data points, usable for further processing in a CAD system.

The basis for our reconstruction is the decomposition of the scattered points into a 3D strip tree data structure. Our data structure is an extension of the original strip tree presented in [3]. It is similar to a quadtree, except that each node in our tree represents a bounding box whose orientation is defined by the best-fit plane approximating the data points inside the box.

2 Re-orienting Bounding Boxes

Given a set of points (x,y,z), a bounding box is defined as a box that contains all the points. Typically a bounding box is a parallelpiped and is oriented with respect to the axes of the global coordinate system. To obtain a tighter bounding box, i.e. , one with smaller volume, the bounding box should be re-oriented by using the three principal directions of the underlying point set. The re-orientation process consists of the following steps:

1. Given a set of points (x_i,y_i,z_i), compute the best-fit plane for the point set. A coordinate frame in 3D space consists of three basis vectors and an origin. (We use principle component analysis (PCA) to determine the basis vectors see, [1].)

2. The normal vector of the best-fit plane defines the ordinate direction in the local frame. We express all points associated with a box in terms of its local coordinate frame, using the average of the points as the local origin.

An example of this is shown in Figure 1.

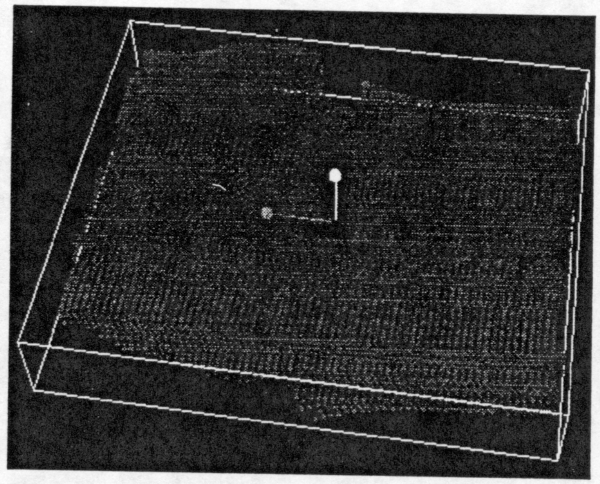

Figure 1: A "tight bounding box" for a set of points.

3 Our Generalized Strip Tree

Our strip tree is a quadtree in 3D space whose nodes are bounding boxes. Given a set of points, the root of the tree is a single, properly re-oriented bounding box that contains all points. A node in the tree is subdivided by first dividing its bounding box into four sub-boxes and then re-orienting these sub-boxes relative to the points inside them. As the strip tree is refined, the bounding boxes of the leaf nodes provide "tighter approximations" to the underlying points.

The subdivision process of a box consists of three steps:

1. First, we determine a subdivision point within the bounding box around which to form the four sub-boxes. This subdivision point is defined as the actual data point closest to the center of the bounding box. This point is used to divide the bounding box into four subboxes, all having parallel faces.

2. Once the four sub-boxes have been determined, the points associated with the original box are now associated with the four sub-boxes. The sub-boxes are reoriented using the procedure outlined in Section 2.

3. The *box points* for the new re-oriented sub-boxes are computed. The box points are either the four midpoints of the four box edges parallel to the normal vector of the best-fit plane or four original points closest to the midpoints of these edges.

The subdivision process is illustrated in Figures 1 - 3.

Figure 1 shows the root bounding box. Figure 2 shows the strip tree after one subdivision step. Figure 3 shows the strip tree after four subdivision steps.

4 Progressively Fitting Surfaces

Once the data points have been approximated well enough, the strip tree approximation is used to fit surfaces that approximate the data points. The approximation process starts from the next-to-last level in the strip tree and proceeds upward towards the root of the tree. The approximation process for a given node in the tree consists of the following steps:

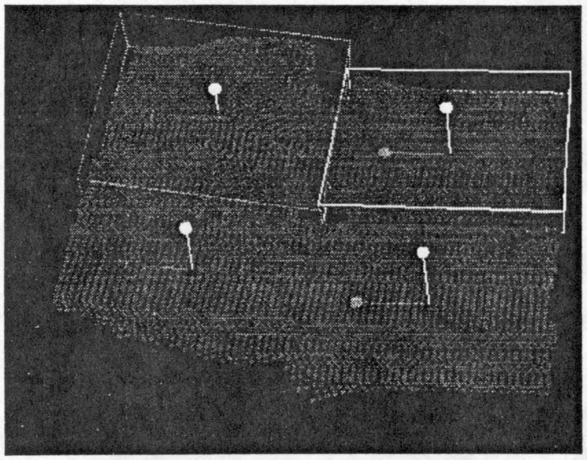

Figure 2: Strip tree after one subdivision step.

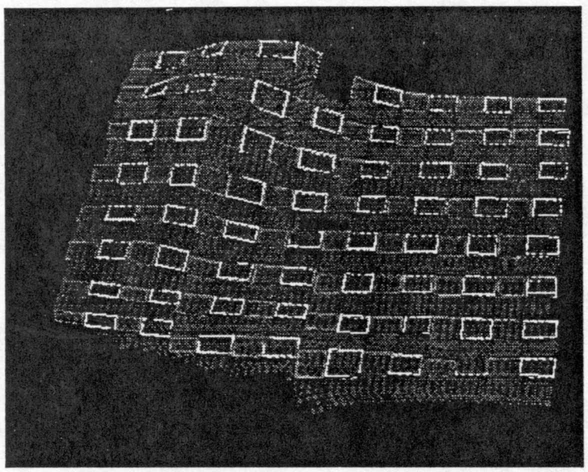

Figure 3: Strip tree after four subdivision steps.

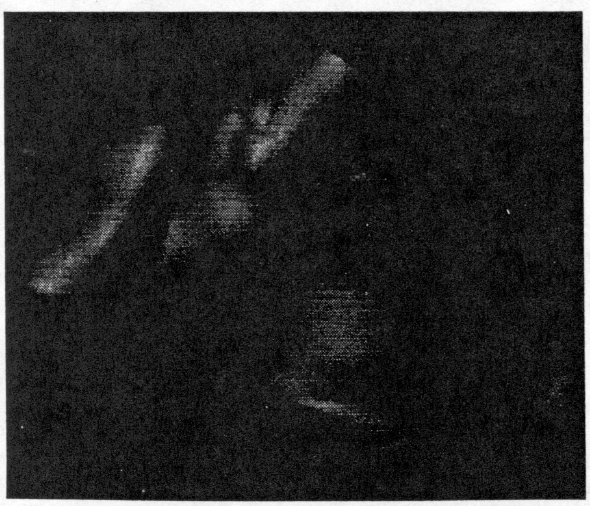

Figure 4: Final approximating surface for strip tree shown in Figure 3. All surface patches have been blended together.

1. If the node's children have not yet been approximated, a single bi-quadratic Bezier patch is fitted to the box points of the node's children. The 3x3 control points used to determine the patch are obtained from the four box points and by averages of interior points. This surface is degree-elevated to a bi-cubic Bezier patch.

2. If the node's children already have been approximated, the four underlying surfaces are blended together to form a new, C^1-continuous B-Spline surface that describes the four Bezier patches of its children.

In the curve case, this process consists of three steps:

1. The two curves become C^0-continuous by forcing the last control point of the first curve and the first control point of the second curve equal.

2. The two curves are combined into one curve by turning the double control point at the joint into a single control point and by removing one interior knot.

3. The curve becomes C^1-continuous by removing one of the interior knots and the corresponding control point. The knot removal algorithm we are utilizing is described in [2].

5 Results

We have tested our algorithm on a scattered data point set originating from a car door, provided by Alias-Wavefront. Figure 4 shows the final B-Spline surface (Gouraud shaded) obtained by blending all surfaces for the strip tree shown in Figure 3 .

6 Conclusions and Future work

We have presented a new method for reconstructing surfaces from scattered points. Our algorithm introduces a generalization of the strip tree data structure that is used to approximate the given points initially. This initial approximation is then used to construct a set of surfaces that approximates the actual data points. The strip tree needs to be enhanced to allow for a non-uniform subdivision of the data in regions with more data points or more complicated behavior. We also need to refine our error estimate. The error is currently approximated from the relative height of the bounding boxes and the error resulting from the knot removal process. These error metrics need to be integrated into a single measurement.

References

[1] J. Edward Jackson. *A User's Guide to Principal Components.* A Wiley-Interscience Publication, John Wiley and Sonc, Inc., New York, 1991.

[2] Les Piegl and Wayne Tiller. *The NURBS Book, 2nd Edition.* Springer-Verlag, New York, 1996.

[3] Hanan Samet. *The Design and Analysis of Spatial Data Structures.* Addison-Wesley, Reading, Massachusetts, 1990.