Eurographics Symposium on Geometry Processing 2009
Marc Alexa and Michael Kazhdan
(Guest Editors)

Volume 28 (2009), Number 5

Recovering Structure from *r*-Sampled Objects

O. Aichholzer[†], F. Aurenhammer[‡], B. Kornberger[‡], S. Plantinga[§], G. Rote[¶], A. Sturm[||], G. Vegter[¶]

Abstract

For a surface $\mathcal{F}$ in 3-space that is represented by a set S of sample points, we construct a coarse approximating polytope P that uses a subset of S as its vertices and preserves the topology of $\mathcal{F}$. In contrast to surface reconstruction we do not use all the sample points, but we try to use as few points as possible. Such a polytope P is useful as a 'seed polytope' for starting an incremental refinement procedure to generate better and better approximations of $\mathcal{F}$ based on interpolating subdivision surfaces or e.g. Bézier patches.
Our algorithm starts from an r-sample S of $\mathcal{F}$. Based on S, a set of surface covering balls with maximal radii is calculated such that the topology is retained. From the weighted α-shape of a proper subset of these highly overlapping surface balls we get the desired polytope. As there is a rather large range for the possible radii for the surface balls, the method can be used to construct triangular surfaces from point clouds in a scalable manner. We also briefly sketch how to combine parts of our algorithm with existing medial axis algorithms for balls, in order to compute stable medial axis approximations with scalable level of detail.

1. Introduction

This paper deals with recovering structural information for a 3-dimensional object that is represented by a sample point cloud. More specifically, given an object $\mathcal{O}$ in 3-space and an r-sample S of its boundary, we want to find an approximating polytope P that uses a subset of the points in S as its vertices and preserves the topology of $\mathcal{O}$. Our goal is, on the one hand, to use as few points of S as possible and, on the other, to get a flexible approximation whose level of detail can be tuned from coarse to fine. We also (briefly) address the problem of finding piecewise linear approximations of the medial axis of $\mathcal{O}$. Motivation for studying these problems is based on open problems in object simplification and surface reconstruction, two fundamental tasks in several areas of computer science, like geometric modeling, computer graphics, and computational geometry.

[†] Institute for Software Technology, Graz University of Technology, Austria, {oaich,bkorn}@ist.tugraz.at
[‡] Institute for Theoretical Computer Science, Graz University of Technology, Austria, auren@igi.tugraz.at
[§] University of Groningen, Department of Mathematics and Computing Science, {simon,gert}@cs.rug.nl
[¶] Institut für Informatik, Freie Universität Berlin, Germany, {rote, sturm}@inf.fu-berlin.de

The main support structure we use is an approximation of the object in question with a union of balls. In the context of object simplification, this approach is used for many purposes, e.g. collision detection [Hub96], shape matching [SS04], and shape interpolation [RF96], to name a few. Regarding surface reconstruction, approximating objects with balls also plays a major role, see for example the power crust algorithm [ACK01], related work [AB99, AK00, AK01] and also [CL08], naming again only a few.

In our approach, which is similar to work in [CL08], we build a union of so-called *surface balls*, centered at the points in our r-sample S on the surface $\mathcal{F}$ of $\mathcal{O}$, whose radii adapt to the local feature size of $\mathcal{F}$. The desired approximating polytope P is then extracted from the weighted alpha shape [Ede95] of a carefully chosen subset of these balls. In contrast to [CL08], where prior knowledge of the local feature size of $\mathcal{F}$ is assumed, we obtain an estimation of this function from the data, by using distances to poles [AB99] (certain vertices of the Voronoi diagram for S). Using a tailored technique of pruning the surface balls, we obtain a coarse-to-fine approximation of $\mathcal{F}$ by polytopes. This is the first result that uses, from a practical point of view, approximations of local feature size and medial axis to obtain locally adaptive reconstructions of an unknown surface.

The polytopes we construct are topologically correct reconstructions of $\mathcal{F}$. Thus our results differ from existing multi-scale surface reconstruction techniques in [NSW08, CL08, CCSL09, GO08] where topological filtering occurs. At the coarsest level, the polytope we obtain is what we call 'seed polytope', as it provides not only a coarse approximation of $\mathcal{F}$ but also a mapping of the non-used sample points in S to the vertices of the polytope. Such a mapping is needed for incrementally generating approximations of $\mathcal{F}$ based on interpolating subdivision surfaces or Bézier patches. We stress that the intended purpose of the seed polytope is not primarily in approximating $\mathcal{F}$ but rather in serving as a (topologically correct and small) starting structure for subsequent approximations by patches. We thus do not try to keep the approximation error small for the seed polytope itself, and use this additional freedom to keep the polytope small. In a previous related approach [BPR*07], point clouds in convex position are approximated by spherical patches.

Strongly related to the surface reconstruction is the medial axis approximation; we refer to [ABE07] for a recent survey paper on medial axes and their algorithmic construction. In this area, many algorithms are based on unions of balls as well, for example [BO04, GMP07, YBM04]. We briefly describe how a variant of our approach, now for balls centered at poles instead on sample points, combines with an existing medial axis algorithm for balls [AK01] to an efficient and stable medial axis approximation algorithm for general objects. It is known that sufficiently dense r-samples lead to topologically correct medial axis approximations; see [AK00] and, for a result more general than for poles, [AB03].

2. Definitions and notation

Throughout this paper, let $\mathcal{O}$ denote the original solid object and let $\mathcal{F}$ denote its surface. The following definitions are standard.

Definition 1

- The *medial axis transform* of $\mathcal{O}$ is the (infinite) set of maximal balls that avoid $\mathcal{O}$, where maximality is with respect to inclusion. The set of the centers of these balls forms the *medial axis* of $\mathcal{O}$. The surface $\mathcal{F}$ splits the medial axis in an *inner medial axis* and an *outer medial axis*.
- The *local feature size* $\mathrm{lfs}(x)$ of a point $x \in \mathcal{F}$ is the minimum distance from x to any point on the medial axis of $\mathcal{O}$.
- A finite point set $S \subset \mathcal{F}$ is an *r-sample* of $\mathcal{F}$ if every point $x \in \mathcal{F}$ has at least one point of S within distance $r \cdot \mathrm{lfs}(x)$ [AB99].

In this paper, we will assume that S is an r-sample of $\mathcal{F}$ for $r = 0.08$.

For each sample point $s \in S$, we define two vertices of the Voronoi diagram of S as the *poles* of s, see [AB99]: the *inner pole* is the vertex of the Voronoi cell of s farthest away from s and in the interior of $\mathcal{O}$, and the *outer pole* is the farthest one from s and outside $\mathcal{O}$. For the inner pole p of each site s we consider the ball with center p and radius $\|p-s\|$. We refer to the set of these polar balls as the *(inner) discrete medial axis transform* $\mathrm{DMAT_{in}}$. Analogously, we generate a set of outer polar balls and denote it by $\mathrm{DMAT_{out}}$.

Definition 2

- The *discrete medial axis* DM_{in} (DM_{out}) is the medial axis of the union of polar balls in the sets $\mathrm{DMAT_{in}}$ ($\mathrm{DMAT_{out}}$).
- The *discrete local feature size* $\widetilde{\mathrm{lfs}}(x)$ of a point $x \in \mathcal{F}$ is the minimum distance from x to $DM_{in} \cup DM_{out}$.
- The *pole distance* $\hat{D}(x)$ of a point x is the distance to the nearest pole.

We will see that $\hat{D}$ is a good estimate of $\widetilde{\mathrm{lfs}}$ (Corollary 5.5), as well as an upper bound on the true local feature size (Lemma 5.1). In practice, $\hat{D}$ is easier to compute than $\widetilde{\mathrm{lfs}}$, and the true local feature size is not computable at all.

The *weighted* α*-shape* is the dual shape of a union of balls [Ede95]. It is a simplicial complex whose vertices are the centers of the balls, and which is homotopy-equivalent to the union of balls. We will refer to the weighted α-shape of $\mathrm{DMAT_{in}}$ as $\mathcal{A}_{in}$ and to the one of $\mathrm{DMAT_{out}}$ as $\mathcal{A}_{out}$.

Proposition 2.1 [AK01] Let $\mathcal{A}_{in}$ and $\mathcal{A}_{out}$ be the weighted α-shapes of $\mathrm{DMAT_{in}}$,$\mathrm{DMAT_{out}}$. Then we have

$$DM_{in} \cup DM_{out} \subseteq \mathcal{A}_{in} \cup \mathcal{A}_{out}.$$

3. Our approach

The flowchart in Figure 1 gives an overview of the work flow for the three tasks considered in this paper: Computing a seed polytope, a scalable surface reconstruction, and the medial axis.

In all cases we start with an r-sample S of the object $\mathcal{O}$ as input and compute from it the two discrete medial axis transforms $\mathrm{DMAT_{in}}$ and $\mathrm{DMAT_{out}}$.

These sets serve two purposes: For seed polytopes and scalable surface reconstruction we use them in order to estimate bounds on the local feature size of the sample points. For medial axis approximation, we use a pruned version of $\mathrm{DMAT_{in}}$ with slightly enlarged radii, representing the object $\mathcal{O}$ in a compact and faithful way.

The union of surface balls. A *surface ball* is a ball with center at a sample point $s \in S$. For seed polytopes, our goal is to represent the surface $\mathcal{F}$ of $\mathcal{O}$ in a topological correct way with as few faces as possible. We try to make the surface balls as large as possible, while guaranteeing correct topology of the the union $U(B_{\mathcal{F}})$ of the set $B_{\mathcal{F}}$ of surface balls. A subsequent pruning step will throw away some of these balls whenever the sample is denser than necessary. For surface

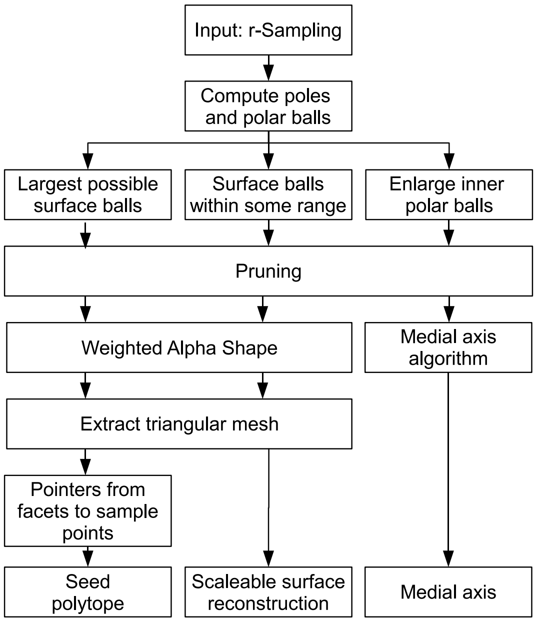

Figure 1: *Work flow*

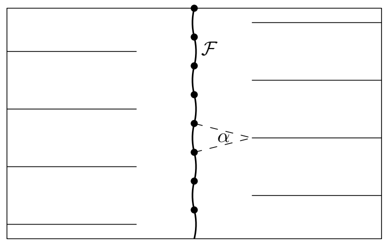

Figure 2: *A wiggly curve $\mathcal{F}$ with a point sample on a straight line. (Adapted from [AK00].)*

reconstruction, we will output meshes of scalable complexity. The only modification necessary to reach this goal is to choose surface balls with smaller radii.

Pruning. To decide which balls to keep, we solve a combinatorial problem. We (virtually) shrink the balls in $B_{\mathcal{F}}$ and compute a minimal subset $B'_{\mathcal{F}}$ of $B_{\mathcal{F}}$ such that the shrunk balls cover the sample S. This is a *set covering problem*, which is solved by a heuristic. The advantage of this approach is that the selection of the pruned subset proceeds now in a purely combinatorial manner, without regard to geometry and topology. The radii of the shrunk balls are chosen in such a way that covering of S by a subset of shrunk balls guarantees that the original, unshrunk, surface balls cover the surface $\mathcal{F}$, and moreover, their union represents the topology of $\mathcal{F}$ correctly.

The polyhedral approximation. Finally we compute the weighted α-shape of $B'_{\mathcal{F}}$, which has the same topology as $\mathcal{F}$ and which gives the desired seed polytope. The vertices of the weighted α-shape are points in S, because the centers of the balls in $B'_{\mathcal{F}}$ have been chosen from S. We use the power diagram of $B'_{\mathcal{F}}$ to find out which vertex of the polytope each sample point $s \in S$ belongs to and provide a list of pointers representing this relation.

Medial axis approximation. The medial axis algorithm of Amenta et al. [AK01] could be used to compute the medial axis of the union of the balls in $\mathrm{DMAT}_{\mathrm{in}}$. However, medial axes are in general unstable because of their disproportional response to even small perturbations on the object surface. Therefore, and also due to noise and numerical inaccuracies, $\mathrm{DMAT}_{\mathrm{in}}$ might contain balls far from a reasonably pruned approximation of the medial axis of $\mathcal{O}$ in practice, as small details might be (correctly) approximated that nevertheless are not needed in the application. Moreover, because S is a dense r-sample, the centers of the balls in $\mathrm{DMAT}_{\mathrm{in}}$ sample the medial axis in a much too dense way. We provide an adequate input for the medial axis algorithm [AK01] by reducing the number of balls significantly and thereby stabilizing $\mathrm{DMAT}_{\mathrm{in}}$.

This is done by adding a small distance ε to the radii of the balls in $\mathrm{DMAT}_{\mathrm{in}}$. Thus we get an enlarged set $\mathrm{DMAT}'_{\mathrm{in}}$ which we use to compute a covering matrix. Our set covering algorithm finds a small subset $\mathrm{DMAT}''_{\mathrm{in}}$ of $\mathrm{DMAT}'_{\mathrm{in}}$ which covers all sample points (but not necessarily $\mathcal{F}$). The goal of stabilization of $\mathrm{DMAT}_{\mathrm{in}}$ is implicitly reached because the set covering algorithm favors balls covering many sample points (which have their center near the medial axis and are therefore usually larger than unstable ones). The degree of simplification (and thus the level of detail of the approximated medial axis) is scalable by the choice of ε.

No implementation of the algorithm developed in [AK01] was available and so we have implemented it using CGAL [CGA]. We obtain—in combination with our pruning technique—stable and efficient medial axes. In practice, the approach works even for poorly sampled inputs which do not meet the r-sampling condition at all; see a companion paper [AAHK09]. Of course, no theoretical guarantees can be given in that case.

Obtaining the local feature size. A distinguishing feature of our problem setting is that we cannot get a lower estimate on the local feature size. Figure 2 shows a section of a curve $\mathcal{F}$ that consists of alternating short circular arcs. The horizontal lines are part of the medial axis. The points of the r-sample S are aligned vertically. By reducing the angle α, such an example can be built for any $r > 0$. The algorithm sees only these samples. Thus, to the algorithm, this input is indistinguishable from a very densely oversampled straight line.

4. Technical results

In order to generate adequate sets of polar balls and surface balls (in both cases, the topology must be maintained), we need to derive certain information concerning the local feature size of the sampled object. The present and the subsequent section are devoted to this issue. We obtain several new properties of r-sampled objects for suitable values of r.

Let $\mathcal{M}_{\text{in}}$ and $\mathcal{M}_{\text{out}}$ denote the inner and the outer medial axis of the given object $\mathcal{O}$, respectively. We start by bounding the distance of poles to the respective parts of the medial axis—a result crucial for bounding the radii of surface balls in Section 5.

Theorem 4.1 For an r-sample S, let p be an inner (resp. outer) pole of a sample point $s \in S$, and denote with B_p the inner (outer) polar ball of s, with radius R_p. The distance from p to $\mathcal{M}_{\text{in}}$ ($\mathcal{M}_{\text{out}}$) is at most $O(r) \cdot R_p$.

In the limit, when the sampling density approaches zero, poles and the medial axis coincide, as has already been shown by Amenta et al. [ACK01, Theorem 35]. In contrast to this result, we give an explicit quantitative analysis in terms of r. Results similar to Theorem 4.1 have been shown (see e. g. [ACK01, Lemma 34], on which Theorem 35 is based, or [BC01, Proposition 16]). However, we could not use these results, since they hold only when the angle between the two closest surface points to a given point on $\mathcal{M}_{\text{in}}$ ($\mathcal{M}_{\text{out}}$) is not too small, (These points form the γ-*medial axis*.)

Proof The idea of the proof is to turn the polar ball B_p into a medial ball, while not moving its center too much. The proof is based on several technical lemmas which are given subsequently. We proceed in three steps, see Figure 3:

1. While keeping the center of B_p fixed we shrink the radius of B_p until the ball becomes empty, touching the surface $\mathcal{F}$ of $\mathcal{O}$ at some point x_0. By Lemma 4.2 below, the difference Δ_1 between the new radius and the original radius R_p is at most $\Delta_1 = O(r^2) \cdot R_p$.
2. We expand the shrunken ball from the touching point x_0 by moving its center in the direction $\overrightarrow{x_0 p}$ until either

   (2a) the ball has the original radius R_p of B_p, or
   (2b) the ball touches the surface at another point. If this occurs we have found a point of $\mathcal{M}_{\text{in}}$ within distance Δ_1, and we are done.

3. In case (2a), we "roll" the new ball B'_p (with radius R_p) on the surface. More precisely, let K_1 be the component of $B_p \cap \mathcal{F}$ which contains x_0. Consider the balls of radius R_p that are tangent to $\mathcal{F}$ in a point of K_1 and lie on the same side of $\mathcal{F}$ as p. The locus of the centers of these balls is the inner parallel surface $\bar{F}$ of K_1. We claim that the rolling ball touches another point of $\mathcal{F}$, and therefore $\bar{F}$ contains a point of $\mathcal{M}_{\text{in}}$.

   We prove this by contradiction. Let us suppose that the ball can roll on K_1 without ever touching a second point of $\mathcal{F}$. K_1 cuts B_p into two parts: B^+ containing p, and the rest B^-. By Lemma 4.4 below, B^+ is completely covered by the tangent balls of K_1. Since by assumption these balls never hit another point of $\mathcal{F}$, it follows that K_1 is the only component of $\mathcal{F} \cap B_p$. Let $s \in B_p$ be the sample point whose pole is p. This point must lie on K_1 and therefore we can roll the empty tangent ball of radius R_p to s. The radius R_M of the medial ball at s is therefore at least R_p. On the other hand, each point of the medial axis is contained in the Voronoi cell of the nearest sample point, therefore $\|p - s\| = R_p \geq R_M$. This implies $R_p = R_M$ and the tangent ball at s has its center on $\mathcal{M}_{\text{in}}$, and we are done. We remark that this last case can actually never arise, since $R_p > R_M$ unless the medial axis branches and the ball touches $\mathcal{F}$ in several points.
   We have established that $\bar{F}$ contains a point m_x of $\mathcal{M}_{\text{in}}$ which is the center of a medial ball with radius R_p touching K_1 in x. We know by Lemma 4.3a that the angle $\gamma = \angle m_x x p$ is at most $3r + O(r)$. Thus, $\|p - m_x\| \leq R_p \cdot (3r + O(r^2))$.

□

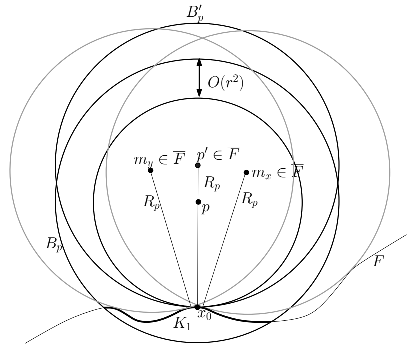

Figure 3: *After shrinking and expanding the ball B_p we roll the new ball B'_p on K_1 (e.g. the gray ball).*

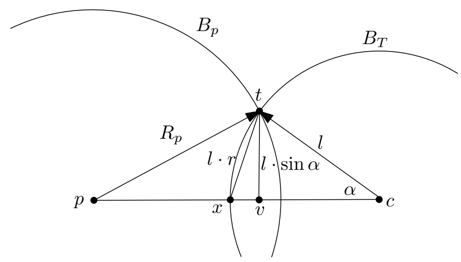

Figure 4: *Deepest penetration into B_p*

In the following, we will assume that p is an inner pole. (The situation is symmetric for outer poles.)

Lemma 4.2 Let p be a pole with polar radius R_p. The surface

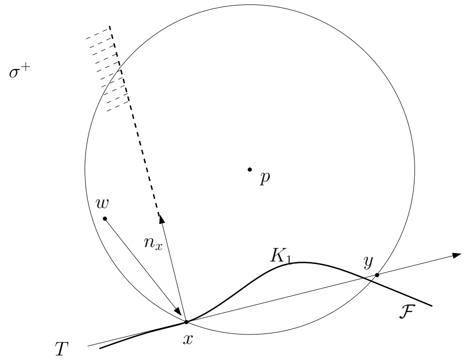

Figure 5: *The tangent balls of K_1 cover B^+*

$\mathcal{F}$ cannot get closer to p than

$$R_p\left(\sqrt{1-4(r^2-\tfrac{r^4}{4})}-r^2\right) \geq R_p\left(1-3r^2-O(r^4)\right).$$

For an r-sample with $r = 0.08$ the distance between the center p of a polar ball with radius R_p and $\mathcal{F}$ is larger than $0.9807 \cdot R_p$.

Proof Let x be the point on $\mathcal{F}$ closest to p. Let B_T be an empty outer ball tangent to x with center c and radius $l = \mathrm{lfs}(x)$. By the sampling condition, there must be a sample t within distance rl of x. t lies outside the balls B_p and B_T and therefore the distance from x to the circle $\partial B_p \cap \partial B_T$ is at most $r \cdot l$ (see Figure 4). Thus, the angle $\alpha = \angle cpt$ is bounded by $\sin\frac{\alpha}{2} \leq \frac{r}{2}$. For fixed l and R_p, the point x is closest to p when α is maximized. We thus analyze the situation for $\sin\frac{\alpha}{2} = \frac{r}{2}$:

$$\sin\alpha = 2\sin\frac{\alpha}{2}\cos\frac{\alpha}{2} \leq 2\cdot\frac{r}{2}\sqrt{1-\frac{r^2}{4}} = \sqrt{r^2-\frac{r^4}{4}}$$

$$\|v-p\| = \sqrt{R_p^2-(l\cdot\sin\alpha)^2} = \sqrt{R_p^2-l^2\cdot(r^2-\tfrac{r^4}{4})}$$

$$\|v-x\| = \sqrt{(l\cdot r)^2-(l\cdot\sin\alpha)^2} = \sqrt{(l\cdot r)^2-l^2\cdot(r^2-\tfrac{r^4}{4})} = \frac{l\cdot r^2}{2}$$

$$\|x-p\| \geq \|v-p\|-\|v-x\| \sqrt{R_p^2-l^2\cdot(r^2-\tfrac{r^4}{4})}-\frac{l\cdot r^2}{2}$$

The inner polar ball B_p contains a point of $\mathcal{M}_{\mathrm{in}}$ ([ACK01, Corollary 13]), therefore $l \leq 2R_p$. It follows that the distance between p and $\mathcal{F}$ is at least

$$\sqrt{R_p^2-4\cdot R_p^2\cdot(r^2-\tfrac{r^4}{4})}-R_p\cdot r^2 = R_p\cdot\left(\sqrt{1-4\cdot(r^2-\tfrac{r^4}{4})}-r^2\right),$$

as claimed in the lemma. □

Lemma 4.3 Let x be a surface point x inside a polar ball B_p with center p.

a) The angle γ between $\overrightarrow{xp}$ and the surface normal at x is bounded by $3r + O(r^2) = O(r)$.
b) (The penetration bound) The distance from x to the boundary of B_p is is bounded by $\frac{3}{2}\,\mathrm{lfs}(x)(r^2+O(r^3))$.

Part b of the lemma is similar to Lemma 4.2, except that the penetration of the surface point x into the pole ball B_p is measured in terms of $\mathrm{lfs}(x)$, and not in terms of the radius of B_p.

The proof of Lemma 4.3 is omitted for lack of space.

To complete the proof of Theorem 4.1, we still need to show that the tangent balls of K_1 cover all parts of B^+. Recall that K_1 cuts B_p in two parts: B^+ containing p, and the rest B^-.

Lemma 4.4 The tangent balls of K_1 completely cover B^+.

Proof Let $w \in B^+$ and let x be the closest point of K_1. We claim that the tangent ball at x covers w. If x lies in the interior of K_1, then wx is perpendicular to $\mathcal{F}$, and the claim is obvious. Let us assume that x is at the boundary of K_1, that is $B_p \cap \mathcal{F}$ (see Figure 5). Assume that the surface normal n_x does not go through p; otherwise it is obvious that w is covered. Consider the plane σ through n_x and through the point p. Figure 5 shows the projection on this plane. Locally around x, $\mathcal{F}$ is approximated by the tangent plane T and $B_p \cap \mathcal{F}$ is the halfspace of T that projects onto the ray xy in Figure 5. It follows that x can only be the point of K_1 closest to w, if w lies in the plane σ and in the closed halfplane σ^+ of σ which is bounded by n_x and does not contain p. □

5. Construction of balls

5.1. Polar balls

For the set $\mathrm{DMAT_{in}}$ of inner polar balls, it is well known [AK00] that the union of the balls in this set is homeomorphic to the original object $\mathcal{O}$. Recall that each ball in $\mathrm{DMAT_{in}}$ is the circumball of a Delaunay tetrahedron and therefore has at least four points of S on its boundary and no such point in its interior. From $\mathrm{DMAT_{in}}$ we generate a set $\mathrm{DMAT'_{in}}$ of slightly enlarged balls which are still centered on S. Such a ball typically covers tens or even hundreds of points of S. In a subsequent set covering step, this redundancy in covering will be eliminated, and thereby only a small and stable subset of $\mathrm{DMAT'_{in}}$ will be kept. We have to ensure, for the goal of topologically correct medial axis approximation, that the union of $\mathrm{DMAT_{in}}$ and the union of $\mathrm{DMAT'_{in}}$ are topologically equivalent. Using the lower bound on the discrete local feature size of sample points developed in Lemma 5.4 below, it is easy to check whether $\mathrm{DMAT'_{in}} \cap \mathcal{A}_{\mathrm{out}} = \emptyset$.

5.2. Surface balls

In order to maintain correct topology of the piecewise linear surface reconstruction, the surface balls we generate have to

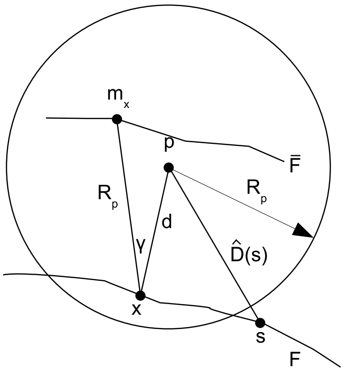

Figure 6: *Distance from pole p to the medial axis point m_x*

be large enough such that their union does not only cover S but also $\mathcal{F}$ and, on the other hand, these balls avoid the medial axis of the union of the balls in $\mathrm{DMAT}_{\mathrm{in}}$ and $\mathrm{DMAT}_{\mathrm{out}}$. The above restrictions limit the possible radii to a certain range. Maximizing the radii within this range will lead to a coarse result (which is desirable for seed polytopes), while minimizing the radii of the surface balls will lead to a faithful and detailed representation of the object. The choice of the radii determines the degree by which the surface balls are pruned in a subsequent set covering step.

5.2.1. Lower bound on the radii

To ensure that $\mathcal{F}$ is completely covered by surface balls we choose the radii of the surface balls such that they cover at least the intersection of their site's Voronoi cells with $\mathcal{F}$. For a point s in an r-sample, this intersection is covered by a sphere around s whose radius is $\rho \geq \frac{r}{1-r}\cdot \mathrm{lfs}(s)$, see [AB99], and so the surface balls need to have at least that radius. As $\mathrm{lfs}(s)$ is unknown, we need to estimate it in terms of the distance $\hat{D}(s)$ between s and the nearest among the poles of all sample points. Using Lemma 5.1 below, we get

$$\mathrm{lfs}(s) \leq 1.2802\cdot \hat{D}(s)$$

and so we must choose the radius ρ of a surface ball around s to be at least

$$\rho \geq \tfrac{r}{1-r}\cdot 1.2802\cdot \hat{D}(s).$$

The distance $\hat{D}(s)$ can be calculated relatively easily using a spatial search structure.

Lemma 5.1 Let $s\in S$ be a point of an r-sample S with $r \leq 0.08$, and let $\hat{D}(s) = \|s-p\|$ denote its distance to the nearest pole p. Then

$$\mathrm{lfs}(s) \leq 1.2802\cdot \hat{D}(s).$$

Proof The local feature size of s cannot be larger than $\hat{D}(s)$ plus the distance from p to the medial axis. To bound the latter distance for a specific value of r, we revisit the cases developed in Theorem 4.1 (and we use the notation introduced there). If case (2a) occurs we know that $\bar{F}$ contains a point $m_x \in \mathcal{M}_{\mathrm{in}}$ ($\mathcal{M}_{\mathrm{out}}$); see Figure 6. By Lemma 4.3a, the maximum angle between the touching point $x\in K_1$ of the medial ball centered at m_x and p is $\gamma = \angle m_x x p < 14.99°$ if $r \leq 0.08$. By Lemma 4.2,

$$d = \|x-p\| \geq \left(\sqrt{1-4(r^2-\tfrac{r^4}{4})} - r^2\right)\cdot R_p > 0.9807\cdot R_p.$$

Therefore

$$\|p-m_x\| \leq 2\cdot R_p\cdot \sin(\frac{\gamma}{2}) + (1-0.9807)R_p < 0.2802\cdot R_p$$

which is at most $0.2802\cdot\hat{D}(s)$ because s lies outside the polar ball centered at p. Otherwise, case (2b) occurs and by Lemma 4.2, p is not farther from $\mathcal{M}_{\mathrm{in}}$ ($\mathcal{M}_{\mathrm{out}}$) than

$$R_p\cdot(1-\sqrt{1-4\cdot(r^2-\tfrac{r^4}{4})}+r^2) < 0.0193\cdot R_p.$$

The lemma follows. □

5.2.2. Upper bound on the radii

To prevent surface balls from "different" parts of $\mathcal{F}$ from intersecting we want to ensure that they don't reach the discrete medial axis $\underline{DM}_{\mathrm{in}}$ (resp. DM_{out}). Thus, the discrete local feature size $\widetilde{\mathrm{lfs}}(s)$ is an upper bound on the radius that we can use. We will replace $\widetilde{\mathrm{lfs}}(s)$ by a smaller value, that is easier to compute, see Proposition 2.1.

Consequently, the minimum distance from s to any of the two weighted α-shapes is a lower bound on $\widetilde{\mathrm{lfs}}(s)$. Computing $\mathcal{A}_{\mathrm{in}}$ and $\mathcal{A}_{\mathrm{out}}$ and determining the minimum distance directly would consume too much time and memory, however. We show how to estimate this distance, again using the distance $\hat{D}(s)$ to the nearest pole to s.

Lemma 5.2 Let s be a sample point, and let v be a point with the following properties

- v lies in the Voronoi cell of s.
- v is not in the interior of the polar ball around the pole p of s that lies on the same side of $\mathcal{F}$ as v.

Then

(a) $\|v-s\| = O(r)\cdot\mathrm{lfs}(s)$. In particular, for $r = 0.08$, the distance to s is at most $0.123\cdot\mathrm{lfs}(s)$.
(b) The distance from v to the closest point $\bar{v}$ on the surface is $O(r^2\,\mathrm{lfs}(s)) = O(r^2\,\mathrm{lfs}(\bar{v}))$. For $r = 0.08$, the distance $\|v-\bar{v}\|$ is at most $0.0355\cdot\mathrm{lfs}(s) \leq 0.0424\cdot\mathrm{lfs}(\bar{v})$.

Lemma 5.3 Let pq be an edge of the weighted α-shape $\mathcal{A}_{\mathrm{in}}$ ($\mathcal{A}_{\mathrm{out}}$). Then the exterior angle of intersection between the polar balls B_q, B_p around p and q is at least $120°$.

Based on the preceding lemmas, it is possible to derive the following bound on $\widetilde{\mathrm{lfs}}(s)$.

Lemma 5.4 If m is a point on an edge pq of $\mathrm{DMAT}_{\mathrm{in}}$ (or in a triangle pqr of $\mathrm{DMAT_{in}}$) and v is outside or on the boundary of $U(\mathrm{DMAT_{in}})$ then

$$\|m-v\| \geq 0.817 \cdot \min\{\|p-v\|, \|q-v\|\},$$

(or $\|m-v\| \geq 0.817 \cdot \min\{\|p-v\|, \|q-v\|, \|r-v\|\}$, respectively).

The proofs for these lemmas are given in the appendix.

Corollary 5.5 Let $s \in S$ be a sample point, and let $\hat{D}(s)$ be its distance to the nearest pole. Then

$$\hat{D}(s) \geq \widetilde{\mathrm{lfs}}(s) \geq 0.817 \cdot \hat{D}(s).$$

Proof Since the poles are part of the discrete medial axis, the inequality $\widetilde{\mathrm{lfs}}(s) \leq D(s)$ is obvious. For the other direction, we bound $\widetilde{\mathrm{lfs}}$ by the distance from v to the weighted α-shape $\mathcal{A}$ of the polar balls, which contains the discrete medial axis. The proof of the lower bound on the ratio

$$\frac{\widetilde{\mathrm{lfs}}(v)}{D} = \frac{\|v-m\|}{D} \geq \max\left\{\frac{\|v-m\|}{\|v-p\|}, \frac{\|v-m\|}{\|v-q\|}\right\},$$

follows from Lemma 5.4. □

5.3. Topological Correctness

To show that the union $U(B_F)$ of surface balls is homotopy-equivalent to the surface $\mathcal{F}$, we follow the standard approach of using a fibration (a partition of $U(B_F)$ into a continuous family of curves, each intersecting $\mathcal{F}$ in a single point) and moving the boundaries of $U(B_F)$ along the fibers towards F.

The usual fibration by surface normals does not work since the medial axis might be closer than it appears from looking at the sample points, see Figure 2. Instead we use the fibers of the union $U(\mathrm{DMAT_{in}})$ of all polar balls. It is known that this union is homotopy-equivalent to $\mathcal{O}$, and its boundary is homotopy-equivalent to $\mathcal{F}$ [AK00].

The boundary of the union $U(\mathrm{DMAT_{in}})$ is not smooth, but still, it is in a certain sense "smooth from the inside" (it has no convex edges or vertices) and has therefore a reasonable fibration connecting the boundary to its inner medial axis $\mathrm{DMAT_{in}}$, see Figure 7. We concentrate on the inner discrete medial axis $\mathrm{DMAT_{in}}$; the outer discrete medial axis $\mathrm{DMAT_{out}}$ is treated analogously. The fibers are line segments that partition $U(\mathrm{DMAT_{in}}) \setminus DM_{\mathrm{in}}$, and they run from a surface point v on the boundary to a point m on the inner discrete medial axis DM_{in}. In three dimensions, there are three types of fibers: from a point v on a spherical patch of the boundary to a vertex m of the medial axis; from a point v on a circular edge formed as the intersection of two spheres to a point m on an edge of the medial axis; and from a vertex v of the boundary, formed as the intersection of three (or more) spheres to a point m on a face of the medial axis. Our proof treats all three cases uniformly.

We take the radius of the surface balls as $\rho\hat{D}(s)$ where the factor ρ can be chosen in the interval

$$\rho_{\min} = 0.24 \leq \rho \leq \rho_{\max} = 0.56. \qquad (1)$$

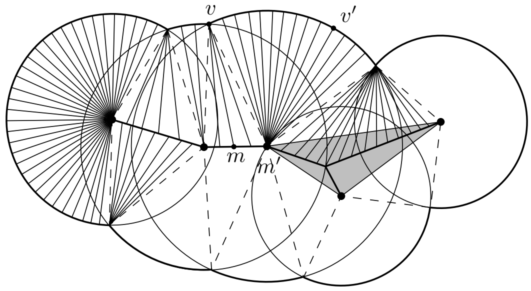

Figure 7: *Part of the fibration which is used to show isotopy. The shaded area is the weighted α-shape.*

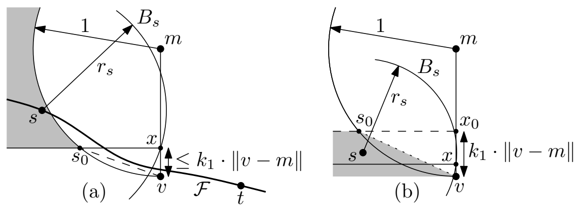

Figure 8: *A ball B_s that intersects the fiber vm improperly*

The upper bound ensures that the surface balls do not intersect the discrete medial axis, and the lower bound ensures that they are large enough to cover the surface completely. The bounds are stricter that would be required to reach only these two goals, since we also want to achieve ensure topological correctness of the union $U(B_F)$ of surface balls:

Lemma 5.6 If ρ is chosen in the interval (1), every fiber from a point v on the boundary of $U(\mathrm{DMAT_{in}})$ to a point m on the medial axis of $U(\mathrm{DMAT_{in}})$ starts in the union $U(B_F)$ of surface balls and intersects the boundary of $U(B_F)$ precisely once.

The lemma implies that the boundary of $U(B_F)$ can be continuously deformed along the fibers into the boundary of $U(\mathrm{DMAT_{in}})$, and thus the two boundaries are homotopy-equivalent. The boundary of $U(\mathrm{DMAT_{in}})$ is already known to be homotopy-equivalent to $\mathcal{F}$, and thus, the correct topology is established.

Proof For simplicity we prove the bound for $\rho = 0.3$. The calculation for general ρ is slightly more involved.
Let B_s be a surface ball around a sample point s such that the segment vm enters B_s in a point x, see Figure 8a. We will show that this does not lead to a violation of the lemma, because the segment vx is covered by the union of surface balls. We assume without loss of generality that vm is vertical and $\|m-v\| = 1$. We first show that x must have distance $\|x-v\| \leq k_1$ for $k_1 = 0.074$.
Suppose that this is not true. The medial ball of radius 1 around m is inside the union of balls, and hence it does not

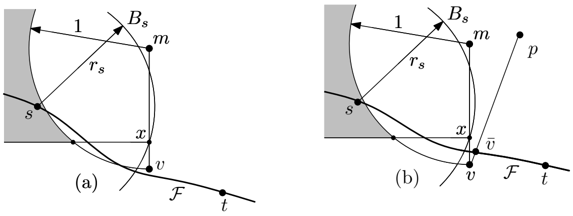

Figure 9: *A ball B_s that intersects the fiber vm improperly, v lies either inside $\mathcal{F}$ (a) or outside $\mathcal{F}$ (b)*

contain s: $||s-m|| \geq 1$. We claim that this implies

$$||s-x|| > 0.37 \cdot ||s-m||. \tag{2}$$

We know that s must lie outside the ball of radius 1 around m; s must also lie above the horizontal line through x. Thus, s is restricted to the shaded area in the figure. The ratio $||s-x||/||s-m||$ is minimized when x is as low as possible ($||x-v|| = k_1$) and s is at the lower right corner s_0 of this area. Here we have $||s-x||^2 + (1-k_1)^2 = 1$, from which one can compute $||s-x||/||s-m|| = ||s-x|| > 0.37$.

On the other hand, since $m \in \mathrm{DMAT_{in}} \subseteq \mathcal{A}_{\mathrm{in}}$, we have by definition $||s-m|| \geq \widetilde{\mathrm{lfs}}(s) \geq 0.817\hat{D}(s)$, by Lemma 5.4. Thus, the radius r_s of B_s is $r_s = ||s-x|| \leq \rho\hat{D}(s) \leq \rho/0.817 \cdot ||s-m|| < 0.368 \cdot ||s-m||$, contradicting (2).

Let us denote the extreme positions of s and x in the above analysis by s_0 and x_0.We have established that x and s lie below horizontal line s_0x_0, see see Figure 8b. For an arbitrary x and s we now claim

$$\frac{||s-x||}{||x-v||} \geq \frac{||s_0-x_0||}{||x_0-v||} \geq 5. \tag{3}$$

We know that s must always lie higher than x, For a fixed point x, we can rotate s around x until it lies at the same height as x, without changing the above ratio, So we can assume that s and x lie at the same height, with $||x-v|| \leq k_1$. The sample s cannot lie in the polar ball around m, and in particular, s must lie below the dotted line segment. The claim (3) follows.

Now to complete the proof we will show that the segment vx is covered by a surface ball, namely by the ball around the surface sample t closest to v. We are done if we can show that the radius r_t of this ball is at least $||t-v|| + ||v-x||$:

$$r_t = \rho\hat{D}(t) \geq ||t-v|| + ||v-x|| \tag{4}$$

This implies that $r_t \geq ||t-v||$ and $r_t \geq ||t-x||$ (by the triangle inequality), and thus ensures that the whole segment vx is covered. It establishes also that the starting point v of the fiber is covered, irrespective of whether another ball B_s intersects vm "in an improper way".

First we show that there is a sample point t with

$$||t-v|| \leq 0.123 \cdot \mathrm{lfs}(t) \tag{5}$$

We distinguish two cases:

(a) v lies inside $\mathcal{F}$ (on the same side as m), see Figure 9(a). Let t be the sample point closest to v. The point v satisfies the assumptions of Lemma 5.2 with respect to t: By definition, v lies in the Voronoi cell of t. Moreover, v lies in none of the polar balls around the vertices of $\mathrm{DMAT_{in}}$. Thus, by Lemma 5.2a, $||t-v|| \leq 0.123 \cdot \mathrm{lfs}(t)$.

(b) v lies outside $\mathcal{F}$, see Figure 9(b). By Lemma 5.4, there is a pole p in $\mathrm{DMAT_{in}}$ such that

$$||p-v|| \leq \frac{1}{0.817} \cdot ||m-v|| \leq 1.224 \cdot ||m-v||$$

The segment vp must intersect $\mathcal{F}$ in some point $\bar{v}$. Lemma 4.3b limits the penetration of the surface point $\bar{v}$ into the ball B_p:

$$||\bar{v}-v|| \leq (3/2 \cdot r^2 + O(r^3)) \cdot \mathrm{lfs}(\bar{v}).$$

In particular, for $r = 0.08$,

$$||\bar{v}-v|| \leq 0.0114 \cdot \mathrm{lfs}(\bar{v}).$$

The nearest sample point t from $\bar{v}$ is less than $r \cdot \mathrm{lfs}(t)$ away:

$$||\bar{v}-t|| \leq r \cdot \mathrm{lfs}(t)$$

The Lipschitz condition yields

$$\mathrm{lfs}(\bar{v}) \leq \mathrm{lfs}(t) + ||\bar{v}-t|| \leq (1+r) \cdot \mathrm{lfs}(t).$$

Therefore we get:

$$\begin{aligned} ||t-v|| &\leq ||v-\bar{v}|| + ||\bar{v}-t|| \\ &\leq 0.0114 \cdot \mathrm{lfs}(\bar{v}) + r \cdot \mathrm{lfs}(t) \\ &\leq 0.0114 \cdot (1+r)\,\mathrm{lfs}(t) + r \cdot \mathrm{lfs}(t) \\ &\leq 0.093\,\mathrm{lfs}(t) \leq 0.123\,\mathrm{lfs}(t) \end{aligned}$$

proving (5).

We have, by Lipschitz continuity, and using (3),

$$\begin{aligned} \hat{D}(t) &\geq \hat{D}(s) - ||s-x|| - ||x-v|| - ||v-t|| \\ &\geq ||s-x||/\rho - ||s-x|| - ||x-v|| - ||v-t|| \\ &\geq 5(1/\rho - 1)||x-v|| - ||x-v|| - ||v-t|| \\ &> 10.6 \cdot ||x-v|| - ||v-t|| \end{aligned} \tag{6}$$

By (5) and Lemma 5.1, we have $||v-t|| \leq 0.123 \cdot \mathrm{lfs}(t) \leq 0.123 \cdot 1.2802 \cdot \hat{D}(t) < 0.1575\hat{D}(t)$ and hence

$$\hat{D}(t) > 6.3 \cdot ||v-t|| \tag{7}$$

Multiplying (6) by 0.095, (7) by 0.175, and adding them together yields

$$0.27\hat{D}(t) \geq ||x-v|| + ||v-t||, \tag{8}$$

implying (4). □

6. Pruning by set covering

If we have a sample that is much denser than required by our conditions, we will get a correct "surface reconstruction", but we would like to obtain a coarser approximation to reduce the data, while maintaining topological correctness. We will therefore only use a subset of the surface balls.

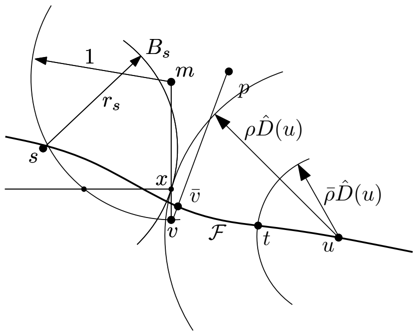

Figure 10: *The segment vm is covered by the enlarged ball around u.*

We establish a condition that is easy to check and guarantees the correct topology: As before, we use balls of radius $\rho\hat{D}(u)$ around surface points u; for each ball we also consider a shrunk copy of radius $\bar{\rho}\hat{D}(u)$, where $\bar{\rho} = 0.03 < \rho$. We can then prove the following statement.

Theorem 6.1 If the shrunk balls around the points u of a subset $S' \subseteq S$ cover all sample points S, then the union of the original balls (of radius $\rho\hat{D}(u)$) around these points is homotopy-equivalent to $\mathcal{F}$.

Proof The proof proceeds via the statement of Lemma 5.6. In that proof, we have established the existence of a sample point t that is close enough to v such that the ball around t covers the segment vx. This is extended to the present setting as follows: we can now no longer be sure that the ball around t is used, but there must be a (shrunk) ball around some sample point u that covers t. Then the (original) ball around u is large enough to guarantee that it reaches vx.
We know, by the pruning condition, that the covering contains a ball of radius $\rho\hat{D}(u)$ around a sample point u such that the shrunk ball with radius $\bar{\rho}\hat{D}(u)$ covers t:

$$\|u-t\| \leq \bar{\rho}\hat{D}(u)$$

From this, together with the above bound (8) on $\|t-x\|$, we obtain

$$\|u-x\| \leq \|u-t\|+\|t-x\| \leq \bar{\rho}\hat{D}(u)+(\rho-\bar{\rho})\hat{D}(u) = \rho\hat{D}(u),$$

and thus the ball B_u covers x. ☐

We try to select a minimum subset of surface balls whose shrunk copies cover the whole sample. This is an instance of the (in general NP-hard) set covering problem. In [AAH*07] and [AAHK09] a combination of exact and heuristic methods is described which yields not only an approximate solution but also a lower bound on the optimal solution, and in our setting the gap between them is typically quite small.

To get the input data for the set covering problem, the information about the sample points covered by each ball, we use a simple spatial search structure, e.g. a kd-tree.

The lemma remains true if the shrinking factor 0.03 is replaced by a smaller number. This parameter allows us to scale the algorithm to different levels of coarseness or refinement of the approximation. If the shrinking factor approaches 0, each shrunk ball will contain no sample points except its center, and thus the full sample will be used.

The small radius $0.03 \cdot \hat{D}$ that we have proved may not seem very impressive, but it must be seen in relation with the sampling constant $r = 0.08$. Thus, balls will start to be eliminated as soon at the actual sampling density exceeds the required minimum by a factor of about 4–5 (in terms of the sampling radius).

The same approach works for approximating the medial axis. Here we start with an enlarged set of *polar balls* DMAT'_{in}, and produce an (almost) minimum subset DMAT''_{in} whose union covers S.

7. Experimental data

Due to lack of space, we only include two examples showing the output produced by our implementations, one for surface reconstruction and one for medial axis approximation.

Figure 11 illustrates how different choices of radii for surface balls lead to different levels of detail in the approximating polyhedral surface mesh. The initial point cloud for this 'double torus' model consists of 85237 points. Due to the effect of pruning, the mesh for the big ring is more and more coarsened, whereas the necessary details are preserved for the small ring. The running times for these computations (for a single threaded application on a Core2 Duo E6700 CPU) are shown in Table 1. Filtered floating point arithmetic has been used.

Figure	11abc	11def	11ghi
Surface balls	55s	55s	55s
Pruning	-	35s	159s
# Remaining balls	85237	4198	549
Weighted α-shape	217s	7s	1s

Table 1: *Runtimes for the double torus model in Figure 11*

We have implemented the medial axis algorithm for balls in [AK01] with CGAL [CGA] and have used it to compute the exact medial axis of the union of the balls in the set DMAT''_{in}. The output is a topologically correct approximation of the medial axis of the original object. The level of simplification is tuned by the parameter ε which specifies how much to grow the radii before the pruning. Figure 12 (model provided by the AIM@SHAPE Repository [AIM]) shows four pruned medial axis transforms and medial axes, computed from a set of 39779 polar balls using different values of ε. Table 2 shows the elapsed runtimes (in seconds) on the same computing platform as before.

The observed runtimes are practical for moderately large

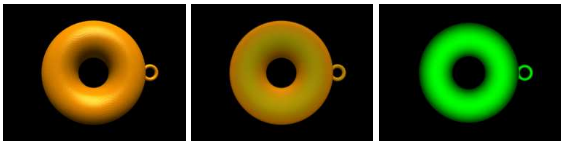

(a) Without pruning: 85237 balls
(b) Transparent
(c) Mesh on top of 85237 vertices

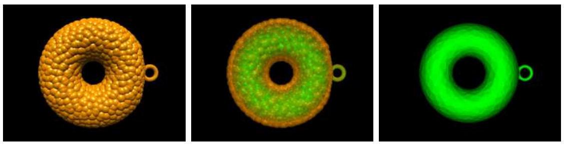

(d) After moderate pruning: 4198 balls
(e) Transparent
(f) Mesh on top of 4198 vertices

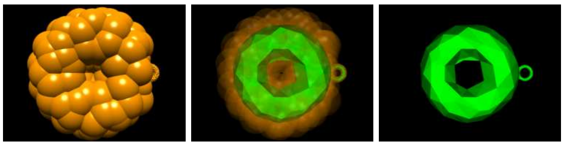

(g) After heavy pruning: 549 balls
(h) Transparent
(i) Mesh on top of 549 vertices

Figure 11: *Double torus reconstruction*

Figure	12(a)	12(b)	12(c)	12(d)
Polar balls	87.1s	87.1s	87.1s	87.1s
Pruning	151.2s	207.5s	289.3s	340.9s
Medial axis	152.2s	25.7s	4.2s	1.3s

Table 2: *Runtimes for the medial axes in Figure 12*

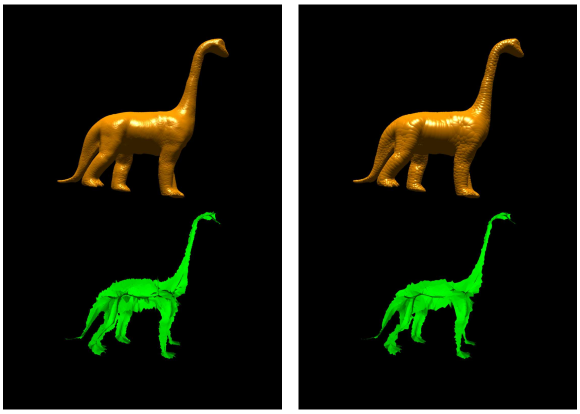

(a) 5917 Balls
(b) 1692 Balls

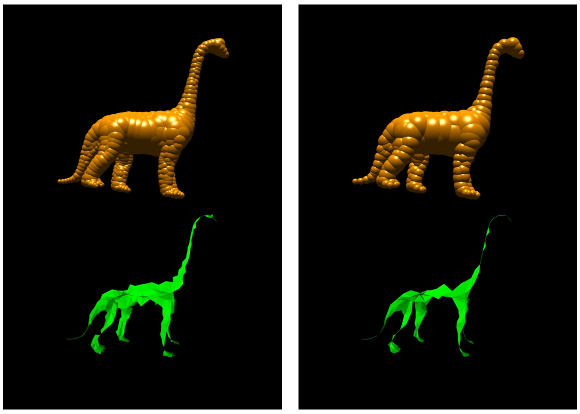

(c) 375 Balls
(d) 170 Balls

Figure 12: *Pruned polar balls and their medial axes*

data sets, but naturally cannot compete with mesh reconstruction methods that do not come with a topological guarantee (see e.g. [KBH06]) or with medial axis algorithms which are not scalable [SFM07]. Still, our approach compares well with mesh reconstruction methods with guarantee; see e.g. [DGH01]. The strength of our method lies in combining topological correctness with scalability.

Acknowledgements. This work has been partially supported by the FWF Joint Research Program 'Industrial Geometry' S9205-N12, and by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-006413 (ACS–Algorithms for Complex Shapes).

References

[AAH*07] Aichholzer O., Aurenhammer F., Hackl T., Kornberger B., Peternell M., Pottmann H.: Approximating boundary-triangulated objects with balls. In *Proc. 23rd European Workshop on Computational Geometry* (Graz, 2007), pp. 130–133.

[AAHK09] Aichholzer O., Aurenhammer F., Hackl T., Kornberger B.: Scalable piecewise linear approximations of 3d medial axes. Manuscript, 2009.

[AB99] Amenta N., Bern M.: Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry 22* (1999), 481–504.

[AB03] Attali D., Boissonnat J.-D.: Complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete & Computational Geometry 30* (2003), 437–452.

[ABE07] Attali D., Boissonnat J.-D., Edelsbrunner H.: Stability and computation of medial axes—a state-of-the-art report. In *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration* (2007), Müller T., Hamann B., Russell B., (Eds.), Springer Series on Mathematics and Visualization.

[ACK01] Amenta N., Choi S., Kolluri R. K.: The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications 19* (2001), 127–153.

[AIM] Aim@shape shape repository. `http://shapes.aim-at-shape.net/`.

[AK00] Amenta N., Kolluri R.: Accurate and efficient unions of balls. In *Proc. 16th Ann. Symp. Computational Geometry* (Hong Kong, 2000), ACM, pp. 119–128.

[AK01] Amenta N., Kolluri R.: The medial axis of a union of balls. *Computational Geometry: Theory and Applications 20* (2001), 25–37.

[BC01] BOISSONNAT J.-D., CAZALS F.: Natural neighbor coordinates for points on a surface. *Computational Geometry: Theory and Applications 19* (2001), 155–173.

[BO04] BRADSHAW G., O'SULLIVAN C.: Adaptive medial-axis approximation for sphere-tree construction. *ACM Transactions on Graphics 23* (2004), 1–26.

[BPR*07] BUCHIN K., PLANTINGA S., ROTE G., STURM A., VEGTER G.: Convex approximation by spherical patches. In *Proc. 23rd European Workshop on Computational Geometry* (Graz, 2007), pp. 26–29.

[CCSL09] CHAZALS F., COHEN-STEINER D., LIEUTIER A.: A sampling theory for compact sets in Euclidean spaces. *Discrete & Computational Geometry 41* (2009), 461–479.

[CGA] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org/.

[CL08] CHAZAL F., LIEUTIER A.: Smooth manifold reconstruction from noisy and non-uniform approximation with guarantees. *Computational Geometry: Theory and Applications 40* (2008), 156–170.

[DGH01] DEY T. K., GIESEN J., HUDSON J.: Delaunay based shape reconstruction from large data. In *Proc. IEEE Symp. Parallel and Large-Data Visualization and Graphics* (San Diego, 2001), pp. 19–27.

[Ede95] EDELSBRUNNER H.: The union of balls and its dual shape. *Discrete & Computational Geometry 13* (1995), 415–440.

[GMP07] GIESEN J., MIKLOS B., PAULY M.: Medial approximation of planar shapes from union of balls: A simpler and more robust algorithm. In *Prof. 19th Canad. Conf. Comput. Geom. (CCCG)* (2007), pp. 105–108.

[GO08] GUIBAS L. J., OUDOT S. Y.: Reconstruction using witness complexes. *Discrete & Computational Geometry 40*, 3 (2008), 325–356.

[Hub96] HUBBARD P.: Approximating polyhedra with spheres for time-critical collision detection. *ACM Transactions on Graphics 15* (1996), 179–210.

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Eurographis Symposium on Geometry Processing* (2006), pp. 61–70.

[NSW08] NIYOGI P., SMALE S., WEINBERGER S.: Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry 39* (2008), 419–441.

[RF96] RANJAN V., FOURNIER A.: Matching and interpolation of shapes using unions of circles. *Computer Graphics Forum 15* (1996), 129–142.

[SFM07] SUD A., FOSKEY M., MANOCHA D.: Homotopy-preserving medial axis simplification. *Int. J. Computational Geometry & Applications 17* (2007), 423–451.

[SS04] SCHARF A., SHAMIR A.: Feature-sensitive 3d shape matching. In *Proc. Computer Graphics International (CGI'04)* (2004), pp. 1530–1552.

[YBM04] YANG Y., BROCK O., MOLL R.: Efficient and robust computation of an approximated medial axis. In *Proc. 9th ACM Symp. Solid Modeling and Applications* (2004), pp. 15–24.

Appendix A: Proofs of technical lemmas

Lemma A.1 (Lemma 5.2) Let s be a sample point, and let v be a point with the following properties

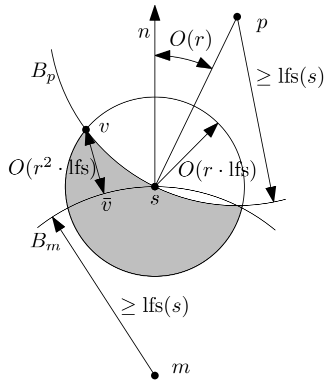

Figure 13: *A point v that is not covered by the polar ball must lie close to the surface.*

- v lies in the Voronoi cell of s.
- v is not in the interior of the polar ball around the pole p of s that lies on the same side of $\mathcal{F}$ as v.

Then

(a) $\|v-s\| = O(r)\cdot \mathrm{lfs}(s)$. In particular, for $r = 0.08$, the distance to s is at most $0.123\cdot \mathrm{lfs}(s)$.

(b) The distance from v to the closest point $\bar{v}$ on the surface is $O(r^2\,\mathrm{lfs}(s)) = O(r^2\,\mathrm{lfs}(\bar{v}))$. For $r = 0.08$, the distance $\|v-\bar{v}\|$ is at most $0.0355\cdot \mathrm{lfs}(s) \leq 0.0424\cdot \mathrm{lfs}(\bar{v})$.

Proof We perform the calculation for $r = 0.08$, and only indicate the asymptotic dependence on r. We will first show part (a). Let p be the pole of s on the same side of the surface as v. If $\|v-s\| > kr\cdot \mathrm{lfs}(s)$ for $k = 1.536$, the angle between sv and the surface normal is at most $\arcsin\frac{1}{k(1-r)} + \arcsin\frac{r}{1-r} < 47.2^\circ$, see [AB99, Lemma 4]. Similarly, the angle between the normal and sp is at most $2\arcsin\frac{r}{1-r} < 12.8^\circ$. In total the angle vsp is less than 60°. Since $\|v-s\| \leq \|p-s\|$, by the definition of the pole, it follows that v must be contained in the polar ball around p, whose radius is $\|p-s\|$, a contradiction. We thus conclude that v is contained in a ball of radius

$$kr\cdot \mathrm{lfs}(s) \leq 0.123\cdot \mathrm{lfs}(s) \quad (= O(r\,\mathrm{lfs}(s)))$$

around s. Since v avoids the polar ball B_p around p, it lies in the shaded region indicated in Figure 13. The direction sp of the polar ball deviates at most $2\arcsin\frac{r}{1-r} < 12.8^\circ$ $(= O(r))$ from the normal direction n at s. Thus the "highest" possible position of v is as indicated in the figure. We know that the surface must pass above the opposite medial ball P_m of s, and thus we can estimate the distance from v to the surface and prove (b). A straightforward calculation gives the bound $\|v-\bar{v}\| \leq 0.0355\,\mathrm{lfs}(s)$ $(= O(r^2\,\mathrm{lfs}(s)))$. By the Lipschitz condition,

$$0.0355\,\mathrm{lfs}(s) \leq \frac{0.0355}{1-0.123-0.0355}\,\mathrm{lfs}(\bar{v}) \leq 0.0424\cdot \mathrm{lfs}(\bar{v})$$

is obtained. □

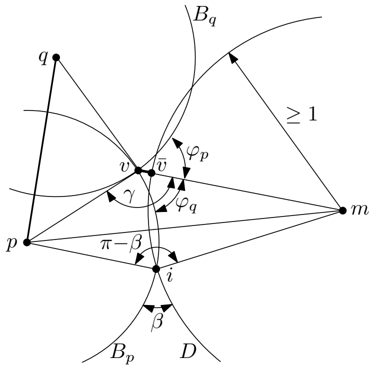

Figure 14: *Schematic figure of an intersection of two polar balls such that their intersection point v is not covered by the union of polar balls.*

Lemma A.2 (Lemma 5.3) Let pq be an edge of the weighted α-shape $\mathcal{A}_{\text{in}}$ ($\mathcal{A}_{\text{out}}$). Then the exterior angle of intersection between the polar balls B_q, B_p around p and q is at least 120°.

Proof Since pq is an edge of the weighted α-shape, there is a point v on the intersection of the boundaries of the two polar balls B_p and B_q which is not covered by any other polar ball, see Figure 14. Therefore, the neighborhood of v contains points outside all polar balls and, by Lemma 5.2. v is close to $\mathcal{F}$: For the closest surface point $\bar{v}$ we have

$$d = \|v-\bar{v}\| \leq 0.0424 \cdot \operatorname{lfs}(\bar{v}).$$

Without loss of generality, we assume $\operatorname{lfs}(\bar{v}) = 1$. Consider the medial ball B of $\bar{v}$ on the opposite site, with center m and radius $\|\bar{v}-m\| \leq \operatorname{lfs}(\bar{v}) = 1$. By [ACK01, Lemma 17], a polar ball B_p or B_q intersects a medial ball D on the opposite site at angle $\beta \leq 2\arcsin 2r$. Let us focus on one ball B_p and the angle ϕ_p between this ball and the surface normal vm. The other ball is treated in the same way, and the total exterior angle is then $\phi_p + \phi_q$.
We have $\phi_p = \gamma - \pi$, where $\gamma = \angle pvm$. To get an upper bound on ϕ_p (or on γ), let us fix the angle γ and try to find circles B_p and D that are consistent with this situation. We have the following constraints:

(i) $1 = \operatorname{lfs}(\bar{v}) \geq \|\bar{v}-m\|$;
(ii) $d := \|v-\bar{v}\| \leq 0.0424 \cdot \operatorname{lfs}(\bar{v}) \leq 0.0424$;
(iii) The intersection angle between B_p and D is $\beta \leq 2\arcsin 2r$.

This gives us a distance $\|c-v\| = 1+d$, using the triangle inequality we get $\|q-v\| = 1-d$. For the triangle qcv only the segment qc is of unknown length. We consider also a second triangle, formed by the points q, c and one intersection point i of the medial ball with the polar ball B_q. Again only the distance of the segment qc is unknown. From the

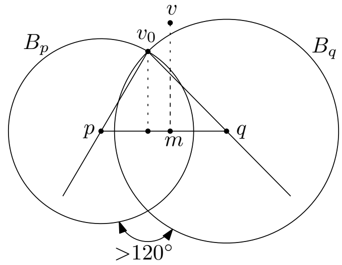

Figure 15: *The distance from the sample point s to the weighted α-shape*

triangles we get the following equations:

$$\cos\beta = \frac{1+(1-d)^2-\|c-q\|^2}{2(1-d)}, \qquad \cos\gamma = \frac{(1+d)^2+(1-d)^2-\|c-v\|^2}{2(1-d)(1+d)},$$

for $\beta = \angle cvq = \pi - \beta = \pi - 2\arcsin 2r$, $\gamma = \angle qic$, $d = 0.0355$. Solving these equations for γ gives an angle $\varphi = 2\cdot(\gamma - \pi/2) > 120^\circ$. $\square$

Lemma A.3 (Lemma 5.4) If m is a point on an edge pq of DMAT_{in} (or in a triangle pqr of DMAT_{in}) and v is outside or on the boundary of $U(\text{DMAT}_{\text{in}})$ then

$$\|m-v\| \geq 0.817 \cdot \min\{\|p-v\|, \|q-v\|\},$$

(or $\|m-v\| \geq 0.817 \cdot \min\{\|p-v\|, \|q-v\|, \|r-v\|\}$, respectively).

Proof We first consider the case when m lies on an *edge* pq,as illustrated in Figure 15. Let m' be the point on pq that is closest to v. If m' is one of the endpoints p or q, we are done:

$$\|m-v\| \geq \|m'-v\| = \min\{\|p-v\|, \|q-v\|\}.$$

Otherwise we know that $m'-v$ is perpendicular to pq. We know from Lemma A.2 that the intersection of the two polar balls B_p and B_q cannot be too thin: their angle of intersection is at least 120°. For fixed balls B_p and B_q, the angles and hence the ratios are minimized when s lies on the intersection between the balls (the point v_0 in the figure).
Now keeping v_0 fixed at the intersection and considering a variation of the balls B_p and B_q, maintaining $\min\{\|v-p\|, \|v-q\|\}$, it is clear that the distance from v to the edge pq is minimized when the angle $\angle pvq$ is at its upper bound of 60° and the two distances are equal: $\|v-p\| = \|v-q\|$. Then the ratio $\|v-v\|/\|v-p\| = \cos 30^\circ > 0.866$.
Now consider the case when m lies in a *triangle* pqr. If the point m' on pqr that is closest to v lies on an edge or at a vertex of the triangle, we have reduced the problem to the previous case. Otherwise we know that $m'-v$ is perpendicular to pqr. The remaining argument is similar as in the case of an edge: The extreme situation is a triangular pyramid with equal angles $\angle pvq = \angle qvr = \angle rvp = 60^\circ$ at the apex m and equal sides $\|p-v\| = \|q-v\| = \|r-v\|$. The ratio between the height of this pyramid and the length $\|p-v\|$ is $\sqrt{(1+2\cos 60^\circ)/3} > 0.817$. $\square$