The Power Crust, Unions of Balls, and the Medial Axis Transform

Nina Amenta Sunghee Choi Ravi Krishna Kolluri

March 2, 2001

Abstract

The medial axis transform (or MAT) is a representation of an object as an infinite union of balls. We consider approximating the MAT of a three-dimensional object, and its complement, with a finite union of balls. Using this approximate MAT we define a new piecewise-linear approximation to the object surface, which we call the *power crust*.

We assume that we are given as input a sufficiently dense sample of points from the object surface. We select a subset of the Voronoi balls of the sample, the *polar balls*, as the union of balls representation. We bound the geometric error of the union, and of the corresponding power crust, and show that both representations are topologically correct as well. Thus, our results provide a new algorithm for surface reconstruction from sample points. By construction, the power crust is always the boundary of a polyhedral solid, so we avoid the polygonization, hole-filling or manifold extraction steps used in previous algorithms.

The union of balls representation and the power crust have corresponding piecewise-linear dual representations, which in some sense approximate the medial axis. We show a geometric relationship between these duals and the medial axis by proving that, as the sampling density goes to infinity, the set of *poles*, the centers of the polar balls, converges to the medial axis

1 Introduction

The input to the *surface reconstruction problem* is a set S of sample points from the surface W of a three-dimensional object, and the output should be a piecewise-linear approximation of W. Surface reconstruction arises in a variety of contexts, and it has recently become important in computer graphics because of the development of laser range scanners and other technologies for collecting sets of sample points from the surfaces of real objects.

Our approach to surface reconstruction, in a nutshell, is first to use the sample points to approximate the *medial axis transform* (or *MAT*) of the object, and then to produce the piecewise-linear surface approximation from the approximate MAT. See Figure 1 for a two-dimensional example.

The MAT is a representation of the object as the infinite union of its maximal internal balls. As our approximation, we use the *polar balls*, a subset of the Voronoi balls of S. The polar balls belong to two sets, one more or less filling up the inside of the object, and the other the outside. These two sets approximate the MAT of the object, and the MAT of its complement, respectively. When the sample S is sufficiently dense it is easy to distinguish the inner from the outer poles; Section 9 contains our algorithm. In a subsequent paper, we will describe additional heuristics and a very robust implementation.

Our main innovation lies in the following algorithm for converting these unions of balls into a surface representation. We compute a weighted Voronoi diagram, the *power diagram*, of the polar balls. The power diagram divides space into polyhedral cells, each cell consisting of the points in $\mathbb{R}^3$ closest to a particular ball, under a convenient distance function, the *power distance*. The boundary separating the cells belonging to inner polar balls from the cells belonging to outer polar balls is a piecewise-linear surface, which is our output, the *power crust*.

We also use the power diagram to define the adjacencies of the polar ball centers (the *poles*). Subsets of inner (resp. outer) poles whose power diagram cells share a face are connected with a dual weighted Delaunay face. These faces form a simplicial complex, the *power shape*, analogous to the medial axis.

We prove a variety of bounds on the quality of our approximations, under

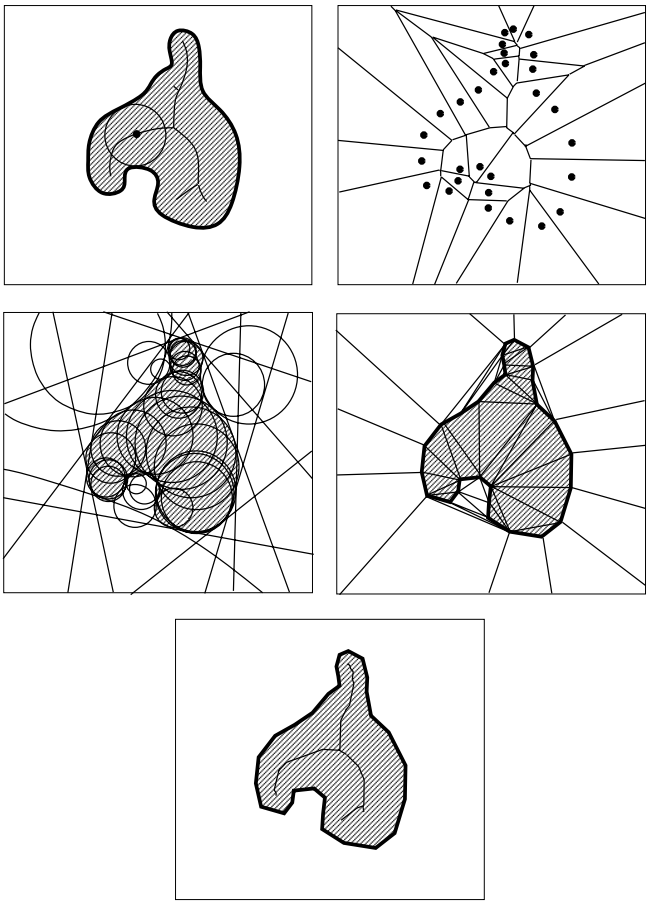

Figure 1: Two-dimensional example of power crust construction. Upper left, an object (shaded) with its medial axis; one maximal interior ball is shown. The medial axis is the union of the centers of these maximal interior balls. Upper right, the Voronoi diagram of a sample of points from the object boundary. In two dimensions, we select all Voronoi vertices as poles; in three dimensions we select only certain ones near the medial axis. Middle left, the sets of inner (shaded) and outer polar balls. Outer polar balls with centers at infinity degenerate to halfspaces on the convex hull. Middle right, the power diagram cells of the poles. In two dimensions this is the same as the Delaunay triangulation of the samples, but not in three dimensions. Bottom, the power crust and the inner portion of the power shape.

the assumption that the input sample is sufficiently dense. We show that the power crust, and the surfaces of the two unions of polar balls, inner and outer, are all close to the surface of the original object, that their surface normals are close, and that they interpolate the samples. These geometric bounds allow us to show not only that the power crust is homeomorphic to the original object surface, but also that it's interior is homeomorphic to the solid object itself. This in turn implies that the power shape, like the medial axis, is homotopy equivalent to the original object. We characterize the geometric accuracy of the power shape by showing that the set of poles converges to the medial axis as the sampling density goes to infinity.

2 Related work

Computer graphics

The clean abstraction of the problem of reconstruction from unorganized points was introduced to the computer graphics community by Hoppe et al. [23]. They proposed an algorithm which locally estimates the *signed distance function*, the function on $\mathbb{R}^3$ which returns the distance from the closest point on the surface; the distance is negative at interior points of the object. They use as an estimate the distance to the closest point in the input sample. The output surface is a polygonization of the zero set of the estimated signed distance function.

Curless and Levoy [14] gave a really effective algorithm which represents the distance function on a voxel grid. To save space, they store only the part of the grid near the input sample. This allows them to handle very large and noisy data sets, so that their algorithm can be applied to combinations of many laser range scans. Because they only approximate part of the distance function, they need a post-processing step for hole-filling. They used silhouettes of the object to constrain hole-filling, which works well except in indentations, where there might be data but no silhouettes; our algorithm, which returns a solid interpolating all of the data, even within indentations, has an advantage here.

Like the algorithms above, our algorithm can be described in terms of the

signed distance function. The medial axis sketches the "ridges" of the signed distance function, the points at which the direction to the closest surface point changes discontinuously. Thus estimating the medial axis is a way of representing the signed distance function on the entire space, in about the same amount of storage as the input data itself.

Computational geometry

The surface reconstruction problem has received a lot of recent attention in the computational geometry community. There have been several algorithms for reconstructing curves [4, 6, 15, 19, 21] including algorithms which handle plane curves with boundaries [16] and curves with sharp corners [1, 20], and an algorithm for space curves with a strong topological guarantee [27] similar to ours. In three dimensions, Amenta and Bern [3] gave an algorithm which selects a subset of the Delaunay triangles of S as the output surface. They defined a sampling condition (which we use, see Section 5), under which they proved that their output surface is close to that of the original object. They also defined the poles, which are at the heart of our algorithm. Their algorithm selects a set of candidate triangles from the Delaunay triangulation, and then selects the output manifold from the candidate set. This manifold extraction step fails when the sampling condition is not met, a serious drawback in practice. Amenta, Choi, Dey and Leekha [2] gave a similar, but simpler algorithm, with a much simpler proof, and also showed that the output surface is homeomorphic to the original object surface. They describe a manifold extraction heuristic which seems to work well. Boissonnat and Cazals [10] avoid the manifold extraction difficulty by proposing an algorithm which reconstructs a smooth surface interpolating the sample points. As part of their theoretical analysis they independently proved a version of the theorem (see Section 8) that the set of poles converges to the medial axis as the sampling density goes to infinity. Computing the smooth surface is time consuming compared to the Voronoi diagram computation.

A key feature that differentiates our algorithm is that in addition to being simple and providing theoretical guarantees, we also guarantee that our output is the boundary of a three-dimensional solid, *irrespective of the sampling density*. This not only avoids the manifold extraction problem, but makes the algorithm quite robust in practice.

Our algorithm is perhaps most similar to an old algorithm of Boissonnat [9], which labels a subset of the Delaunay tetrahedra of the input sample as the interior of the solid. We avoid difficulties in labeling by using the power diagram instead of the Delaunay tetrahedra.

Another algorithm based on Delaunay triangulation is the α-shape algorithm of Edelsbrunner and Mücke [18]. This algorithm selects candidate Delaunay triangles based on the radius of their smallest empty circumspheres. We use many of the beautiful ideas developed in the context of α-shapes, although in a different way. In particular, the relationship between power diagrams and unions of balls was developed by Edelsbrunner [17], and the power shape is almost, but not exactly, the same as the weighted α-shape of the polar balls.

Bernardini et al. [8] have also given an algorithm based, conceptually, on α-shapes, while *avoiding* the computation of the Delaunay triangulation. This allows them to apply the algorithm to very large data sets. The Delaunay triangulation is the expensive step in the construction of the power crust as well. It would be very interesting to find a power crust algorithm which similarly avoids computing the Delaunay triangulation.

Medial axis approximation

Another distinguishing feature of our algorithm is that it generates a discrete approximation of the MAT, the power shape, which is a useful alternative representation of the object. Applications for the MAT have been proposed in a wide variety of contexts, but particularly in three dimensions it has generally failed to lead to practical algorithms. One problem is that the MAT is hard to compute exactly. The computation of the exact medial axis for simple polyhedra has been demonstrated only recently [13]. For more complicated shapes, approximation probably continues to be more appropriate. Attali and Montanvert [7] and others [26] have proposed approximating the medial axis using the Voronoi diagram. This approach is sometimes justified by a reference to [22], which argues, incorrectly, that the set of three-dimensional Voronoi vertices converges to the true medial axis as the sampling density goes to infinity. Since the set of poles *does* converge to the medial axis, we believe that the power shape is a better MAT approximation.

The set of interior polar balls is a good approximation of the object as a

union of balls, which is a also a useful shape approximation. Hubbard [24] promotes the use of unions of balls for collision detection, guided by the observation that detecting the intersection of two balls is much easier than detecting intersections of two other primitives like triangles or polyhedra. He constructs a hierarchical representation, using increasingly simple unions of balls, and gives convincing experimental evidence that this hierarchy is more efficient in practice than others. Hubbard's experience shows that the success of the approach depends on the quality of the shape approximation. He finds that the set of Voronoi balls is superior to a larger and less accurate set of balls derived from a quad-tree; we believe that the set of polar balls should be better still.

Finite unions of balls or discrete medial axis transforms have also been proposed as a representation for deformable objects. Rajan and Fournier [25] use a union of balls for interpolating between shapes. Teichman and Teller [28] use a discrete medial axis as a skeleton in a semi-automatic system for animating arbitrary computer models. Both papers again begin with the set of Voronoi balls and use a heuristic clean-up phase, and again, we believe that the polar balls would be a better starting point. Cheng, Edelsbrunner, Fu and Lam [11] do morphing in two dimensions with *skin surfaces*, which are smooth surfaces based on unions of balls. Our work can be seen as a step toward converting an arbitrary polygonal surface into a provably accurate skin surface.

3 Geometry

In this section we formally introduce the geometric structures we will use, and describe some of their known, although perhaps not widely known, properties.

Surfaces and balls

All our ideas are based on the relationships between surfaces and balls.

Let W be the closed, bounded two-dimensional surface of an object $\mathcal{W}$ in $\mathbb{R}^3$. To avoid having to deal with points at infinity, we assume that surface $\mathcal{W}$

is contained in an open, bounded region $\mathcal{Q}$. W divides $\mathcal{Q}$ into two bounded open solids, the *inside* and the *outside* of $\mathcal{W}$. Hence, we allow both the inside and the outside to be disconnected. For our theoretical arguments, we will assume that W is not only closed but smooth, by which we mean C^1-continuous.

A (Euclidean) ball $B_{c,\rho}$ has a center c and radius ρ. In the context of power diagrams, a ball is often equivalently represented by a *weighted point* with position c and weight ρ^2. We will also need the concept of a point c with negative weight $-\rho^2$, equivalent to a ball $B_{c,i\rho}$ (with $i = \sqrt{-1}$). A point with weight zero (ie a ball with radius zero) is *unweighted*.

Medial axis transform

We say ball $B = B_{c,\rho} \subset \mathcal{Q}$ is *empty* (with respect to W) if the interior of B contains no point of W. A *medial ball* is a maximal empty ball; that is, it is completely contained in no other empty ball. The center of a medial ball is either a point with more than one closest point on W, or a center of curvature of W.

Definition: The *medial axis transform* of surface W is the set of medial balls. The set of centers of the medial balls form the *medial axis* M of W.

We could equivalently define the medial axis as the closure of the set of all points with more than one closest point of W. Notice that either way the medial axis includes both a part inside of W (the *inner medial axis*) and a part outside of W (the *outer medial axis*). Also note that since we define the medial axis as a locus of centers of balls contained in $\mathcal{Q}$, the medial axis is always bounded.

Barring degeneracies, the medial axis of a two dimensional surface in a region $\mathcal{Q}$ is another two dimensional surface. The medial axis is usually not a manifold, but if W is piecewise smooth and in general position its medial axis will consist of piecewise smooth two-dimensional patches, with singular edges and vertices contained in two and three patches respectively.

The medial axis is *homotopy equivalent* to the complement of W, $\mathcal{Q} - W$ [12]. This is a way of saying that the two solid shapes, $\mathcal{W}$ and its complement, and the medial axis have the same holes, tunnels and connected components,

even though they generally differ in dimension. This can be shown by giving a continuous *deformation retraction* of $\mathcal{Q} - W$ onto M, defining a continuous motion which moves every point away from its closest surface point [12].

Power diagrams

The duality just described between the surface W and its medial axis M is akin to the relationship between a finite union of balls and its α-shape, which in turn is related to *power diagrams*, a species of weighted Voronoi diagram. Since we will use power diagrams extensively we review them in some detail.

Definition: The *power distance* between two weighted points c_1, ρ_1^2 and c_2, ρ_2^2 is $d^2(c_1, c_2) - \rho_1^2 - \rho_2^2$.

Note that either ρ_1 or ρ_2 might be imaginary. Equivalently,

Definition: The *power distance* between two balls $B_1 = B_{c_1,\rho_1}$ and $B_2 = B_{c_2,\rho_2}$ is $d_{pow}(B_1, B_2) = d^2(c_1, c_2) - \rho_1^2 - \rho_2^2$.

For example, let $B = B_{c,\rho}$ be a ball with real weight ρ, and let x be a point with weight zero (equivalently, a ball of radius zero). If x is on the boundary of B, then $d_{pow}(B, x) = 0$; if x is inside B, then $d_{pow}(B, x) < 0$ and if x is outside of B then $d_{pow}(B, x) > 0$.

The motivation behind the definition of power distance is that computing the induced weighted Voronoi diagrams is easy.

Definition: The *power diagram* $Pow(\mathcal{B})$ of a set of balls is the weighted Voronoi diagram which assigns an (unweighted) point $x = B_{x,0}$ in space to the cell of the ball B which minimizes $d_{pow}(B, x)$.

Very conveniently, programs which compute the (unweighted) d-dimensional Voronoi diagram by computing a convex hull in dimension $(d + 1)$ (the standard approach in dimensions three and higher) can be easily converted to compute power diagrams as well.

The two-dimensional faces separating the cells of a three-dimensional power diagram - the sets of points in space with two "closest" samples - are subsets of two-dimensional planes. If the balls corresponding to the two weighted points determining a face intersect, then the face is a subset of the plane

containing the circle in which the boundaries of the two balls intersect.

Just as the regular unweighted Voronoi diagram defines the Delaunay triangulation, the power diagram defines a dual *weighted Delaunay triangulation*, sometimes called a *regular triangulation* (since not all triangulations of a set S of samples are regular.)

Definition: A face f of the *weighted Delaunay triangulation* $WDT(\mathcal{B})$ of a set of weighted points (equivilantly, a set of balls) is the simplex formed by the convex hull of the set $\mathcal{B}_h$ of weighted points inducing a face h of $Pow(\mathcal{B})$. We say that faces f and h are *duals*.

In the usual unweighted Voronoi diagram, the maximal empty ball centered at an arbitrary point x is incident to the samples which induce the face of the Voronoi diagram containing x. Similarly in the weighted case we can describe the power diagram face containing a point x using a ball centered at x, as follows. We begin with a definition due to Edelsbrunner [17].

Definition: Two weighted points are *orthogonal* if the power distance between them is zero.

We can give a geometric interpretation of orthogonality.

Definition: The boundaries of two balls B_1 and B_2 intersect in a circle C. We say that B_1 and B_2 *meet at angle* α, where α is the angle between their tangent planes at any point on C; equivalently, $\alpha = \pi - \beta$, where β is the angle between the normal vectors to B_1 and B_2 at any point of C. Considering Figure 2 and using the Pythagorean theorem, we see that two orthogonal *positive* balls meet at an angle of $\pi/2$, which implies that the center of either ball is outside the other. When a ball with negative weight $B_1 = B_{c_1, i\rho_1}$ is orthogonal to a positive ball $B_2 = B_{c_2,\rho_2}$, c_1 lies inside B_2 and the two balls intersect in a great circle of B_1. Two negatively weighted points cannot be orthogonal.

Observation 1 *Let x be a point, and let $\mathcal{B}$ be a set of balls, and let ball B be element of $\mathcal{B}$ which minimizes $d_{pow}(B, x)$. Consider the ball $B_{x,w}$ orthogonal to B, and let $B_{x,w'}$ be the ball orthogonal to any other ball $B' \in \mathcal{B}$. Then $w \leq w'$.*

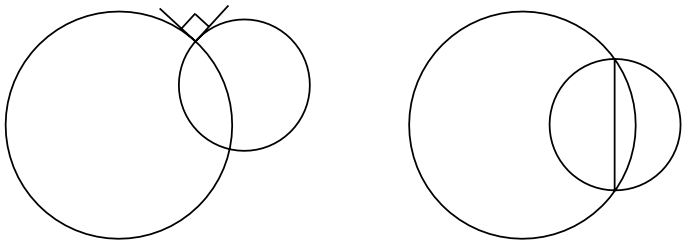

Figure 2: On the left, two orthogonal positively weighted balls meet at a right angle. On the right, the larger ball is positively weighted and the smaller is negatively weighted.

Proof: We know that $d_{pow}(B, x)) \leq d_{pow}(B', x)$, and

$$0 = d_{pow}(B, B_{x,w}) = d_{pow}(B, x) - w^2$$

so that

$$0 = d_{pow}(B', x) - (w')^2 \leq d_{pow}(B', x) - w^2$$

□

For example, a point x on a two-dimensional face of $Pow(\mathcal{B})$, with positive power distance to the two balls B_1, B_2 inducing the face, will be the center of a ball $B_{x,w}$ which intersects both B_1 and B_2 orthogonally, and meets every other ball corresponding to a point in $\mathcal{B}$ at an angle of less than $\pi/2$.

The following Lemma may be obvious to experts. It highlights the role of the power diagram in the well-known duality between the Voronoi diagram and the Delaunay triangulation. In order to handle infinite Voronoi faces cleanly, we assume that the Voronoi diagram includes a vertex at infinity, to which all infinite Voronoi edges and faces are adjacent.

Lemma 2 *Let S be a set of points in general position in $\mathbb{R}^3$ and let $\mathcal{B}$ be the set of Voronoi balls centered at the Voronoi vertices of S. $Pow(\mathcal{B})$ is the Delaunay triangulation $DT(S)$.*

Proof: Each point $p \in S$ lies on the boundary of each ball corresponding to a vertex of the Voronoi cell of p, and outside of all the other Voronoi balls.

Consider a k-face f of $DT(S)$, let $\mathcal{V}_f$ be the set of Voronoi vertices corresponding to d-simplicies in $DT(S)$ containing f, and let $\mathcal{B}_f$ be the set of their Voronoi balls. Face f is the convex hull of a set S_f of $(k+1)$ points of S. The points in S_f all lie on the boundaries of the Voronoi balls in $\mathcal{B}_f$. Each point $x \in S_f$ thus has $d_{pow}(B_i, x) = 0$, for all $B_i \in \mathcal{B}_f$, and cannot lie in the interior of any other balls in $\mathcal{B}$, and hence belongs to the face f' of $Pow(\mathcal{B})$ determined by $\mathcal{B}_f$ (which must exist). This face f' is convex, and is a subset of the affine hull of S_f, so that $f \subseteq f'$. Since this is true of every face f in $DT(S)$, we have $DT(S) = Pow(\mathcal{B})$.
□

4 Our constructions

We now define our construction of the union of balls representation and the power crust. Let S be a sufficiently dense sample of points from a smooth surface W; we shall define "sufficiently dense" in the following section. Again, to avoid dealing with infinity, we add a set Z of eight points, the vertices of a large box surrounding W, to S, so that all the Voronoi vertices of each sample in S are finite points. Amenta and Bern [3] made the following definition.

Definition: The *poles* p_1, p_2 of a sample $s \in S$ are the two vertices of its Voronoi cell farthest from s, one on either side of the surface. The Voronoi balls $B_{p_1,\rho_1}, B_{p_2,\rho_2}$ are *polar balls*, with $\rho_i = d(p_i, s)$.

Amenta and Bern [3] show that both poles of s are found correctly by the following procedure, assuming that S is sufficiently dense in the technical sense described below. Select the Voronoi vertex of s farthest from s as the first pole p_1. From among those Voronoi vertices v of s such that the angle $\angle vsp_1 > \pi/2$, select the farthest as the second pole p_2.

The intuition behind the definition of poles is that the polar balls approximate medial balls. Let $\mathcal{P}$ be the set of poles. The surface W divides the set of poles into the set $\mathcal{P}_I$ of *inside poles* and the set $\mathcal{P}_O$ of *outside poles*. The corresponding sets of polar balls are $\mathcal{B}_I$ and $\mathcal{B}_O$.

Definition: Let $\mathcal{U}_I = \bigcup \mathcal{B}_I$ be the union of Voronoi balls centered at inside

poles, and $\mathcal{U}_O = \bigcup \mathcal{B}_O$ be the union of Voronoi balls centered at outside poles. Let $U_I = \partial \mathcal{U}_I$ and $U_O = \partial \mathcal{U}_O$ be the boundaries of these unions.

Observation 3 *Every sample $s \in S$ lies on both U_I and U_O.*

We will show in Section 6 that both U_I and U_O form good approximations of W when S is sufficiently dense.

Now consider the power diagram $Pow(\mathcal{B}_I \cup \mathcal{B}_O)$. Some cells of this power diagram belong to balls in $\mathcal{B}_I$, and others to balls in $\mathcal{B}_O$. (Unlike the power diagram of a general set of balls, every input ball ends up with a cell in $Pow(\mathcal{B}_I \cup \mathcal{B}_O)$.) The collection of two-dimensional faces induced by one inside and one outside polar ball separate the part of the domain $\mathcal{Q}$ belonging to the inside balls from the part belonging to outside balls.

Definition: The *power crust* of S is the set of faces in $Pow(\mathcal{B}_I \cup \mathcal{B}_O)$ separating cells belonging to inside polar balls from cells belonging to outside polar balls.

Observation 4 *Every sample lies on the power crust.*

Observation 5 *The power crust is the (possibly non-regular) boundary of a three-dimensional solid.*

We show that the power crust is also a good approximation of W, in Section 7.

Dual shapes

Both the union of balls and the power crust have *dual shapes*, skeletal representations by simplicial complexes. The dual shapes can be considered discrete analogs to the medial axis. Some partial geometric results can be found in Section 8.

Edelsbrunner [17] defined the *dual shape* of a union of balls (also known as the *weighted α-shape*), for which he demonstrated an elegant correspondance with the structure of the union.

Definition: The dual shape of a union of balls $\mathcal{U} = \cup\mathcal{B}$ is a simplicial complex. The centers of a subset $\mathcal{B}_f \subseteq \mathcal{B}$ are connected by a simplex whenever the power cells of all balls in $\mathcal{B}_f$ have a point x in common, such that $x \in \mathcal{U}$.

Edelsbrunner proved the following [17], establishing a topological analogy between the dual shape and the medial axis.

Theorem 6 *The dual shape of a union of balls is homotopy equivalent to the union.*

He in fact defines a deformation retraction, very similar to that which establishes that the medial axis is homotopy equivalent to the object [12].

We make a similar definition of the dual shape of a subset of power diagram faces (eg. the power crust). (We abuse notation by writing $Pow(\mathcal{B})$ for the set of faces of $Pow(\mathcal{B})$.)

Definition: Let $\mathcal{B}$ be a set of balls, and let Y be a set of closed faces selected from $Pow(\mathcal{B})$. The *dual shape* of Y is the union of the dual faces of of every face in $Pow(\mathcal{B}) - Y$.

Definition: The *power shape* of S is the dual shape of the power crust.

Note that the dual shape is the dual, in the standard computational geometry definition, of the *complement* of the power crust, not of the power crust itself.

Again using techniques borrowed from Edelsbrunner, we show that when Y is a surface selected from the 2-faces of $Pow(\mathcal{B})$, the dual shape of Y is analgous to the medial axis in the following sense.

Theorem 7 *Let Y be a set of $(d-1)$-dimensional faces, together with all their subfaces, selected from $Pow(\mathcal{B})$. The dual shape is homotopy equivalent to $\bigcup(Pow(\mathcal{B}) - Y)$.*

Proof: The d-dimensional cells of $Pow(\mathcal{B}) - Y$ form a family of convex sets. The *nerve* of a family of convex sets is a simplicial complex, with a vertex for every convex set and a simplex connecting every subset of convex sets which have a common intersection. The Nerve Theorem states that the nerve of a family of convex sets is homotopy equivalent to their union. The dual shape is a geometric realization of the nerve of the d-dimensional cells in

$Pow(\mathcal{B}) - Y$, since any tuple of weighted points (that is, ball centers) which induce a face in the power diagram either induces a face of Y, in which case the corresponding convex sets fail to intersect, or a face of the dual shape, in which case the corresponding convex sets do intersect.
□

5 Sampling condition

Before we get into the proofs that the unions of polar balls and the power crust are geometrically accurate, we need to define what we mean by a "sufficiently dense" sample S. We use the following definitions and lemmata from recent papers on surface reconstruction [3],[4].

Definition: The *Local Feature Size* at a point $w \in W$, written $LFS(w)$, is the distance from w to the nearest point of the medial axis of W.

Intuitively, LFS is small where two parts of the surface pass close together, since they are separated by the medial axis. The medial axis is also close to the surface where the curvature is high. We use the LFS function to define the sampling density we require to produce a good surface reconstruction.

Definition: $S \subseteq W$ is an *r-sample* if the distance from any point $w \in W$ to its closest sample in S is at most a constant fraction r times $LFS(w)$.

Sampling assumption: *We assume that S is a r-sample from W and $r \leq 0.1$.*

The usefulness of this assumption depends on LFS being well behaved. The following lemma says that the LFS function is Lipschitz.

Lemma 8 (Amenta and Bern [3]) *For any two points p and q on W, $|LFS(p) - LFS(q)| \leq d(p,q)$.*

Observation 9 *If d(u,s)= O(r)LFS(u) then d(u,s)=O(r)LFS(s) as well, for $r < 1$.*

The following lemma is a Lipschitz condition on the surface normal with respect to LFS.

Lemma 10 (Amenta and Bern [3]) *For any two points p and q on W with $d(p,q) \leq \rho \min\{LFS(p), LFS(q)\}$, for any $\rho < 1/3$, the angle between the normals to W at p and q is at most $\rho/(1-3\rho)$.*

We need to state one more key lemma, which will be useful in our proofs later on. Informally, the idea is that when S is sufficiently dense, the Voronoi cell of every sample $s \in S$ is long and skinny and roughly perpendicular to the surface. The way we quantify this is to say that, given a sample s and a point v in its Voronoi region, the angle between the vector from s to v and the surface normal at s has to be small (linear in r) when v is far away from s (as a function of LFS).

For convenience, we define $r' = r/(1-r) = O(r)$.

Lemma 11 (Amenta and Bern [3]) *Let s be a sample point from an r-sample S. Let v be any point in $Vor(s)$ such that $d(v,s) \geq \kappa LFS(s)$ for $\kappa > r'$. Let α be the angle between the vector $s\vec{v}$ and the surface normal $\vec{n}$ at s. Then $\alpha \leq \arcsin(r'/\kappa) + \arcsin r'$.*

Conversely, if the angle is large, then point v has to be close to s. Specifically, if $\alpha \geq \arcsin(r'/\kappa) + \arcsin r'$, then $d(v,s) \leq \kappa LFS(s)$. Rearranging things, we get

Corollary 12 *For any v such that $\alpha > \arcsin r'$, we have $d(v,s) \leq \kappa LFS(s)$ with*

$$\kappa = \frac{r'}{\sin(\alpha - \arcsin r')}$$

The Voronoi cell of a sample $s \in W$ must contain the point m of the inside (outside) medial axis for which s is a closest surface point. Since m is at least distance $LFS(s)$ from s, while the inside (outside) pole p of s is at least as far away, the angle between the vector to m and the vector to p is at most $2\arcsin r'$ by Lemma 11. The polar ball B_p centered at p is at least as large as the medial ball centered at m, so that m has to fall inside B_p, whenever $2\arcsin r' < \pi/3$.

Corollary 13 *Every polar ball contains a point of the medial axis, when $r < 1/3$.*

6 Unions of polar balls

We will now show that, under the sampling assumption, first, that the boundary U_I of the inner polar balls and the boundary U_O of the union of the outer outer polar balls, are both close to W. Second, we will establish that their surface normals agree with those of W, and third, that both of them are homeomorphic to W.

Shallow intersections

First, an observation, illustrated by Figure 14 below.

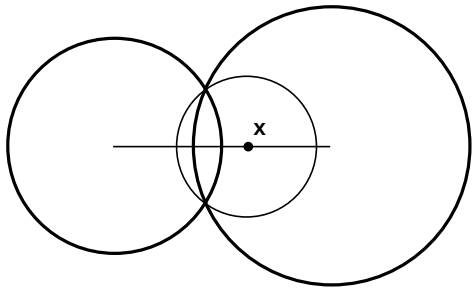

Figure 3:

Observation 14 *Let B_I and B_O be two intersecting balls, and let x be a point on the segment connecting them. Any ball centered at x and containing a point outside of both B_I and B_O also completely contains $B_I \cap B_O$.*

The main idea in all the proofs is that inside and outside balls cannot intersect each other deeply. We say this in three different ways in the lemmata below. We measure the depth of the intersection by the angle α at which the balls intersect, as in Figure 4, below.

The first version of the lemma deals with the special case in which the two balls are the inner and outer polar balls of the same sample s, for which we can get the best bound.

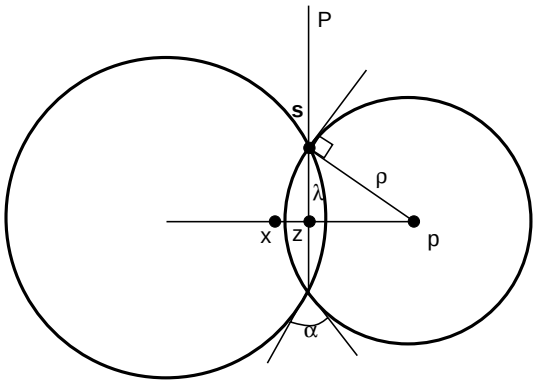

Figure 4: An inside and outside ball can intersect only at a small angle α.

Lemma 15 *The two polar balls of a sample s intersect at an angle of $O(r)LFS(s)/\rho$, where ρ is the radius of the smaller polar ball.*

Proof: Without loss of generality let the inner polar ball $B_{p,\rho}$ be smaller than the outer polar ball B_O. The line segment between p and the center of B_O intersects the surface in at least one point x. Since $B_{p,\rho}$ and B_O cannot contain samples, s is the nearest sample to x (Observation 14) and $d(x,s) \leq rLFS(x)$.

Let z be the center of the circle C in which the boundaries of B_I and B_O intersect, and let λ be the radius of C, as in Figure 4. We have $\lambda \leq d(x,s)$, and so, using Observation 9, $\lambda \leq O(r)LFS(s)$. The angle between P and the tangent plane to $B_{p,\rho}$ at s is the same as $\angle zps = \arcsin(O(r)LFS(s)/\rho)$. Since $LFS(s) \leq \rho$, for small enough r this is $O(r)LFS(s)/\rho$. The angle between P and the tangent plane to B_O is no greater, so $\alpha = O(r)LFS(s)/\rho$.
□

Now we show that in the general case, any pair consisting of an inner and an outer polar ball must intersect shallowly.

Lemma 16 *Let B_I be an inside polar ball and B_O be an outside polar ball. B_I and B_O intersect at an angle of at most $2\arcsin 3r = O(r)$.*

Proof: Consider the line segment connecting c_I and c_O, the centers of B_I and B_O. Since c_I and c_O lie on opposite sides of W, this segment crosses W in at least one point x.

Let $B_{c,\rho}$ be the smaller of the two balls B_I and B_O. If $x \in B_{c,\rho}$, we have $LFS(x) \leq 2\rho$, since the polar ball $B_{c,\rho}$ contains a point of the medial axis (Corollary 13).

Otherwise x is in the larger of the two balls, but not in the smaller, as in Figure 4. Let c be the center of the smaller ball, and again define z and λ as in Figure 4. By Corollary 13, we have $LFS(x) \leq d(x,c)+\rho = d(x,z)+d(z,c)+\rho$. But the distance from x to the nearest sample is at least

$$\sqrt{\lambda^2 + d^2(x,z)} = \sqrt{\rho^2 - d^2(z,c) + d^2(x,z)}$$

So the r-sampling requirement means that

$$\sqrt{\rho^2 - d^2(z,c) + d^2(x,z)} \leq r[\rho + d(x,z) + d(z,c)]$$

Since $d(z,c) \leq \rho$, we can simplify to

$$d(x,z) \leq 2r'\rho$$

which, for $r \leq 1/3$, means that x is very close to $B_{c,\rho}$, and $LFS(x) \leq 3\rho$.

Since the distance from x to the nearest sample is at least λ and at most $3r\rho$, we know that $\lambda \leq 3r\rho$. The angle between the plane P containing C and a tangent plane on $B_{c,\rho}$ at any point on C is thus at most $\arcsin 3r$, the angle between the plane P containing C and the tangent plane of the larger ball is smaller, and the two balls meet at an angle of at most $2\arcsin 3r$.
□

The third lemma shows that a similar fact holds when one of the balls is a medial, rather than a polar, ball.

Lemma 17 *Let B_p be an inside (outside) polar ball and let B_m be an outside (inside) medial ball. The angle at which B_p and B_m intersect is at most $2\arcsin 2r = O(r)$.*

Proof: Again we consider the line segment connecting p and m, the centers of B_p and B_m, which crosses W in at least one point x, which is in B_p but not in B_m (since the interior of any medial ball is empty of points of the surface).

We have $LFS(x) \leq 2\rho_p$, since B_p contains a point of the medial axis. When $2\rho_p \leq \rho_m$, we use this bound to show that the balls intersect at an angle of at most $2 \arcsin 2r$, as in the proof of Lemma 16.

Otherwise, since m itself is a point of the medial axis, we have $LFS(x) \leq d(x,m) = d(x,z) + d(z,m)$. Again, the distance from x to the nearest sample is at least

$$a = \sqrt{\lambda^2 + d^2(x,z)} = \sqrt{\rho_m^2 - d^2(z,m) + d^2(x,z)}$$

So the r-sampling requirement means that

$$\sqrt{\rho_m^2 - d^2(z,m) + d^2(x,z)} \leq r[d(x,z) + d(z,m)]$$

Since $d(z,c) \leq \rho_m$, we can simplify to

$$(1-r)d(x,z) \leq r\rho_m$$

which, for $r \leq 1/2$, means that $LFS(x) \leq 2\rho_m$. We use this bound to show that the angle between the two balls is most $2 \arcsin 2r$, again as in Lemma 16.
□

Proximity

We now turn to the proof that the union boundaries U_I and U_O approximate the surface W. We can immediately infer from Lemma 16 that the surface W cannot penetrate too far into the interior of either union, as a function of the radii of the balls forming the unions. We extend this to a stronger bound in terms of LFS, which could be much smaller than the radius of either medial ball at a surface point x.

Lemma 18 *Let u be a point in the Voronoi cell of s but not in the interior of either polar ball at s. The distance from u to s is $O(r)LFS(s)$, for small enough r.*

Proof: We assume without loss of generality that $LFS(s) = 1$. Let p_1 be the pole farther from s. If $\angle usp_1 \leq \pi/2$, we let $p = p_1$, otherwise we consider $p = p_2$, the pole nearer to s. We let $B_{p,\rho}$ be the polar ball centered at p. In either case $d(u, s) \leq \rho$, because of the way in which the poles were chosen. Let θ be the angle between vectors $\vec{su}$ and $\vec{sp}$. Since u is outside the polar ball, $d(s, u) \geq 2\rho \cos\theta$.

Since $d(s, u) \leq \rho$, we have $\theta \geq \frac{\pi}{3} > 3 \arcsin r'$, for small enough values of r. Let $\vec{n}$ represent the normal at s. We find $\angle \vec{n}\ \vec{sp} < 2 \arcsin r'$ by Lemma 11. So $\angle \vec{n}\ \vec{su} > \pi/3 - 2 \arcsin r' > \arcsin r'$. From Corollary 12 it follows that, for any point u in the Voronoi cell of s,

$$d(u, s) \leq \frac{r'}{(\sin(\theta - 3 \arcsin r'))}$$

Since $\theta \geq \frac{\pi}{3}$, the angle, $(\theta - 3 \arcsin r') \geq \frac{\pi}{6}$, again for small enough values of r. Thus $d(u, s) \leq 2r'$. Since we assumed $LFS(s) = 1$, the lemma follows. □

Corollary 19 *Any point u which does not lie in the interior of either $\mathcal{U}_I$ or $\mathcal{U}_O$ is within distance $O(r)LFS(s)$ of its closest sample s.*

It remains to bound the distance from any point on the boundary of one union and in the interior of the other to the surface.

Lemma 20 *For a point u contained in both $\mathcal{U}_I$ and $\mathcal{U}_O$, the distance to the closest sample s is $O(r)LFS(s)$.*

Proof: Point u is contained in an inner ball B_I and an outer ball B_O. The line joining the centers of B_O and B_I intersects the surface at some point x. Let s_x be the closest sample to x and let s be the closest sample to u; see Figure 6. A ball centered at x, and with radius $d(x, s_x)$, must also contain u (Observation 14). This and the r-sampling condition give a bound on $d(x, u)$.

$$d(x, u) \leq d(x, s_x) = O(r)LFS(x)$$

Hence

$$d(u, s) \leq d(u, s_x) \leq d(u, x) + d(x, s_x) = O(r)LFS(x)$$

By Observation 9, $d(u,s) = O(r)LFS(s)$.
□

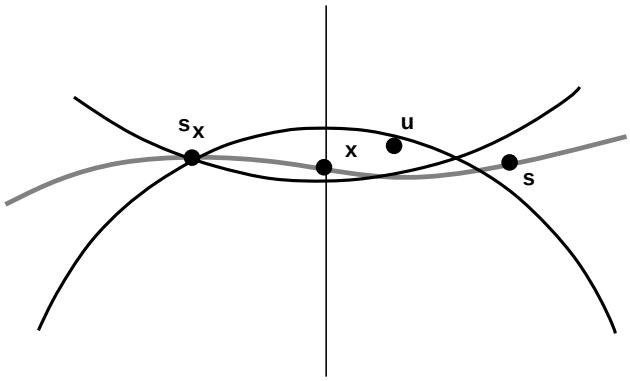

Figure 5: The point u is closer to x than s_x, which is outside both the polar balls.

We use the two lemmata above to show that the two union boundaries $U_I = \partial\mathcal{U}_I$ and $U_O = \partial\mathcal{U}_O$ have to be close to the surface.

Theorem 21 *The distance from a point $u \in U_I$ or $u \in U_O$ to its closest point on the surface $x \in W$ is $O(r)LFS(x)$.*

Proof: Let s be the closest sample to u. Assume without loss of generality that u is on the boundary U_I. The either $u \in \mathcal{U}_I$ and $u \in \mathcal{U}_O$, so that $d(u,s) = O(r)LFS(s)$ by Lemma 20, or u is in the interior of neither $\mathcal{U}_I$ or $\mathcal{U}_O$, so that $d(u,s) = O(r)LFS(s)$ by Corollary 19. The point x is at least as close to u as s is, and hence $d(x,u) = O(r)LFS(s)$ and $d(x,s) = O(r)LFS(s)$. The result follows from Observation 9.
□

Lemmata 18 and 20 imply that most of the domain $\mathcal{Q}$ lies in either the union of inner balls or the union of outer balls, and only points very near the surface W might lie in both unions, or in neither.

Definition: The *tubular neighborhood* around surface W is the set of points

within distance $O(r)LFS(x)$ of a point $x \in W$. Figure 6 illustrates the tubular neighborhood.

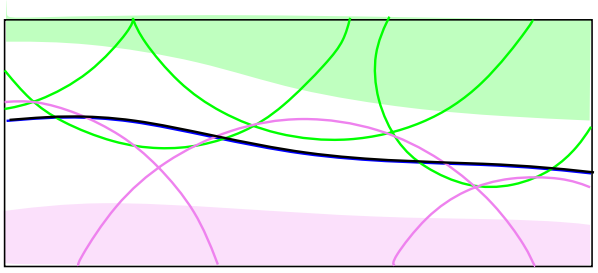

Figure 6: The boundaries of the unions of balls $\mathcal{U}_I$ and $\mathcal{U}_O$ must lie close to the surface W. Specifically, the boundaries are contained in the *tubular neighborhood*, defined as the set u of points such that the distance from u to the closest point $x \in W$ is at most $O(r)$ times the distance from x to the medial axis.

Normals

Now we show that the normals on the union boundaries U_I and U_O are also close to the normals of nearby points of the surface W, approaching the correct normal at a rate proportional to $\sqrt{r}$ as $r \to 0$.

Observation 22 *Let $B = B_{c,\rho}$ be a polar ball, at distance at most k from a point $x \in W$. Then $\rho \geq \frac{LFS(x)-k}{2}$.*

This follows because B is a polar ball, so it contains a point of the medial axis, by Corollary 13, while the nearest point of the medial axis to x is at distance $LFS(x)$.

Lemma 23 *Let u be a point such that the distance to the nearest surface point $x \in W$ is at most $O(r)LFS(x)$. Let $B_{c,\rho}$ be an inner (resp. outer) polar ball containing u. Then the angle, in radians, between the inner (resp. outer) surface normal at x and the vector $\vec{uc}$ is $O(\sqrt{r})$.*

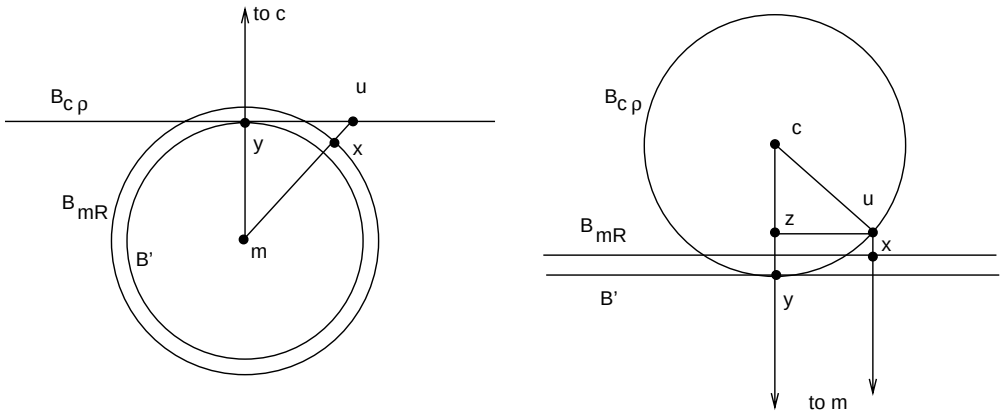

Figure 7: Since B cannot intersect B_M very deeply, and $d(u,x)$ has to be small, the angles at m and c cannot be very large.

Proof: Let $B_{m,R}$ be the medial ball at x on the opposite side of the surface from c. Since x is the nearest surface point to u, the vector $\vec{xu}$ is normal to the surface at x, and m, x and u are collinear. So we can write the angle we are interested in as $\alpha = \angle ucm + \angle umc$. We begin by bounding $\angle umc$. Without loss of generality, assume $LFS(x) = 1$.

Since $B_{c,\rho}$ and $B_{m,R}$ cannot intersect at x at an angle greater than $2 \arcsin 2r$ (Lemma 17), the thickness of the lune in which they intersect is at most a factor of $O(r^2)$ times the smaller of the two radii. Let B' be the ball centered at m and touching this lune, as in Figure 7.

Angle $\beta = \angle cmu$ will depend on the ratio of the two radii R and ρ. The following argument establishes that β is maximized when $R = 1$ and $\rho = \infty$, as in Figure 7 on the left. Since β decreases as u moves towards the center c, we assume u is on the boundary of $B_{c,\rho}$. For any fixed ρ, increasing R makes β smaller, so we assume $R = LFS(x) = 1$, its minimum value since $B_{m,R}$ is a medial ball at x. For any fixed R, increasing ρ makes β larger, so we assume that $B_{c,\rho}$ is infinitely large.

Since in this situation $B_{m,R}$ is the smaller ball, the radius of B' is $R(1-O(r^2))$. Let y be the point at which segment c, m intersects B'. The distance

$$d(u,y) = \sqrt{d^2(m,y) - d^2(m,u)} = O(\sqrt{(1+O(r))^2 - (1-O(r^2))^2}) = O(\sqrt{r})$$

We get $\angle ucm = \arcsin\left(\frac{d(u,y)}{d(m,u)}\right) = O(\sqrt{r})$.

We use a similar argument to bound $\gamma = \angle ucm$. Again we can assume that u is on the boundary of $B_{c,\rho}$. For any fixed ρ, increasing R increases γ, and for any fixed R, increasing ρ decreases γ, so in contrast to the previous situation, we let ρ take on its minimum value of $(1-d(u,x))/2 = \theta(1)$ (Observation 22), and let R become infinitely large. This worst case is shown on the right in Figure 7. Here $d(z,y)$ is at most the thickness of the lune added to $d(u,x)$, that is, $O(r^2) + O(r) = O(r)$. Distance $d(x,u) = \sqrt{d(c,u)^2 - d(c,z)^2} = O\left(\sqrt{1-(1-r)^2}\right) = O(\sqrt{r})$. This finally gives us $\angle ucm = O(\sqrt{r})$, completing the $O(\sqrt{r})$ bound on α.

□

Theorem 24 *Let u be a point on U_I (resp. U_O), and let $x \in W$ be the closest surface point to u. The difference between the outer (resp. inner) normal n_u (where it is defined) to the union boundary at u and the outer (resp. inner) surface normal n_x at x is $O(\sqrt{r})$ radians.*

Proof: Point u is contained in the tubular neighborhood, and the distance $d(u,x) = O(r)LFS(x)$ (Theorem 21). If n_u is defined, then u is contained in the surface of exactly one ball and n_u is the vector pointing towards the ball center, so we can apply Lemma 23.

□

Homeomorphism

We use these geometric theorems to show that the surface of either U_I or U_O is homeomorphic to the actual surface W. We'll do this using a natural map from U to W.

Definition: Let $\mu : R^3 \to W$ map each point $q \in R^3$ to the closest point of W.

Lemma 25 *Let U be either U_I or U_O. The restriction of μ to U defines a homeomorphism from U to W.*

Proof: We consider U_I; the argument for U_O is identical. Since U_I and W are both compact, it suffices to show that μ defines a continuous, one-to-one and onto function. The discontinuities of μ are the points of the medial axis. From Theorem 21, every point of U_I is within distance $O(r)LFS(x)$ from some point $x \in W$, whereas every point of the medial axis is at least $LFS(x)$ from the nearest point $x \in W$. Thus μ is continuous on U_I.

Now we show that μ is one-to-one. For any $u \in U_I$, let $x = \mu(u)$ and let $n(x)$ be the normal to W at x. Orient the line $l(x)$ through x with direction $n(x)$ according to the orientation of W at x. Any point on U_I such that $\mu(u) = x$ must lie on $l(x)$; let u' be the outer-most such point. By Theorem 21, $d(u', x) = O(r)LFS(x)$.

Let $B_{c,\rho}$ be the ball in U_I with u on its boundary. Let α be the angle between $\vec{uc}$ and the surface normal $n(x)$. By Theorem 24, $\alpha = O(\sqrt{r})$. Meanwhile $\rho = \Omega(LFS(x))$, by Observation 22.

Point u' is at most $O(r)LFS(x)$ from u, while the portion of $l(x)$ extending outward from u lies in the interior of $B_{c,\rho}$ for distance at least $2\rho\cos\alpha = O(LFS(x))$. So every point of $l(x)$ farther out from u but closer than $O(r)LFS(x)$ from x lies in $B_{c,\rho}$; in other words, u' has to be identical to u.

Finally, we need to establish that $\mu(U)$ is onto W. Since μ maps U, a closed and bounded surface, continuously onto W, $\mu(U)$ must consist of some subset of the closed, bounded connected components of W. But since every connected component of W contains samples of S, and $\mu(s) = s$ for $s \in S$, $\mu(U)$ must consist of all the connected components of W.
□

7 The power crust

It seems natural that since $\mathcal{U}_I$ and $\mathcal{U}_O$ are accurate representations of $\mathcal{W}$ and its complement, that the power crust that they induce is also an accurate representation of W. We establish this formally in this section.

Proximity

The fact that the power crust is close to W is actually immediate from our results so far. Since any point on a face separating an inside from an outside cell is contained in either both of their Voronoi balls or in no Voronoi ball at all, Theorem 21 implies the following.

Corollary 26 *Any point u on a face of the power crust lies within $O(r)LFS(x)$ of some point $x \in W$.*

Notice that although a point u on the power crust might be nearest to inner (outer) polar ball B, in Euclidean distance, it might belong to the power cell of some other inner (outer) ball B' which is nowhere near B. Our proof that the power crust is homeomorphic to the original surface hinges on showing that B and B' cannot, in fact, be too far apart.

Observation 27 *Let p be a point in the tubular neighborhood, and let s be the sample nearest p. Then $d(p, s) = O(r)LFS(s)$.*

Let $x \in W$ be the closest point on the surface to p. The Observation above follows since the distance $d(p, s)$ is at most distance $d(p, x) + d(x, s')$, where s' is the sample nearest x, using Observation 9.

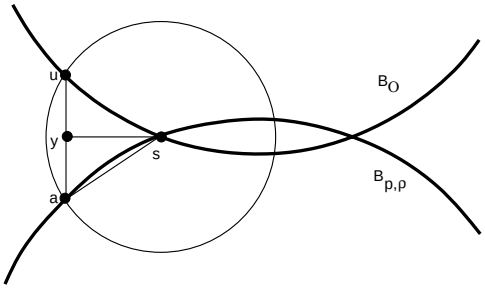

Figure 8:

Lemma 28 *Let u be a point in the tubular neighborhood outside of any polar ball, let $x \in W$ be the nearest surface point to u, and let s be the closest sample to x. Let $B_{p,\rho}$ be the smaller of the two polar balls at s. Then $d(u, B_{p,\rho}) = O(r^2)LFS^2(s)/\rho$.*

Proof: Since u is in the tubular neighborhood, $d(u,x) = O(r)LFS(x)$, and $d(x,s) \leq rLFS(x)$. So by Observation 27, $d(u,s) = O(r)LFS(s)$, that is, u is contained in a ball of radius $O(r)LFS(s)$ centered at s, as in Figure 8. The distance from u to $B_{p,\rho}$ will be maximized when a) the two polar balls intersect in as large an angle as possible (which is $O(r)LFS(s)/\rho$, by Lemma 15) and b) the radius of B_O is as small as possible (which is ρ).

From Figure 8, we have $d(u, B_{p,\rho}) \leq d(u,a)$. The length of the chord sa is $O(r)LFS(s)$, so the angle between the chord and the tangent plane to $B_{p,\rho}$ at s is $\arcsin[O(r)LFS(s)/2\rho] = O(r)LFS(s)/\rho$. So the total angle $\angle ysa = O(r)LFS(s)/\rho$ as well.

This gives $d(u,a) = O(r)LFS(s)\sin[O(r)LFS(s)/\rho]$, and hence $d(u, B_{p,\rho}) = O(r^2)LFS^2(s)/\rho$, for small enough r.

Lemma 29 *Let u be a point in the tubular neighborhood, and let p be the inner (outer) pole at minimum power distance to u, with polar ball $B_{p,\rho}$. Let $x \in W$ be the nearest surface point to u and let s be the nearest sample to x. Let $B_{c,\mu}$ be the smaller of the two polar balls at s. If $u \notin B_{p,\rho}$, then $d(u, B_{p,\rho}) = O(r)LFS(s)$, for small enough r.*

Proof: If u is inside $B_{c,\mu}$, then it is inside $B_{p,\rho}$, and the lemma is trivial.

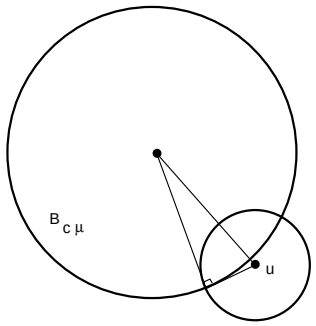

Figure 9:

Otherwise, we claim that the radius λ of the ball B_u centered at u and orthogonal to $B_{c,\mu}$ is at most $O(r)LFS(s)$, for small enough r. Since this ball must also intersect $B_{p,\rho}$ (Observation 1), the Lemma follows.

To establish the claim, assume without loss of generality that $LFS(s) = 1$, so that $LFS(x) = O(1)$ (Observation 9). By Lemma 28, $d(u, B_{p,\rho}) \leq k(r^2/\mu)$, for some constant k. We have, as in Figure 9,

$$\lambda = \sqrt{\mu^2 + 2(kr^2/\mu)\mu + k^2r^4/\mu^2 - \mu^2}$$

$$= \sqrt{2kr^2 + k^2r^4/\mu^2} = O(r)$$

□

Homeomorphism

From Lemma 23 and Lemma 29 we get the following observation, which we need to establish the homeomorphism between the power crust and W.

Observation 30 *Let u be a point in the tubular neighborhood, and let p be the inner (outer) pole at minimum power distance to u, with polar ball B_p. Let $x \in W$ be the surface point closest to u with surface normal n_x. The vector $\vec{uc}$ forms an angle of at most $\alpha = \pi/6$ with n_x, for small enough r.*

The set of points in the tubular neighborhood whose closest point on W is x forms a line segment g, perpendicular to the surface at x. Note that when we take a point u in the tubular neighborhood to its nearest point $x \in W$, it travels along the segment g corresponding to x.

Lemma 31 *The segment g normal to the surface at a point $x \in W$ and passing through the tubular neighborhood intersects the power crust exactly once.*

Proof: Consider the function $f_I(u)$ which returns the minimum power distance to any pole $p \in P_I$. The level sets of f_I are piecewise-quadratic surfaces formed by patches of spheres of equal power distance centered at the poles. The restriction of f_I to the segment g is a piecewise quadratic function. We claim that this function is monotonically decreasing as u goes from the outer

end of g to the inner end, since, by Observation 30, a point x moving inwards on g is always moving at an angle of at most $\pi/6$ from the the vector from the nearest pole p to x, and any angle less than $\pi/2$ would suffice to ensure that x is moving away from p.

Similarly, the function f_O is monotonically increasing on g. So f_I and f_O are equal at exactly one point, at a face of the power diagram separating the cells of an inside and an outside pole.
□

Theorem 32 *There is a continuous deformation of $\mathcal{Q}$ taking the power crust into W.*

Proof: Let Y be the power crust. We define a deformation of all of the domain $\mathcal{Q}$ which takes Y into W, and hence the interior of Y into the interior of W and the exterior of Y into the exterior of W. Specifically, we define a continuous parameterized map $f_t : \mathcal{Q} \to \mathcal{Q}$, for $t \in [0, 1]$, such that at any time t, f_t is a continuous, one-to-one and onto map, and such that at time $t = 0$, $f_0(Y) = Y$, and at time $t = 1$, $f_1(Y) = W$.

The power crust is strictly contained in the tubular neighborhood around W (Lemma 26). Outside of the tubular neighborhood, we define f_t to be the identity, at every time t. By Lemma 31, the segment g normal to W at a point $x \in W$ and passing through the tubular neighborhood intersects the power crust exactly once, in a point $y \in Y$. By the definition of the tubular neighborhood, g intersects W only in x. Let g_i and g_o be the inner and outer endpoints of g. We define $f_t(y) = tx + (1 - t)y$, and we let f_t linearly map the segments g_i, y to $g_i, f_t(y)$ and y, g_o to $f_t(y), g_o$.
□

8 Medial axis approximation

Both topologically and geometrically, we can show that the power shape is a good appoximation to the medial axis in some basic ways.

Theorem 33 *The power shape is homotopy equivalent to* $\mathcal{Q} - W$.

Proof: Theorem 7 established that the power shape is homotopy equivalent to $\bigcup(Pow(\mathcal{B}) - Y)$, where $Pow(\mathcal{B})$ is the set of faces of the power diagram and Y is the power crust. Since $\bigcup Pow(\mathcal{B}) = \mathcal{Q}$, this means that the power shape is homotopy equivalent to $\mathcal{Q} - Y$. The space homeomorphism of Theorem 32 shows that $\mathcal{Q} - Y$ is homeomorphic to $\mathcal{Q} - W$.
□

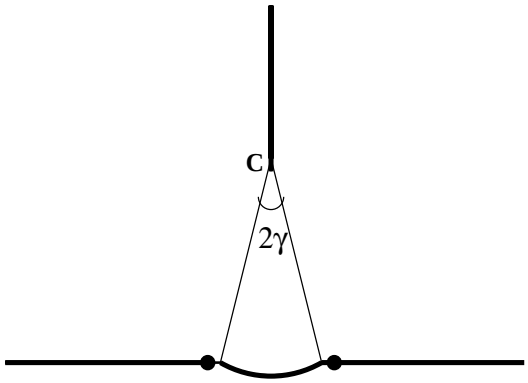

Figure 10: A small bump on the surface induces a long "hair" on the medial axis without having to contain any samples. Here, c is the endpoint of the "hair", and r is about 1/2, so that neither of the samples lies on the bump.

In addition to this topological equivalence, we show that the set P_I of poles converges, geometrically, to the true medial axis of $\mathcal{W}$ as the sampling density goes to infinity. In contrast to our previous results, we cannot guarantee that every medial axis is adequately approximated by an r-sample for a specific value of r such as 0.1. This is because, as in Figure 10, for any finite value of r, we can construct a a very small, shallow bump on the surface W, inducing a "hair" on the medial axis but without requiring samples on the bump. Note, however, that we have to choose the angle γ to be small with respect to r. This motivates the following definition.

Definition: A medial axis point c belongs to the γ-medial axis of $\mathcal{W}$ when at least two points $u_1, u_2 \in W$ on the boundary of the medial ball centered

at c form an angle $\angle u_1 c u_2 > 2\gamma$.

Interestingly, the γ-medial axis can be disconnected.

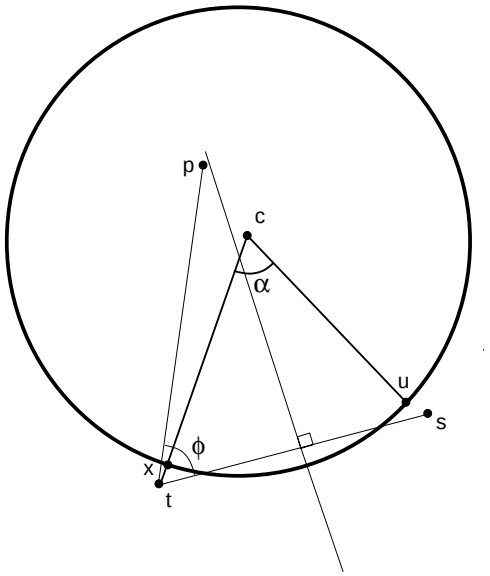

Figure 11: Since p is in the Voronoi cell of t, it has to be on the same side of the bisector of ts as t.

Lemma 34 *Let $B_{c,\rho}$ be a medial ball such that c belongs to the inner (outer) γ-medial axis, for some fixed γ, with $\gamma = \Omega(r^{1/3})$. Let t be the nearest sample to c. Then the distance from c to the inner (outer) pole p of t is $O(\rho r^{2/3})$.*

Proof: Without loss of generality let c be a point on the inner γ-medial axis. Let t be the closest sample to c, and let u_1, u_2 be two surface points on the boundary of $B_{c,\rho}$ such that $\angle u_1 c u_2 \geq 2\gamma$. Let α be the maximum of angles $\angle t c u_1$ and $\angle t c u_2$, so that $\alpha \geq \gamma$. Let $u \in \{u_1, u_2\}$ be the one realizing this maximum angle and let s be u's closest sample; see Figure 11.

From the sampling criterion we have that $d(u,s) \leq rLFS(s) \leq r\rho$. Let x be the point at which segment ct intersects the medial ball. Since $\angle xcu = \alpha$, $d(x,u) = 2\rho \sin(\alpha/2)$. Also, $d(c,t) \leq d(c,s) \leq \rho(1+r)$, so $d(x,t) \leq r\rho$. We conclude that $d(t,s) \leq 2\rho(r + \sin\alpha/2)$.

Next, we will bound the angle $\phi = \angle pts \leq \beta + \psi + \epsilon$, where $\beta = \angle ptc$, $\psi = \angle utc$, and $\epsilon = \angle uts$.

Point c lies in the Voronoi cell of t, and $d(c,t) \geq LFS(t)$. So from Lemma 11, both $\angle \vec{n}\ \vec{tc}$ and $\angle \vec{n}\ \vec{tp}$ are at most $2 \arcsin r'$, where $\vec{n}$ is the surface normal at t. So $\beta = \angle ptc \leq 4 \arcsin r' = O(r)$.

Since x and t are collinear, $d(c,t) \geq d(c,x)$, and both are much greater than $d(x,u)$, we have $\psi = \angle utc \leq \angle uxc = \pi/2 - \alpha/2$.

We have $d(s,u) \leq \rho r$, and both $d(t,s)$ and $d(t,u)$ are at least $2\sin(\frac{\alpha}{2} - r)$, so that angle $\epsilon = \angle uts \leq \arcsin(\frac{r}{2\sin(\frac{\alpha}{2}-r)}) = O(r^{2/3})$. This completes our upper bound on ϕ.

Since p is t's pole, p is closer to t than it is to s. By intersecting the cone at t around ts at angle ϕ with the plane equidistant from t and s (see Figure 11), we can get an upper bound on $d(t,p)$:

$$d(t,p) \leq \frac{d(t,s)/2}{\sin(\frac{\pi}{2} - \phi)} \leq \frac{\rho(r + \sin\frac{\alpha}{2})}{\sin(\frac{\alpha}{2} - \beta - \epsilon)}$$

$$= O\left(\rho \frac{r + r^{1/3}}{r^{1/3} - r - r^{2/3}}\right) = O(\rho(1 + r^{2/3}))$$

Let q be the point at which the circle centered at t and passing through c intersects the segment tp. Since $\beta = \angle ctp = O(r)$, we get $d(c,q) = O(\rho r)$. Also, since $d(t,p) - d(t,q) = d(t,p) - d(t,c)$, we have $d(q,p) = O(\rho r^{2/3})$ and finally $d(p,c) = O(\rho r^{2/3})$.
□

Note that the value $\gamma = \Omega(r^{1/3})$ in the theorem above is not crucial; a similar statement could be made for any $\gamma = o(r^{1/2})$, with an appropriate modification of the bound.

Now we apply this bound to make a precise statement to the effect that the set of poles converges to the medial axis as $r \to 0$.

Theorem 35 *Consider a sequence of samples $S_0, S_1, \cdots$ from W, with the property that r_i converges to 0 as $i \to \infty$. The set of inner (outer) poles of S_i converges to the inner (outer) medial axis of $\mathcal{W}$.*

Proof: Lemma 34 shows that for every point c on the γ-medial axis, $\gamma > 0$, and any fixed radius $\epsilon > 0$, there is some finite i such that, for all $j > i$,

there is a pole of S_j within distance ϵ of c. A point c on the 0-medial axis (a center of curvature of W) belongs to the closure of the γ-medial axis, with $\gamma > 0$, so that, again, for any $\epsilon > 0$ there is a sufficiently small γ such that there is a point c' of the γ-medial axis within distance $\epsilon/2$ of c, and a finite i such that for any $j > i$ there is a pole of S_j within distance $\epsilon/2$ of c'. This shows that in the limit the set of poles contains the medial axis.

We now argue that in the limit the medial axis contains the set of poles. First, we assiciate a value γ_x with each point $x \in W$. Point x is associated with an inner and an outer medial axis point, belonging respectively to the γ_i- and γ_o- medial axes; let γ_x be the minimum of γ_i, γ_o. Let the subset of W with $\gamma_x \geq \gamma$ be the γ-surface.

Now fix $\gamma > 0$. The Voronoi cell of any sample s in the γ-surface contains the interior medial axis point c corresponding to s, which belongs to the γ-medial axis. So for any ϵ there is some i such that small enough so that, if $s \in S_j$ for any $j > i$, c is within distance ϵ if the interior pole of s, by Lemma 34. □

9 Theoretical algorithm

Even when the surface W is unknown, it is possible to correctly construct the power crust given an r-sample for small enough r. The difficulty of course is in determining which are the inside, and which are the outside, poles. We know that the polar balls of an inner and an outer pole can only intersect shallowly. If we could determine that two inner (outer) polar balls which induce a face of the power diagram must intersect deeply, then we could assign all power diagram two-faces corresponding to shallowly intersecting pairs of balls to the power crust, giving an algorithm analogous to that of Attali [6] in $\mathbb{R}^2$. Unfortunately, we could not establish that adjacent inner (outer) polar balls intersect deeply. Instead, we have the following.

Lemma 36 *Two inside (resp. outside) polar balls inducing a face f intersecting the tubular neighborhood meet at an angle of at least $2\pi/3$, for small enough r.*

Proof: Let p be any point on the f inside the tubular neighborhood, and let c_1, c_2 be the centers of the two inside (resp. outside) polar balls inducing the face. Since p is in the tubular neighborhood, and c_1 and c_2 are the poles with (equal) minimum power distance to p, we can apply Observation 30. Thus the angle between the surface normal $n(x)$ at the point $x \in W$ closest to p and either pc_1 or pc_2 is at most $\pi/6$, so $\angle c_1 p c_2$ is at most $\pi/3$.
□

This leads to the following algorithm to label each pole as either outside (O') or inside (I').

Input: An r-sample S from a closed, bounded smooth surface W.
Output: The power crust of S.

Step 1: Construct the Delaunay triangulation of S, find the Voronoi vertices, and select two poles for each sample. Let $\mathcal{B}_P$ be the set of polar balls.

Step 2: Construct the power diagram $Pow(\mathcal{B}_P)$.

Step 3: Select a sample on the convex hull of S.
Label its infinite outer pole with O' and the opposite inner pole I'.
Insert both poles in a queue.

Step 4: While the queue is non-empty:
Remove a labeled pole p from the queue, and examine
each unlabeled neighbor q of p in $Pow(\mathcal{B}_P)$.
If the Voronoi ball surrounding q intersects the Voronoi ball of p
at an angle of more than $\pi/4$:
Give q the same label as p and insert it in the queue.
For each sample s such that q is a pole of s,
if the pole q' opposite q at s is unlabeled:
Give q' the opposite label from
q and insert q' into the queue.

Step 5: Output the faces of $Pow(\mathcal{B}_P)$ separating the cells of
one pole labeled I' and one pole labeled O' as the power crust.

To prove that this algorithm is correct, we need to show that the sets I and O, corresponding to the inside an outside of W, are identical to the sets I' and O'.

Lemma 37 *No pole in I receives label O' and no pole in O receives label I'.*

Proof: Let q be the first mislabeled pole, and let p be the pole from whose label that of q was determined. Either p and q should have opposite labels but they meet at an angle of more than $\pi/2$, or p and q should have the same label but they are opposite poles of the same sample s. The first case is impossible by Lemma 16, and the second is impossible because the two poles of any sample always should have opposite labels.
□

Lemma 38 *Every pole receives a label.*

Proof: We consider a pole $p \in I$, Every ball in I has at least one point on the power crust, since each sample s such that p is a pole of s appears on the power crust.

By Lemma 36 we know that every power crust edge is contained in two balls which intersect deeply (they meet at an angle of at least $2\pi/3$). Therefore if any pole q in the same connected component of $\mathcal{U}_I$ receives label I', then p will eventually as well.

Each connected component of either I or O eventually gets at least one labeled pole. Assume not; consider some component that remains unlabeled; we claim that there must be a sample on this component. If this is true, we are done, because a label will be propogated across this sample.

The claim must be true; otherwise, consider any point x on the boundary of that component. The line segment connecting x to its nearest sample s must cross the medial axis, so that the distance $d(x, s) \geq LFS(x)$, a contradiction.
□

10 The Anti-Crust

We conclude with a brief comparison of the dual shapes described in this paper with the *anti-crust*, the dual shape corresponding to crust constructions such as [2],[3],[4],[5],[6],[16], and [21].

Definition: Let T be a triangulated manifold (possibly with boundary) selected from the Delaunay triangulation of a set S of surface samples, which we shall generically call a *crust*. The *anti-crust* A is the set of Voronoi faces of S whose duals do not belong to T.

Observation 39 *The dimension of every face of A is at most two.*

In this, the anti-crust is more like the medial axis than the power shape is.

Observation 40 *Every Voronoi vertex of S is a vertex of A.*

This is the main failing of the anti-crust, as discussed below.

Definition: Let $\mathcal{B}_V$ be the set of Voronoi balls of S.

From Observation 2 above, we see that T is a subset of $Pow(\mathcal{B}_V)$.

Observation 41 *A is the dual shape of T in $Pow(\mathcal{B}_V)$.*

Theorem 7 therefore implies the following.

Corollary 42 *When T is homeomorphic to W, the anti-crust A is homotopy equivalent to $\mathcal{Q} - W$.*

This is one sense in which A is a good approximation to the medial axis M; $\mathcal{Q} - W$ is homotopy equivalent to both. Another sense in which it is good is functional: the surface approximation can be recovered from it.

Observation 43 *The surface approximation T can be recovered from A by computing $Pow(\mathcal{B}_V)$ and selecting the faces not dual to faces of A.*

Again, this follows from Observation 2. Thus, we can think of A as the approximate MAT dual to T, an approximate surface representation.

In the limit, as the sampling density becomes infinite, $T \to W$ (eg. [3],[2]). One would like to conclude that $A \to M$ as well, but this is *not* in fact true. The problem is that even an arbitrarily dense sample can produce Voronoi vertices very close to W and far from M: whenever four samples adjacent on the surface and determining a Voronoi vertex v are nearly co-circular, v might be anywhere on the line perpendicular to the circle. This introduces unwanted "hairs" on the anticrust A, purely due to quantization, which do not correspond to any feature of M.

It must be admitted, however, that this theoretical difficulty has not had much effect in practice. Equally erroneous "hairs" are caused by small errors in the sample positions.

11 Open Questions

We were unable to resolve the following conjecture.

Conjecture 44 *The power crust faces are exactly those for which the two polar balls determining the face intersect in a lune of at most $O(r)$ degrees.*

This is the criterion used in Attali's (and Gold's) two-dimensional surface reconstruction algorithm. If the conjecture is true, then the generalization of their algorithms to three dimensions produce the power crust.

References

[1] E. Althaus and K. Mehlhorn. Polynomial time TSP-based curve reconstruction. *Symposium on Discrete Algorithms (SODA)*, pages 686-695, (2000).

[2] N. Amenta, S. Choi, T. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. To appear in the *International Journal of Computational Geometry and its Applications*, special issue

on papers from the *16th ACM Symposium on Computational Geometry*, (2000).

[3] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, **22**, pages 481–504, (1999). An extended abstract appeared in the *14th ACM Symposium on Computational Geometry*, 1998, 39–48.

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

[5] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Siggraph* 1998.

[6] D. Attali. r-Regular shape reconstruction from unorganized points. *Computational Geometry: Theory and Applications*, **10**:4, pages 239–247 (1998). An extended abstract appeared in *Proc. 13th ACM Symp. Computational Geometry*, 1997, pages 248–253.

[7] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons, *Computer Vision and Image Understanding*, 67(3) (1997), pp. 261–273.

[8] F. Bernardini, J. Mittleman, H. Rushmeir, C. Silva and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Vision and Computer Graphics.* **5**:4, (1999). Also IBM Tech. Report RC21463(96842).

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

[10] J-D. Boissonnat and F. Cazals. Natural coordinates of points on a surface. *Proceedings of the 16th Annual ACM Symposium on Computational Geometry*, pages 223–232, (2000).

[11] S-W. Cheng, H. Edelsbrunner, P. Fu, and K-P. Lam. Design and analysis of planar shape deformation. *Proceedings of the 14th Annual ACM Symposium on Computational Geometry* pages 29–38, (1998).

[12] H.I. Choi, S.W.Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics*, **181**:1 (1997), pages 57–88.

[13] T. Culver, J. Keyser, and D. Manocha. Accurate computation of the medial axis of a polyhedron, *Proc. ACM Solid Modeling and Applications*, pages 179–190, (1999). Also *UNC Technical Report* TR98-038.

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

[15] T. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. 10th Annual ACM-SIAM Symp. on Discrete Algorithms*, 1999, pages S893–4.

[16] T. Dey, K. Mehlhorn and E. Ramos. Curve reconstruction: connecting the dots with good reason. *Comput. Geom. Theory Appl.*, **15** (2000), 229–244. A previous version appeared in *SoCG '99*.

[17] H. Edelsbrunner. The union of balls and its dual shape. *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*, pp, 218-231, (1993).

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

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

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

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

[22] 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.

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

[24] P. Hubbard. Approximating polyhedra with spheres for time-critical collision detection. *ACM Transactions on Graphics*, 15(3), pp. 179–210, (1996).

[25] V. Ranjan and A. Fournier. Matching and interpolation of shapes using unions of circles. *Computer Graphics Forum*, 15(3), pp. 129–142, (1996).

[26] D. Sheehy, C. Armstrong and D. Robinson. Shape description by medial axis construction. *IEEE Transactions on Visualization and Computer Graphics*, 2(1), pp. 62–72, (1996).

[27] T. Sakkalis and Ch. Charitos. Approximating curves via α-shapes. *Graphical Models and Image Processing* **61**, pp. 165–179 (1999).

[28] M. Teichman and S. Teller. Assisted articulation of closed polygonal models. *Proc. 9th Eurographics Workshop on Animation and Simulation*, (1998).