Surface Reconstruction by Voronoi Filtering

Nina Amenta* Marshall Bern†

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 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 has been mostly heuristic. Only recently have researchers started publishing algorithms for the two-dimensional problem with provable properties.

Hoppe et al. [20, 21, 22] brought the reconstruction problem to the attention of the computer graphics community. Their algorithm computes an approximating surface—not interpolating but close. The 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 similar algorithm by Curless and Levoy [13] is tuned for data samples collected by a laser range scanner, but could be applied to the general

*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

SCG 98 Minneapolis Minnesota USA

reconstruction problem. Their 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 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. [15, 16]. 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. [8] follow α-shape-based reconstruction with a clean-up phase to resolve sharp dihedral angles. Edelsbrunner and Raindrop Geomagic [14] are continuing to develop α-shape-based reconstruction along with proprietary extensions.

An early algorithm due to Boissonnat [10] 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 [19] 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 [18] prove that the Euclidean minimum spanning tree can be used to reconstruct uniformly sampled curves in the plane. Bernardini and Bajaj [7] 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 graph in which edges are defined by empty regions between vertices. Our previous paper showed that both the crust and the β-skeleton [23] (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).

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) Pick the vertex p^- of $Vor(s)$ farthest from s with negative projection onto sp^+.
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 2.2θ for the other vertices of T).
6. (Trimming) Orient triangles and poles (inside and outside) consistently, and extract a piecewise-linear manifold without sharp dihedral angles.

Figure 1. The surface reconstruction algorithm.

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 vertices of the Voronoi cell $Vor(s)$ farthest from s, one on each side of the surface. Since the algorithm does not know the surface, only the sample points, it chooses the poles by the method given in step 2 in Figure 1. Lemma 2, parts (a) and (c), in Section 5 imply that that this method is correct for well-sampled surfaces. Voronoi cells of sample points on the convex hull of S are unbounded in the outwards direction; for such a point the outwards pole simply represents a direction in which the cell is unbounded. Such a pole is used only to help find the sample point's opposite pole and is not included in the eventual Delaunay triangulation. 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 2.2θ. Angles are unsigned angles in the range $[0, \pi/2]$. As stated in Theorem 4, 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 [16].

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.

Define a *sharp* edge to be an edge for which the cyclic order of triangle sides alternates between outside and inside. (As we show in the proof of Theorem 4, a sharp edge indeed has all triangles within a small wedge. Notice that an edge bounding only a single triangle is necessarily a sharp edge.) Step 6 trims off pockets by greedily removing triangles with *sharp* edges. Now the remaining triangles form a "quilted" surface, in which each edge borders 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 start with some definitions. Figure 2 gives an example of the medial axis in $\mathbb{R}^2$; we allow F to have more than one connected component.

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.*

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.

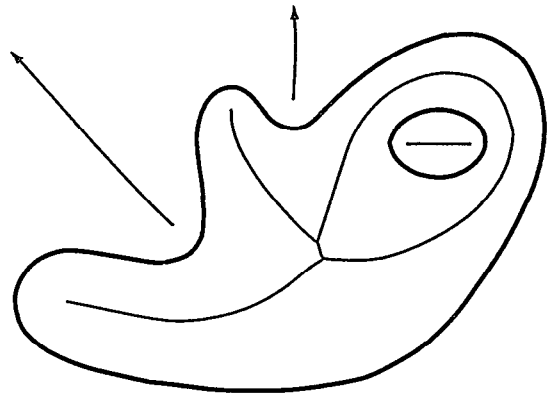

Figure 2. The medial axis of a smooth curve.

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.*

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 theorem states the theoretical guarantees for the three-dimensional (raw) crust.

Theorem 3. *(a) If S is an r-sample for $r \leq .1$, then the crust includes all the good triangles. (b) 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. This step is indeed necessary for this guarantee, as the raw crust sometimes includes small skinny triangles with deviant surface normals. For example, the insides of the sausages shown on the left in Figure 11 have a sort of "washboard" texture due to small deviant triangles lining the inside curves. Finally, Step 6 adds the guarantee of topological equivalence.

Theorem 4. *Assume S is an r-sample and set $\theta = 3r$. (a) 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 $O(r)$ radians. (b) For sufficiently small r, the trimmed θ-crust is homeomorphic to F.*

5 Proofs

In this section we sketch proofs of the theoretical guarantees. We start with some terminology. 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 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: Lipschitz conditions for the $LFS(p)$ function and for the direction of surface normals (a function from F to the two-dimensional sphere). We use $d(p,q)$ to denote the Euclidean distance from p to q. Angles are measured in radians.

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

Proof: For part (a), $LFS(p) \geq LFS(q) - d(p,q)$, 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, $LFS(p) - d(p,q) \leq LFS(q)$.

For part (b), let us parametrize 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 $d(p,q) \max_t |n'(t)|$.

There are tangent balls of radius $LFS(f(t))$ on either side of F at $f(t)$, so $|n'(t)|$ is at most the rate at which the normal turns on one of these tangent balls. Referring to Figure 3, we see that

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

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

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

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, with long direction nearly normal to the surface F. Together, Lemma 2(a) and (c) below show that the vector from a sample point to a pole gives 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.

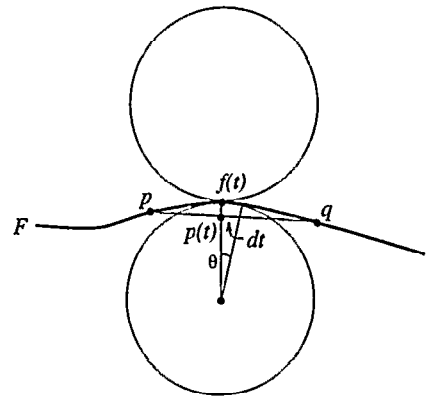

Figure 3. Bounding $|n'(t)|$ in terms of the radius $LFS(f(t))$ and $d(f(t),p(t))$.

Lemma 2. *(a) On either side of F, the distance from s to its pole is at least $LFS(s)$. (b) The intersection of $Vor(s)$ and F is contained in a ball of radius $\frac{r}{1-r}LFS(s)$. (c) Let v be a vertex of $Vor(s)$ such that $d(v,s) \geq LFS(s)$. The angle at s between the vector to v and the normal to the surface is at most $2\arcsin\frac{r}{2-2r}$.*

Proof: (a) On either side of F at s, the tangent ball of radius $LFS(s)$ is empty, since it is contained in a medial ball at s. The center c of such a ball lies within $Vor(s)$. So either $Vor(s)$ has a vertex at distance at least $LFS(s)$, or the pole on that side lies at infinity.

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

(c) 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 and B_v each have radius at least $LFS(s)$, $d(q,s) \geq 2\sin(\phi/2)LFS(s)$. Since S is an r-sample, $d(q,s)$ must be less than $\frac{r}{1-r}LFS(s)$. Combining the last two inequalities, $2\sin(\phi/2) \leq \frac{r}{1-r}$, or $\phi \leq 2\arcsin\frac{r}{2-2r}$. ∎

Lemma 3. *Let T be a good triangle and s its vertex with largest angle. (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 $LFS(s)$

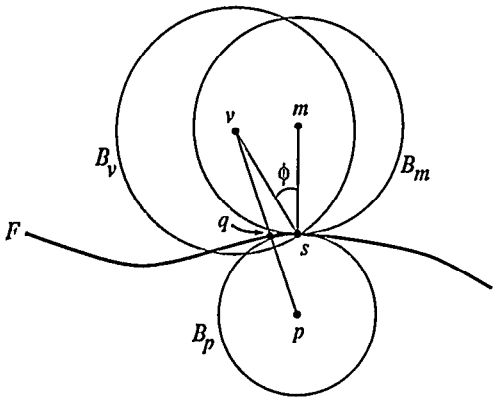

Figure 4. The vector from s to a distant Voronoi vertex such as a pole must be nearly normal to the surface. B_p, B_m and B_v all have radius at least $LFS(s)$.

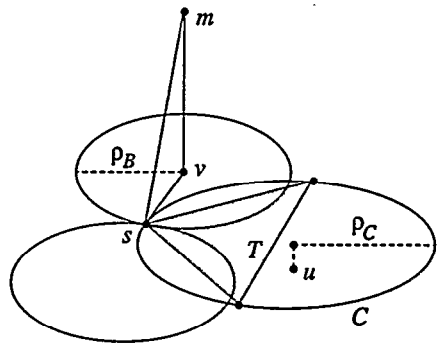

Figure 5. Bounding the angle between the normal to the triangle and the normal to the surface at s.

tangent to F at s on either side of F. These balls intersect the plane of T in two disks of the same radius, which we shall denote ρ_B. Since the balls are empty of sample points, the disks cannot contain the other two vertices of T. The other two vertices are at most distance $\sqrt{3}\rho_C$ away from s, which in turn implies that $\rho_B \leq \sqrt{3}\rho_C$.

We can rewrite 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 2(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 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 1(b), the angle between the normals to F at s and s' is at most $2r/(1-7r)$. ∎

We are now ready to sketch a proof of Theorem 2: the good triangles form a polyhedron homeomorphic to F. The proof uses a result of Edelsbrunner and Shah [17].

Proof of Theorem 2: It suffices [17] 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 2(b) shows that $Vor(s) \cap F$ is small, fitting inside a ball B of radius $\frac{r}{1-r} LFS(s)$. Lemma 1(b) then shows that F is nearly horizontal throughout $Vor(s) \cap F$, more precisely, the normal to F is nowhere farther than $r/(1-3r) \leq .15$ radians from vertical.

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 3(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 the perpendicular bisector of s and another sample point s'. If f intersects F, then some side of f must pierce F, and since such an edge can form an angle no greater than .9 radians with vertical, f itself lies within .9 radians of vertical. Now consider a single connected component C of $f \cap F$. We use the curved segment C to divide $\mathbb{R}^3$ into two pieces. Let H be the set of points p in $\mathbb{R}^3 \setminus C$ 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 thus a union of wedges with vertices on C. We assert that all points of $F \cap B_s$ lie either in C or in H. If F crossed the boundary of H other than at C, then there must be a point of F with normal more than .2 radians from the vertical. Similarly, we assert that all points of f lie in either C or $\mathbb{R}^3 \setminus H$. 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. Any point within this strip can be connected to C by a line within .9 radians of vertical. Altogether, we can conclude that C is the only connected component of $f \cap F$, and of course is a topological 1-ball.

Finally we consider $Vor(s)$ itself, the case $k = 3$. Consider any connected component C of the intersection of F and the Voronoi cell. We mimic the argument from the case $k = 2$, again dividing $\mathbb{R}^3$ into two pieces using the angle formed by a shortest line segment to C. Again H will be a union of wedges with vertices on C. Except at C itself, $F \cap B_s$ must lie inside the wedges (closer to horizontal), whereas $Vor(s)$ must lie outside the wedges. ∎

Next we give a proof of Theorem 3(a): the 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.

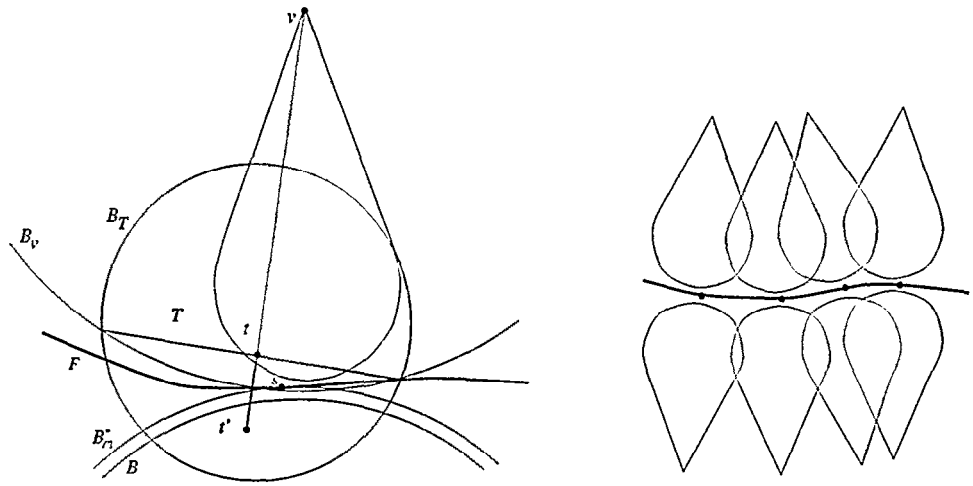

Figure 6. (a) The Delaunay ball B_T of a triangle intersecting the spindle must contain a big patch of surface F. (b) Spindles of sample points fuse so that all triangles must lie close to F.

Proof of Theorem 3(a): 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 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 the polar ball B_v must be contained in B'_u, and that B_v must contain a point of the medial axis, thereby giving a contradiction. Let B_m be the medial ball with center m, tangent to F at s on the same side of F as v; this ball has radius at least $LFS(s)$. By Lemma 2(c), $\angle msv$ measures at most $2\arcsin(r/2)$, which is less than .12 for $r \leq .1$. An easy 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 $2rLFS(u)$ 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 3(b). 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 ball of radius $LFS(s)$ tangent to F at s, on the same side of F as v, and let B_m^- be the tangent ball of radius $LFS(s)$ on the opposite side of F from v. Surface F must lie between these two balls, since these balls are contained in medial balls at s. Let B be the ball concentric with B_m^- with radius $(1-r)LFS(s)$, as shown in Figure 6(a). Notice that Lemma 2(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 $\{ t \in B_v \mid$ segment tt' intersects $B \}$, that is, all points in B_v whose reflection lies in or beyond B.*

The spindle is shaded in Figure 6(a). Our plan is to confine large crust triangles between the union of spindles on each side of F as shown in Figure 6(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 r.

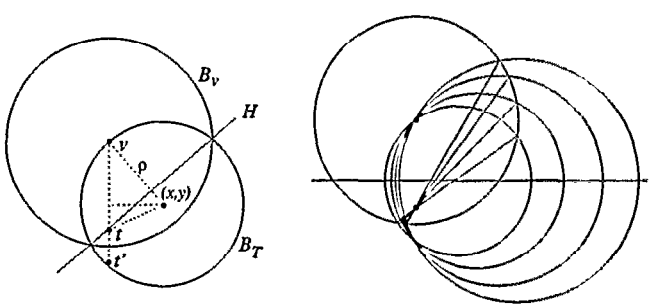

Figure 7. (a) B_T must contain reflection point t'. (b) The family of possible B_T circles.

Lemma 4. *Assume $r \leq .06$. No crust triangle T whose Delaunay ball B_T has radius greater than $5rLFS(s)$ can penetrate the spindle of s.*

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. We may assume that t lies on H, since this is the worst case for our assertion.

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, as shown in Figure 7(a). We may assume that ball B_T passes through v, for this situation is again the worst case.

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 7(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 $t' = (0,-y_t)$.

Thus if the original point t lies within the spindle of s, then B_T must intersect B. Since F is confined between B_m^- and B_m^+, this statement implies that B_T intersects F. In fact, a short calculation using the assumption that the radius of B_T is greater than $5rLFS(s)$ reveals that B_T must intersect F in a patch large enough that it must contain a sample point in its interior, a contradiction to B_T being a Delaunay ball. ∎

The next lemma shows that spindles have flat bottoms. In this lemma we assume that B and B_v have equal radius. It is not hard to confirm that this assumption is worst case: a larger B_v just gives a larger, flatter spindle.

Lemma 5. *Assume that B and B_v are unit balls, and that the distance between them is at most $\delta \le .06$. Let t be a point outside B and outside the spindle induced by B in B_v. Let p be the closest point on B to t. If $|\angle omp|$, the measure of $\angle omp$ in radians, is less than .20, then $d(t,p) \le \delta + |\angle omp|$.*

Proof: Assume v has coordinates $(0,1)$. The worst case for the lemma occurs when δ assumes its maximum value, as larger δ means a higher and narrower spindle, thereby maximizing $d(t,p)$ relative to $\delta + |\angle omp|$. So assume m has coordinates $(0,-1.06)$.

Draw the .20-radian ray with origin m and the .32-radian ray with origin v as shown in Figure 8. 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 as stated in the lemma, and $|\angle omp| = .20$. The distance from x to m is less than 1.252, so $d(t,p) - \delta \le .192 \le |\angle omp|$. 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 3(b): 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 opposite medial axis ball, and the concentric ball with radius reduced by $rLFS(s)$ as in Figure 9. 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

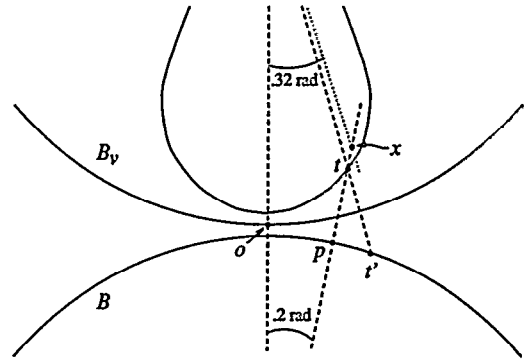

Figure 8. The spindle curves gradually, so t must be close to B.

be the closest sample point to o', since B_m^- and B_v are both empty. Hence by Lemma 2(b), $d(s,o') \le rLFS(s)/(1-r)$. Since B_v has radius at least that of B_m^-, $d(s,o) \le d(s,o')$.

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 $d(t,q) \le d(p,t)$. By an argument analogous to that used for o', $d(s,p') \le rLFS(s)/(1-r)$, and so by the triangle inequality, $d(o,p') \le 2rLFS(s)/(1-r)$. So $\angle omp' \le 2\arcsin(r/(1-r))$, which for $r \le .06$, is less than .20 radians. The set-up satisfies the hypotheses of Lemma 5, only with radii scaled by $(1-r)LFS(s)$.

By Lemma 4, t must lie between the spindle and B_m. Applying Lemma 5,

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

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

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

Finally, $d(s,q) \le rLFS(s)/(1-r)$, so

$$LFS(q) \ge (1-2r)LFS(s)/(1-r),$$

and hence $5rLFS(q) \ge d(t,p) \ge d(t,q)$. ∎

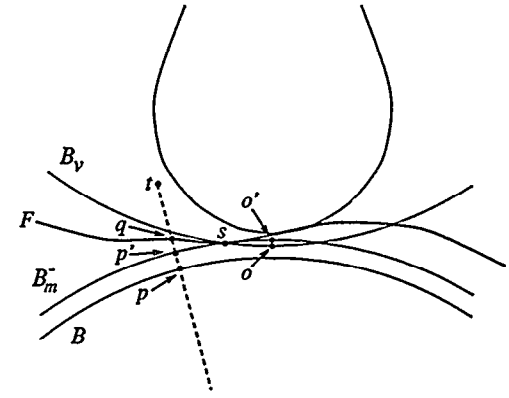

Figure 9. Crust point t must be near surface point q.

Let T be a triangle of the θ-crust, with $\theta = 3r$, and let t be a point on T. Theorem 4, part (a), states that for the angle between the normal to T and the normal to F at the point $p \in F$ closest to t measures $O(r)$ radians. Part (b) states that for sufficiently small r, the trimmed θ-crust is homeomorphic to F.

Proof of Theorem 4: We first prove that the (untrimmed) θ-crust contains all the good triangles, so that we do not lose the guarantees of the raw crust. Theorem 3(a) shows that the crust contains all the good triangles, so we need only 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 3(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 2(c), the angle between the pole vector at s and the normal to F at s measures at most $2\arcsin(r/2)$. Combining these two bounds, the angle between the normal to T and either pole vector at s must be less than $3r = \theta$. Similarly, Lemmas 3(b) and 2(c) combine to show that the angle between the normal to T and the pole angle at any other vertex of T is at most $2\arcsin(r/2) + 2r/(1-7r) + \arcsin(\sqrt{3}r/(1-r))$ radians, which is less than $6.6r = 2.2\theta$.

We next prove 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 $6.6r$ with the vector to the pole. By Lemma 2(c), the pole vector deviates from the normal to F by at most $2\arcsin(r/2)$, which is less than $1.1r$ 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 3(b), $d(t,p) \leq 5rLFS(p)$. We next prove that the normal to T does not deviate by more than $O(r)$ radians from the normal to F at p. If t is less than distance $LFS(s)/4$ from a vertex s of T, then combining the bounds from Lemmas 1(b) and 3(b) establishes the result. If t is not close to any vertex of T, then we use an argument related to our spindle argument above. Let B_T be the Delaunay ball of T, and s and s' be two vertices of T at least $LFS(s)/3$ apart. Since B_T is fairly large with respect to both $LFS(s)$ and $LFS(s')$, the vectors from s and s' to the center v of B_T deviate only slightly (linearly in r) from the normals to F at s and s'. Since these surface normals in turn deviate only slightly (again linearly in r) from the normal to T, as r shrinks the vectors from s and s' to v become nearly parallel, By picking r sufficiently small, we force the boundary of B_T to be very close to t and parallel to T. Now consider the medial balls, which have radius at least $LFS(p)$, on either side of F at p. As in the proof of Lemma 4, F is confined between these medial balls. If the normal to F at p were not nearly parallel to the normal to T, then F would be forced to penetrate B_T in a large patch (a constant times $LFS(p)$). F cannot avoid B_T since its curvature is bounded by Lemma 1(b). Thus we can conclude that Theorem 4(a) holds.

For part (b), we must 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 2(c), the vector from s to one of its poles is within .2 radians of vertical. Since the sum of .5 and .2 is bounded below $\pi/2$, 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.

Sharp edges are exactly those edges at which a walk on the triangles can cross from an outside side to an inside side without piercing a triangle. After all triangles with sharp edges have been removed, all walks along the remaining triangles must run along either only inside or only outside sides.

Consider the mapping that takes each point of space to its closest point on F. We claim that this mapping is a homeomorphism. By Theorem 4(a) each triangle is nearly parallel to F, so the map is one-to-one on each triangle. Because the triangles are consistently oriented, points on two different triangles cannot map to the same point on F. ∎

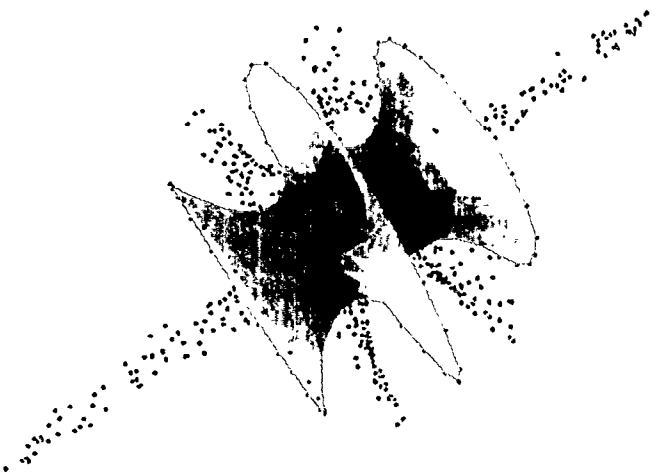

Figure 10. A reconstructed minimal surface along with the poles of sample points. The crust contains exactly the original triangles. (Sample points courtesy of Hugues Hoppe)

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 [12] for Delaunay triangulation, and *Geomview* [24] 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 cocircularity 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 10 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.

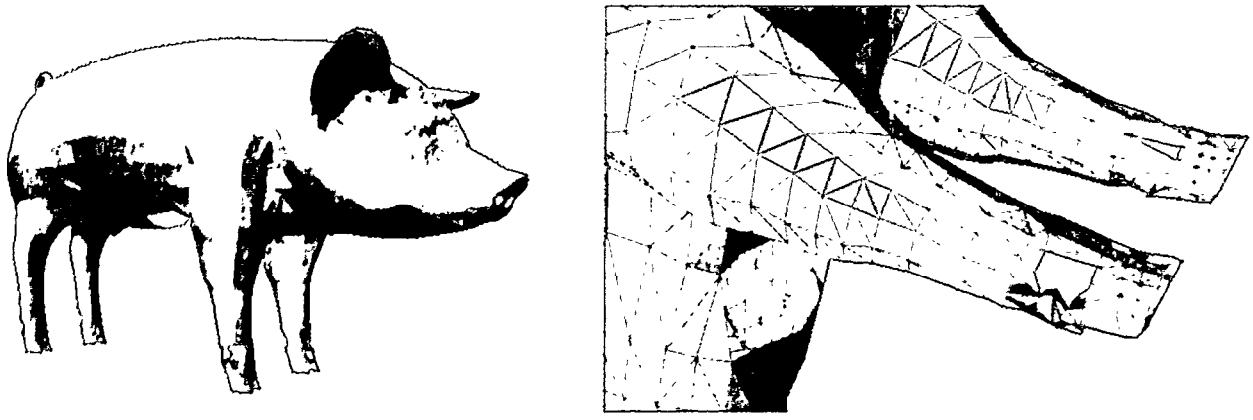

Figure 12. (a) The pig sample set contains 3511 points. (b) A close-up of the front feet shows an effect of undersampling. (Sample points courtesy of Tim Baker)

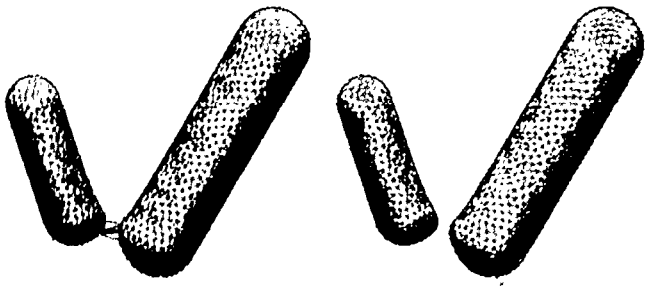

Figure 11. The raw crust contains some extra triangles linking the sausages; this defect is corrected by step 5. (Sample points courtesy of Paul Heckbert)

Figure 11 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 12 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 [11]. 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 [11], 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 3 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, David Goldberg, David Eppstein, and Manolis Kamvysselis for helpful suggestions, 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] M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T.-S. Tan. Edge insertion for optimal triangulations. *Disc. and Comp. Geometry* 10 (1992) 47–65.

[7] F. Bernardini and C. Bajaj. Sampling and reconstructing manifolds using α-shapes, *9th Canadian Conference on Computational Geometry*, 1997, 193–198.

[8] 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).

[9] F. Bernardini, C. Bajaj, J. Chen, D. Schikore. A triangulation-based object reconstruction method. *Proc. 13th ACM Symp. Computational Geometry*, 1997, 481–484.

[10] J-D. Boissonnat. Geometric structures for three-dimensional shape reconstruction. *ACM Trans. Graphics* 3 (1984) 266–286.

[11] C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. *Proc. 5th International Conf. on Computer Vision*, 1995, 494–499.

[12] K. Clarkson. *Hull*: a program for convex hulls. http://cm.bell-labs.com/netlib/voronoi/hull.html.

[13] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proc. SIGGRAPH '96*, 1996, 303–312.

[14] H. Edelsbrunner. Surface reconstruction by wrapping finite sets in space. Tech. Rept. 96-001, Raindrop Geomagic, Inc., 1996.

[15] 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.

[16] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* 13 (1994) 43–72.

[17] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc. 10th ACM Symp. Computational Geometry*, 1994, 285–292.

[18] L. H. de Figueiredo and J. de Miranda Gomes. Computational morphology of curves. *Visual Computer* 11 (1995) 105–112.

[19] 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.

[20] 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

[21] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proc. SIGGRAPH '92*, 1992, 71–78.

[22] H. Hoppe, T. DeRose, T. Duchamp, H. Jin, J. McDonald, and W. Stuetzle. Piecewise smooth surface reconstruction. *Proc. SIGGRAPH '94*, 1994, 19–26.

[23] D. G. Kirkpatrick, J. D. and Radke. A framework for computational morphology. *Computational Geometry* (G. Toussaint, ed.), Elsevier, pp. 217-248.

[24] S. Levy, T. Munzner, and M. Phillips. Geomview. http://www.geom.umn.edu/software/download/geomview.html

[25] S. Lodha. Scattered Data Techniques for Surfaces. To appear in *Geometry Detection, Estimation and Synthesis for Scientific Visualization*, Academic Press.

[26] 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.

[27] R. C. Veltkamp. *Closed object boundaries from scattered points*. LNCS Vol. 885, Springer, 1994.