A Simple Algorithm for Homeomorphic Surface Reconstruction

Nina Amenta [*] Sunghee Choi [†] Tamal K. Dey [‡] Naveen Leekha [§]

Abstract

The problem of computing a piecewise linear approximation to a surface from a set of sample points is important in solid modeling, computer graphics and computer vision. A recent algorithm [1] using the Voronoi diagram of the sample points gave a guarantee on the distance of the output surface from the original sampled surface assuming the sample was 'sufficiently dense'. We give a similar algorithm, simplifying the computation and the proof of the geometric guarantee. In addition, we guarantee that our output surface is homeomorphic to the original surface; to our knowledge this is the first such topological guarantee for this problem.

1 Introduction

A number of applications in CAD, computer graphics, computer vision and mathematical modeling involve the computation of a piecewise linear approximation to a surface from a set of sample points. The point set can be generated by a laser range scanner, manually with a contact probe digitizer, using medical images like CT or MRI scans, or in any other way. In a specific application, the input may contain additional information such as estimated surface normals, which can be quite useful; for example, see [8]. But, as observed by Hoppe et al. [16], solutions to the general surface reconstruction problem can provide a baseline for solving and analyzing specialized problems. The two dimensional version of the problem, namely curve reconstruction in the plane, has received a lot of recent attention. Several algorithms [2, 4, 5, 9, 10, 14, 15, 18] with various theoretical guarantees have been proposed.

The three-dimensional problem has been addressed by researchers in computer graphics and computer vision. Hoppe et. al [16] presented an algorithm in which the surface is represented by the *zero set* of a signed distance function. Curless and Levoy [8] give a very effective algorithm using the same basic idea, carefully using estimates of error and surface normals in laser range data. Edelsbrunner reports success with a proprietary commercial program [11]. The α-shapes algorithm, as described in [12] by Edelsbrunner and Mücke is useful for reconstructing surfaces from uniform sample sets. There is no published analysis showing a relationship between the original object surface and the output of any of these algorithms.

*Dept. of Computer Science, U. of Texas, Austin TX 78712. e-mail: amenta@cs.utexas.edu, supported by NSF grant CCR-9731977

†Dept. of Computer Science, U. of Texas, Austin, TX 78712. e-mail: sunghee@cs.utexas.edu, supported by NSF grant CCR-9731977

‡Dept. of Computer and Information Science, Ohio State U., Columbus, OH 43210. e-mail: tamaldey@cis.ohio-state.edu

§Dept. of Computer and Information Science, Ohio State U., Columbus, OH 43210. e-mail: leekha@cis.ohio-state.edu

Clearly, it is not possible to compute a surface that is faithful to the topology and geometry of the original unless the sampling is sufficiently dense, so any such analysis must include some assumption about the sampling density. Amenta and Bern [1] assumed that the distance between samples is proportional to the distance to the *medial axis*, and presented a surface reconstruction algorithm based on Voronoi diagrams. They proved that the output of their algorithm, the *crust*, is close to the surface S, under the assumption that S is a smooth (twice-differentiable) 2-manifold without boundary and that the sampling meets their assumption. Their algorithm uses two passes of Voronoi diagram computation, and also two postprocessing steps, called *normal filtering* and *trimming*. In this paper, we give a simpler, single-pass Voronoi-based algorithm, and eliminate the normal filtering step. Amenta and Bern [1] did not prove that the crust is homeomorphic to S. In this paper, we present the first such proof.

Our algorithm is based on the following structural theorem. Let T be a set of triangles satisfying three conditions, namely, I. T contains all triangles whose dual Voronoi edges intersect S, II. each triangle in T is small, that is, the radius of its circumcircle is much smaller than the distance to the medial axis at its vertices, and III. all triangles in T are "flat", that is, the triangle normals make small angles with the surface normals at their vertices. Assuming again that S is smooth and the sampling is sufficiently dense, condition I ensures that T contains a piecewise-linear manifold homeomorphic to S. Using conditions II and III we show that *any* piecewise-linear 2-manifold N extracted from T which spans all the sample points and for which every adjacent pair of triangles meets at an obtuse angle must be homeomorphic to S.

We compute T by filtering triangles from the Delaunay triangulation as follows. Let p be a sample point and let e be a Voronoi edge in the Voronoi cell of p. We can estimate the surface normal at p by the vector from p to its *pole*, as in [1]. We then determine if e has a point x, where px makes an angle close to $\pi/2$ with an estimated normal at the sample point p. If this condition is satisfied for all three Voronoi cells adjacent to e, its dual is included in the candidate set T. We prove that T satisfies conditions I, II and III and an acceptable piecewise-linear manifold N can be selected from T by the *manifold extraction* step of [1]. Thus, in theory, our algorithm produces a piecewise-linear manifold. Unfortunately in the usual case in which the required sampling assumptions are *not* met, this manifold extraction step is not particularly robust. Instead of implementing the manifold extraction step as described in [1], we implement a heuristic which seems to give reasonable results in practice.

After some definitions and preliminaries in Section 2, we will describe the algorithm in detail in Section 3; our implementation is described later in Section 7. We prove in Section 4 that T satisfies conditions I, II and III, and in Section 5 we derive some additional geometric consequences of these conditions. In Section 6 we show a homeomorphism between the output and the original surface. Sections 5 and 6 will therefore establish that the output of our algorithm is both geometrically and topologically correct.

2 Definitions and Preliminaries

We assume that surface S is a smooth manifold without boundary, embedded in $\mathbf{R}^3$. We adopt the following definition of sampling density from [1].

Medial axis and ϵ-sampling

The *medial axis* of a surface S in $\mathbf{R}^3$ is the closure of the set of points which have more than one closest point on S. The *local feature size*, $f(p)$, at point $p \in S$ is the least distance of p to the medial axis. The maximal balls tangent to S at p are centered on points of the medial axis; call these the *medial balls* at p. Notice that $f(p)$

is not necessarily same as the radius of the *medial balls* at p. A very useful property of $f(\cdot)$ is that it is 1-Lipschitz, that is, $f(p) \le f(q) + |pq|$ for any two points p, q on S. A point set P is called an ϵ-sample of a surface S if every point $p \in S$ has a sample within distance $\epsilon f(p)$.

Main theorem

Given these definitions, we can state formally the main theorem of this paper.

Theorem 1 *Let P be an ϵ-sample for a smooth surface S, with $\epsilon \le 0.08$. Our algorithm computes a piecewise-linear 2-manifold N homeomorphic to S, such that any point on N is at most $\frac{1.3\epsilon f(x)}{1-\epsilon}$ from some point $x \in S$.*

Restricted Delaunay triangulation

We assume that the input sample $P \in \mathbf{R}^3$ is in general position; in practice most Delaunay triangulation codes simulate general position, so this is not unreasonable. Let D_P and V_P denote the Delaunay triangulation and the Voronoi diagram of P. A Voronoi cell $V_p \in V_P$ for each point $p \in P$ is defined as the set of points $x \in \mathbf{R}^3$ such that $|px| \le |qx|$ for any $q \in P$ and $q \neq p$. The Delaunay triangulation has an edge pq if and only if V_p, V_q share a face, has a triangle pqr if and only if V_p, V_q and V_r share an edge, and a tetrahedron $pqrs$ if and only if V_p, V_q, V_r and V_s share a Voronoi vertex.

Consider the restriction of V_P to the surface S. This defines the *restricted Voronoi diagram* $V_{P,S}$, with *restricted Voronoi cells* $V_{p,S} = V_p \cap S$. The dual of these restricted Voronoi cells defines the *restricted Delaunay triangulation* $D_{P,S}$. Specifically, an edge pq is in $D_{P,S}$ if and only if $V_{p,S} \cap V_{q,S}$ is nonempty; a triangle pqr is in $D_{P,S}$ if and only if $V_{p,S} \cap V_{q,S} \cap V_{r,S}$ is nonempty. Assuming that S does not pass through a Voronoi vertex in general position there is no tetrahedron in $D_{P,S}$. Edelsbrunner and Shah [13] showed that the underlying space of $D_{P,S}$ is homeomorphic to S if the following *closed ball property* holds: each $V_{p,S}$ is a topological 2-ball, each nonempty pairwise intersection $V_{p,S} \cap V_{q,S}$ is a topological 1-ball, and each nonempty triple intersection $V_{p,S} \cap V_{q,S} \cap V_{r,S}$ is a single point, that is, a 0-ball. Amenta and Bern used this result to show that if P is a ϵ-sample of S with $\epsilon \le 0.1$, then $V_{P,S}$ satisfies the closed ball property, and hence D_P contains the set $D_{P,S}$ of triangles forming a piecewise-linear manifold homeomorphic to S.

Conditions for homeomorphism

Our algorithm selects a set of *candidate triangles* T that satisfy the following three conditions.

I. RESTRICTED DELAUNAY CONDITION. The set of triangles includes the restricted Delaunay triangles.

II. SMALL TRIANGLE CONDITION. The circumcircle of each triangle $t \in T$ is small; specifically, its radius is $O(\epsilon)f(p)$, where p is any vertex of t.

III. FLAT TRIANGLE CONDITION. The normal to each $t \in T$ makes a small, $O(\epsilon)$, angle with the surface normal at the vertex p, where p is the vertex with the largest interior angle in t.

3 Algorithm

Our algorithm selects the candidate triangles using *co-cones* at each sample point, and then (at least in theory) extracts a piecewise-linear manifold from T.

Co-cones

The normal to S at each sample point is estimated using "poles", which were introduced in [1]. For each Voronoi cell V_p, the Voronoi vertex farthest from the sample point p is taken as a pole. The line through p and its pole is

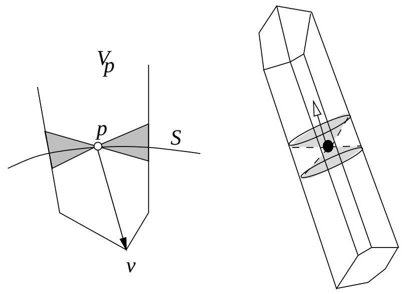

Figure 1: The co-cone for a sample in two dimensions (left), and three dimensions (right). In the left the co-cone is shaded, in the right its boundary is shaded.

almost normal to S and is called the *estimated normal line* at p; see Figure 1. For an angle θ, we define a cone-complement – the *co-cone* at p – as the complement of the double cone with apex p making an angle of $\pi/2-\theta$ with the estimated normal line at p. We determine the set of Voronoi edges in V_p that intersect the co-cones at all three of the sample points inducing the edge. The dual triangles of these edges form our candidate set T. We will argue that T satisfies conditions I, II and III for $\theta \leq \pi/8$.

Note that this is equivalent to the definition of T mentioned in the introduction. Any Voronoi edge e of a sample p which intersects the co-cone at p must contain a point x such that the angle between the estimated normal line and the vector px is at least $\pi/2-\theta$.

Computing T is absolutely straightforward. We first compute the Delaunay triangulation of P, and the Voronoi vertices dual to every tetrahedron, and we find the pole v of every sample p. Denote any ray from p to a point $y \in V_p$ as $\vec{y}$. Let e be an edge in the Voronoi cell V_p, and let w_1, w_2 be its two endpoints. We compute $\angle \vec{w_1}\vec{v}$ and $\angle \vec{w_2}\vec{v}$ and check if the range of angles determined by these two angles intersects the desired range $[\pi/2-\theta, \pi/2+\theta]$. If it does, we mark e. We include a Delaunay triangle t in T if its dual edge e is marked by all three Voronoi cells adjacent to e.

Manifold extraction

For completeness, we review the manifold extraction step of the crust algorithm [1]. First, we delete all triangles incident to *sharp edges.* An edge is called *sharp* if the angle between any two consecutive triangles around the edge is more than $3\pi/2$. An edge with a single incident triangle is also sharp. Next, we extract the outer boundary N of the set of triangles by a depth-first walk along the outer boundary of each of its connected components. As mentioned earlier, we use a heuristic to implement the manifold extraction to deal with practical data that may not satisfy the sampling condition required by our theory.

4 Conditions

We use the following three lemmas from [1]. The first two establish that the vector from a sample to its pole estimates the normal at the sample.

Lemma 2 *Let y be any point in V_p such that $|py| \geq \delta f(p)$ for $\delta > 0$. The acute angle between $\vec{y}$ and n_p is less than* $\sin^{-1}\frac{\epsilon}{\delta(1-\epsilon)} + \sin^{-1}\frac{\epsilon}{1-\epsilon}$.

Using Lemma 2 and the fact that $|pv| > f(p)$ (recall that v is the pole of p), we can bound the deviation of $\vec{v}$ from the surface normal n_p.

Lemma 3 *The acute angle between n_p and $\vec{v}$ is less than* $2\sin^{-1}\frac{\epsilon}{1-\epsilon}$.

We also use the following lemma [1], which establishes that, within a bounded region around p, the surface normal is also a Lipschitz function.

Lemma 4 *Let p, q be two points on S so that $|pq| < \rho \min\{f(p), f(q)\}$ with $\rho < 1/3$. Then the angle between n_p and n_q is at most $\frac{\rho}{1-3\rho}$ radians.*

Restricted Delaunay condition

Condition I requires the restricted Delaunay triangles to be in T. We begin with a technical observation, which says that the line segment connecting two points close together on S must be nearly parallel to the surface.

Observation 5 *A line segment connecting two points $x, x' \in S$, such that the distance $|x, x'| \leq c\epsilon f(x)$, with $c \leq \sqrt{2}$, makes an acute angle with the surface normal n_x at x of at least $\pi/2 - sin\frac{c\epsilon}{2}$.*

This follows from the fact that x' must lie outside the two tangent balls of radius $f(x)$ at x.

Lemma 6 *Let y be any point in the restricted Voronoi cell $V_{p,S}$. The acute angle between n_p and $\vec{y}$ is larger than $\pi/2 - \epsilon$, for $\epsilon < 0.1$.*

Proof. The distance $|yp| \leq \epsilon f(y)$, since $y \in V_{p,S}$ and P is an ϵ-sample. By the Lipschitz condition $f(y) \leq f(p) + |py|$ giving $f(y) \leq \frac{f(p)}{1-\epsilon}$, and hence $|py| \leq \epsilon f(y) \leq \frac{\epsilon}{1-\epsilon} f(p)$. We can therefore apply Observation 5. ◻

We can now prove that T satisfies Condition I.

Theorem 7 *All restricted Delaunay triangles are in T, for $\epsilon \leq 0.1$ and $\theta \leq \pi/8$.*

Proof. Let e be the dual edge of a restricted Delaunay triangle. Consider the point $y = e \cap S$. We have $y \in V_{p,S}$ for the three points $p \in P$ determining e. For each such p, the acute angle between n_p and $\vec{y}$ is larger than $\pi/2 - \epsilon$ by Lemma 6. Therefore $\angle \vec{y}\vec{v} \in [\pi/2 - \epsilon - \alpha, \pi/2 + \epsilon + \alpha]$, where α is the acute angle between $\vec{v}$ and n_p. Plugging in the upper bound on α from Lemma 3 we find that $\alpha + \epsilon < \pi/8$, so $\angle \vec{y}\vec{v} \in I$. ◻

Small triangle condition

We now show that T meets Condition II. Looking at the contrapositive of Lemma 2 and plugging in a value of $\delta = \frac{1.3\epsilon}{1-\epsilon}$ we obtain the following.

Corollary 8 *Let x be any point in V_p so that the acute angle between $\vec{x}$ and n_p is at least $\pi/2 - \theta - 2\sin^{-1}\frac{\epsilon}{1-\epsilon}$. Then $|px| < \frac{1.3\epsilon}{1-\epsilon} f(p)$, for $\theta = \pi/8$ and $\epsilon \leq 0.08$.*

Proof. If the acute angle between $\vec{x}$ and n_p is at least $\alpha = \sin^{-1}\frac{\epsilon}{\delta(1-\epsilon)} + \sin^{-1}\frac{\epsilon}{1-\epsilon}$, then $|px| < \delta f(p)$ according to Lemma 2. With $\delta = \frac{1.3\epsilon}{1-\epsilon}$ we have

$$\alpha = \sin^{-1}\frac{1}{1.3} + \sin^{-1}\frac{\epsilon}{1-\epsilon}$$

which is less than $\pi/2 - \theta - 2\sin^{-1}\frac{\epsilon}{1-\epsilon}$ for $\theta = \pi/8$ and $\epsilon \leq 0.08$. ◻

Lemma 9 *Let p be a vertex of any triangle $t \in T$. The radius of the smallest Delaunay ball of t is at most $\frac{1.3\epsilon}{1-\epsilon} f(p)$ for $\epsilon \leq 0.08$ and $\theta = \pi/8$.*

Proof. Let e be the dual edge of t and p any vertex of t. By our choice of e, there is a point $x \in e$ so that $\vec{x}$ makes an angle in the range $I = [\pi/2 - \theta, \pi/2 + \theta]$ with $\vec{v}$. Taking into account the angle between $\vec{v}$ and n_p we conclude that this ray makes an acute angle more than $\pi/2 - \theta - 2\sin^{-1}\frac{\epsilon}{1-\epsilon}$ with n_p. From Corollary 8, $|px| < \frac{1.3\epsilon}{1-\epsilon} f(p)$. ◻

Theorem 10 *Let r denote the radius of the circumcircle of any triangle $t \in T$. Then, for each vertex p of t, $r \leq \frac{1.3\epsilon}{1-\epsilon} f(p)$ for $\epsilon \leq 0.08$ and $\theta = \pi/8$.*

Proof. The radius of the smallest Delaunay ball, bounded in Lemma 9, is an upper bound on the radius of the circumcircle of t, which is centered at the intersection of the line containing e with the plane containing t. ◻

Flat triangle condition

Here we show that T meets Condition III.

Theorem 11 *The normal to any triangle $t \in T$ makes an acute angle of no more than $\alpha + \sin^{-1}(\frac{2}{\sqrt{3}} \sin 2\alpha)$ with n_p where p is the vertex at the largest interior angle of t, and $\alpha \leq sin^{-1}\frac{1.3\epsilon}{1-\epsilon}$, $\epsilon \leq 0.08$.*

Proof. Consider the medial balls M_1 and M_2 touching S at p with the centers on the medial axis. Let D be the ball with the circumcircle of t as a diameter; refer to Figure 2. The radius r of D is equal to the radius of the circumcircle of t. Denote the circles of intersection of D with M_1 and M_2 as C_1 and C_2 respectively. The normal to S at p passes through m, the center of M_1. This normal makes an angle less than α with the normals to the planes of C_1 and C_2, where

$$\begin{aligned} \alpha &\leq \sin^{-1} r/|pm| \\ &\leq \sin^{-1} \frac{1.3\epsilon}{1-\epsilon} \end{aligned}$$

since $|pm| \geq f(p)$ by the definition of f and $r \leq \frac{1.3\epsilon}{1-\epsilon} f(p)$ by Theorem 10. This angle bound also applies to the plane of C_2, which implies that the planes of C_1 and C_2 make a wedge, say W, with an acute dihedral angle no more than 2α.

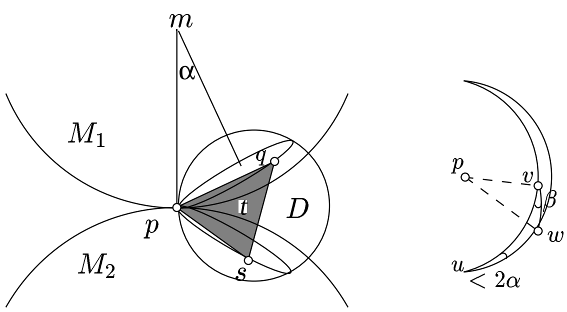

Figure 2: Normal to a small triangle and the normal to S at the vertex with the largest face angle.

The vertices other two vertices q, s of t cannot lie inside M_1 or M_2. This implies that t lies completely in the wedge W. Consider a cone at p inside the wedge W formed by the three planes; π_t, the plane of t, π_1, the plane of C_1 and π_2, the plane of C_2. A unit sphere centered around p intersects the cone in a spherical triangle uvw, where u, v and w are the points of intersections of the lines $\pi_1 \cap \pi_2$, $\pi_t \cap \pi_1$ and $\pi_t \cap \pi_2$ respectively with the unit sphere. See the picture on right in Figure 2. Without the loss of generality, assume that the angle $\angle uvw \leq \angle uwv$. We have the following facts. The arc length of wv, denoted $|wv|$, is at least $\pi/3$ since p subtends the largest angle in t and t lies completely in the wedge W. The spherical angle $\angle vuw$ is less than or equal to 2α. We are interested in the spherical angle $\beta = \angle uvw$ which is also the acute dihedral angle between the planes of t and C_1. By standard *sine laws* in spherical geometry, we have $\sin\beta = \sin|uw| \frac{\sin\angle vuw}{\sin|wv|} \leq \sin|uw| \frac{\sin 2\alpha}{\sin|wv|}$. If $\pi/3 \leq |wv| \leq 2\pi/3$, we have $\beta \leq \sin^{-1}\frac{2}{\sqrt{3}}\sin 2\alpha$. For the range $2\pi/3 < |wv| < \pi$, we use the fact that $|uw| + |wv| \leq \pi$ since $\angle vuw \leq 2\alpha < \pi/2$ for sufficiently small ϵ. So, in this case $\frac{\sin|uw|}{\sin|wv|} < 1$. Thus, $\beta \leq \sin^{-1}\frac{2}{\sqrt{3}}\sin 2\alpha$.

The normals to t and S at p make an acute angle at most $\alpha + \beta$ proving the lemma. ◻

The upper bound on the angle between the normal to t and n_p provided by this lemma is 0.4 radians; and the angle is $O(\epsilon)$.

5 Geometric consequences

The preceding lemmas have told us a great deal about T. We know [1] that the restricted Delaunay triangulation is a piecewise-linear surface homeomorphic to S, when $\epsilon \leq 0.1$. Condition I ensures that T contains the restricted Delaunay triangulation.

Triangle interiors

Conditions II and III relate properties of each triangle $t \in T$ to the value of $f(\cdot)$ and the surface normal direction, respectively, at its vertices. But since the triangles are small, we can use the Lipshcitz properties to show that similar properties hold at any point q in the interior of t. To define these properties, we map q to the nearest surface point. Let $\mu : \mathbf{R}^3 \to S$ map each point $q \in \mathbf{R}^3$ to the closest point of S. The restriction of μ to T is a well-defined function $\mu : T \to S$, Since if some point q had more than one closest point on the surface, q would be a point of the medial axis; but by Lemma 12 every point $q \in T$ is within $\frac{1.3\epsilon}{1-\epsilon} f(p)$ of a triangle vertex $p \in S$.

Lemma 12 *Let q be any point on a triangle $t \in T$. The distance between q and and the point $x = \mu(t)$ is at most $0.11f(x)$, for $\epsilon \leq 0.08$.*

Proof. The circumcircle of t is small; the distance from q to the vertex p of t with largest angle is at most $2\delta\, f(p)$, with $\delta = \frac{1.3\epsilon}{1-\epsilon} \leq .12$, by Theorem 10. Since there is a sample, namely, a vertex of t within $\delta f(p)$ from q, we have $|qx| \leq \delta f(p)$. We are interested in expressing this as a function of $f(x)$, so we need an upper bound on $|px|$.

The triangle vertex p has to lie outside the tangent ball at x, while, since x is the nearest surface point to q, q must lie on the segment between x and the center of this tangent ball. For any fixed $|pq|$, these facts imply that $|px|$ is maximized when the angle pqx is a right angle. Thus, $|px| \leq \sqrt{5}\delta f(p) \leq 0.27f(p)$ for $\epsilon \leq 0.08$. This implies that $f(p) \leq 1.37f(x)$ by Lipschitz property of $f(\cdot)$, giving $|qx| \leq 0.165f(x)$. ◻

With a little more work, we can also show that the triangle normal agrees with the surface normal at the surface point closest to t.

Lemma 13 *Let q be a point on triangle $t \in T$, and let n_x be the surface normal at $x = \mu(q)$. The acute angle between n_x and and the normal to t is at most $42°$ for $\epsilon \leq 0.08$. Also, the acute angle between n_x and the surface normal n_p at the vertex p of t with largest angle is at most $19°$.*

Proof. Applying Lemma 4, and taking $\rho = 0.165$, shows that the angle between n_x and n_p is less than $19°$. The angle between the triangle normal of t and n_p is less than $23°$ for $\epsilon \leq 0.08$ (Theorem 11). Thus, the triangle normal and n_x make an angle of at most $42°$. ◻

Sharp edges

The manifold extraction step selects a piecewise-linear manifold from T. It begins by recursively removing any triangle in T adjacent to a *sharp* edge; recall that a sharp edge is one for which the angle between two adjacent triangles, in the circular order around the edge, is greater than $3\pi/2$. Let T' be the remaining set of triangles. The following lemma shows that none of the restricted Delaunay triangles are removed, so that T' is guaranteed to contain a piecewise-linear manifold homeomorphic to S.

Lemma 14 *No restricted Delaunay triangle has a sharp edge, for $\epsilon \leq 0.08$.*

Proof. Let t and t' be adjacent triangles in the restricted Delaunay triangulation, let e be their shared edge, and let $p \in e$ be any of their shared vertices. Since t and t' belong to the restricted Delaunay triangulation, they have circumspheres B and B', respectively, centered at points v, v' of S.

The boundaries of the circumspheres B and B' intersect in a circle C contained in a plane H, with $e \subset H$. H separates t and t', since the third vertex of each triangle must lie on the boundary of its circumsphere, and $B \subseteq B'$ on one side of H, while on the other $B' \subseteq B$. The line through v, v' is perpendicular to H, and the distance $|vv'| \leq 2\epsilon/(1-\epsilon)f(v)$ (using the sampling condition). So segment v, v' forms an angle of at least $\pi/2 - sin\frac{\epsilon}{1-\epsilon}$ with n_v (Observation 5).

This normal differs, in turn, from n_p by at most $\frac{\epsilon}{1-4\epsilon}$ (Lemma 4), so H is nearly parallel to n_p, at an angle of at most .118 radians. The normals of both t and t' differ from the surface normal at p by at most .4 radians (Theorem 11).

Thus we have t on one side of H, t' on the other, and the smaller angle between H and either triangle is at least $\pi/2 - .518$ radians. Hence the smaller angle between t and t' is at least 2.62 radians, and e is not sharp. $\square$

6 Homeomorphism

In this section, we will show a homeomorphism between S and any piecewise-linear surface made up of candidate triangles from T with two additional properties. The piecewise-linear manifold N selected by the manifold extraction step of our algorithm does in fact have these properties, thus completing the proof of Theorem 1.

Additional properties

A pair of triangles $t_1, t_2 \in N$ are *adjacent* if they share at least one common vertex v. Since the normals to all triangles sharing v differ from the surface normal at v by at most $42°$ (Lemma 13), and that normal in turn differs from the vector to the pole at v by less than $5°$ (Lemma 3), we can orient the triangles sharing v, arbitrarily but consistently calling the normal facing the pole the *inside* normal and the normal facing away from the pole the *outside* normal. Let α be the angle between the two inside normals of t_1, t_2. We define the angle at which the two triangles meet at v to be $\pi - \alpha$.

PROPERTY I: Every two adjacent triangles in N meet at their common vertex at an angle of greater than $\pi/2$.

Requiring this property excludes manifolds which contain sharp folds and, for instance, flat tunnels. Since the triangles of T are all nearly perpendicular to the surface normals at their vertices, and the manifold extraction step eliminates triangles adjacent to sharp edges, N has this property.

PROPERTY II: Every sample in P is a vertex of N.

Lemma 14 ensures that T' contains the restricted Delaunay triangulation, which contains a triangle adjacent to every sample in P. Lemma 16, below, ensures that at least one triangle must be selected for each sample by the manifold extraction step. This implies that N has the second property as well.

Homeomorphism proof

We define the homeomorphism explicitly, using the function $\mu : N \to S$, as defined above. The function μ defines a homeomorphism between N and S if it is continuous, one-to-one and onto. Our approach will be first to show that μ is well-behaved on the samples themselves, and then show that this behavior continues in the interior of each triangle of N.

Lemma 15 *The restriction of μ to N is a continuous function $\mu : N \to S$.*

Proof. By Theorem 10 every point $q \in N$ is within $\frac{1.3\epsilon}{1-\epsilon} f(p)$ of a triangle vertex $p \in S$. Function μ is continuous except at the medial axis of S, so that since N is continuous and avoids the medial axis, μ is continuous on N. $\square$

Lemma 16 *Let p be a sample and let m be the center of a medial ball M tangent to the surface at p. No candidate triangle intersects the interior of the segment pm.*

Proof. In order to intersect segment pm, a candidate triangle t would have to intersect M, and so would the smallest Delaunay ball D of t. Let H be the plane of the circle where the boundaries of M and D intersect. We show that H separates the interior of pm and t.

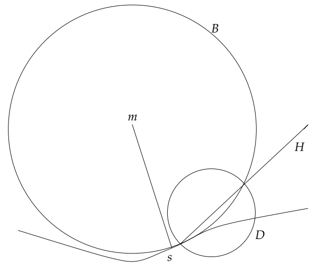

Figure 3: Proof of Lemma 16.

On one side of H, M is contained in D, and on the other, D is contained in M. Since the vertices of t lie on S and hence not in the interior of M, t has to lie in the open halfspace, call it H^+, in which D is outside M. Since D is Delaunay, p cannot lie in the interior of D; but since p lies on the boundary of M, it therefore cannot lie in H^+. We claim that $m \notin H^+$ either. (see Figure 3.) Since $m \in M$, if it lay in H^+ then m would have to be contained in D. Since m is a point of the medial axis, this would mean that the radius of D would be at least $1/2\ f(p')$ for any vertex p' of t, contradicting, by Lemma 9, the assertion that t is a candidate triangle. Therefore p, m and hence the segment pm cannot lie in H^+, and H separates t and pm. ◻

Since any point q such that $\mu(q) = p$ lies on such an open segment pm, we have the following.

Corollary 17 *The function μ is one-to-one from N to every sample p.*

In what follows, we will show that μ is indeed one-to-one on all of N.

Our proof proceeds in three short steps. We show that μ induces a homeomorphism on each triangle, then on each pair of adjacent triangles, and finally on N as a whole.

Lemma 18 *Let U be a region contained within one triangle $t \in N$ or in adjacent triangles of N. The function μ defines a homeomorphism between U and $\mu(U) \subset S$.*

Proof. We know that μ is well-defined and continuous on U, so it only remains to show that it is one-to-one. First, we prove that if U is in one triangle t, μ is one-to-one. For a point $q \in t$, the vector $\vec{n_q}$ from $\mu(q)$ to q is perpendicular to the surface at $\mu(q)$; since S is smooth the direction of $\vec{n_q}$ is unique and well defined. If there was some $y \in t$ with $\mu(y) = \mu(q)$, then q, $\mu(q)$ and y would all be colinear and t itself would have to contain the line segment between q and y, contradicting Lemma 13, which says that the normal of t is nearly parallel to $\vec{n_q}$.

Now, we consider the case in which U is contained in more than one triangle. Let q and y be two points in U such that $\mu(q) = \mu(y) = x$, and let v be a common vertex of the triangles that contain U. Since μ is one-to-one in one triangle, q and y must lie in the two distinct triangles t_q and t_y. Let n_x be the surface normal at x. The line l through x with direction n_x pierces the patch U at least twice; if y and q are not adjacent intersections along l, redefine q so that this is true ($\mu(q) = x$ for any intersection q of l with U). Now consider the orientation of the patch U according to the direction to the pole at v. Either l passes from inside to outside and back to inside when crossing y and q, or from outside to inside and back to outside.

The acute angles between the triangle normals of t_q, t_y and n_x are at most $42°$ (Lemma 13), that is, the triangles are stabbed nearly perpendicularly by n_x. But since the orientation of U is opposite at the two intersections, the angle between the two *oriented* triangle normals is at least $\pi - 84°$, meaning that t_q and t_y must meet at v at an acute angle of at most $84°$. This would contradict PROPERTY I, which is that t_q and t_y meet at v at an obtuse angle. Hence there are no two points in y, q with $\mu(q) = \mu(y)$. ◻

We finish the theorem using a theorem from topology.

Theorem 19 *The mapping μ defines a homeomorphism from the triangulation N to the surface S for $\epsilon \leq 0.08$.*

Proof. Let $S' \subset S$ be $\mu(N)$. We first show that (N, μ) is a *covering space* of S'. Informally,

(N, μ) is a covering space for S' if function μ maps N smoothly onto S', with no folds or other singularities; see Massey [17], Chapter 5. Showing that (N, μ) is a covering space is weaker than showing that μ defines a homeomorphism, since, for instance, it does not preclude several connected components of N mapping onto the same component of S', or more interesting behavior, such as a torus wrapping twice around another torus to form a *double covering*.

Formally, the (N, μ) is a covering space of S' if, for every $x \in S'$, there is a path-connected *elementary neighborhood* V_x around x such that each path-connected component of $\mu^{-1}(V_x)$ is mapped homeomorphically onto V_x by μ.

To construct such an elementary neighborhood, note that the set of points $|\mu^{-1}(x)|$ corresponding to a point $x \in S'$ is non-zero and finite, since μ is one-to-one on each triangle of N and there are only a finite number of triangles. For each point $q \in \mu^{-1}(x)$, we choose an open neighborhood U_q of around q, homeomorphic to a disk and small enough so that U_q is contained only in triangles that contain q.

We claim that μ maps each U_q homeomorphically onto $\mu(U_q)$. This is because it is continuous, it is onto $\mu(U_q)$ by definition, and, since any two points x and y in U_q are in adjacent triangles, it is one-to-one by Lemma 18.

Let $U'(x) = \cap_{q \in \mu^{-1}(x)} \mu(U_q)$, the intersection of the maps of each of the U_q. $U'(x)$ is the intersection of a finite number of open neighborhoods, each containing x, so we can find an open disk V_x around x. V_x is path connected, and each component of $\mu^{-1}(V_x)$ is a subset of some U_q and hence is mapped homeomorphically onto V_x by μ. Thus (N, μ) is a covering space for S'.

We now show that μ defines a homeomorphism between N and S'. Since N is onto S' by definition, we need only show that μ is one-to-one. Consider one connected component G of S'. A theorem of algebraic topology (see eg. Massey [17], Chapter 5 Lemma 3.4) says that when (N, μ) is a covering space of S', the sets

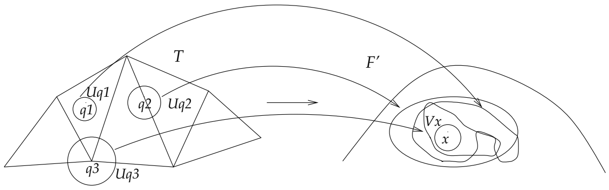

Figure 4: Proof of Theorem 19.

$\mu^{-1}(x)$ for all $x \in G$ have the same cardinality. We now use Corollary 17, that μ is one-to-one at every sample. Since each connected component of S contains some samples, it must be the case that μ is everywhere one-to-one, and N and S' are homeomorphic.

Finally, we show that $S' = S$. We must have S' to be closed and compact since N is. So S' cannot include part of a connected component of S, and hence S' must consist of a subset of the connected components of S. Since every connected component of S contains a sample p (actually many samples), and $\mu(p) = p$, all components of S belong to S', $S' = S$, and N and S are homeomorphic. ◻

7 Implementation

We have implemented the algorithm and tested it on several data sets. We faced a serious implementation difficulty with the manifold extraction step. The manifold extraction step depend heavily on the assumptions that the surface is smooth, has no boundaries, and that the sampling is sufficiently dense. In practice the data do not satisfy these assumptions. The manifold extraction step deletes edges at boundaries, since they are sharp edges. Recursively deleting boundary edges can, in the worst case, delete the entire output of the co-cone filtering step. Even when the surface has no boundaries, the output of the co-cone step generally has holes which are produced due to noise and undersampling, mostly in non-smooth regions. To prevent deleting too many triangles, we use a heuristic which we call *UmbrellaCheck*.

objects	points	triangles	time(sec.)
FOOT set T	20021	40,341	153
CLUB set T	16864	33,692	122
FOOT surface N		40,004	234
CLUB surface N		33,670	189

Table 1: Experimental data.

We say that a vertex v has an umbrella if there exists a set of triangles incident to v which form a topological disk and no two consecutive triangles around the disk meet at a dihedral angle less than $\frac{\pi}{2}$ or more than $\frac{3\pi}{2}$. *UmbrellaCheck* determines if a vertex v has an umbrella or not. This is done by considering the process of deleting triangles adjacent to sharp edges only on the triangles incident to v. We recursively mark all triangles adjacent to v and to a sharp edge as "deleted". If there v has no umbrella, all triangles adjacent to v will be marked. Otherwise, an unmarked triangle incident to v remains, and we conclude that v has an umbrella. In the manifold extraction step, we actually delete a triangle incident with a sharp edge only if all of its three vertices have umbrellas.

We implemented the algorithm in $C++$ using the well known `qhull` code for Delaunay triangulation. We show outputs for two data sets FOOT and CLUB in Figure 5. The outputs are shown before (set T) and after (surface N) the manifold extraction step. There is not much visible difference between the two outputs though the number of triangles and running times in Table 1 indicate the difference. We also provide zoomed pictures of the FOOT near the ankle for the set T and the surface N. One can notice some slivers with one missing triangle in the picture for T. The remaining three triangles form almost a square with a hanging triangle on top of it. They disappear in the picture for the surface N. The surface N is computed correctly almost everywhere except at the boundaries and near sharp features, as expected.

Our running times, measured on a SUN machine with the 300Mhz processor and 256 MB memory, are faster than those reported in [3] for the crust algorithm. For example, the foot took 153 seconds for extracting the set T, for which the crust algorithm required 15 minutes on a SGI Onyx machine with 512 MB memory. The difference can be explained by two factors; first, this algorithm requires only one Delaunay triangulation step, and second, the implementation of [3] used the exact-arithmetic Delaunay triangulation program `hull`, which we have observed to be about four times slower than `qhull` on these inputs.

8 Conclusions

The main advantages of our algorithm over the original crust algorithm [1] are: (i) it requires only one Voronoi diagram computation as opposed to two; (ii) it collects a set of triangles from the Delaunay triangulation by checking a single simple condition; (iii) the proofs are simpler; and (iv) we can give a topological guarantee on the output.

Our theory is supported by the output of our program on some reasonably large data sets. We should note, however, that in practice many surfaces have sharp corners and boundaries, and that sets of sample points are often noisy and fail to meet our sampling condition, so that our theoretical results do not guarantee good reconstruction in practice.

Important goals that remain in this area are to correctly reconstruct surfaces with sharp edges, corners, and boundaries, to develop reconstruction algorithms that gracefully handle noise, and to find more efficient algorithms that avoid computing the Delaunay triangulation of all the input samples.

Acknowledgement

The authors thank J. Havlicek and R. Wenger for giving insightful comments on the paper.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.*, **22**, (1999), 481–504.

[2] N. Amenta, M. Bern and D. Eppstein. The crust and the β-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, **60** (1998), 125-135.

[3] N. Amenta, M. Bern and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *SIGGRAPH 98*, (1998), 415-421.

[4] D. Attali. r-regular shape reconstruction from unorganized points. *Proc. 13th ACM Sympos. Comput. Geom.*, (1997), 248–253.

[5] F. Bernardini and C. L. Bajaj. Sampling and reconstructing manifolds using α-shapes. *Proc. 9th Canadian Conf. Comput. Geom.*, (1997), 193–198.

[6] J. D. Boissonnat. Shape reconstruction from planar cross-sections. *Computer Vision, Graphics, and Image Processing* **44** (1988), 1–29.

[7] J. D. Boissonnat and F. Cazals. Smooth shape reconstruction. To appear in *Proc. 16th ACM Sympos. Comput. Geom.*, 2000.

[8] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *SIGGRAPH 96*, (1996), 303-312.

[9] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. ACM-SIAM Sympos. Discr. Algorithms*, (1999), 893–894.

[10] T. K. Dey, K. Mehlhorn and E. A. Ramos. Curve reconstruction: connecting dots with good reason. *Comput. Geom. Theory Appl.*, **15** (2000), 229–244.

[11] H. Edelsbrunner. Shape reconstruction with Delaunay complex. *LNCS 1380, LATIN'98:Theoretical Informatics* (1998), 119–132.

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

[13] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc. 10th ACM Sympos. Comput. Geom.*, (1994), 285–292.

[14] J. Giesen. Curve reconstruction, the TSP, and Menger's theorem on length. *Proc. 15th ACM Sympos. Comput. Geom.*, (1999), 207–216.

[15] C. Gold. Crust and anti-crust: a one-step boundary and skeleton extraction algorithm. *Proc. 15th ACM Sympos. Comput. Geom.*, (1999), 189–196.

[16] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stuetzle. Surface reconstruction from unorganized points. *SIGGRAPH 92*, (1992), 71-78.

[17] W. S. Massey. *Algebraic Topology: An Introduction*, Springer-Verlag, Graduate texts in Mathematics 56, 1967.

[18] M. Melkemi. $\mathcal{A}$-shapes and their derivatives. *Proc. 13th ACM Sympos. Comput. Geom.*, (1997), 367–369.

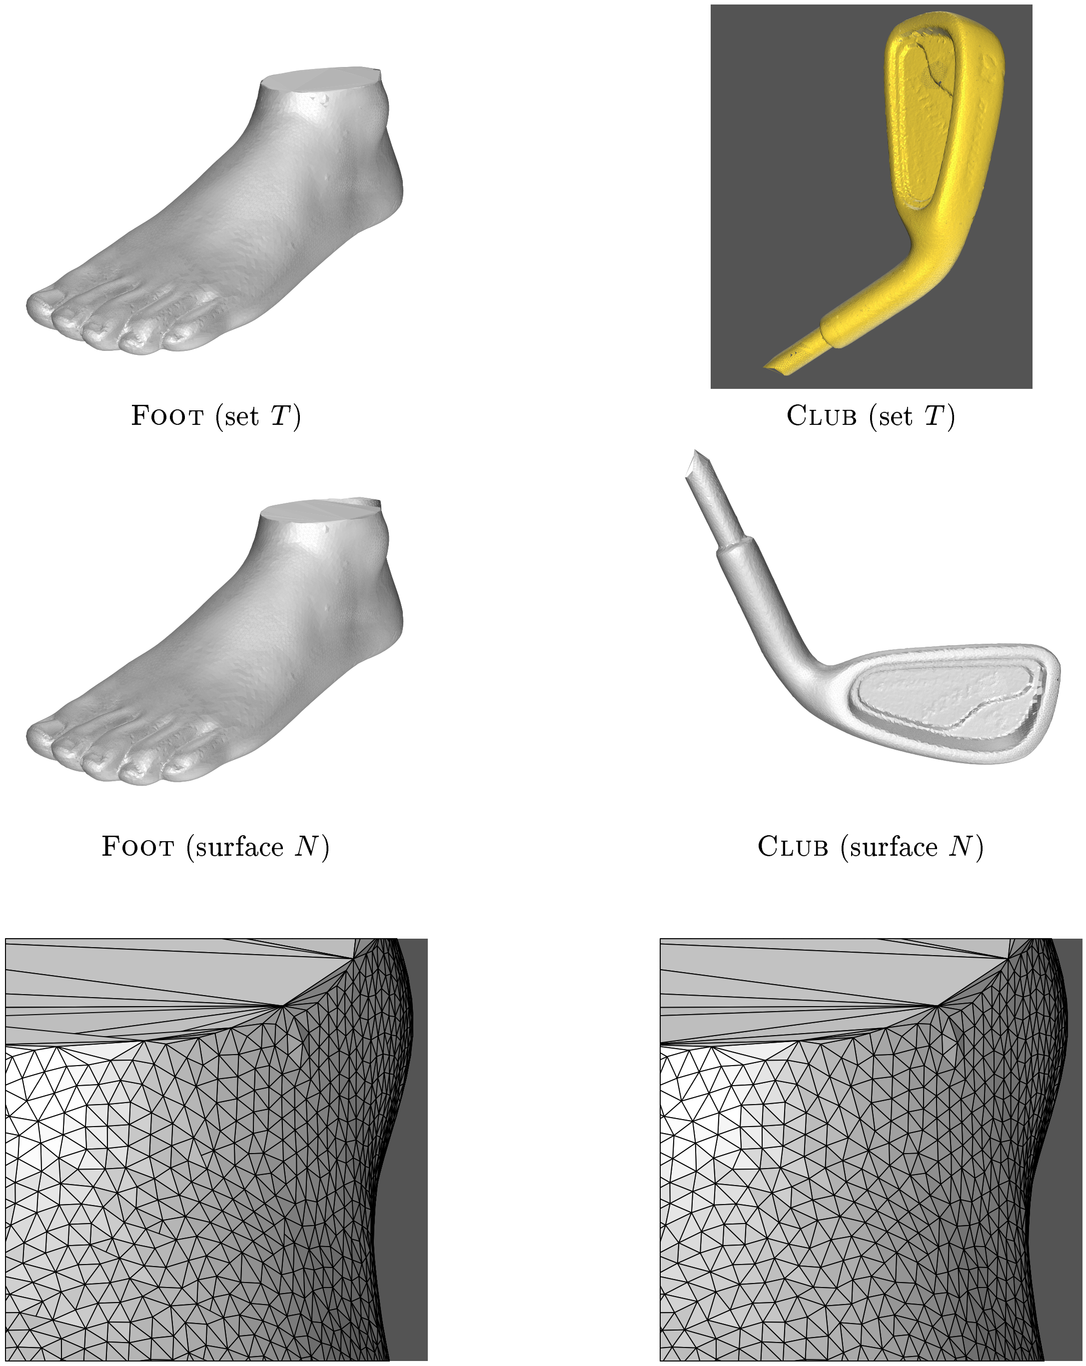

FOOT (set T)

CLUB (set T)

FOOT (surface N)

CLUB (surface N)

Zoomed T for FOOT

Zoomed N for FOOT

Figure 5: Experimental results