

The Crust and the β -Skeleton: Combinatorial Curve Reconstruction

Nina Amenta* Marshall Bern[†] David Eppstein[‡]

December 2, 1997

Abstract

We construct a graph on a planar point set, which captures its shape in the following sense: if a smooth curve is sampled densely enough, the graph on the samples is a polygonalization of the curve, with no extraneous edges. The required sampling density varies with the *Local Feature Size* on the curve, so that areas of less detail can be sampled less densely. We give two different graphs that, in this sense, reconstruct smooth curves: a simple new construction which we call the *crust*, and the β -skeleton, using a specific value of β .

1 Introduction

There are many situations in which a set of sample points lying on or near a surface is used to reconstruct a polygonal approximation to the surface. In the plane, this problem becomes a sort of unlabeled version of connect-the-dots: we are given a set of points and asked to connect them into the most likely polygonal curve. We show that under fairly generous and well-defined sampling conditions either of two proximity-based graphs defined on the set of points is guaranteed to reconstruct a smooth curve. These two graphs are the *crust*, which we define below, and the β -skeleton, defined ten years ago by Kirkpatrick and Radke [KR88], with an appropriately chosen value of β .

Figure 1 shows an example of a point set and its crust. The points were chosen by hand. Notice that fewer samples are required on the goose's back than on its head and foot.

*Computer Sciences, University of Texas, Austin, TX 78712. Most of this work was done at Xerox PARC, partially supported by the NSF grant CCR-9404113.

[†]Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA.

[‡]Department of Information and Computer Science, University of California, Irvine, CA 92697. Work supported in part by NSF grant CCR-9258355 and by matching funds from Xerox Corp, and performed in part while visiting Xerox PARC.




Figure 1: A point set and its *crust*.

The reconstruction of curves in the plane is important in computer vision. Simple *edge detectors* select image pixels which are likely to belong to edges, often delimiting the boundaries of objects. Grouping these pixels into likely curves is an area of active research. Extension of our ideas to three dimensions would be useful for constructing three-dimensional models from laser range data, stereo measurements, and medical images.

2 Definitions

In this paper, we will consider closed, compact, twice-differentiable 1-manifolds, without boundary, embedded in the plane; we shall call such a manifold a *smooth curve*. According to our definition, then, a smooth curve can have several connected components, but no endpoints, branches or self-intersections. Let F be a smooth curve and $S \subset F$ a finite set of sample points on F .

Definition: A *polygonal reconstruction* of F from S is a graph that connects every pair of samples adjacent along F , and no others.

Clearly no algorithm can reconstruct *any* curve from *any* set of samples; we need some condition on the quality of S . Our condition will be that the distance from any point p on F to the nearest sample $s \in S$ is at most a constant factor r times the *Local Feature Size* at p , which we define to be the distance from p to the *medial axis* of F (see section 4). This condition has the nice property that less detailed sections of the curve do not have to be sampled as densely.

Given a sample S from a smooth curve which meets the sampling condition for an appropriately small value of r , we show that a polygonal reconstruction is given

by either of the two graphs defined below.

We shall say that a disk B *touches* an object x if the intersection $B \cap x$ is a subset of the boundary of B (that is, we mean that B “just touches” x). We say that B is *empty* of points in x if its interior contains no points of x . The graph definitions are both related to the Voronoi diagram and Delaunay triangulation of S ¹, and we shall refer to the following well-known property:

Empty Circle Property: Two points in S determine an edge of the Delaunay triangulation if there is a disk B , empty of points in S , which touches them both.

We now define the graphs we will use for reconstruction.

Definition: Let S be a finite set of points in the plane, and let V be the vertices of the Voronoi diagram of S . Let S' be the union $S \cup V$, and consider the Delaunay triangulation of S' . An edge of the Delaunay triangulation of S' belongs to the *crust* of S if both of its endpoints belong to S .

An alternate definition can be given using the Empty Circle Property:

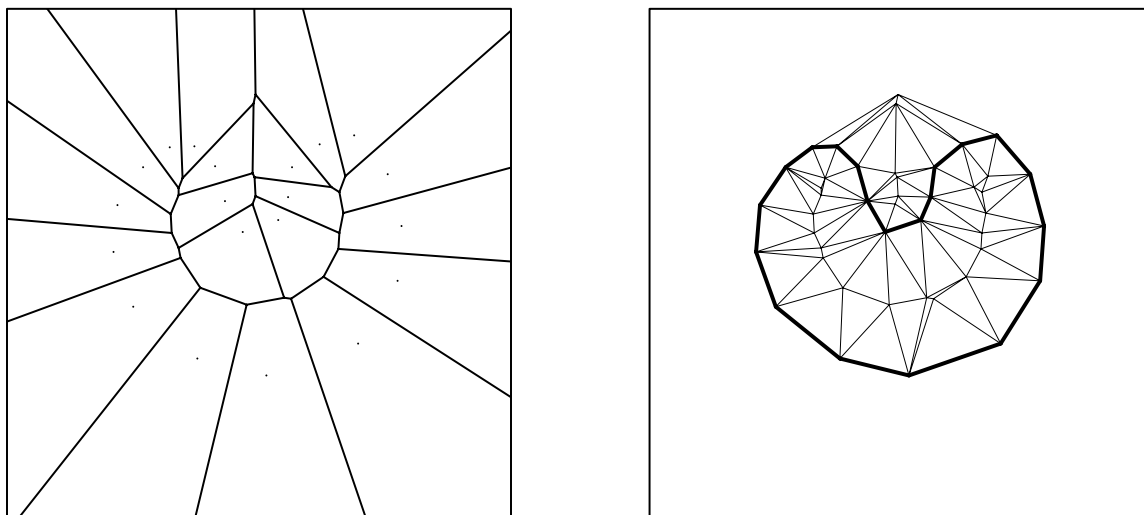


Figure 2: A Voronoi diagram of a point set S , and the Delaunay triangulation of $S \cup V$, with the crust highlighted.

Alternate Definition: Let S be a finite set of points in the plane, and let V be the vertices of the Voronoi diagram of S . An edge between points $s_1, s_2 \in S$ belongs to the *crust* of S if there is a disk, empty of points in $S \cup V$, touching s_1 and s_2 .

The intuition behind the definition of the crust is that the vertices V of the Voronoi diagram of S approximate the medial axis of F , and the Voronoi disks of $S' = S \cup V$ approximate empty circles between F and its medial axis. Note that if an edge

¹For more on the Voronoi diagram and Delaunay triangulation see any of the standard computational geometry texts, eg. [PS85],[E87] or [O94].

between two points of S belongs to the Delaunay triangulation of S' it certainly belongs to the Delaunay triangulation of S , and hence the crust of S is a subset of the Delaunay triangulation of S .

We now review the definition of the β -skeleton. An edge is present in the β -skeleton if the following *forbidden region* is empty of points of S . Let $\beta \geq 1$ be a constant.

Definition: Let s_1, s_2 be a pair of points in the plane, at distance $d(s_1, s_2)$. The *forbidden region* of s_1, s_2 is the union of the two disks of radius $\beta d(s_1, s_2)/2$ touching s_1 and s_2 .

Examples of the forbidden region for different values of β are shown in Figure 3. Reasonable definitions for $\beta \leq 1$ can also be made; see [KR88].

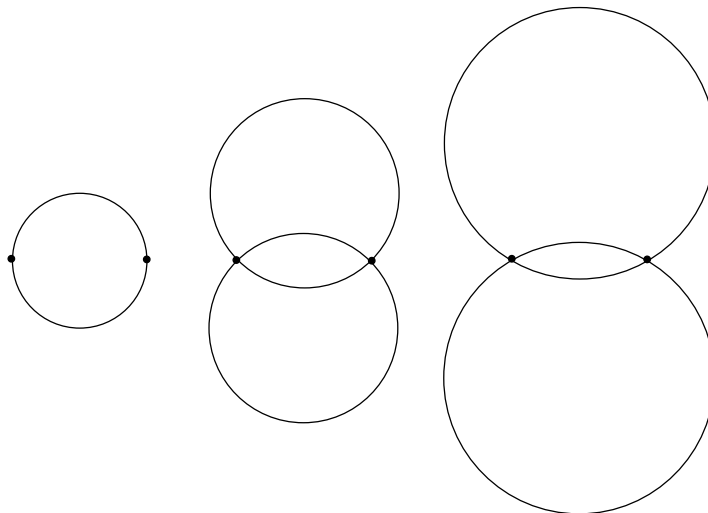


Figure 3: Forbidden regions for $\beta = 1, 3/2, 2$

Definition: Let S be a finite set of points in the plane. An edge between $s_1, s_2 \in S$ belongs to the β -skeleton of S if the forbidden region of s_1, s_2 is empty.

The β -skeleton, like the crust, is a subset of the Delaunay triangulation. Values of β which are either too large or too small require denser sampling, and hence smaller values of r , to guarantee reconstruction. The largest value of r for which we can guarantee reconstruction corresponds to a value of $\beta = 1.70$.

Both the crust and the β -skeleton are very easy to compute, given a good program for the Delaunay triangulation and Voronoi diagram of points in the plane, for which see [Sh], among others. To compute the crust, one computes the Voronoi diagram of S , combines S with the set V of Voronoi vertices to make $S' = S + V$, computes the Delaunay triangulation of S' , and finally selects from the result all those edges whose two endpoints lie in S . For the β -skeleton, one computes the Delaunay triangulation of S and then selects each edge e for which the circumcircles of the adjacent triangles are centered on opposite sides of e and both have radius greater than $\beta/2$ times the

length of e . In either case the running time is bounded by the time required to compute the Voronoi diagram and Delaunay triangulation, which is $O(n \log n)$, for $n = |S|$.

3 Previous Work

Our work draws on a variety of sources. The closest line of research concerns shape recognition for computer vision. The emphasis there is on the closely related problem of estimating the medial axis from a set of boundary points. Brandt and Algazi [BA92] showed that the Delaunay triangulation of a sufficiently dense set of samples contains a reconstruction of the boundary as a subset of its edges (a slightly weaker version of our Theorem 12). Robinson, Colchester, Griffin and Hawkes [RCGH92] propose selecting the boundary reconstruction edges by comparing the length of dual Voronoi and Delaunay edges; our paper essentially gives two equally easy and provably better filtering algorithms. Ogniewicz [O94] studies the computation of an approximate medial axis from a densely sampled boundary, and uses the approximate medial axis to produce successively simpler representations of the boundary. Similar ideas were used by O'Rourke, Booth and Washington [O84], who proposed reconstructing simple closed polygons in the plane from a set of points by choosing a subset of the Delaunay triangulation so as to optimize the approximate medial axis of the resulting polygon.

A successful earlier computational geometric approach to defining the shape of a set of points is the α -shape, introduced by Edelsbrunner, Kirkpatrick and Seidel [EKS83], and studied extensively by Edelsbrunner and others. The α -shape is a simplicial complex defined on a set of points in arbitrary dimension d ; each $k \leq d+1$ points are connected into a $(k-1)$ -simplex if they touch an empty ball of radius $-1/\alpha$. The α -shape tends to work well for sample points which are evenly distributed in the interior of an object, and has proved particularly useful for modeling molecules. But α -shapes are often unsatisfying for reconstructing surfaces; the user needs to find the correct value of the threshold α , and the same α has to apply to the whole data set.

In this paper we continue the study of the β -skeleton, which was defined by Kirkpatrick and Radke [KR88]. Up until now it has been assumed that the parameter β , like α , needs to be found by the user. For our reconstruction problem, we give a value for β which is guaranteed to work when S meets the sampling condition.

The γ -neighborhood graph, introduced by Veltkamp [Vel92], is a generalization of the β -skeleton in which the two forbidden disks may have different radii. We believe that results similar to ours can be proved for a suitably defined family of γ -neighborhood graphs, in which the angle between the two circles at the point of intersection (see Observation 16) is fixed at an optimal value, probably a bit more than $\pi/2$.

We have recently become aware of two concurrent independent research efforts related to ours. Attali [A97] proves that uniformly sampled curves can be reconstructed by (essentially) the above-mentioned family of γ -neighborhood graphs. She requires

the sampling density be everywhere great enough to resolve the finest detail of the curve. Our results are better in that they allow the sampling density to vary along with the level of detail. Melkemi [Mel97] defines an \mathcal{A} -shape on a set S of points as follows: let S' be the union of S with an arbitrary set of points \mathcal{A} . An edge of the Delaunay triangulation of S' belongs to the \mathcal{A} -shape if both of its endpoints belong to S . Our crust is an \mathcal{A} -shape for which \mathcal{A} is the set of Voronoi vertices. \mathcal{A} -shapes for other choices of \mathcal{A} may also have interesting provable properties.

4 The Medial Axis

In this section we review the definition of the *medial axis* [B67], and prove some useful lemmata about it. The medial axis can be thought of as the Voronoi diagram generalized to an infinite set of input points.

Definition: The *medial axis* of a curve F is closure of the set of points in the plane which have two or more closest points in F .

Figure 4 shows the medial axis of a smooth curve. Note that we include components

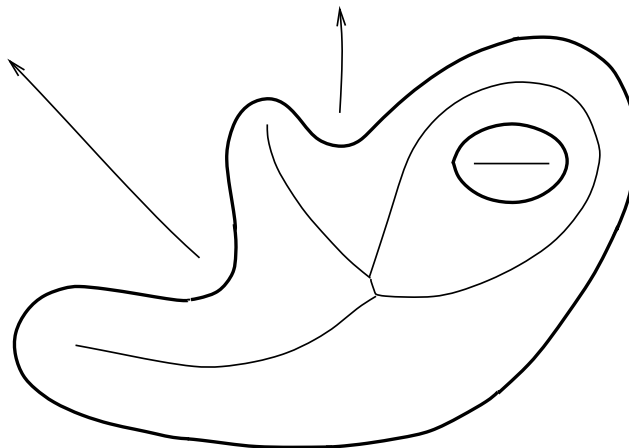


Figure 4: The light curves are the medial axis of the heavy curves.

of the medial axis on either side of the curve, so that some components of the medial axis may extend to infinity. Note also that since we define the medial axis to be a closed set, it includes the centers of all empty *osculating disks* (the empty disks tangent to F with matching curvature), which are its limit points. The medial axis can be defined similarly for a $(d - 1)$ -dimensional surface in \mathbb{R}^d .

Many of our arguments will be based on the following topological Lemma. Note that it concerns two distinct kinds of disks, circular Euclidian disks and topological 1-disks, that is, curve segments.

Lemma 1 *Any (Euclidean) disk containing at least two points of a smooth curve in the plane either intersects the curve in a topological 1-disk or contains a point of the*

medial axis (or both).

Proof: If $B \cap F$ is a topological disk there is nothing to prove, so assume that $B \cap F$ is not a topological disk. If some connected component C of $B \cap F$ is a closed loop, forming a Jordan curve in the interior of B , then there is a connected component of the medial axis interior to C which is entirely contained in B , and we are done.

Otherwise $B \cap F$ consists of two or more connected components of F . Let c be the center of B , and let p be the nearest point on F to c . If p is not unique, then c is a point of the medial axis and we are done. Otherwise p lies in a unique connected component f_p of $B \cap F$. Consider the point q nearest c in some other connected component f_q . Any point x on the line segment cq is closer to q than to any point outside B , so the closest point of F to x is always some point on one of the connected components of $B \cap F$. Since at one end of the segment the nearest connected component is f_p , and at the other it is f_q , at some point x the nearest connected component must change. Point x has two nearest points on two distinct connected components, and so is a point of the medial axis. \square

A *Voronoi disk* of a finite set S of points is a maximal empty disk centered at a Voronoi vertex of S . Each Voronoi disk has at least three points of S on its boundary, and none in its interior.

Lemma 2 *In the plane, any Voronoi disk of a finite set $S \subset F$, where F is a smooth curve, must contain a point of the medial axis of F .*

Proof: Assume first that in the neighborhood of one of the samples $s \in S$ on the boundary of B , $F - s$ is contained completely in B . Then either $B \cap F$ is entirely contained in the boundary of B and the center of B is a point of the medial axis, or shrinking B around s will produce a smaller disk B_ρ , contained in B , with $B_\rho \cap F$ consisting of at least two connected components. By Lemma 1 B_ρ contains a point of the medial axis. If there is no such s , then the intersection of F with B already consists of at least two connected components, and B contains a point of the medial axis by Lemma 1. \square

Note: This lemma does not hold in dimension three; an arbitrarily dense sample S on a smooth surface F can have very small Voronoi balls centered on the surface F itself (or anywhere else), which are very far from the medial axis. Such a ball can be constructed as follows: select a point p on the surface F , $p \notin S$. Construct a small ball B around p , empty of samples, and add four new samples to S on the intersection $B \cap F$. Such examples arise naturally with grid-like sample sets.

5 Sampling

In this section we define our sampling condition. Our condition is based on a *Local Feature Size* function, which in some sense quantifies the local “level of detail” at a point on smooth curve. Local Feature Size functions are used in the computational geometry literature on mesh generation; the term was first used, to the best of our knowledge, by Ruppert [R92] (with a similar definition).

Definition: The *Local Feature Size*, $LFS(p)$, of a point $p \in F$ is the Euclidean distance from p to the closest point m on the medial axis.

The segment of length $LFS(p)$ between a point $p \in F$ and the closest point m on the medial axis of F is perpendicular to the medial axis, not to F , as in figure 5.

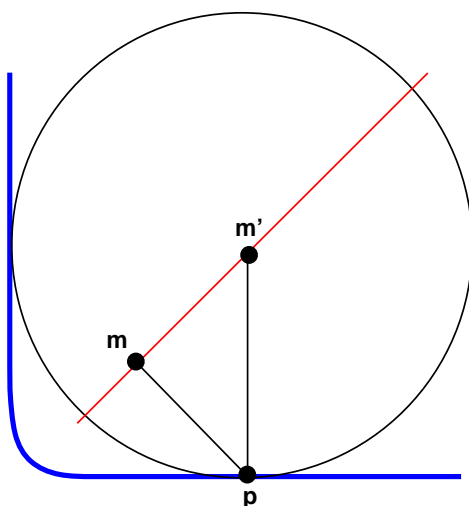


Figure 5: $LFS(p)$ is the distance $d(p, m)$, *not* the perpendicular distance $d(p, m')$ to the center of the largest empty tangent ball at p .

Notice that, because it uses the medial axis, this definition of Local Feature Size depends on both the curvature at p and the proximity of nearby features.

We can now define the sampling condition we will require for curve reconstruction in terms of the *LFS* function.

Definition: F is *r-sampled* by a set of sample points S if every $p \in F$ is within distance $r LFS(p)$ of a sample $s \in S$.

We shall be concerned with values of $r \leq 1$.

Armed with this definition of Local Feature Size, we can clarify the intuition that a small enough disk intersects a curve in a topological 1-disk. The following are corollaries of Lemma 1.

Corollary 3 *A disk containing a point $p \in F$, with diameter at most $LFS(p)$, intersects F in a topological disk.*

Proof: Consider the contrapositive: any disk B containing p that does *not* intersect F in a topological disk contains a point m of the medial axis, by Lemma 1. The nearest point to p on the medial axis is at distance $LFS(p)$ from p , so $d(p, m) \geq LFS(p)$. Since B contains the segment (p, m) , its diameter is greater than $LFS(p)$. \square

Corollary 4 *A disk centered at a point $p \in F$, with radius at most $LFS(p)$, intersects F in a topological disk.*

Proof: Similar to Corollary 3. \square

The following objects were defined by Chew [C93], from whom we borrow the idea of polygonalizing a curve using empty disks centered on the boundary. We take responsibility for the names.

Definition: A *Curve Voronoi Disk* is a maximal disk, empty of sample points, centered at a point of the curve. A *Curve Voronoi Vertex* is the center of a Curve Voronoi Disk.

Note that a *Curve Voronoi Vertex* is the restriction of an edge of the Voronoi diagram of S to the curve F .

Corollary 5 *A Curve Voronoi Disk on an r -sampled smooth curve F , $r \leq 1$, intersects F in a topological disk.*

Proof: Follows from Corollary 4. \square

For large r , it is possible for there to be a set S of points that r -samples two topologically different curves, as in Figure 6. The sample points are placed at the vertices of two regular octagons, positioned so that two adjacent pairs of vertices form a square. The points 1-sample two different curves, one having a single connected component and the other two.

Observation 6 *Let S be a set of points in the plane. There may not be a unique graph on S that is the polygonal reconstruction of a smooth curve r -sampled by S , for $r \geq 1$.*

For considerably smaller r , we shall show that there is only one possible reconstruction, and our graphs find it.

6 Flatness

Considering the definition of the medial axis, and referring back to Figure 5, we observe the following:

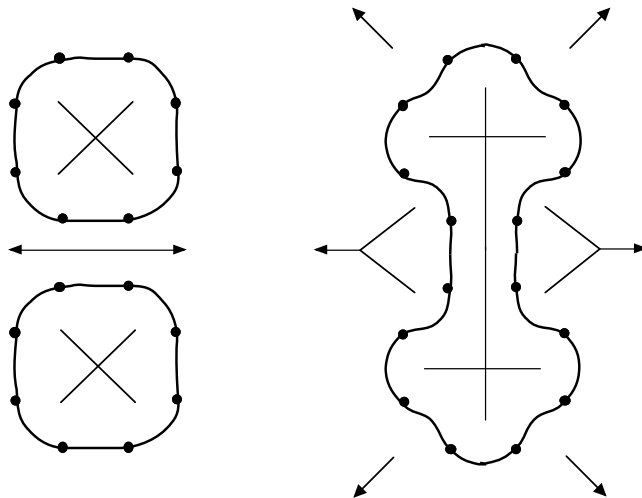


Figure 6: The 16 points 1-sample both heavy curves. The light lines are the medial axes.

Lemma 7 *A disk tangent to a smooth curve F at a point p with radius at most $LFS(p)$ contains no points of F in its interior.*

Proof: The perpendicular distance from p to the point m' on the medial axis that is the center of the largest empty tangent disk at p is at least $LFS(p)$. The tangent disk of radius $LFS(p)$ at p must therefore be contained in the largest tangent disk and hence is also empty. \square

We use this lemma to show quantitatively that the intersection of a smooth curve with a small enough disk is not only a topological disk but also rather flat. The calculations will be based on simple geometric facts about the angles and points labeled in Figure 7. Roughly speaking, we can think of s as a sample and p as an adjacent Curve Voronoi Vertex. Let r be the distance from s to p , and let the distance from s to c , and the distance from p to c , equal one.

Observation 8 *It is easy to verify the following:*

- i. The length of segment (s, x) is $\sin(\gamma)$.*
- ii. $r = d(s, p) = 2 \sin(\gamma/2)$, so $\gamma = 2 \arcsin(r/2)$.*
- iii. The angle $\alpha = \gamma/2 = \arcsin(r/2)$.*
- iv. The angle between the tangent line L at p and the segment (s, p) is $\alpha = \arcsin(r/2)$.*

Lemma 9 *For an r -sampled curve in the plane, $r < 1$, the angle formed at a Curve Voronoi Vertex between two adjacent samples is at least $\pi - 2 \arcsin(r/2)$.*

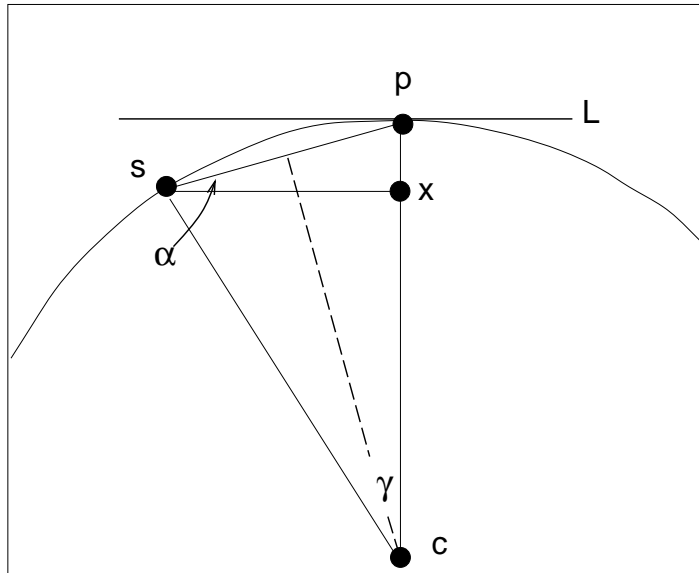


Figure 7: Line L is tangent to the circle at p . $d(c, p) = d(c, s) = 1$ and $d(s, p) = r$.

Proof: Let p , in Figure 7, be the Curve Voronoi Vertex, and let the disk B centered at c be a tangent disk of radius $LFS(p)$, which we assume without loss of generality to be equal to one. The curve F does not intersect the interior of B , so the sharpest angle is achieved when the adjacent sample points lie on the boundary of B at distance r from p , as does s in the figure. The angle formed at p is then $\pi - 2\alpha = \pi - 2 \arcsin(r/2)$ (Observation 8). \square

A very similar argument shows

Lemma 10 *For an r -sampled curve in the plane, $r < 1$, the angle spanned by three adjacent samples is least $\pi - 4 \arcsin(r/2)$.*

7 Polygonal Reconstruction

We now begin our study of curve reconstruction by showing that for a densely r -sampled curve, the Delaunay triangulation of the samples contains, as a subset of its edges, a polygonal reconstruction of the curve.

Lemma 11 *Let F be an r -sampled smooth curve in the plane, $r \leq 1$. There is a Curve Voronoi Disk touching each pair of adjacent samples.*

Proof: Let s_1, s_2 be two samples adjacent along F . The interval of F between s_1 and s_2 crosses the bisector of s_1, s_2 at least once, so let p be one such crossing point. Let B be the maximal disk centered at p which has no sample in its interior. If s_1 and s_2 are on the boundary of B , then B is a Curve Voronoi Disk touching s_1 and s_2 .

Otherwise the maximality of B implies that B touches some third sample s_i . Since p lies between s_1 and s_2 on F , s_i does not lie between s_1 and s_2 on F , p is inside B , s_1 and s_2 are outside B , and B touches s_i , B must intersect F in at least two connected components. In that case B must contain a point of the medial axis, by Lemma 1, and the radius of B is greater than $LFS(p)$, by the definition of Local Feature Size. Since there is no sample within distance $LFS(p)$ of p , this contradicts the assumption that F is r -sampled, with $r \leq 1$. \square

Theorem 12 *Let F be an r -sampled smooth curve in the plane, $r < 1$. The Delaunay triangulation of the set S of samples contains an edge between every adjacent pair of samples.*

Proof: Implied immediately by Lemma 11 and the Empty Circle Property. \square

Note: Brandt and Algazi [BA92] also show that adjacent points on a densely sampled curve are separated by a Voronoi edge (the dual statement of Theorem 12). Let d^* be the minimum, over all points $p \in F$, of $LFS(p)$. Their sampling condition is that every point p must have a sample within distance d^* .

The polygonal reconstruction is close to the curve in the following sense:

Theorem 13 *The distance from a point p on an r -sampled smooth curve F to some point on the polygonal reconstruction of the samples is at most $(r^2/2)LFS(p)$.*

Proof: Let p be the point of F between two samples s_1 and s_2 which is farthest from the reconstruction. Assuming without loss of generality that $LFS(p) = 1$, then the distance from p to the nearer of the two samples, say s_1 , is at most r . Since the curve is smooth and p is maximally distant from the segment (s_1, s_2) , the tangent at p is parallel to (s_1, s_2) . The disk of radius one tangent to the curve at p is empty of sample points, so the maximal distance from p to (s_1, s_2) is achieved when s_1 lies on the surface of the disk, at distance r from p , once again as in Figure 7. The distance $d(p, x)$ there is $r \sin \alpha = r \sin \gamma/2 = r \sin(\arcsin(r/2))$ (Observation 8). \square

Note: The distance from the reconstruction to F is, like the required sampling density, scale invariant; the reconstruction in areas of less detail, which are sampled less densely, can be farther away from the curve. Theorem 13 implies that to obtain a reconstruction that is everywhere within a constant distance d of F , every point p on F should have a sample within distance $\sqrt{2d LFS(p)}$.

In the following sections we give criteria for selecting the edges of the polygonal reconstruction from the Delaunay triangulation.

8 The Crust

We now prove that for small enough r , the crust edges fall exactly between adjacent vertices. First we show that all the desired edges belong to the crust, and then that

no undesired edges do.

Theorem 14 *The crust of an r -sampled smooth curve, $r < 0.40$, contains an edge between every pair of adjacent samples.*

Proof: An edge appears in the crust if and only if there is a circle touching its endpoints which is empty of both sample points and Voronoi vertices. We claim that this is true of every Curve Voronoi Disk on an r -sampled smooth curve. There is a Curve Voronoi Disk touching every pair of adjacent vertices (Lemma 11), so this claim establishes the theorem.

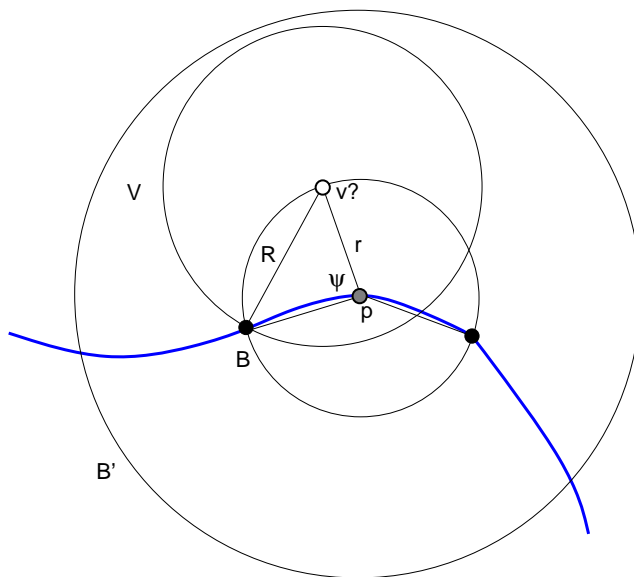


Figure 8: The construction of Theorem 14.

Let B be a Curve Voronoi Disk centered at p . By definition, B cannot contain a sample point. To see that B cannot contain a Voronoi vertex, consider Figure 8. The point v is a Voronoi vertex which, we assume for the purpose of contradiction, falls within B . We assume, once again without loss of generality, that $LFS(p) = 1$.

Since v is a Voronoi vertex, the radius R of the Voronoi circle V around v is at most the distance to the nearer of the two samples inducing p . This Voronoi circle must contain a point of the medial axis (Lemma 2). On the other hand, the disk B' of radius $LFS(p) = 1$ around p cannot contain a point of the medial axis, by the definition of Local Feature Size.

We now choose r so that V lies entirely within B' , establishing the contradiction. Any point in V is at most distance $r + R$ from p , and R is maximized when v lies on the boundary of B . In this case R is length of the base of an isosceles triangle whose other two edges have length r . Since the curve is pretty flat at p (Lemma 9), the angle ψ at p opposite the base is at most $\frac{1}{2}(\pi + 2 \arcsin(r/2))$, and $R/2 \leq r \sin \psi/2$. So we want

$$r + 2r \sin\left(\frac{\pi + 2 \arcsin(r/2)}{4}\right) \leq 1$$

The parenthetical quantity is less than $\pi/2$ for r in the interval $[0, 1]$, so the left-hand side is increasing in that interval. Choosing $r \leq 0.40$ satisfies the inequality. \square

Theorem 15 *The crust of an r -sampled smooth curve does not contain any edge between non-adjacent vertices, for $r < 0.252$.*

Proof: We need to show that there is no circle, empty of both Voronoi and sample points, touching any two non-adjacent sample points s and t . We assume, for the purpose of contradiction, that there is such a circle B , as in Figure 9.

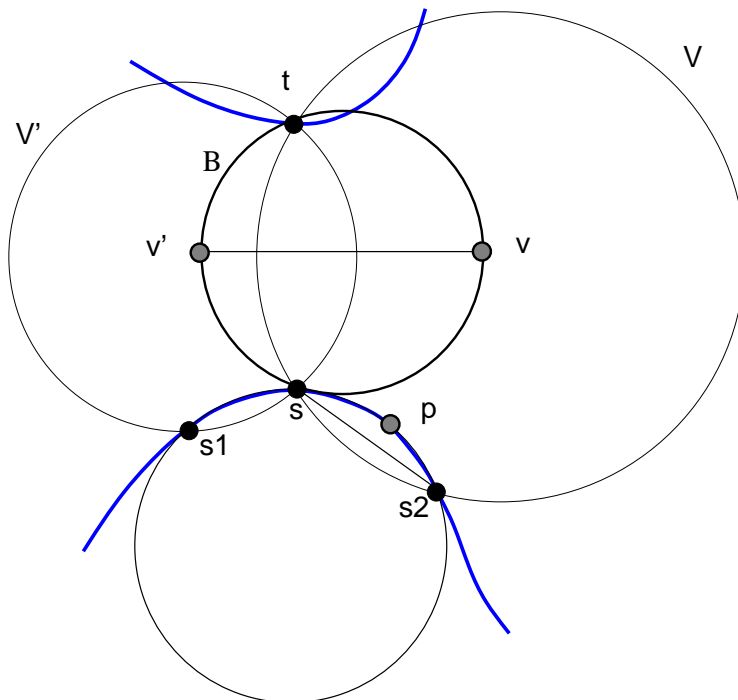


Figure 9: The construction of the contradiction in Lemma 15.

Consider the two circles V, V' touching s and t and centered at the points v, v' at which B intersects the perpendicular bisector of s and t .

We claim that if B is empty of Voronoi points, then V and V' are empty of sample points. For if one of them were nonempty, it would contain a sample s' determining a minimal circle touching s, t and s' , empty of all other samples and hence inducing a Voronoi vertex inside of B .

Consider for a moment Figure 10, which is a closeup of the situation at s . The angle ω between the tangents to the circles V, V' at s is equal to $\pi/2$ (since the lower half-circle of B , containing s , is the locus of points which form a right angle with v and v' , and the tangents are perpendicular to (v, s) and (v', s)).

The angle $\angle(s_1, s, s_2)$ is at least $\pi - 4 \arcsin(r/2)$ (Lemma 10). Without loss of generality let V be the circle such that the angle ψ between the tangent to V at s

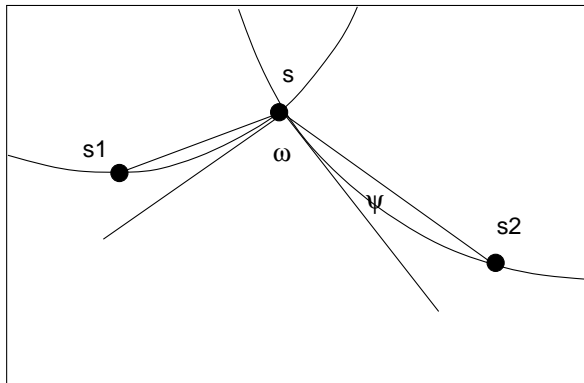


Figure 10: The angle ψ between the tangent and the chord is greater than the corresponding angle on the other side.

and the chord (s, s_1) or (s, s_2) is greater than the corresponding angle on the other side. Then $\psi \geq 1/2(\pi - 4 \arcsin(r/2) - \pi/2) = \pi/4 - 2 \arcsin(r/2)$. If we assume, once again with no loss of generality, that the radius of V is equal to one, this bound on ψ implies (Observation 8) that $d(s, s_2) \geq 2 \sin(\pi/4 - 2 \arcsin(r/2))$.

There is a Curve Voronoi Vertex p between s and s_2 (Lemma 11), and hence, since the curve is r -sampled, $\sin(\pi/4 - 2 \arcsin(r/2)) \leq r LFS(p)$.

We now give an upper bound for $LFS(p)$, so as to derive a contradiction. The samples s and t are on two different connected components of the intersection $F \cap V$, so V contains a point of the medial axis (Lemma 1). The point p lies in V , which has radius one, so $LFS(p) \leq 2$. Thus

$$\sin(\pi/4 - 2 \arcsin(r/2)) \leq 2r$$

The right hand side is increasing in r , while the left hand side is decreasing in r in the range $[0, 1]$. Choosing $0 < r \leq 0.252$ violates the inequality, producing a contradiction. \square

9 The β -Skeleton

In this section, we show that with an appropriately chosen value of β , the β -skeleton of the samples on an r -sampled smooth curve forms a polygonal reconstruction of the curve. To simplify our calculations, we shall sometimes define the forbidden region of an edge in terms of the angle between the two forbidden circles, rather than the length of the edge (see Figure 11).

Observation 16 *Let s_1, s_2 be a pair of points in the plane, let $\beta \geq 1$, and let $\phi = \arcsin 1/\beta$. The tangents to the two disks of radius $d(s_1, s_2)\beta/2$ touching s_1 and s_2 form an angle of 2ϕ at s_1 and s_2 .*

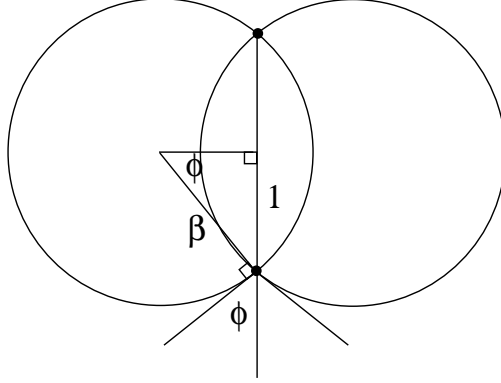


Figure 11: Forbidden regions can be defined by β or ϕ

Lemma 17 *Let s_1, s_2, s_3 be three successive samples on an r -sampled smooth curve. When $\phi > 4 \arcsin(r/2)$, s_3 cannot lie in the forbidden region of the edge (s_1, s_2) .*

Proof: If we choose ϕ so that the angle $\angle(s_1, s_2, s_3) > \pi - \phi$, then s_3 cannot lie in the forbidden region of (s_1, s_2) . Since the curve is r -sampled, $\angle(s_1, s_2, s_3)$ is at least $\pi - 4 \arcsin(r/2)$ (Lemma 10). \square

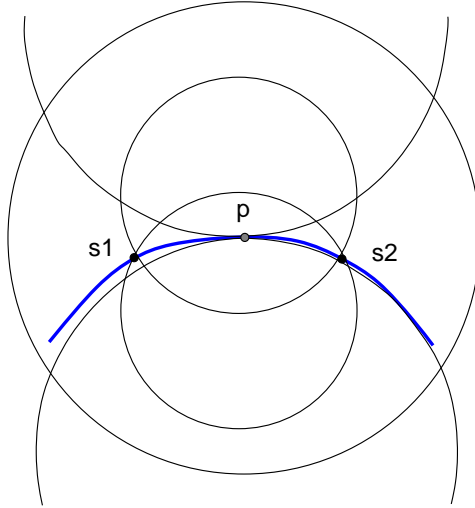


Figure 12: The construction of Lemma 18

Lemma 18 *The forbidden region of an edge between two adjacent samples on an r -sampled smooth curve cannot contain a point of the medial axis, when $\phi > \arcsin(2 \sin(2 \arcsin(r/2)))$.*

Proof: Let s_1 and s_2 be adjacent samples. Let p be a Curve Voronoi Vertex between s_1 and s_2 (Lemma 11). We assume without loss of generality that $LFS(p) = 1$.

We begin by choosing ϕ so that the radius R of the circles defining the forbidden region of (s_1, s_2) is at most $1/2$. $R = d(s_1, s_2)/(2 \sin \phi)$. And, since $LFS(p) = 1$, $d(s_1, s_2) \leq 2 \sin(2 \arcsin(r/2))$ (Observation 8), so we choose $\phi > \arcsin(2 \sin(2 \arcsin(r/2)))$.

The samples s_1 and s_2 must lie outside the interior of the two tangent circles of radius one at p , so there is a circle of radius at least one, touching p , s_1 and s_2 . Since $R \leq 1$, and the forbidden disks also touch s_1 and s_2 , p must lie in the interior of both of the forbidden disks, as in Figure 12. Since $R \leq 1/2$, the forbidden region lies entirely within the circle of radius one around p , which by the definition of the Local Feature Size does not contain a point of the medial axis. \square

Lemma 19 *The β -skeleton of an r -sampled smooth curve does not contain an edge between any pair of non-adjacent samples, when $\phi < \arccos(2r) - 2\arcsin(r/2)$.*

Proof: The proof of this theorem is similar to that of Theorem 15, so this presentation is somewhat sketchy. We let $s, t \in S$ be two samples, not adjacent on F , and replace the circles V and V' in that proof with the forbidden circles B and B' of the potential edge (s, t) , as in Figure 13.

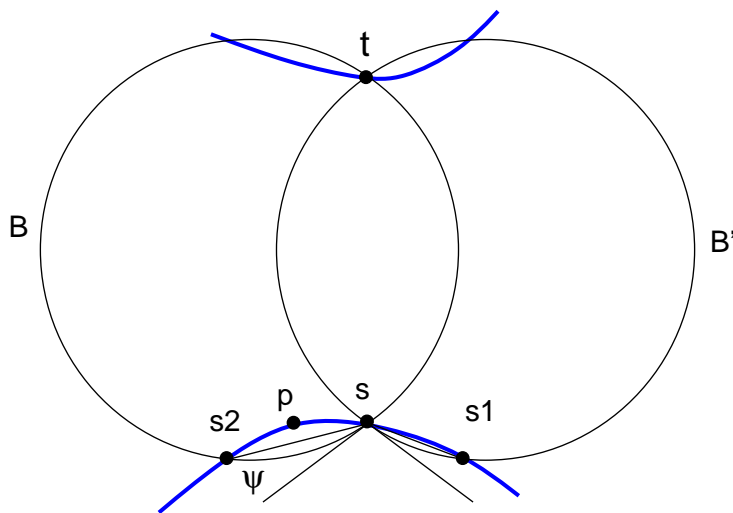


Figure 13: The forbidden region is drawn horizontally

The angle between the tangents of B and B' at s is 2ϕ . The angle $\angle(s_1, s, s_2)$ is at least $\pi - 4\arcsin(r/2)$ (Lemma 10). On one side of s , without loss of generality the side of s_2 , the angle ψ between the chord (s, s_2) and the tangent to B at s is at least $\pi/2 - 2\arcsin(r/2) - \phi$. Assuming without loss of generality that the radius of B is equal to one, we find (Observation 8) that the length $d(s, s_2) \geq 2\sin(\pi/2 - 2\arcsin(r/2) - \phi) = 2\cos(2\arcsin(r/2) + \phi)$, and that at the Surface Voronoi Vertex p between s and s_2 , $r LFS(p) \geq \cos(2\arcsin(r/2) + \phi)$.

Again, since B has radius one and intersects the curve F in two connected components, $LFS(p) \leq 2$ (Lemma 1 and the definition of $LFS(p)$) so we have:

$$\cos(2\arcsin(r/2) + \phi) \leq 2r$$

To produce a contradiction, we choose ϕ so as to violate this inequality:

$$\phi < \arccos(2r) - 2\arcsin(r/2)$$

□

Theorem 20 *Let S r -sample a smooth curve, with $r \leq 0.297$. The β -skeleton of S contains exactly the edges between adjacent vertices, for $\beta = 1.70$.*

Proof: Lemma 19 established that the β -skeleton contains no edges between non-adjacent vertices for

$$\phi < \arccos(2r) - 2 \arcsin(r/2) \quad (1)$$

Let s_1, s_2 be two adjacent samples, and let s_0 and s_3 be the other samples adjacent to s_1, s_2 respectively. There would fail to be an edge between s_1 and s_2 if some third sample fell into the forbidden region. Lemma 17 implies that neither s_0 nor s_3 can lie in the forbidden region when

$$\phi > 4 \arcsin(r/2) \quad (2)$$

If some other sample s_i lay in the forbidden region, but s_0 and s_3 did not, that would imply that one of the forbidden circles intersects F in at least two connected components (one containing s_1 and the other containing s_i), and hence must contain a point of the medial axis (Lemma 1). This cannot occur when

$$\phi > \arcsin 2 \sin 2 \arcsin(r/2) \quad (3)$$

(Lemma 18). These three functions are plotted in Figure 14. All three inequalities are satisfied in the shaded region. There is a feasible choice of ϕ for any $r < 0.297$.

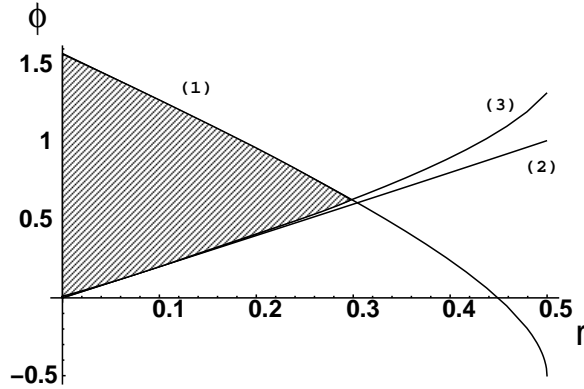


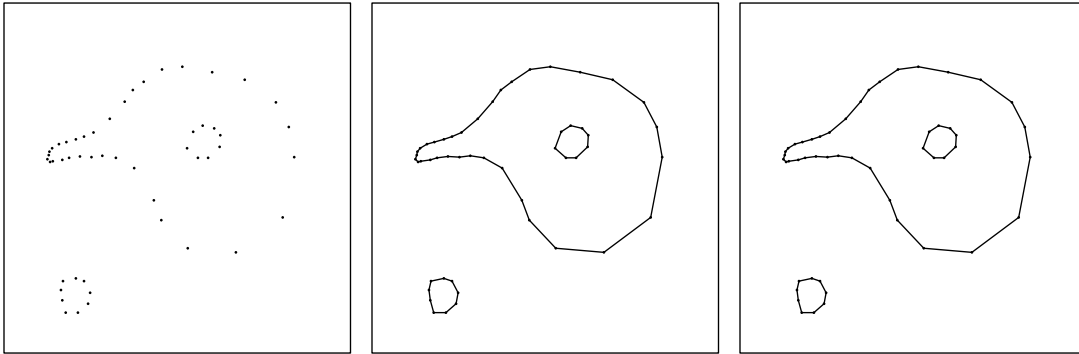
Figure 14: The three functions from Theorem 18, plotted in Mathematica and annotated with idraw.

The value of ϕ which allows the sparsest sampling, maximizing r , is roughly $\phi = 0.637$, which corresponds to $\beta = 1.70$. □

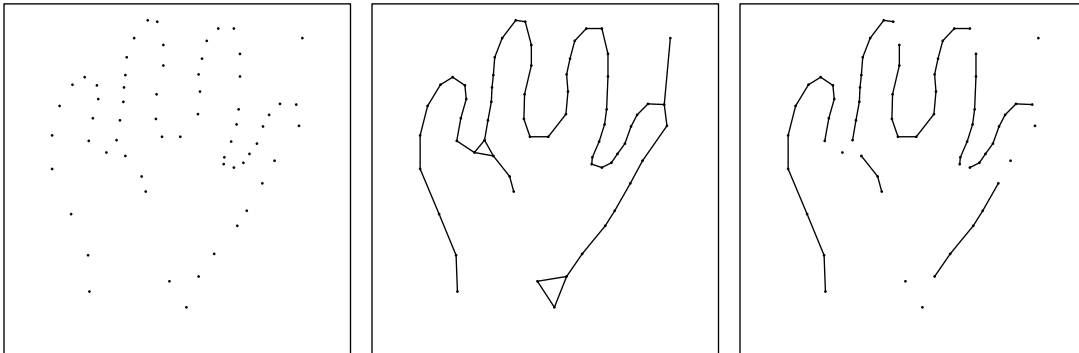
Note that this value for β is not likely to be the optimal one; it is just the value corresponding to the largest r for which our somewhat crude bounds allow us to prove that an r -sampling yields a correct reconstruction.

10 Implementation and Examples

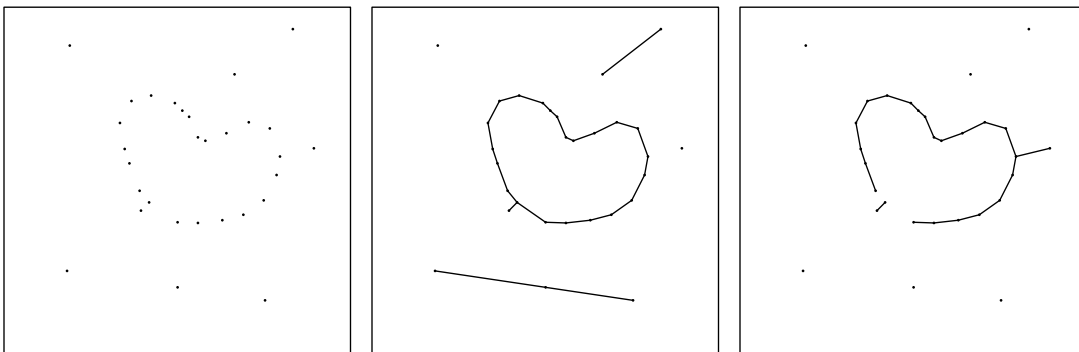
We implemented the two reconstruction algorithms using the Delaunay triangulation and Voronoi diagram programs in Shewchuk's `Triangle` package [Sh], generating data with an interactive Java front end. Here are some examples of point sets, with their crusts in the center and their β -skeletons on the right.



The crust and the β -skeleton are identical on any point set S which is a 0.253-sample of a smooth curve F , since they both produce a correct reconstruction. They are often identical in practice for larger values of r ; in the example above, $r > 1/2$.



On more sparsely sampled curves, the crust is usually more liberal in adding edges. Note that β -skeletons and crusts can contain vertices of degree one or degree three. Vertices of degree four or greater cannot occur in crusts, while the maximum degree in a β -skeleton depends on the choice of β .



A curve can be reconstructed fairly well in the presence of sparse added noise. Notice the unusual occurrence, in this last example, of an edge in the β -skeleton which is not in the crust.

11 Conclusion and Open Questions

We can summarize our main results as follows. Let S be an r -sample from a smooth curve F . For $r \leq 1$, the Delaunay triangulation of S contains the polygonal reconstruction of F . For $r \leq 0.40$, the crust of S contains the polygonal reconstruction of F . For $r \leq 0.279$, the β -skeleton of S is the polygonal reconstruction of F , and for $r \leq 0.252$ the crust of S is the polygonal reconstruction of F .

The minimum required sample density that we can show for the β -skeleton is somewhat better than the density that we can show for the crust. The crust tends to err on the side of adding edges, which can be useful. But the β -skeleton could be biased towards adding edges, at the cost of increasing the required sampling density, by tuning the parameter β .

The main open question is the polygonal reconstruction of two-dimensional surfaces in \mathbb{R}^3 . This is an important problem in graphics, and a series of Siggraph papers have presented effective practical algorithms [HDDMS92], [TL94], [BBX95], [CL97]. Neither of our planar graphs gives a polygonal reconstruction when generalized to \mathbb{R}^3 in a straightforward way, although it seems possible that either idea could be elaborated into a working algorithm.

Many questions remain about two-dimensional reconstruction. There should be results on the quality of the reconstruction of curves with branches and endpoints. There are probably versions of our theorems that do not require smoothness, but only that any angles be bounded away from zero by a function of r . It should be possible to prove something about the quality of the reconstruction in the presence of small errors in sample positions and of additive noise.

Better lower bounds would also be interesting. None of our constants are tight, and they are far from the lower bound $r \leq 1$ of Observation 6. The comparison is not really fair here, since our graphs also reconstruct some curves with branches and endpoints. An algorithm that produced only reconstructions of smooth closed curves could perhaps get by with a larger value of r .

The work in [RCGH92], [BA92], [O94], dealt with the polygonal analog of the medial axis, consisting of those edges of the Voronoi diagram of S whose dual Delaunay edges do not belong to the polygonal reconstruction of the boundary; see Figure 15. One can think of this graph as the *anti-crust*. Our bounds on the quality of the polygonal reconstruction of the boundary should imply something about the quality of the anti-crust as a reconstruction of the medial axis.

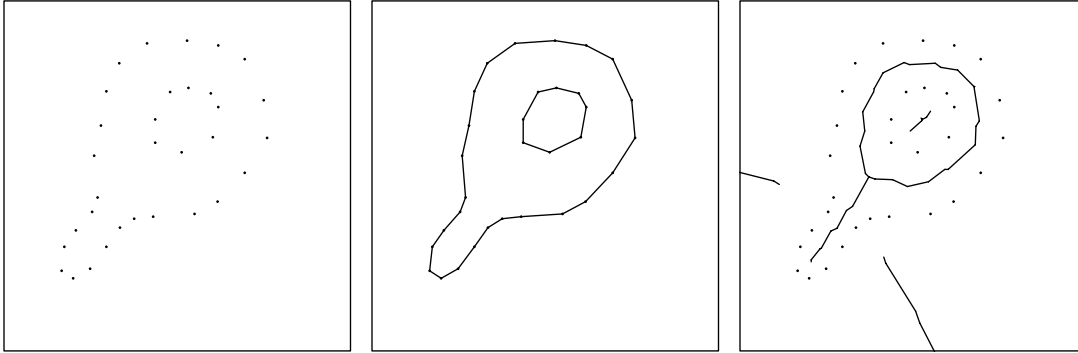


Figure 15: A point set, its crust, and the corresponding polygonal analog of the medial axis.

Frequently piecewise-linear reconstruction is only a step towards smooth reconstruction. Since the LFS gives an upper bound on the curvature, it should be possible to reconstruct F with spline rather than line segments in such a way as to improve Theorem 13. This might give a near-minimal representation of F which does not sacrifice any important features.

References

- [A97] Attali, Dominique. R -regular shape reconstruction from unorganized points, *Proceedings of the ACM Symposium on Computational Geometry*, (1997), pp. 248-253.
- [BBX95] Bajaj, C.L., Bernadini, F., and Xu, G., Automatic reconstruction of surfaces and scalar fields from 3D scans, *SIGGRAPH '95 Proceedings*, (1995), pp. 109-118.
- [B67] Blum, H., A transformation for extracting new descriptors of shape, *Models for the Perception of Speech and Visual Form* (W. Walthen-Dunn, ed.) MIT Press, (1967) pp. 362-380.
- [BA92] Brandt, J. and Algazi V.R., Continuous skeleton computation by Voronoi diagram, *Computer Vision, Graphics and Image Processing* **55** (3), (1992), pp. 329-338.
- [C93] Chew, L.P., Guaranteed-quality mesh generation for curved surfaces, *Proceedings of the ACM Symposium on Computational Geometry*, (1993), pp. 274-280.
- [CL97] Curless, B. and Levoy, M. A volumetric method for building complex models from range images, *Siggraph '96*, (1997) pp. 303-312.

- [E87] Edelsbrunner H., *Algorithms in Combinatorial Geometry*, Springer Verlag, (1997).
- [EKS83] Edelsbrunner, H., Kirkpatrick D.G., and Seidel, R., On the shape of a set of points in the plane, *IEEE Transactions on Information Theory* **29** (4) (1983), pp. 551-559.
- [HDDMS92] Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., and Stuetzle, W., Surface reconstruction from unorganized points, *SIGGRAPH '92 Proceedings*, (1992), pp. 71-78.
- [KR88] Kirkpatrick, D.G., and Radke J.D., A framework for computational morphology, *Computational Geometry* (G. Toussaint, ed.), North-Holland (1998), pp. 217-248.
- [O94] Ogniewicz, R.L., Skeleton-space: A multiscale shape description combining region and boundary information, *Proceedings of Computer Vision and Pattern Recognition*, (1994), pp. 746-751.
- [Mel97] Melkemi, Mahmoud, \mathcal{A} -shapes and their derivatives, *Proceedings of the ACM Symposium on Computational Geometry*, (1997), pp. 367-369.
- [O94] Joseph O'Rourke, *Computational Geometry in \mathbb{C}* , (1994), Cambridge University Press.
- [O84] O'Rourke, J., Booth, H. and Washington, R., Connect-the-dots: a new heuristic, *Computer Vision, Graphics and Image Processing*, **39**, (1984), pp. 258-266.
- [PS85] Preparata, F. and Shamos, M.I., *Computational Geometry: An Introduction*, Springer Verlag, (1995).
- [RCGH92] Robinson, G.P., Colchester, A.C.F., Griffin, L.D., and Hawkes, D.J., Integrated skeleton and boundary shape representation for medical image interpretation, *Proceedings of the European Conference on Computer Vision*, (1992), pp. 725-729.
- [R92] Ruppert, J., A new and simple algorithm for quality two-dimensional mesh generation, *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, (1993), pp. 83-92.
- [Sh] Shewchuk, J.R., Triangle. <http://www.cs.cmu.edu/~quake/triangle.html>.
- [TL94] Turk, G., and Levoy, M., Zippered polygon meshes from range images, *SIGGRAPH '94 Proceedings*, (1994), pp. 311-318.
- [Vel92] Veltkamp, R.C., The γ -neighborhood graph, *Computational Geometry: Theory and Applications* **1:4**, (1992), pp. 227-246.