UC Davis

IDAV Publications

Title

Surface Reconstruction by Voronoi Filtering

Permalink

<https://escholarship.org/uc/item/4rc7z0d1>

Journal

Discrete and Computational Geometry, 22

Authors

Amenta, Nina
Bern, Marshall

Publication Date

1999

Peer reviewed

Surface Reconstruction by Voronoi Filtering

Nina Amenta*

Marshall Bern[†]

August 25, 1998

Abstract

We give a simple combinatorial algorithm that computes a piecewise-linear approximation of a smooth surface from a finite set of sample points. The algorithm uses Voronoi vertices to remove triangles from the Delaunay triangulation. We prove the algorithm correct by showing that for densely sampled surfaces, where density depends on “local feature size”, the output is topologically valid and convergent (both pointwise and in surface normals) to the original surface. We describe an implementation of the algorithm and show example outputs.

1 Introduction

The problem of reconstructing a surface from scattered sample points arises in many applications such as computer graphics, medical imaging, and cartography. In this paper we consider the specific reconstruction problem in which the input is a set of sample points S drawn from a smooth two-dimensional manifold F embedded in three dimensions, and the desired output is a triangular mesh with vertex set equal to S that faithfully represents F . We give a “provably correct” combinatorial algorithm for this problem. That is, we give a condition on the input sample points, such that if the condition is met the algorithm gives guaranteed results: a triangular mesh of the same topology as the surface F , with position and surface normals within a small error tolerance. The algorithm relies on the well-known constructions of the Delaunay triangulation and the Voronoi diagram.

This paper is an extension of previous work by Amenta, Bern, and Eppstein [1] on reconstructing curves in two dimensions. Our previous work defined a planar graph on the sample points called the “crust”. The crust is the set of edges in the Delaunay triangulation of the sample points that can be enclosed by circles empty not only of sample points, but also of Voronoi vertices. The crust comes with a guarantee: if the curve is “well-sampled”, then the crust contains exactly the edges between sample points adjacent on the curve. Our notion of well-sampled, which involves the medial axis of the curve, is sensitive to the local geometry. Hence our algorithm, unlike other algorithms for this problem, allows highly nonuniform sampling, dense in detailed areas yet sparse in featureless areas. Any provably

*Computer Sciences, University of Texas, Austin, TX 78712. Work performed in part while at Xerox PARC, partially supported by NSF grant CCR-9404113.

[†]Xerox Palo Alto Research Center, 3333 Coyote Hill Rd., Palo Alto, CA 94304

correct algorithm must impose some sampling density requirement, similar to the Nyquist limit in spectral analysis.

The extension to three dimensions in this paper requires both new algorithmic ideas and new proof techniques. Most notably the algorithm uses only a subset of the Voronoi vertices to remove Delaunay triangles. The algorithm picks only two Voronoi vertices—called *poles*—per sample point: the farthest vertices of the point’s cell on each side of the surface. With this modification, the straightforward generalization of our two-dimensional algorithm now works. Delaunay triangles with circumspheres empty of poles give a piecewise-linear surface pointwise convergent to F . The poles, however, also enable further filtering on the basis of triangle normals. Adding this filtering gives a piecewise-linear surface that converges to F both pointwise and in surface normals (and hence in area). We believe that poles may be applicable to other algorithms as well, perhaps whenever one wishes to estimate a surface normal or tangent plane.

This paper is organized as follows. Section 2 describes previous work on surface reconstruction. Section 3 gives our algorithm. Section 4 states our theoretical guarantees, and Section 5 sketches their proofs. Section 6 shows some example outputs.

2 Previous Work

Previous work on the reconstruction problem falls into two camps: computer graphics and computational geometry. The algorithms in use in computer graphics typically compute an approximating surface, that is, a surface passing close by, rather than exactly through, the original sample points. The algorithms devised by computational geometers typically restrict attention to surfaces on the original sample points, usually a carefully chosen subset of the Delaunay triangulation. Only recently have computational geometers started publishing algorithms with provable properties, and until this current paper these algorithms with guarantees applied only to reconstructing curves in two dimensions.

The first and most widely known reconstruction algorithm in the computer graphics community is the work of Hoppe et al. [19, 20, 21]. This algorithm estimates a tangent plane at each sample using the k nearest neighbors, and uses the distance to the plane of the closest sample point as a signed distance function. The zero set of this function is then contoured by a continuous piecewise-linear surface using the marching cubes algorithm. A later algorithm by Curless and Levoy [12] is designed for data samples collected by a laser range scanner. This algorithm sums anisotropically weighted contributions from the samples to compute a signed distance function, which is then discretized on voxels to eliminate the marching cubes step. These two computer graphics algorithms are quite successful in practice, but have no provable guarantees. Indeed there exist arbitrarily dense sets of samples, for example ones with almost collinear nearest neighbor sets, for which the algorithm of Hoppe et al. would fail.

The most famous computational geometry construction for associating a polyhedral shape with an unorganized set of points is the α -shape of Edelsbrunner et al. [14, 15]. Like our reconstructed surface, the α -shape is a subcomplex of the Delaunay triangulation. A Delaunay simplex (edge, face, etc.) belongs to the α -shape of S if its circumsphere has radius at most α . The major drawback of using α -shapes for surface reconstruction is that the optimal value of α depends on the sampling density, which often varies over different

parts of the surface. For uniformly sampled surfaces, however, α -shapes are workable. Bernardini et al. [7] follow α -shape-based reconstruction with a clean-up phase to resolve sharp dihedral angles. Edelsbrunner and Raindrop Geomagic [13] are continuing to develop α -shape-based reconstruction along with proprietary extensions.

An early algorithm due to Boissonnat [9] is related to ours. He proposed a “sculpting” heuristic for selecting a subset of Delaunay tetrahedra to represent the interior of an object. The heuristic is motivated by the observation that “typical” Delaunay tetrahedra have circumspheres approximating maximal empty balls centered at points of the medial axis; our algorithm relies on this same observation. Boissonnat’s algorithm, however, overlooks the fact that even dense sample sets can give Delaunay tetrahedra with circumspheres that are arbitrarily far from the medial axis; indeed it is this second observation which motivates our definition of poles. Goldak, Yu, Knight and Dong [18] made a similar oversight, asserting incorrectly that the Voronoi diagram vertices asymptotically approach the medial axis as the sampling density goes to infinity.

Finally, for the two-dimensional problem there are a few recent algorithms with provable guarantees. Figueiredo and Miranda Gomes [17] prove that the Euclidean minimum spanning tree can be used to reconstruct uniformly sampled curves in the plane. Bernardini and Bajaj [6] prove that α -shapes also reconstruct uniformly sampled curves in the plane. Attali [3] gives yet another provably correct reconstruction algorithm for uniformly sampled curves in the plane, using a subgraph of the Delaunay triangulation in which each edge is included or excluded according to the angle between the circumcircles on either side. Our previous paper showed that both the crust and the β -skeleton [22] (another empty-region planar graph) correctly reconstruct curves even with nonuniform sampling. Our two-dimensional results [1] are in this way strictly stronger than those of the other authors.

3 Description of the Algorithm

We start by describing the algorithm of Amenta et al. [1] for the problem of reconstructing curves in $\mathbb{R}^2$. Let F be a smooth (twice differentiable) curve embedded in $\mathbb{R}^2$, and S be a set of sample points from F . Let V denote the vertices of the Voronoi diagram of S . The *crust* of S contains exactly the edges of the Delaunay triangulation of $S \cup V$ with both endpoints from S . Saying this another way, the crust contains exactly those Delaunay edges around which it is possible to draw a circle empty of Voronoi vertices. In our earlier paper, we proved that if S is a sufficiently dense sample, this simple algorithm constructs a polygonal approximation of F (Theorem 1 in Section 4 below).

The straightforward generalization of this algorithm fails for the task of reconstructing a smooth two-dimensional manifold embedded in three dimensions. The problem is that vertices of the Voronoi diagram may fall very close to the surface, thereby punching holes in the crust. For example, the Voronoi center of a sliver can lie arbitrarily close to the surface F . A *sliver* is a tetrahedron with bad aspect ratio yet a reasonably small circumradius to shortest edge ratio, such as the tetrahedron formed by four nearly equally spaced vertices around the equator of a sphere.

The fix is to consider only the *poles*. The poles of a sample point s are the two farthest vertices of its Voronoi cell, one on each side of the surface. Since the algorithm does not know the surface, only the sample points, it chooses the poles by first choosing the farthest

1. Compute the Voronoi diagram of the sample points S .
2. For each sample point s :
 - (a) If s does not lie on the convex hull of S , let p^+ be the vertex of $Vor(s)$ farthest from s .
 - (b) If s does lie on the convex hull of S , let p^+ be a point at “infinite distance” outside the convex hull with the direction of sp^+ equal to the average of the outward normals of hull faces meeting at s .
 - (c) Among all vertices p of $Vor(s)$ such that $\angle p^+sp$ measures more than $\pi/2$, choose the farthest from s to be p^- .
3. Let P denote all poles p^+ and p^- , except those p^+ ’s at infinite distance. Compute the Delaunay triangulation of $S \cup P$.
4. (Voronoi Filtering) Keep only those triangles in which all three vertices are sample points.
5. (Filtering by Normal) Remove each triangle T for which the normal to T and the vector to the p^+ pole at a vertex of T form too large an angle (greater than θ for the largest-angle vertex of T , greater than $3\theta/2$ for the other vertices of T).
6. (Trimming) Orient triangles and poles (inside and outside) consistently, and extract a piecewise-linear manifold without boundary.

Figure 1. The surface reconstruction algorithm.

Voronoi vertex regardless of direction (or a fictional pole at “infinity” in the case of an unbounded Voronoi cell), and then choosing the farthest in the opposite half-space. See step 2 in Figure 1. Lemma 5 in Section 5 shows that this method is indeed correct for well-sampled surfaces. Denoting the poles by P , we define the *crust* of S to be the triangles of the Delaunay triangulation of $S \cup P$, all of whose vertices are members of S .

Steps 1–4 compute the crust (sometimes called the *raw crust* to distinguish it from the more finished versions). The crust has a relatively weak theoretical guarantee: it is pointwise convergent to F as the sampling density increases. Steps 5 and 6 are “postprocessing” steps that produce an output with a stronger guarantee: convergence both pointwise and in surface normals, and topological equivalence.

Step 5 removes triangles based on the directions of their surface normals. Let T be a triangle of the crust and let s be its vertex of maximum angle. Step 5 removes T if the angle between the normal to T and the vector from any one of T ’s vertices to its first-chosen pole is too large. The definition of “too large” depends on which vertex of T is under consideration: for the vertex with largest angle, too large means greater than an input parameter θ , and for the other two vertices it means greater than $3\theta/2$. Angles are unsigned angles in the range $[0, \pi/2]$. As stated in Theorem 5, the choice of θ is connected with the sampling density. If the user of our algorithm does not have an estimate of the sampling density (the parameter r in Definition 3 below), then the user can slowly decrease θ , backing off when holes start to appear in the surface, similar to choosing a surface from the spectrum of α -shapes [15].

Step 6 ensures that the reconstructed surface has the topology of the original surface; before this final step, the computed surface will resemble the original surface geometrically,

but may have some extra triangles enclosing small bubbles and pockets. The problem once again is slivers: all four faces of a flat sliver may make it past steps 4 and 5.

Step 6 first orients all triangles. Start with any sample point s on the convex hull of S . Call the direction to p^+ at s the *outside* and the direction to p^- the *inside*. Pick any triangle T incident to s , and define the outside side of T to be the one visible from points on the sp^+ ray. Orient the poles of the other vertices of T to agree with this assignment. Orient each triangle sharing a vertex with T so that they agree on the orientations of their shared poles, and continue by breadth-first search until all poles and triangles have been oriented. Our Theorem 5, below, guarantees that this orientation is consistent.

In a triangulated piecewise-linear two-dimensional manifold, two triangles meet at each edge, with outside sides together and inside sides together. Define a *sharp* edge to be an edge which has a dihedral angle greater than $3\pi/2$ between a successive pair of incident triangles in the cyclic order around the edge. In other words, a sharp edge has all its triangles within a small wedge. We consider an edge bounding only one triangle to have a dihedral of 2π , so such an edge is necessarily sharp.

Step 6 trims off pockets by greedily removing triangles with sharp edges. Now the remaining triangles form a “quilted” surface, in which each edge bounds at least two triangles, with consistent orientations. Finally, Step 6 extracts the outside of this quilted surface by a breadth-first search on triangles.

4 Theoretical Guarantees

What sets our algorithm apart from previous algorithms are its theoretical guarantees. We begin with the required sampling density, which is defined with respect to the medial axis.

Definition 1. *The medial axis of a manifold F embedded in $\mathbb{R}^d$ is the closure of the set of points in $\mathbb{R}^d$ with more than one nearest neighbor on F .*

Figure 2 gives an example of the medial axis in $\mathbb{R}^2$; in $\mathbb{R}^3$, the medial axis is generally a two-dimensional surface. Note that we allow the surface F to have more than one connected component.

Definition 2. *The local feature size $LFS(p)$ at a point p on F is the Euclidean distance from p to (the nearest point of) the medial axis.*

Definition 3. *Set $S \subset F$ is an r -sample of F if no point p on F is farther than $r \cdot LFS(p)$ from a point of S .*

Notice that the notion of r -sample does not assume any global—or even local—uniformity. Further notice that to prove an algorithm correct, we must place some condition on the set of sample points S , or else the original surface could be any surface passing through S . Our paper on curve reconstruction [1] proved the following theoretical guarantee.

Theorem 1 (Amenta et al. [1]). *If S is an r -sample of a curve in $\mathbb{R}^2$ for $r \leq .40$, then the crust includes all the edges between pairs of sample points adjacent along F . If S is an r -sample for $r \leq .25$, then the crust includes exactly those edges.*

Figure 2. The medial axis of a smooth curve.

To state our results for the three-dimensional problem, we must define a generalization of adjacency. Consider the Voronoi diagram of the sample points S . This Voronoi diagram induces a cell decomposition on surface F called the *restricted Voronoi diagram*: the boundaries of the cells on F are simply the intersections of F with the three-dimensional Voronoi cell boundaries. We call a triangle with vertices from S a *good triangle* if it is dual to a vertex of the restricted Voronoi diagram; good triangles are necessarily Delaunay triangles. Our first three-dimensional result shows that good triangles deserve their name. To our knowledge, our proof of this result is the first proof that the three-dimensional Delaunay triangulation of a sufficiently dense set of samples contains a piecewise-linear surface homeomorphic to F .

Theorem 2. *If S is an r -sample of F for $r \leq .1$, then the good triangles form a polyhedron homeomorphic to F .*

Our next two theorems state the theoretical guarantees for the three-dimensional (raw) crust.

Theorem 3. *If S is an r -sample for $r \leq .1$, then the crust includes all the good triangles.*

Theorem 4. *If S is an r -sample for $r \leq .06$, then the crust lies within a fattened surface formed by placing a ball of radius $5rLFS(q)$ around each point $q \in F$.*

Step 5 adds another guarantee: convergence in surface normals. The raw crust sometimes includes small skinny triangles with surface normals that deviate significantly from the surface normals. For example, the insides of the sausages shown on the left in Figure 13 have a sort of “washboard” texture. Convergence in surface normal ensures that the area of the trimmed θ -crust converges to that of the surface, and we use it in the proof of Theorem 6.

Theorem 5. *Assume S is an r -sample and set $\theta = 4r$. Let T be a triangle of the θ -crust and t a point on T . The angle between the normal to T and the normal to F at the point $p \in F$ closest to t measures $O(\sqrt{r})$ radians.*

Finally, the trimming or “manifold extraction” step, Step 6, adds the guarantee of topological equivalence.

Theorem 6. *Assume S is an r -sample and set $\theta = 4r$. For sufficiently small r , the trimmed θ -crust is homeomorphic to F .*

5 Proofs

In this section we give the proofs of the theoretical guarantees. We begin with some definitions. At each point $p \in F$, there are two tangent *medial balls* centered at points of the medial axis. The vectors from p to the centers of its medial balls are normal to F , and F does not intersect the interiors of the medial balls. Since $LFS(p)$ is at most the radius of the smaller medial ball, F is also confined between the two tangent balls of radius $LFS(p)$. We call these the *big tangent balls* at p (this is somewhat misleading since in general the medial balls at p are bigger); we will use the big tangent balls to bound the curvature of F in terms of $LFS(p)$. We call a maximal empty ball centered at a Voronoi vertex a *Voronoi ball*, and the Voronoi ball centered at a pole a *polar ball*.

Our first lemma is rather basic: a Lipschitz condition for the $LFS(p)$ function. We use $d(p, q)$ to denote the Euclidean distance from p to q . Angles are measured in radians.

Lemma 1. *For any two points p and q on F , $|LFS(p) - LFS(q)| \leq d(p, q)$.*

Proof: $LFS(p) \geq LFS(q) - d(p, q)$, since the ball of radius $LFS(q)$ around q contains the ball of radius $LFS(q) - d(p, q)$ around p and contains no point of the medial axis. Similarly, $LFS(q) \geq LFS(p) - d(p, q)$. ■

Our second lemma is a sort of Lipschitz condition for the direction of surface normals, which can be regarded as a function from F to the two-dimensional sphere.

Lemma 2. *For any two points p and q on F with $d(p, q) \leq \rho \min\{LFS(p), LFS(q)\}$, for any $\rho < 1/3$, the angle between the normals to F at p and q is at most $\rho/(1 - 3\rho)$.*

Proof: Let us parameterize the line segment pq by length. Let $p(t)$ denote the point on pq with parameter value t and let $f(t)$ denote the nearest point to $p(t)$ on the surface F . In other words, $f(t)$ is the point at which an expanding sphere centered at $p(t)$ first touches F . Point $f(t)$ is unique, because otherwise $p(t)$ would be a point of the medial axis, contradicting $d(p, q) \leq \rho LFS(p)$.

Let $n(t)$ denote the unit normal to F at $f(t)$, and $|n'(t)|$ the magnitude of the derivative with respect to t , that is, the rate at which the normal turns as t grows. The change in normal between p and q is at most $\int_{pq} |n'(t)| dt$, which is at most $d(p, q) \max_t |n'(t)|$.

The surface F passes between the big tangent balls of radius $LFS(f(t))$ at $f(t)$, so the greater of the two principal curvatures at $f(t)$ is no more than the curvature of these tangent balls. The rate at which the normal changes with $f(t)$ is at most the greater principal curvature, and hence $|n'(t)|$ is at most the rate at which the normal turns (as a function of t) on one of these tangent balls. Referring to Figure 3, we see that

$$dt \geq (LFS(f(t)) - d(f(t), p(t))) \cdot \sin \theta.$$

Figure 3. Bounding $|n'(t)|$ in terms of the radius $LFS(f(t))$ and $d(f(t), p(t))$.

Now $\sin \theta$ approaches θ as θ goes to zero, so

$$|n'(t)| = \lim_{\theta \rightarrow 0} \theta/dt \leq 1/(LFS(f(t)) - d(f(t), p(t))).$$

We have that

$$d(f(t), p(t)) \leq d(p(t), p) \leq \rho LFS(p)$$

and

$$d(f(t), p) \leq d(f(t), p(t)) + d(p(t), p) \leq 2\rho LFS(p),$$

so by Lemma 1, $LFS(f(t)) \geq (1 - 2\rho)LFS(p)$. Altogether we obtain $\max_t |n'(t)| \leq 1/((1 - 3\rho)LFS(p))$, which yields the lemma. ■

We next show that the cells of the Voronoi diagram of S are long and skinny. We let $Vor(s)$ denote the closure of the Voronoi cell of s , that is, all points at least as close to s as to any other sample point. We ignore the degenerate case that $Vor(s)$ is unbounded on both sides of F .

Lemma 3. *Let s be a sample point from an r -sample S .*

- (a) *On either side of F at s , some point of $Vor(s)$ has distance at least $LFS(s)$ from s .*
- (b) *The intersection of $Vor(s)$ and F is contained in a ball of radius $\frac{r}{1-r}LFS(s)$ about s .*

Proof: On either side of F at s , the center c of the big tangent ball of radius $LFS(s)$ lies within $Vor(s)$, and hence (a) holds. For part (b), let $p \in Vor(s) \cap F$. Since s is the closest sample point to p , $d(p, s) \leq rLFS(p) \leq r(LFS(s) + d(p, s))$ by Lemma 1. So $d(p, s) \leq \frac{r}{1-r}LFS(s)$. ■

The next lemma makes precise the idea that these long skinny Voronoi cells are perpendicular to the surface.

Lemma 4. *Let s be a sample point from an r -sample S . Let v be any point in $Vor(s)$ such that $d(v, s) \geq \nu LFS(s)$ for $\nu > 0$. The angle at s between the vector to v and the normal to the surface (oriented in the same direction) is at most $\arcsin \frac{r}{\nu(1-r)} + \arcsin \frac{r}{1-r}$.*

Figure 4. The vector from s to a distant Voronoi vertex such as a pole must be nearly normal to the surface.

Proof: Let B_v be the Voronoi ball centered on v . Let B_m be the medial ball touching s on the same side of the surface F , and let m be its center. Let ϕ be the angle between the segments sv and sm , that is, the angle referred to in the lemma. Let B_p be the ball of radius $LFS(s)$, tangent to F at s , but lying on the opposite side of F from B_m ; let p be the center of B_p . The surface F passes between B_m and B_p at s , and does not intersect the interior of either of them, as shown in Figure 4.

Since p and v lie on opposite sides of F , line segment pv must intersect F at least once. Let q be the intersection point closest to p . No sample point can lie in either B_p or B_v , so the nearest sample point to q must be s . Since B_p has radius $LFS(s)$, $d(q, s) \geq \sin(\alpha)LFS(s)$, where α is the angle $\angle spq$. We are interested in angle $\angle vsm$, which is $\phi = \alpha + \beta$. Since B_v has radius at least $\nu LFS(s)$, $d(q, s) \geq \nu \sin(\beta)LFS(s)$, where β is the angle $\angle svq$. Since S is an r -sample, $d(q, s)$ must be less than $\frac{r}{1-r}LFS(s)$. Combining the inequalities, we obtain $\alpha \leq \arcsin \frac{r}{1-r}$ and $\beta \leq \arcsin \frac{r}{\nu(1-r)}$, which together give the bound on ϕ . ■

Together Lemmas 3(a) and 4 show that the vector from a sample point to its first pole p^+ is a good approximation to the surface normal. This observation may have wider applicability than to our own surface reconstruction algorithm; for example, the Voronoi diagram and the poles could be used to obtain provably reliable estimates of tangent planes in the algorithm of Hoppe et al.

Our next lemma shows that we do indeed correctly select the second pole p^- . Recall that p^- is defined to be the farthest Voronoi vertex from s on the opposite side of the surface from p^+ .

Lemma 5. *Let s be a sample point from an r -sample S with $r \leq 1/3$. The second pole p^- of s is the farthest Voronoi vertex v of s such that the vector sv has negative dot product with sp^+ .*

Proof: By Lemma 3(a), $d(s, p^-) \geq LFS(s)$, so by Lemma 4 the angle between sp^+ and sp^- is at least $\pi - 4 \arcsin(r/(1-r))$, so $sp^- \cdot sp^+ < 0$. Lemma 4 also shows that for any Voronoi vertex v on the same side of F as p^+ , with $d(s, v) \geq LFS(s)$, the angle between

sv and sp^+ is at most $4 \arcsin \frac{r}{1-r} \leq \pi/4$. Hence any v farther from s than p^- must have $sv \cdot sp^+ > 0$. ■

Our next lemma bounds the angle between the normal to a good triangle and the surface normals at its vertices.

Lemma 6. *Let T be a good triangle and s a vertex of T with angle at least $\pi/3$, and choose $r < 1/7$. (a) The angle between the normal to T and the normal to F at s is at most $\arcsin(\sqrt{3}r/(1-r))$. (b) The angle between the normal to T and the normal to F at any other vertex of T is at most $2r/(1-7r) + \arcsin(\sqrt{3}r/(1-r))$.*

Proof: For part (a), let C be the circumcircle of T and let ρ_C be its radius. Consider the balls of radius $LFS(s)$ tangent to F at s on either side of F . These balls intersect the plane of T in “twin” disks of common radius ρ_B , tangent at point s , as shown in Figure 5. Our first aim is to bound ρ_B in terms of ρ_C .

Since the balls of radius $LFS(s)$ are empty of sample points, the twin disks cannot contain vertices of T . In order to maximize ρ_B relative to ρ_C , we assume that the twin disks pass through the vertices of T and that the angle at s measures exactly $\pi/3$. Now it is not hard to show that ρ_B is maximized exactly when T is equilateral: if we move s away from the midpoint of the arc covered by the twin disks, keeping the twin disks passing through the vertices of T , the radius ρ_B decreases, until s reaches one of the other vertices of T and $\rho_B = \rho_C$. Since the worst-case configuration is equilateral T , we can conclude that $\rho_B \leq \sqrt{3}\rho_C$.

We can bound these radii in terms of $LFS(s)$. Let u denote the restricted Voronoi diagram vertex dual to T . Since u lies on the line through the center of C normal to the plane of C , $\rho_C \leq d(u, s)$. By Lemma 3(b), $d(u, s) \leq \frac{r}{1-r} LFS(s)$, so altogether $\rho_B \leq \sqrt{3}r LFS(s)/(1-r)$.

Now to find the angle between the normal to T and the normal to F at s , we consider one of the big tangent balls B at s . Let m denote the center of B and v denote the center of the disk of radius ρ_B that is the intersection of B with the plane of T , as shown in Figure 5. The segment sm is normal to F at s and the segment mv is normal to T , so the angle we would like to bound is $\angle smv$. The triangle smv is right, with hypotenuse of length $LFS(s)$ and leg opposite $\angle smv$ of length $\rho_B \leq \sqrt{3}r LFS(s)/(1-r)$. Hence $\angle smv$ measures at most $\arcsin(\sqrt{3}r/(1-r))$.

For part (b), let s' be one of the other vertices of T . Since T is a good triangle, s and s' are neighbors in the restricted Voronoi diagram. Let p be a point on the boundary of both restricted Voronoi diagram cells. Then

$$d(p, s) \leq r LFS(p) \leq \frac{r}{1-r} \min\{LFS(s), LFS(s')\}.$$

So $d(s, s') \leq \frac{2r}{1-r} \min\{LFS(s), LFS(s')\}$. By Lemma 2, the angle between the normals to F at s and s' is at most $2r/(1-7r)$ for $r < 1/7$. ■

We need one more lemma for the proof of Theorem 2. This lemma is a topological result concerning the medial axis that may be independently useful.

Lemma 7. *If a ball B intersects surface F in more than one connected component, then B contains a point of the medial axis of F .*

Figure 5. Bounding the angle between the normal to the triangle and the normal to the surface at s .

Proof: Assume $B \cap F$ has more than one connected component. Let c be the center of B and p the nearest point on F to c . If p is not unique, then c is a point of the medial axis and we are done. Let q be the nearest point to c in a connected component of $B \cap F$ that does not contain p . Imagine a point c' moving from c towards q along segment cq . Throughout this journey, c' is closer to q than to any point outside B , so the closest point on F to c' must be some point of $B \cap F$. At the beginning of the journey, the closest point to c' is p and at the end it is q , so at some critical c' the closest point must change connected components. Such a c' is a point of the medial axis. ■

We now give the proof of Theorem 2: the good triangles form a polyhedron homeomorphic to F . The proof relies on the lemmas above along with a result of Edelsbrunner and Shah [16].

Proof of Theorem 2: The theorem of Edelsbrunner and Shah tells us that it suffices to show that S has the following *closed-ball property*: the closure of each k -dimensional face, $1 \leq k \leq 3$, of the Voronoi diagram of S intersects F in either the empty set or in a closed $(k - 1)$ -dimensional topological ball.

Let s be a sample point and $Vor(s)$ its Voronoi cell. Let the direction of the normal to F at s be vertical. Lemma 3(b) shows that $Vor(s) \cap F$ is small, fitting inside a ball B around s of radius $\frac{r}{1-r}LFS(s)$. Now Lemma 7 shows that $F \cap B$ has a single connected component, and Lemma 2 with $\nu = r/(1 - r)$ shows that $F \cap B$ is nearly horizontal, more precisely, the normal to $F \cap B$ is nowhere farther than $r/(1 - 4r) \leq 1/6$ radians from vertical, assuming $r \leq .1$. These statements in turn imply that $F \cap B$ is a topological disk; it cannot have a handle since it is everywhere nearly horizontal, and it cannot have a hole because its boundary is confined to the “low latitudes” of B .

First consider an edge e of $Vor(s)$, that is, the case $k = 1$. If e has nonempty intersection with F , then e is normal to the good triangle T dual to its intersection point. By Lemma 6(b), e must be within $2r/(1 - 7r) + \arcsin(\sqrt{3}r/(1 - r))$ radians from the normal to F at s . For $r \leq .1$, this expression is less than .9, so e is within .9 radians from vertical, and consequently can intersect F only once within B .

Next consider a face f of $Vor(s)$, that is, the case $k = 2$. Face f is contained in a plane h , the perpendicular bisector of s and another sample point s' , where ss' is an edge of a good triangle. Plane h must contain a vector parallel to the normal of T , so again

Lemma 6(b) establishes that the angle between h and the surface normal at s , and hence between f and the surface normal at s , is at most .9 radians when $r \leq .1$.

Consider a single connected component C of $f \cap F$, a nearly horizontal curve drawn across the face f . Let H be the set of points p in $B \setminus C$ such that the line segment from p to its closest point on C forms an angle smaller than .2 radians with horizontal. (Set H is a union of wedges with vertices on the curve C .) We assert that all points of $(F \cap B) \setminus C$ lie in H . We prove this assertion by showing that $F \cap B$ cannot cross the boundary of H . Assume $(F \cap B) \setminus C$ does contain a point p on the boundary of H . Let q be the closest point of C to p . The vertical plane P through p and q intersects $F \cap B$ in a curve. By the Mean Value Theorem there must be a point along this curve at which the tangent forms an angle greater than .2 radians with horizontal; the normal to F at this point must be at least .2 radians from vertical, a contradiction.

We further assert that all points of $f \cap B$ lie outside of H . Face f lies in a plane within .9 radians of vertical, and within a strip on this plane bounded by lines within .9 radians of vertical. All shortest segments from points of f to C lie within this strip, and hence are within .9 radians of vertical. Since f lies outside H and F inside H within B , C must be the only connected component of $f \cap F$, so $f \cap F$ is a topological 1-ball.

Finally consider $Vor(s)$ itself, the case $k = 3$. Consider any connected component C of the intersection of $F \cap B$ and the Voronoi cell. As in case of $k = 2$, let H contain each point that can be connected to its closest point of C by a line segment forming an angle smaller than .2 radians with horizontal. The same argument as above shows that $F \cap B$ cannot cross the boundary of H . Since each point along a face of $Vor(s)$ intersecting F can be connected to its closest point of C by a segment within .9 radians of vertical, the same is true of an interior point of $Vor(s)$. Since $F \cap B$ is confined to one piece of $B \setminus C$ and $Vor(S)$ to another, we can conclude that C is the only connected component of $F \cap Vor(S)$.

Aiming for a contradiction, assume that C is a topological disk with holes. Consider any vertical plane P that meets two components of the boundary of C at angles at least $\pi/2 - 1/6$. (To find such a plane, we could project the two boundary components onto a horizontal plane, and then sweep around a normal to one closed curve in order to find a line meeting each closed curve perpendicularly.) As shown in Figure 6, within plane P the boundary of $Vor(s)$ meets F at an angle larger than 1.07, extends some distance on the other side of F , and then recrosses F again at 1.07 from vertical. (Why 1.07? The face of $Vor(s)$ is within .9 of vertical as above, and .17 $> 1/6$ is added for P 's deviation from perpendicularity with the face.) Since the tangent to $F \cap P$ is everywhere within $1/6$ radians of horizontal, if F recrossed $Vor(s)$ within P , then $P \cap Vor(s)$ would be nonconvex, a contradiction.

Finally C cannot have a handle because it is a piece of the topological disk $F \cap B$. Hence C must itself be a topological disk and we are done. ■

Next we give a proof of Theorem 3: the raw crust contains all the good triangles. The intuition behind this proof is that restricted Voronoi cells are small and poles are far away, so that the ball centered at a vertex u of the restricted Voronoi diagram, passing through the three sample points whose cells meet at u , must be empty of poles.

Proof of Theorem 3: Let T be a triangle dual to a vertex u of the restricted Voronoi diagram. Consider the ball B_u centered on u with boundary passing through the vertices

Figure 6. A vertical cross-section of $F \cap \text{Vor}(s)$ shows the impossibility of a disk with a hole.

of T . Since T is a Delaunay triangle, B_u contains no point of S in its interior. Since S is an r -sample of F for $r < 1$, the radius of B_u is less than $rLFS(u)$. By the definition of LFS , even the larger ball B'_u of radius $LFS(u)$ centered on u cannot contain a point of the medial axis.

Now assume that B_u contains a pole v of a sample point s . We will show that under this assumption, first, that B_v must contain a point of the medial axis, and second, that the polar ball B_v must be contained in B'_u , thereby giving a contradiction. In particular, B_v must contain the center m of the medial ball B_m at s that is on the same side of F as v . Notice that m necessarily lies in $\text{Vor}(s)$ and ball B_m has radius at least $LFS(s)$, while the radius of B_v is at least that of B_m (by Lemma 3). By Lemma 4, $\angle msv$ measures at most $2 \arcsin \frac{r}{1-r}$, which is less than .23 for $r \leq .1$. A calculation shows that B_v must contain the medial axis point m .

Since v lies in B_u , the radius of B_v is no greater than the distance from v to the nearest vertex of T , which is at most $2rLFS(u)$ since S is an r -sample. Since $d(u, v) \leq rLFS(u)$, ball B_v lies entirely within B'_u since $3rLFS(u) \leq LFS(u)$. ■

We now move on to the proof of Theorem 4. Let s be a sample point and v a pole of s . We shall define a forbidden region inside polar ball B_v , which cannot be penetrated by large crust triangles.

Let B_m^+ be the big tangent ball at s , on the same side of F as v , and let B_m^- be the big tangent ball on the other side, with F passing between them. Let B be the ball concentric with B_m^- with radius $(1-a)LFS(s)$, as shown in Figure 7(a); a is a constant that will be chosen later. Notice that Lemma 3(a) shows that the radius of B_v is at least that of B .

Definition 4. The reflection of a point t through B_v is the point t' along ray vt such that line segment tt' is divided into equal halves by the boundary of B_v . The spindle of s is $\{t \in B_v \mid \text{segment } tt' \text{ intersects } B\}$, that is, all points in B_v whose reflection lies in or beyond B .

The spindle is shaded in Figure 7(a). Our plan is to confine large crust triangles between the union of spindles on each side of F as shown in Figure 7(b). (Small crust triangles lie within the fattened surface simply due to their size.) We start by proving two lemmas about spindles: they are indeed forbidden regions, and they have relatively “flat” bottoms, meaning that their width does not shrink with shrinking r .

Lemma 8. No crust triangle T whose Delaunay ball B_T has radius greater than $5rLFS(s)$ can penetrate the spindle of s .

Figure 7. (a) The Delaunay ball B_T of a triangle intersecting the spindle must contain a big patch of surface F . (b) Spindles of sample points fuse so that all triangles must lie close to F .

Proof: Assume t is a point inside B_v on a crust triangle T with Delaunay ball B_T . We first assert that B_T contains the reflection point t' . Let H be the plane containing the intersection of the boundaries of B_v and B_T . Since the vertices of T lie on B_T outside B_v , T must be contained in the closed halfspace bounded by H not containing v . It suffices to prove the lemma for the case in which t lies right on H , as the reflection of any t in the interior of the halfspace lies between H and a reflection of a point on H .

We may also assume that ball B_T passes through v , since if we replace B_T with the ball that touches v and has the same intersection with H , the part of B_T outside B_v shrinks (making things harder for our lemma).

Now consider any plane containing line vt . Balls B_v and B_T intersect this plane in circles and plane H intersects in a line containing the mutual chord of these circles. See Figure 8(a).

Assume w.l.o.g. that the cross-section of B_v is the unit circle with center $v = (0, 1)$. Let $t = (0, y_t)$. Denote the center and radius of B_T 's cross-section by (x, y) and ρ . Since t lies along the mutual chord, it has equal “power distance” to $(0, 1)$ and (x, y) :

$$(1 - y_t)^2 - 1 = x^2 + (y - y_t)^2 - \rho^2.$$

Substituting $(1 - y)^2$ for $\rho^2 - x^2$, we obtain

$$y_t^2 - 2y_t = (y - y_t)^2 - (1 - y)^2,$$

which simplifies to $y = (1 - 2y_t)/(2 - 2y_t)$. Thus the centers of all possible B_T circles lie on the same horizontal line, as shown in Figure 8(b).

Any B_T passes through the reflection of $(0, 1)$ across the horizontal line, the point $(0, (1 - 2y_t)/(1 - y_t) - 1)$. For any value of $y_t < 1$, $(1 - 2y_t)/(1 - y_t) - 1 < -y_t$, so B_T contains $t' = (0, -y_t)$.

Figure 9. The spindle curves gradually, so t must be close to B .

Draw the .20-radian ray with origin m and the .32-radian ray with origin v as shown in Figure 9. The rays intersect at a point x with coordinates about $(.259, .218)$. By computing the distances to the boundaries of B_v and B along ray vx , we can confirm that x lies inside the spindle. Thus the boundary of the spindle lies below x on the .20-radian ray with origin m . Assume t and p are at the extremal positions allowed by the lemma, so that t is on the boundary of the spindle and $|\angle omp| = .20$. The distance from x to m is less than 1.252, so $d(t, p) - \delta \leq .192 \leq |\angle omp|$. Since $d(t, p)$ increases ever more rapidly as $|\angle omp|$ increases, this inequality also applies to points t and p such that $|\angle omp| < .20$ as well. ■

We are now in a position to finish the proof of the theorem: all crust triangles lie within the fattened surface formed by placing a ball of radius $5rLFS(q)$ around each point $q \in F$.

Proof of Theorem 4: Let B_T be the Delaunay ball of the crust triangle containing point t . Let s be the sample point nearest t . If B_T has radius less than $5rLFS(s)$, then there is nothing to prove, since s itself could be the q of the theorem.

So assume B_T has radius at least $5rLFS(s)$. Let B_v , B_m^- , and B be the polar ball of s , the tangent ball of radius $LFS(s)$ on the opposite side of F , and the concentric ball with radius reduced by $rLFS(s)$ as in Figure 10. Let o and o' be the points of lune $B_m^- \cap B_v$ closest to the centers of B_m^- and B_v , respectively. Surface F could pass through the point o' , and if it did, s would necessarily be the closest sample point to o' , since B_m^- and B_v are both empty. Hence by Lemma 3(b), $d(s, o') \leq rLFS(s)/(1-r)$. Since B_v has radius at least that of B_m^- , $d(s, o) \leq d(s, o')$.

Let p and p' be the closest points to t on B and B_m^- , respectively, and let q be the point of F on line pt closest to t . Hence $d(t, q) \leq d(p, t)$. By an argument analogous to that used for o' , $d(s, p') \leq rLFS(s)/(1-r)$, and so by the triangle inequality, $d(o, p') \leq 2rLFS(s)/(1-r)$. So $\angle omp' \leq 2 \arcsin(r/(1-r))$, which for $r \leq .06$, is less than .20 radians. The set-up satisfies the hypotheses of Lemma 9, only with radii scaled by $(1-r)LFS(s)$.

By Lemma 8, t must lie between the spindle and B_m . Applying Lemma 9,

$$d(t, p) \leq rLFS(s) + |\angle omp|(1-r)LFS(s).$$

We now use the fact that $|\angle omp| \leq 2 \arcsin(r/(1-r)) \leq 3r$, to obtain

$$d(t, p) \leq rLFS(s) + 3r(1-r)LFS(s) \leq 4rLFS(s).$$

Figure 10. Crust point t must be near surface point q .

Finally, $d(s, q) \leq rLFS(s)/(1-r)$, so by Lemma 1, $LFS(q) \geq (1-2r)LFS(s)/(1-r)$, and hence $5rLFS(q) \geq d(t, p) \geq d(t, q)$. ■

Let T be a triangle of the θ -crust with $\theta = 4r$, t be a point on T , and p be the closest point to t on F . Theorem 5 states that the angle between the normal to T and the normal to F at p measures $O(\sqrt{r})$ radians.

Proof of Theorem 5: First, we establish the easier claim that at each sample point s , the normals to incident θ -crust triangles do not deviate by more than $O(r)$ radians from the normal to F . This statement follows from the fact that Step 5 of the algorithm removes each triangle around s whose normal forms an angle larger than $6r$ with the vector to the pole. By Lemma 4, the pole vector deviates from the normal to F by at most $\psi = 2 \arcsin(r/(1-r))$, so that $\psi \leq 1.3r$ for $r \leq .06$.

Now let t be any point on a θ -crust triangle T , and let p be the closest point on F to t . By Theorem 4, $d(t, p) \leq 5rLFS(p)$. Let s denote the closest vertex of T to t , C the radius of T 's circumcircle, and ρ the radius of T 's Delaunay ball B_T . If $C \leq \sqrt{r}LFS(s)$ then $d(s, p)$ is $O(\sqrt{r})$, and Theorem 5 follows from Lemma 2 and the bound on ψ .

So assume C and hence ρ is at least $\sqrt{r}LFS(s)$. Let ϕ denote the angle between the normal to F at s and the vector from s to the center v of B_T . Lemma 4 with $\nu = \sqrt{r}$ implies that $\phi \leq 2\sqrt{r}/(1-r)$ radians. Next let δ denote the angle between the normal to T at s and the vector from s to v , as shown in Figure 11. Angle $\delta \leq \phi + \psi$, where ψ , as above, is the angle between the normal to the surface at s and the normal to T . Since $\psi = O(r)$, we can conclude that $\delta \leq 2\sqrt{r}$ for small enough r .

Now $C = \rho \sin \delta$, so $\rho = C/\sin \delta \geq LFS(s)/2$. Thus the assumption that C is large (at least $\sqrt{r}LFS(s)$) shows that ρ must be very large (at least $LFS(s)/2$). We can now return to Lemma 4 with $\nu = 1/2$. This time we obtain an upper bound of $O(r)$ on ϕ and δ , and a lower bound of $\Omega(LFS(s)/\sqrt{r})$ on ρ . (Sadly, we cannot repeat this trick to inflate ρ indefinitely, since ψ remains $O(r)$.)

Notice that since δ is $O(r)$, the plane containing T cuts a small spherical cap on B_T , one subtending solid angle of only $O(r)$. This means that T itself is small with respect to B_T ; the point $t \in T$ can be at most $O(r\rho)$ from a vertex s , bounding (by Lemma 1)

Figure 11. Repeated use of Lemma 4 shows that if triangle T is large B_T must be enormous.

$LFS(t) \leq O(r\rho) + LFS(s)$, which is $O(\sqrt{r}\rho)$. And since t is within $5rLFS(p)$ of p , $LFS(p)$ is $O(\sqrt{r}\rho)$ as well.

Now assume that the normal to F at p deviates from the normal to T by $\Omega(\sqrt{r})$, and consider the big tangent balls of radius $LFS(p)$ at p . The point p is close – within $O(rLFS(p))$ – to the surface of B_T , while the radius of B_T is much larger – $\rho = O(LFS(p)/\sqrt{r})$ – than the radius of the big tangent balls at p . For some small enough value of r , the big tangent balls intersect B_T in circular patches of radius $\Omega(\sqrt{r})LFS(p)$. As in the proof of Lemma 8, F is confined between these two balls, so there must be a similar-size patch of F inside B_T , and hence empty of sample points, which gives a contradiction to S being an r -sample. This contradiction establishes Theorem 5. ■

Finally, Theorem 6 states that for sufficiently small r , the trimmed θ -crust is homeomorphic to F .

Proof of Theorem 6: We first prove that the (untrimmed) θ -crust still contains all the good triangles. Since Theorem 3 shows that the raw crust contains all the good triangles, we only need to show that each good triangle passes the filtering-by-normal step. Let T be a good triangle and s its vertex of maximum angle. By Lemma 6(a), the angle between the normal to T and the normal to F at s measures at most $\arcsin(\sqrt{3}r/(1-r))$ radians. By Lemma 4, the angle between the pole vector at s and the normal to F at s measures at most $2\arcsin(r/(1-r))$. Combining these two bounds, the angle between the normal to T and either pole vector at s must be less than $4r = \theta$. Similarly, Lemmas 6(b) and 4 combine to show that the angle between the normal to T and the pole angle at any other vertex of T is at most $2\arcsin(r/(1-r)) + 2r/(1-7r) + \arcsin(\sqrt{3}r/(1-r))$ radians, which, for small enough r , is less than $6r = 3\theta/2$.

We must now show that the trimming operation (Step 6) produces a set of triangles with the same topology as the good triangles. Let s be a sample point, and assume the normal to F at s is vertical. Step 5 ensures that for $r \leq .06$, all triangles around s remaining after Step 5 have normals within $.5$ radians of vertical. By Lemma 4, the vector from s to one of its poles is within $.2$ radians of vertical. Since $.5 + .2 < \pi/2$, the vertex-to-triangle breadth-first-search in Step 6 orients triangles consistently: the orientations do not depend on the actual search order, and at each vertex they agree with an orientation of F .

Figure 12. A reconstructed minimal surface along with the poles of sample points. The crust contains exactly the original triangles. (Sample points courtesy of Hugues Hoppe)

After all triangles with sharp edges have been removed, all walks along the remaining set of triangles, that do not pierce a triangle, must run along either only inside or only outside sides of triangles. Good triangles cannot have sharp edges, since the dihedral between adjacent good triangles is less than $\pi/2$, and hence are never removed.

Consider the mapping that takes each point of space to its closest point on F . We claim that the restriction of this mapping to the trimmed θ -crust is a homeomorphism. Since the good triangles survived up until the final breadth-first-search, the trimmed θ -crust contains a set of triangles homeomorphic to F and at least one point of the trimmed θ -crust is mapped to each point of F . By Theorem 5 each triangle is nearly parallel to F , so the map is one-to-one on each triangle. And because the triangles are consistently oriented, points on two different triangles cannot map to the same point on F . ■

6 Implementation and Examples

Manolis Kamvysselis, an undergraduate from MIT, implemented steps 1–4 of the crust algorithm during a summer at Xerox PARC. We used Clarkson’s *Hull* program [11] for Delaunay triangulation, and *Geomview* [23] to visualize and print the results. We used vertices from pre-existing polyhedral models as inputs, in order to compare our results with “ground truth”. A companion paper [2] reports on our experimental findings.

The only tricky part of the implementation was the handling of degeneracies and near degeneracies. Our test examples, many of which started from approximately gridded sample points, included numerous quadruples of points supporting slivers. Kamvysselis incorporated an explicit tolerance parameter ϵ ; the circumcenter of quadruples within ϵ of cocircu-

Figure 13. The raw crust contains some extra triangles linking the sausages; this defect is corrected by step 5. (Sample points courtesy of Paul Heckbert)

larity was computed by simply computing the circumcenter of a subset of three. This “hack” did not affect the overall algorithm, as these centers were never poles. Running time was only a little more than the time for two three-dimensional Delaunay triangulations. Notice that the Delaunay triangulation in step 3 involves at most three times the original number of vertices.

Figure 14. (a) The pig sample set contains 3511 points. (b) A close-up of the front feet shows an effect of undersampling. (Sample points courtesy of Tim Baker)

Figure 12 shows an especially advantageous example for our algorithm, a well-spaced point set on a smooth surface. Even though our algorithm is not designed for surfaces with boundary, it achieves perfect reconstruction on this example. Of course, the trimming step should not be used in reconstructing a surface with boundary.

Figure 13 shows an effect of undersampling. (We say we have *undersampled* if the

sample set is not an r -sample for a sufficiently small r .) In this example, the raw crust contains all the good triangles, along with some extra triangles. The extra triangles turn separated sausages into link sausages, and as mentioned above roughen the inside surfaces of the sausages. Both of these defects are corrected by step 5, filtering by normals. Figure 14 shows another effect of undersampling: missing triangles around the chest and hooves. Some sample points are not “opposed” by samples on the other side of these roughly cylindrical surfaces; hence Voronoi cells extend too far and poles filter out some good triangles. An r -sample for a sufficiently small r would be very dense near the hooves, which include some rather sharp corners.

7 Conclusions and Future Work

In this paper we have given an algorithm for reconstructing an interpolating surface from sample points in three dimensions. The algorithm is simple enough to analyze, easy enough to implement, and practical enough for actual use.

Our previous paper [1] gave two provably good algorithms for reconstructing curves in two dimensions, one using Voronoi filtering as in this paper, and the other using the β -skeleton. It is interesting to ask whether the β -skeleton can be generalized to the problem of surface reconstruction. (We know that the most straightforward generalization of the β -skeleton does not work.)

Another interesting question concerns the generalization of Voronoi filtering to higher dimensions. *Manifold learning* is the problem of reconstructing a smooth k -dimensional manifold embedded in $\mathbb{R}^d$. This problem arises in modeling unknown dynamical systems from experimental observations [10]. Even if Voronoi filtering can be generalized to this problem, its running time for the important case in which $k \ll d$ would not be competitive with algorithms that compute triangulations only in k -dimensional subspaces [10], rather than in $\mathbb{R}^d$.

Along with the two theoretical open questions outlined above, there are several quite practical directions for further research on our algorithms. What is the empirical maximum value of r for which our algorithm gives reliable results? We believe that the value of $r \leq .06$ in Theorem 4 is much smaller than necessary. Is the crust useful in simplification and compression of polyhedra? Can the crust be extended to inputs with creases and corners, such as machine parts? Can the crust be modified for the problem of reconstruction from cross-sections, in which the input is more structured than scattered points?

Acknowledgements

We would like to thank Alan Cline, Bob Connelly, Tamal Dey, Herbert Edelsbrunner, David Eppstein, David Goldberg, and Manolis Kamvysselis for helpful conversations, and Ken Clarkson and The Geometry Center for making their software available.

References

- [1] N. Amenta, M. Bern, and D. Eppstein. The crust and the β -skeleton: combinatorial curve reconstruction. To appear in *Graphical Models and Image Processing*.
- [2] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. To appear in *Siggraph* 1998.
- [3] D. Attali. r -Regular shape reconstruction from unorganized points. In *Proc. 13th ACM Symp. Computational Geometry*, 1997, 248–253.
- [4] G. Barequet. Piecewise-linear interpolation between polygonal slices. In *Proc. 10th ACM Symp. Computational Geometry*, 1994, 93–102.
- [5] C. Bajaj, F. Bernardini, and G. Xu. Automatic reconstruction of surfaces and scalar fields from 3D scans. *Proc. SIGGRAPH '95*, 1995, 109–118.
- [6] F. Bernardini and C. Bajaj. Sampling and reconstructing manifolds using α -shapes, *9th Canadian Conference on Computational Geometry*, 1997, 193–198.
- [7] F. Bernardini, C. Bajaj, J. Chen and D. Schikore. Automatic reconstruction of 3D CAD models from digital scans. Technical report CSD-97-012, Purdue University (1997).
- [8] F. Bernardini, C. Bajaj, J. Chen, D. Schikore. A triangulation-based object reconstruction method. *Proc. 13th ACM Symp. Computational Geometry*, 1997, 481–484.
- [9] J-D. Boissonnat. Geometric structures for three-dimensional shape reconstruction. *ACM Trans. Graphics* 3 (1984) 266–286.
- [10] C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. *Proc. 5th International Conf. on Computer Vision*, 1995, 494–499.
- [11] K. Clarkson. *Hull*: a program for convex hulls. <http://cm.bell-labs.com/netlib/voronoi/hull.html>.
- [12] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proc. SIGGRAPH '96*, 1996, 303–312.
- [13] H. Edelsbrunner. Surface reconstruction by wrapping finite sets in space. Tech. Rept. 96-001, Raindrop Geomagic, Inc., 1996.
- [14] H. Edelsbrunner, D.G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. on Information Theory* 29 (1983), 551–559.
- [15] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* 13 (1994) 43–72.
- [16] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc. 10th ACM Symp. Computational Geometry*, 1994, 285–292.
- [17] L. H. de Figueiredo and J. de Miranda Gomes. Computational morphology of curves. *Visual Computer* 11 (1995) 105–112.
- [18] J. Goldak, X. Yu, A. Knight, and L. Dong. Constructing discrete medial axis of 3-D objects. *Int. J. Computational Geometry and its Applications* 1 (1991) 327–339.

- [19] H. Hoppe. *Surface Reconstruction from Unorganized Points*. Ph.D. Thesis, Computer Science and Engineering, U. of Washington, 1994. <http://www.research.microsoft.com/research/graphics/hoppe/thesis/thesis.html>
- [20] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proc. SIGGRAPH '92*, 1992, 71–78.
- [21] H. Hoppe, T. DeRose, T. Duchamp, H. Jin, J. McDonald, and W. Stuetzle. Piecewise smooth surface reconstruction. *Proc. SIGGRAPH '94*, 1994, 19–26.
- [22] D. G. Kirkpatrick, J. D. and Radke. A framework for computational morphology. *Computational Geometry*, G. Toussaint, ed., Elsevier, pp. 217-248.
- [23] S. Levy, T. Munzner, and M. Phillips. Geomview. <http://www.geom.umn.edu/software/download/geomview.html>
- [24] S. Lodha. Scattered Data Techniques for Surfaces. To appear in *Geometry Detection, Estimation and Synthesis for Scientific Visualization*, Academic Press.
- [25] S. Mann, C. Loop, M. Lounsbery, D. Meyers, J. Painter, T. DeRose, and K. Sloan. A survey of parametric scattered data fitting using triangular interpolants. *Curve and Surface Design*, H. Hagen, ed., SIAM, 1992, 145–172.
- [26] R. C. Veltkamp. *Closed object boundaries from scattered points*. LNCS Vol. 885, Springer, 1994.