

Approximating the Medial Axis from the Voronoi Diagram with a Convergence Guarantee

Tamal K. Dey Wulue Zhao

Department of CIS, Ohio State University, Columbus, OH 43210, USA
email: {tamaldey,zhaow}@cis.ohio-state.edu

Abstract

The medial axis of a surface in 3D is the closure of all points that have two or more closest points on the surface. It is an essential geometric structure in a number of applications involving 3D geometric shapes. Since exact computation of the medial axis is difficult in general, efforts continue to improve their approximations. Voronoi diagrams turn out to be useful for this approximation. Although it is known that Voronoi vertices for a sample of points from a curve in 2D approximate its medial axis, similar result does not hold in 3D. Recently, it has been discovered that only a subset of Voronoi vertices converge to the medial axis as sample density approaches infinity. However, most applications need a non-discrete approximation as opposed to a discrete one. To date no known algorithm can compute this approximation straight from the Voronoi diagram with a guarantee of convergence. We present such an algorithm and its convergence analysis in this paper. One salient feature of the algorithm is that it is scale and density independent. Experimental results corroborate our theoretical claims.

Keywords: Medial axis, Geometric Modeling, samples, Voronoi diagram, Delaunay triangulation.

1 Introduction

The medial axis of a shape is defined when the shape is embedded in an Euclidean space and is endowed with a distance function. Informally, it is the set of all points that have more than one closest point on the shape. The medial axis provides a compact representation of the shapes which has been used in a number of applications including image processing [22], computer vision [9, 23], solid modeling [19, 20, 28], mesh generation [25, 26], motion planning [18] and many others [21, 29]. The shapes in this paper are surfaces embedded in three dimensions.

Application demands have prompted research in the computational as well as the mathematical aspects of the medial axis in recent years. As a mathematical structure they are unstable since a small change in shape can cause a relatively large change in its medial axis [17, 30]. They are hard

to compute exactly due to numerical instability associated with their computations. Few algorithms, and only for special classes of shapes, have been designed till date to compute the exact medial axis [12, 20]. Consequently, efforts have been made to approximate the medial axis. For polyhedral input Etzion and Rappoport [16] suggest an approximation method based on octree subdivisions of space. Another scheme considered by many uses a set of sample points on the shape and then approximates the medial axis with the Voronoi diagram of these points [4, 5, 6, 11, 25, 29].

We follow the Voronoi diagram approach. It is particularly suitable for point cloud data, which are increasingly being used for geometric modeling over a wide range of applications. It is known that the Voronoi vertices approximate the medial axis of a curve in 2D. In fact, Schmitt [24] and Brandt [10] show that if the sample density approaches infinity, the Voronoi vertices in this case converge to the medial axis. Unfortunately, the same is not true in three dimensions. Amenta, Bern and Eppstein [2] observe that some Voronoi vertices, the centers of the flat tetrahedra called ‘slivers’, can come close to the surface no matter how dense a sample is. In order to alleviate this problem in the context of surface reconstruction, Amenta and Bern [1] identify some Voronoi vertices called ‘poles’ that remain far from the surface. These poles are the farthest Voronoi vertices from the sample points in their Voronoi cells. Boissonnat and Cazals [7] and Amenta, Choi and Kolluri [4] show that the poles indeed lie close to the medial axis and converge to it as the sample density approaches infinity.

The convergence result of poles to the medial axis is a significant progress in the medial axis approximation in 3D. However, many applications require and often prefer a non-discrete approximation rather than a discrete one. In 2D, Brandt and Algazi [11] achieve this by retaining a subset of Voronoi edges incident to the Voronoi vertices. In 3D, since poles lie close to the medial axis, Amenta, Choi and Kolluri [4] design an algorithm that connects them with a cell complex. They consider the Delaunay balls centering the poles and then compute the medial axis of the boundary of the union of these balls by the weighted Delaunay triangulation of the poles with the radii of the Delaunay balls as weights. This is the first algorithm that produces a continuous approximation of the medial axis with theoretical guarantees in 3D.

However, this method requires a second Voronoi diagram to compute the medial axis and, more importantly, produces noisy medial axis in some cases. Heuristics have been proposed to clean up these noisy medial axes, but these heuristics are not scale independent.

In this paper we propose to approximate the medial axis directly from the Voronoi diagram. Approximating the medial axis straight from the Voronoi diagram in 3D has been proposed in the past. In a nice work, Attali and Lachaud [5] (also see [6]) show how to prune the Voronoi diagram with an angle and length criterion to approximate the medial axis. Although the strategy achieves good results in many cases, the pruning is scale dependent and more seriously depends on the sampling density. Consequently, one needs to fine tune the pruning parameters individually for each data set and it is not clear if these strategies are apt for a data set where the density varies in different parts of the shape.

Our algorithm also uses two criteria to select the Voronoi facets from the Voronoi diagram. But, unlike [5, 6], these two criteria are scale and density independent and the algorithm has a convergence guarantee. We filter Delaunay edges from the Delaunay triangulation of the sample points and then output their dual Voronoi facets as an approximate medial axis. The approximation depends upon a sampling density parameter ε that tends to zero as the sampling density on the surface S approaches infinity. A subset of the medial axis satisfying certain condition that depends on ε is approximated by a point in the output. As ε approaches zero with increasing density, this subset coincides with the complete medial axis meaning that all points in the medial axis are approached by a point in the output. The thresholds used for two criteria to filter the Delaunay edges remain fixed over data sets of different densities. Thus, there is no need for fine tuning the parameter values. Our experiments with different data sets also support our theoretical claims.

The rest of the paper is organized as follows. Section 2 contains preliminaries and definitions that we use later. Section 3 details the two conditions that we use to filter the Delaunay edges and describes the algorithm. In section 4 we prove the guarantees of convergence. Section 5 contains the experimental results and we conclude in section 6.

2 Preliminaries and Definitions

Let P be a point sample from a smooth compact surface $S \subset \mathbb{R}^3$ without boundary. A ball is called *medial* if it meets S only tangentially in at least two points. The medial axis of S is defined as the closure of the set of centers of all medial balls. Each point on S has two medial balls, one touching it from outside and the other touching it from inside. It follows that the line going through a point $p \in S$ and the centers of its medial balls is normal to S at p . See Figure 1 for an illustration in 2D.

Obviously, the medial axis of S can be approximated from a sample P only if it is dense enough to carry information




Figure 1: A curve (dashed), some medial balls (dotted) and the medial axis (solid) in 2D.

about the features of S . Following Amenta and Bern [1] we define the local feature size $f()$ as a function $f : S \rightarrow \mathbb{R}$ where $f(x)$ is the distance of $x \in S$ to the medial axis. Intuitively, $f()$ measures how complicated S is locally. It is known that the function $f()$ is 1-Lipschitz continuous, i.e., $f(p) \leq f(q) + \|p - q\|$ for any two points p, q in S [1]. A sample is an ε -sample if each point $x \in S$ has a sample point within $\varepsilon f(x)$ distance. Generally, in practice, a sample is dense if $\varepsilon < 0.25$ though in most cases theoretical guarantees require much smaller values.

The Voronoi diagram and its dual, the Delaunay triangulation, play a key role in capturing information about shapes. This observation has led to a number of algorithms for the related problem of surface reconstruction which exploit the structures of these diagrams [1, 3, 7, 14, 15]. The Voronoi diagram V_P for a point set $P \in \mathbb{R}^3$ is a cell complex consisting of Voronoi cells $\{V_p\}_{p \in P}$ and their facets, edges and vertices, where $V_p = \{x \in \mathbb{R}^3 \mid \|p - x\| \leq \|q - x\|, \forall q \in P\}$. The dual complex, D_P , called the Delaunay triangulation of P , consists of Delaunay tetrahedra and their incident triangles, edges and vertices. A Delaunay tetrahedron is dual to a Voronoi vertex, a Delaunay triangle is dual to a Voronoi edge, a Delaunay edge is dual to a Voronoi facet and a Delaunay vertex is dual to a Voronoi cell. We say $e = \text{Dual } g$ if e and g are dual to each other in the two diagrams. It is an important result proved by Amenta and Bern [1] that the Voronoi cells are elongated along the normal direction to the surface at the sample points if the sample is sufficiently dense. The definition of *poles* plays an important role in approximating these normals.

Definition 1 The pole p^+ of a sample point p is the farthest Voronoi vertex in the Voronoi cell V_p . If V_p is unbounded, p^+ is taken at infinity. The vector $\mathbf{v}_p = p^+ - p$ is called the pole vector for p and its direction is taken as the average of all directions of infinite edges in case V_p is unbounded, see Figure 2.

It is proved that the pole vector \mathbf{v}_p approximates the normal \mathbf{n}_p to the surface S at p up to orientation [1].

Definition 2 The tangent polygon for a sample point is defined as the polygon in which the plane through p with \mathbf{v}_p as normal intersects V_p . See Figure 2(a) for an illustration.

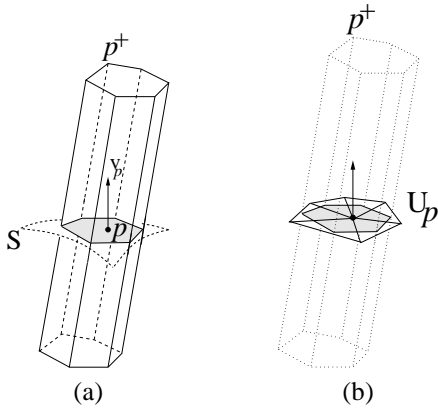


Figure 2: A Voronoi cell V_p . The corresponding pole, pole vector, tangent polygon (a), and the umbrella (b).

Since \mathbf{v}_p approximates \mathbf{n}_p , the tangent polygon approximates the tangent plane at p restricted within V_p . We define a dual structure to the tangent polygon from the Delaunay triangulation D_p .

Definition 3 The umbrella U_p for a sample point p is defined as the topological disc made by the Delaunay triangles incident to p that are dual to the Voronoi edges intersected by the tangent polygon. See Figure 2(b).

The umbrella U_p approximates the surface locally at p . The triangles in the umbrella lie very close to the restricted Delaunay triangles that are dual to the Voronoi edges in V_p intersecting the surface.

Notations. In what follows we use the following notations. The notation $\angle \mathbf{u}, \mathbf{v}$ denotes the acute angle between the lines supporting two vectors \mathbf{u} and \mathbf{v} . The tangent vector going from a point p to q is denoted with \mathbf{t}_{pq} . The normal to a triangle pqr is \mathbf{n}_{pqr} and its circumradius is R_{pqr} .

3 Algorithm

Our aim is to approximate the medial axis with a subset of Voronoi facets and their closures. The closure ClF of a Voronoi facet F is the set of all incident Voronoi edges and vertices of F and F itself. We choose these Voronoi facets as the dual of a set of selected Delaunay edges. This means we need some conditions to filter these Delaunay edges from D_p . Let us examine a medial ball B closely to determine which Delaunay edges we should select. Consider Figure 3. The segment pq makes an angle θ with the tangent plane at p and q where the medial ball touches the surface S . If B touches S in more than two points, let p and q be such that the angle θ is maximal. We associate each medial axis point m , and also the points where B meets S , with such an angle θ , which we call their *medial angle*. The medial axis is approximated by Voronoi facets dual to the Delaunay edges.




Figure 3: A medial axis point m , its medial angle θ and the corresponding medial ball.

These edges are either long edges making large angle with the surface - thus selected by the so-called Angle condition. When the Delaunay edges tend to become parallel to the tangent planes at their endpoints, the Angle condition fails and we resort to a Ratio condition; more precisely the edges used there are those whose length is significantly larger than the circumradii of the umbrella triangles.

3.1 Angle condition

Approximation of the medial angle θ for a medial axis point requires an approximation to the tangent plane at that point. It follows from Lemma 3 that the triangles in the umbrellas necessarily lie flat to S .

Therefore, we take umbrella triangles in U_p for approximating the tangent plane at a sample point p and determine all Delaunay edges pq that make relatively large angle with this tangent plane. The angle between an edge pq and a triangle ptu is measured by the acute angle $\angle \mathbf{n}_{ptu}, \mathbf{t}_{pq}$. Our goal is to capture all edges pq that make more than a threshold angle θ with triangles in U_p , or equivalently make an angle less than $\frac{\pi}{2} - \theta$ with their normals, see Figure 4. We say a Delaunay edge pq satisfies *Angle condition* $[\theta]$ if

$$\max_{ptu \in U_p} \angle \mathbf{n}_{ptu}, \mathbf{t}_{pq} < \frac{\pi}{2} - \theta.$$

The roles of p and q are interchangeable, i.e., if the umbrella triangles in U_q are used in the condition, we say qp satisfies the Angle condition $[\theta]$. Consistent with this notation the edge pq is considered twice in the algorithm, once as pq and another time as qp .

Only Angle condition cannot approximate the medial axis in a density independent manner. If we fix θ for all models, some of the medial axis points with medial angle below θ are not approximated. In that case we cannot hope for convergence in the limit when density approaches infinity.

Figure 5 illustrates the results with ‘only Angle condition’. We varied the value of θ in order to get a good approximation to the medial axis. As expected, larger value of θ produces less facets in the output resulting in undesirable ‘holes’, see the rightmost picture of 3HOLE data in Figure 5. On the other hand, smaller values of θ generate too many facets resulting in undesirable ‘spikes’, see the leftmost picture for the




Figure 5: Results with ‘only Angle condition’.




Figure 4: The angle of interest for the Angle condition $[\theta]$. The Voronoi edge shown with the dotted line is normal to the shaded triangle.

3HOLE data in Figure 5. For each model shown in Figure 5 we attempted to determine a value of θ for which we can obtain an approximation as good as possible. The second row of Figure 5 shows the output of this experiment. The major drawback of the ‘only Angle condition’ approach is that the value of θ for which we obtain good approximation differs from sample to sample. It turns out that the required value of θ gets larger with decreasing sample density. Consequently, we could not find any consistent value for which the approach works for all models we experimented with.

3.2 Ratio condition

Consider the medial ball B as shown in Figure 3. From simple geometry of spheres, we get

$$\|p - q\| = 2\mu \sin \theta$$

where $\mu = \|m - p\|$ is the radius of B . Therefore, if $\theta \gg \varepsilon$, where ε is the sampling density, we have $\|p - q\| \gg 2\mu \sin \varepsilon$. Also, the radius of the umbrella triangles are only of the order of $\mu\varepsilon$ as we show later. This means that the length of pq will be much larger than the radii of the umbrella triangles.

Taking the cue from the above observation, we compare the length of the Delaunay edges with the circumradii of the umbrella triangles. By this we can approximate all medial axis points with medial angles only few times larger than ε . Of course, we will not be able to approximate the medial axis points with medial angle less than ε with this method, but as ε approaches zero, we get the required convergence. We say an edge pq satisfies the *Ratio condition* $[\rho]$ if (see Figure 6)

$$\min_{ptu \in U_p} \frac{\|p - q\|}{R_{ptu}} > \rho.$$

Here, too, the roles of p and q are interchangeable.

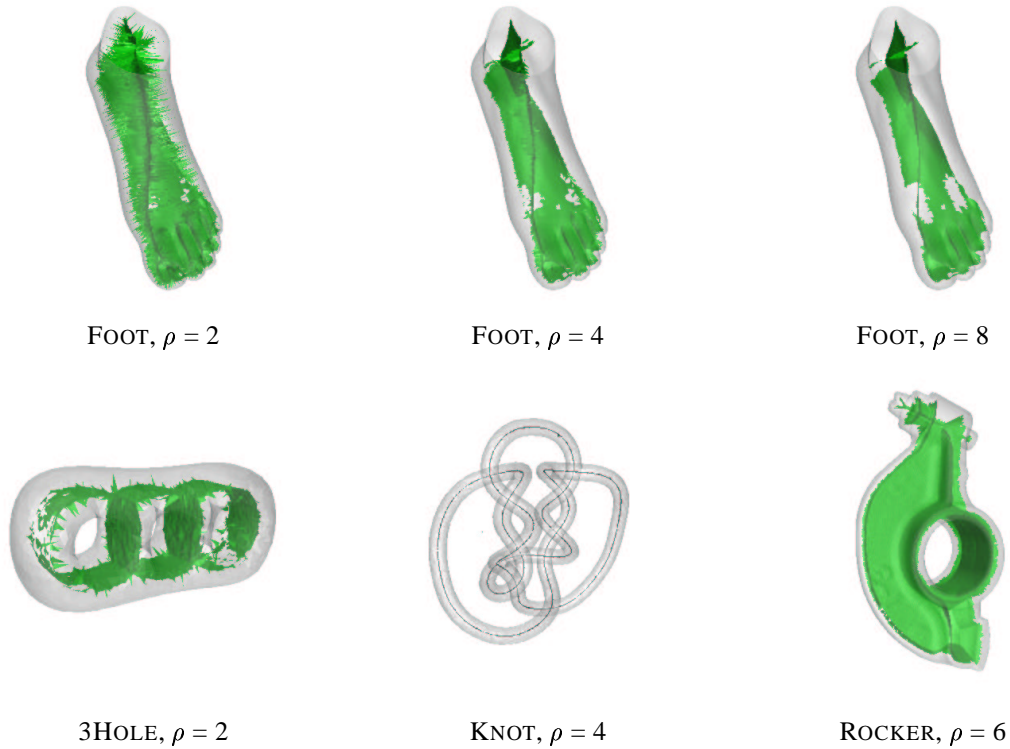


Figure 7: Results with ‘only Ratio condition’.




Figure 6: Radius of interest for the Ratio condition $[\rho]$.

Figure 7 illustrates the effect of ‘only Ratio condition’. As expected, larger value of ρ produces less noisy medial axis, but with ‘holes’ as shown in the rightmost picture for the FOOT data in Figure 7. On the other hand, smaller ρ produces undesirable spurious ‘spikes’ as exhibited by the leftmost picture of the FOOT in Figure 7.

In this case also we could not find a value of ρ for which all samples produce good result. When the sample is less dense, a smaller ρ is needed to obtain an approximation as good as possible. ‘Only Ratio condition’ approach has one more disadvantage. If the sample density is not uniform over the entire surface, one value of ρ cannot capture the medial axis for the entire shape. This is why no value of ρ gave a good result for the FOOT and 3HOLE data though we could find an appropriate ρ for ROCKER and KNOT data which are

mostly uniform.

Although, in theory, for sufficiently dense sample, the ratio condition alone is adequate to filter all necessary Delaunay edges, often this density requirement is not satisfied in practice. Our experiments suggest that both angle and ratio conditions are necessary in practice to keep the thresholds fixed over varying sample density. Experimentally we observe that $\theta = \theta_0 = \frac{\pi}{8}$ and $\rho = \rho_0 = 8$ are appropriate for all reasonably dense data sets. With these two values we enumerate the steps of our algorithm MEDIAL to approximate the medial axis.

```

MEDIAL( $P$ )
1 Compute  $V_P$  and  $D_P$ ;
2  $F = \emptyset$ ;
3 for each  $p \in P$ 
4   Compute  $U_p$ ;
5   for each Delaunay edge  $pq \in U_p$ 
6     if  $pq$  satisfies Angle Condition  $[\theta_0]$ 
7       or Ratio condition  $[\rho_0]$ 
8        $F := F \cup \text{Dual } pq$ 
9   endif
10 endfor
11 output  $CF$ 

```

Remark: As we mentioned earlier, the pole vectors approximate the normals to S at the sample points. Therefore, it is plausible to use them in the angle condition instead of the normals to the umbrella triangles. Similarly, the width of the tangent polygon can replace the circumradii of the umbrella triangles in the ratio condition. However, our experiments suggest that these alternatives do not produce as good result as the stated angle and ratio conditions do. We believe that these alternatives do not eliminate Delaunay edges lying close to the surface very well in practice.

4 Guarantees

In this section we prove that a significant subset of the output of MEDIAL converges to the actual medial axis when ε approaches zero. For a sample point p we define μ_p , δ_p , η_p and β_p as follows. Motivation of their definitions becomes clear when we use them later. Let

μ_p = radius of the larger medial ball at p

$$\delta_p = \max_{pqr \in U_p} \frac{R_{pqr}}{\mu_p}$$

$$\eta_p = \delta_p \left(\frac{\mu_p}{f(p)} \right)$$

$$\beta_p = \arcsin \frac{\eta_p}{1 - 2\eta_p} + \arcsin \left(\frac{2}{\sqrt{3}} \sin(2 \arcsin \frac{\eta_p}{1 - 2\eta_p}) \right) + \frac{2\eta_p}{1 - 6\eta_p}.$$

Our analysis uses the two medial balls touching the surface at p . Their radii are not necessarily equal to the local feature size $f(p)$ which warrants the introduction of μ_p . Both δ_p and η_p measure the local density of the sample around p which may be much smaller than the global density measured by ε . We will see later that $\delta_p = O(\varepsilon) \frac{f(p)}{\mu_p}$. This in turn implies that η_p is $O(\varepsilon)$. The term β_p measures the angle $\angle \mathbf{n}_p, \mathbf{n}_{pqr}$ between the normal at p and the normal to any of its incident umbrella triangle. The $O(\varepsilon)$ bound on η_p also puts an $O(\varepsilon)$ bound on β_p .

4.1 Umbrella triangles

The Ratio and Angle conditions in MEDIAL are based on the assumption that the umbrella triangles lie flat to the surface. We prove this fact now. Similar result has been proved by Amenta, Choi, Dey and Leekha [3]. But, we will need a different form of the result here. For completeness and due to the differences, we include the proofs where necessary.

Recall that R_{pqr} denotes the circumradius of a triangle pqr and \mathbf{n}_{pqr} denotes its normal. Let p be the vertex subtending the largest angle in pqr .

Lemma 1 *If $\frac{R_{pqr}}{f(p)} < 0.1$, the angle $\angle \mathbf{n}_{pqr}, \mathbf{n}_p$ is no more than $\alpha + \sin^{-1}(\frac{2}{\sqrt{3}} \sin 2\alpha)$, where $\alpha \leq \arcsin \frac{R_{pqr}}{f(p)}$.*

PROOF. Consider the medial balls B_1 and B_2 sandwiching S at p with the centers on the medial axis. Let D be the ball with the circumcircle of pqr as a diametric circle; refer to Figure 8. Denote the circles of intersection of D with B_1 and B_2 as C_1 and C_2 respectively.

The line of normal to S at p passes through m , the center of B_1 . This normal makes an angle less than α with the normals to the plane of C_1 , where

$$\begin{aligned} \alpha &\leq \arcsin \frac{R_{pqr}}{\|m - p\|} \\ &\leq \arcsin \frac{R_{pqr}}{f(p)} \end{aligned}$$

This angle bound also applies to the plane of C_2 , which implies that the planes of C_1 and C_2 make a wedge, say W , with an acute dihedral angle no more than 2α .




Figure 8: Normal to a triangle and the normal to S at the vertex with the largest face angle.

The two vertices q, r of pqr cannot lie inside B_1 or B_2 . This implies that pqr lies completely in the wedge W . Since the face angle at p is at least $\frac{\pi}{3}$, the triangle pqr has to lie somewhat flat within the wedge W if α is small enough. With the condition that $\frac{R_{pqr}}{f(p)} < 0.1$, α is small and it is proved in [3] that $\alpha' = \arcsin(\frac{\sin 2\alpha}{\sin \pi/3})$ where α' is the acute angle between the planes of C_1 and pqr . The angle $\angle \mathbf{n}_{pqr}, \mathbf{n}_p$ is at most $\alpha + \alpha'$ proving the theorem. \square

The next lemma is proved by Amenta and Bern [1] which says that surface normals do not differ much if the points are close with respect to local feature sizes.

Lemma 2 *Let p and q be two points in S so that $\|p - q\| \leq \tau f(p)$. Then $\angle \mathbf{n}_p, \mathbf{n}_q \leq \frac{\tau}{1-3\tau}$ if $\tau < \frac{1}{3}$.*

We use the previous two lemmas to derive a bound between the normals to the surface and to the umbrella triangles respectively.

Lemma 3 *Let pqr be any triangle in U_p . We have $\angle \mathbf{n}_{pqr}, \mathbf{n}_p \leq \beta_p$.*

PROOF.

Notice that if p is a vertex subtending the maximum angle in pqr , then we can apply Lemma 1 with $R_{pqr} \leq \delta_p \mu_p$ to get

$$\angle \mathbf{n}_p, \mathbf{n}_{pqr} \leq \arcsin \eta_p + \arcsin \left(\frac{2}{\sqrt{3}} \sin(2 \arcsin \eta_p) \right) \leq \beta_p.$$

If p is not the vertex subtending the maximum angle, without loss of generality assume q be such one. Since pq cannot be larger than the diameter of the circumcircle of pqr , $\|p - q\| \leq 2\delta_p \mu_p = 2\eta_p f(p)$. We have $f(p) \leq \frac{1}{1-2\eta_p} f(q)$ by the Lipschitz property of $f()$. It follows $\|p - q\| \leq \frac{2\eta_p}{1-2\eta_p} f(p)$.

Apply Lemma 2 with $\tau = \frac{2\eta_p}{1-2\eta_p}$ to get

$$\angle \mathbf{n}_p, \mathbf{n}_q \leq \frac{2\eta_p}{1-6\eta_p}.$$

Applying Lemma 1 for q and then using the fact $f(q) \geq (1-2\eta_p)f(p)$ we get that $\angle \mathbf{n}_{pqr}, \mathbf{n}_q$ is no more than

$$\arcsin \frac{\eta_p}{1-2\eta_p} + \arcsin \left(\frac{2}{\sqrt{3}} \sin(2 \arcsin \frac{\eta_p}{1-2\eta_p}) \right).$$

Since $\angle \mathbf{n}_{pqr}, \mathbf{n}_p \leq \angle \mathbf{n}_{pqr}, \mathbf{n}_q + \angle \mathbf{n}_q, \mathbf{n}_p$, the bound as claimed follows. \square

The proof of the next lemma follows from a result in [1].

Lemma 4 *Let pqr be any umbrella triangle in U_p . The circumradius $R_{pqr} \leq \left(\frac{\varepsilon}{1-\varepsilon}\right) \left(\frac{1}{\sin(\pi/2-3 \arcsin \varepsilon/(1-\varepsilon))}\right) f(p)$.*

PROOF. The radius R_{pqr} cannot be more than the distance between p and the farthest point on the tangent polygon of p from it. Any point w in the tangent polygon satisfies $\angle \mathbf{v}_p, \mathbf{t}_{pw} = \frac{\pi}{2}$ where \mathbf{v}_p is the pole vector. Therefore, we can use the reverse implication of the following statement proved in [1]. If $w \in V_p$ is a point so that $\|w - p\| > \tau f(p)$, then

$$\angle \mathbf{v}_p, \mathbf{t}_{pw} \leq \arcsin \frac{\varepsilon}{\tau(1-\varepsilon)} + \arcsin \frac{\varepsilon}{1-\varepsilon}.$$

Also, the pole vector \mathbf{v}_p satisfies $\angle \mathbf{v}_p, \mathbf{n}_p \leq 2 \arcsin \frac{\varepsilon}{1-\varepsilon}$. The reverse implication as referred above gives that $R_{pqr} \leq \|w - p\| \leq \left(\frac{\varepsilon}{1-\varepsilon}\right) \left(\frac{1}{\sin(\pi/2-3 \arcsin \varepsilon/(1-\varepsilon))}\right) f(p)$. \square

It follows from the above lemma that

$$\delta_p = \max_{pqr \in U_p} \frac{R_{pqr}}{\mu_p} \leq O(\varepsilon) \frac{f(p)}{\mu_p}.$$

This implies that both η_p and β_p are $O(\varepsilon)$ only.

Corollary 1 $\delta_p = O(\varepsilon) \frac{f(p)}{\mu_p}$, $\eta_p = O(\varepsilon)$, and $\beta_p \leq 6\eta_p = O(\varepsilon)$ for small ε .

4.2 Analysis

The convergence analysis proceeds in part by showing that each point in a specific subset of the output is within a small distance from a medial axis point. As $\varepsilon \rightarrow 0$ this distance vanishes. Conversely, we also argue that each medial axis point has a nearby point in the subset, the distance between which also vanishes as $\varepsilon \rightarrow 0$.

Some of the points on the sampled surface S have infinitely large medial balls. This poses some difficulty in our analysis.

To prevent this we enclose S and hence its sample within a sufficiently large bounding sphere. This ensures that μ_p has an upper bound for each point p on S . Of course, the bounding sphere changes the medial axis outside S , but we can keep these changes as far away from S as we wish by choosing a sufficiently large bounding sphere. In particular, the medial axis inside S does not change at all with this modification. In the analysis to follow, we assume that the input point set P samples S as well as the bounding sphere. With the bounding sphere assumption, we have $\Delta \leq \frac{f(p)}{\mu_p} \leq 1$ for any point $p \in S$ where $\Delta > 0$ is a constant dependent on S . It is important that, although Δ depends on S , it remains independent of its sampling.

In the analysis we prove convergence for the subset of the medial axis of S that remains unchanged with the bounding sphere assumption. Let M denote this subset of the medial axis which consists of the centers of the medial balls that touch S but not the bounding sphere. We also need to make some adjustments in MEDIAL to accommodate this change. Given the sample points of S , MEDIAL has to first add the sample points of the bounding sphere and then filter only from those Delaunay edges that connect sample points of S . We assume this modification to MEDIAL while carrying out the analysis. However, we do not implement these modifications since we observe that the algorithm without this additional sample points work well in practice.

For the ε -sample P , we will define \mathcal{L}_ε , a subcomplex of the output complex computed by MEDIAL and show that the underlying space L_ε of \mathcal{L}_ε converges to M in the limit $\varepsilon \rightarrow 0$. Let m and m' be the centers of the two medial balls at p and μ, μ' their radii respectively. It is a simple observation that m and m' are contained in V_p [1]. Suppose $w \in V_p$ be a point in the output so that $\mathbf{t}_{pw} \cdot \mathbf{t}_{pm} > 0$. This means w and m lies on the same side of the tangent plane at p . We will show that, if $\|w - p\| \leq \mu$, the distance between w and its closest point on pm is small. This fact is used to show that, if w belongs to a Voronoi facet of V_p which makes large angle with the line of pm , then w must be near to m .

Before we proceed to prove the above fact, we need another geometric property of the umbrella triangles and their circumcircles. For an umbrella triangle pqr , consider the cone on its plane with p as apex and opening angle $\angle qpr$. Let L_{pqr} denote the intersection of this cone with the circumcircle of pqr , see Figure 9. We define the *flower* of p , Fl_p as

$$\text{Fl}_p = \cup_{pqr \in U_p} L_{pqr}.$$

Lemma 5 V_p does not contain any point of the boundary of Fl_p inside.

PROOF. Consider any point x on the boundary of Fl_p as shown in Figure 9. Let x belong to the boundary of the circumcircle of pqr . It is a simple geometric fact of circles that x is closer to either q or r than to p . It means x cannot be inside V_p . \square




Figure 9: L_{pqr} of an umbrella triangle pqr (left), and the flower of p .

The flower of p lies very flat to the surface and cannot intersect the segment mm' at any point other than p . This means that the above lemma implies that Fl_p intersects V_p completely and separates m and m' on its two sides within V_p . In particular, any segment connecting a point $w \in V_p$ with m , where w and m lies on the opposite sides of Fl_p , must intersect Fl_p .

In the next two lemmas let w° denote the closest point to $w \in V_p$ on the line of mm' . Also see Figure 10.




Figure 10: Illustration for Lemma 6.

Lemma 6 If w° lies in the segment mm' , $\|w - w^\circ\| \leq 2 \tan(\arcsin 2\eta_p)\mu_p$.

PROOF. Let m' be the center of the medial ball at p so that w and m' lie on opposite sides of the flower at p within V_p . Consider the segment wm' . Let y be the foot of the perpendicular dropped from p to wm' . Since wm' intersects Fl_p and any point in Fl_p is within $2\delta_p\mu_p$ distance, we must have $\|y - p\| \leq 2\delta_p\mu_p$. Therefore,

$$\angle pm'y \leq \arcsin \frac{\|y - p\|}{\|m' - p\|} \leq \arcsin \frac{2\delta_p\mu_p}{f(p)} = \arcsin 2\eta_p.$$

It follows that

$$\|w - w^\circ\| \leq \|m - m'\| \tan \angle pm'y \leq 2\mu_p \tan(\arcsin 2\eta_p). \quad \square$$

Lemma 7 Let $F = \text{Dual}pq$ be a Voronoi facet where pq satisfies the angle condition $[\theta]$ with $\theta \geq 2\eta_p + \beta_p$. Any point w in F with $\|w - p\| \leq \mu$ is within a distance of $\frac{1}{\sin(\theta - \beta_p)}(2 \tan(\arcsin 2\eta_p))\mu_p$ from m where m and μ are the center and radius of the medial ball at p with $\mathbf{t}_{pm} \cdot \mathbf{t}_{pw} > 0$.

PROOF. Consider the ball B with radius μ around p . The point w necessarily lies inside B since $\|w - p\| \leq \mu$. Therefore, w° lies in the segment pm . We can apply Lemma 6 to assert $\|w - w^\circ\| \leq 2 \tan(\arcsin 2\eta_p)\mu_p$.




Figure 11: Illustration for Lemma 7. The picture in the left is not a possible configuration due to the constraint on α .

Let the plane of F intersect the line of pm at x at an angle α . We have $\alpha > \theta - \beta_p$ since the normal to F makes more than θ angle with the normal of an umbrella triangle (by the Angle condition $[\theta]$) which in turn makes an angle less than β_p with the surface normal at p (Lemma 3). With the requirement that $\theta > 2\eta_p + \beta_p$, we have $\alpha > 2\eta_p$.

Also, the plane of F cannot intersect the segment mm' . This is because the medial balls at p are empty of any other sample point and thus both m and m' belong to V_p . In particular, the segment mm' must be inside the Voronoi cell V_p . The segment wm' necessarily lies inside V_p and intersects the flower of p , say at y . This means $\|y - p\| \leq 2\delta_p\mu_p$. If x lies below m' as shown in the left picture of Figure 11, we have

$$\alpha < \angle pxw < \angle pm'y < \tan \frac{\|y - p\|}{\|m' - p\|} < 2\eta_p.$$

This contradicts the assertion that $\alpha > 2\eta_p$. So, x cannot lie below m' . Instead, it lies above m as shown in the right picture in Figure 11. From the triangle $w w^\circ x$ we have

$$\|w - m\| \leq \|w - x\| \leq \frac{\|w - w^\circ\|}{\sin \alpha} \leq \frac{2 \tan(\arcsin 2\eta_p)}{\sin(\theta - \beta_p)}\mu_p. \quad \square$$

To complete the proof of convergence we need the following lemma which says that for a long Delaunay edge pq there must exist a point $w \in \text{Dual}pq$ which cannot be too far from a medial axis point. This lemma is extracted from a result (Proposition 18) of Boissonnat and Cazals [8]. Although we use slightly different version with different constants and exponents, the proof remains same.

Lemma 8 Let pq be a Delaunay edge with $\|p - q\| > \eta\mu$, where μ is the radius of a medial ball at p with the center m and $\eta \geq \varepsilon^{1/4}$. If the medial angle of p is larger than $\varepsilon^{1/3}$, then all points $w \in \text{Dual}pq$ with $\mathbf{t}_{pm} \cdot \mathbf{t}_{pw} > 0$ satisfy $\|w - m\| = O(\varepsilon^{3/4})\mu$ for sufficiently small $\varepsilon \geq 0$.

Now we define a subcomplex \mathcal{L}_ε from the complex computed by MEDIAL. The definition is motivated by the conditions of Lemma 7 and Lemma 8.

Definition 4 We define $\mathcal{L}_\varepsilon \subseteq V_P$ as $\mathcal{L}_\varepsilon = \{CIF \mid F \in V_P\}$ where $pq = \text{Dual } F$ is selected by MEDIAL either by (i) Angle condition, or (ii) Ratio condition with $\|p - q\| \geq \varepsilon^{1/4}\mu_p$ and the medial angle of p is larger than $\varepsilon^{1/3}$. Here μ is the radius of the medial ball at p with the center m so that $\mathbf{t}_{pm} \cdot \mathbf{t}_{pq} > 0$. Let L_ε denote the underlying space of \mathcal{L}_ε .

Consider a sequence of L_ε with decreasing ε . Theorem 1 and 2 establish that the limit of this sequence converges to M where M is the subset of the medial axis defined by medial balls that touch only S but not the bounding sphere. This shows that a subset of the output of MEDIAL converges to the medial axis though the difference between this subset and the output is small; specifically when $\varepsilon = 0$, this subset coincides with the output.

Theorem 1 $\lim_{\varepsilon \rightarrow 0} L_\varepsilon \subseteq M$.

PROOF. Let $F \in \mathcal{L}_\varepsilon$ be a facet computed by MEDIAL(). First, consider the case when $pq = \text{Dual } F$ is selected by the Angle condition. Let $w \in F$ be any point and μ and m are as defined in Lemma 7. If w is more than μ away from p , we can apply Lemma 8 to conclude $\|w - m\| = O(\varepsilon^{3/4})\mu$ if p has a medial angle $\Omega(\varepsilon^{1/3})$. Otherwise, Lemma 7 applies to assert that $\|w - m\| \leq \frac{2 \tan(\arcsin 2\eta_p)\mu_p}{\sin(\pi/8 - \beta_p)} = O(\varepsilon)\mu_p$. In both cases as ε approaches zero, w reaches m in the limit $\varepsilon \rightarrow 0$.

Next, consider the case when $pq = \text{Dual } F$ is selected by MEDIAL by the Ratio condition. Since $F \in \mathcal{L}_\varepsilon$, all conditions to apply Lemma 8 are satisfied. So, we have $\|w - m\| = O(\varepsilon^{3/4})\mu$. This implies that w reaches m in the limit $\varepsilon \rightarrow 0$. \square

Next we wish to establish the converse of the previous theorem, i.e., the points on the medial axis have a nearby point in the output, L_ε in particular. It turns out that only a subset of M satisfies this guarantee which converges to entire M as ε approaches zero.

Definition 5 The sampled surface S is said to be well behaved in the neighborhood of a point $p \in S$ with respect to $\varepsilon \geq 0$ if the following condition holds: any ball that meets S tangentially at p can be tangent to another point $q \in S$ only if $\|p - q\| \geq 2\varepsilon^{1/4}\mu$, where μ is the radius of the medial ball at p with the center m so that $\mathbf{t}_{pm} \cdot \mathbf{t}_{pq} > 0$.

Roughly speaking, the above definition says that S cannot oscillate around a well behaved point arbitrarily. Let $S_\varepsilon \subseteq S$ be the set of all points where S is well behaved with respect to ε . By definition all points in S_ε have a medial angle more than $\arcsin \varepsilon^{1/4} \geq \varepsilon^{1/3}$ for sufficiently small ε . Next lemma claims that all sample points where S is well behaved must have a long Delaunay edge.

Lemma 9 Let $p \in S_\varepsilon$ be a sample point. For each medial ball at p with the center m and radius μ , there is a Delaunay edge pq with $\mathbf{t}_{pm} \cdot \mathbf{t}_{pq} > 0$ so that $\|p - q\| \geq \varepsilon^{1/4}\mu$ where ε is sufficiently small.

PROOF. Let B be a medial ball touching S at p with radius μ and center m . Grow B keeping it tangent to S at p . Initially, the ball is empty of any other point of S . But, as the growth continues it starts meeting S and eventually meets a sample point q . Certainly, we have $\mathbf{t}_{pm} \cdot \mathbf{t}_{pq} > 0$. Let q belong to a component of $B \cap S$ that originated at $x \in S$ during the growth process. This means B met S tangentially at x . Since S is well behaved at p , we must have $\|p - x\| \geq 2\varepsilon^{1/4}\mu$. Also, applying the ε -sampling condition and some sphere geometry we can show that, for sufficiently small ε , $\|q - x\| \leq 2\varepsilon f(x)$. Now assume $\|p - q\| \leq 2\varepsilon^{1/4}\mu$ because otherwise there is nothing to prove. Using Lipschitz condition on $f(\cdot)$ and the condition $\|q - x\| \leq 2\varepsilon f(x)$, it is easy to establish that $f(x) \leq \frac{f(q)}{1 - 2\varepsilon}$. Similarly, using the condition $\|p - q\| \leq 2\varepsilon^{1/4}\mu \leq \frac{2\varepsilon^{1/4}}{\Delta} f(p)$, we get $f(q) \leq \frac{f(p)}{1 - 2\varepsilon^{1/4}/\Delta} \leq \frac{\mu}{1 - 2\varepsilon^{1/4}/\Delta}$. Therefore,

$$\begin{aligned} \|p - q\| &\geq \|p - x\| - \|q - x\| \\ &\geq 2\varepsilon^{1/4}\mu - \frac{2\varepsilon\mu}{1 - 2\varepsilon^{1/4}/\Delta} \\ &\geq \varepsilon^{1/4}\mu \end{aligned}$$

for sufficiently small ε . \square

We also need the following result proved by Amenta, Choi, Kolluri [4] which says that poles are not very far from the medial axis points inside a Voronoi cell.

Lemma 10 Let m be a medial axis point with nearest sample point p and the medial angle at m is larger than $\varepsilon^{1/3}$. Then there is a pole of p which is only $O(\varepsilon^{2/3})\mu$ distance away from m where μ is the radius of the respective medial ball.

Let $M_\varepsilon \subseteq M$ be the set of medial axis points whose nearest point(s) in S and P lie in S_ε . Since all points in S_ε have medial angles more than $\varepsilon^{1/3}$, so are the points in M_ε . We show that M_ε has a nearby point in L_ε .

Lemma 11 Let $m \in M_\varepsilon$. There exists a point $w \in L_\varepsilon$ so that $\|w - m\| = O(\varepsilon^{1/4})\mu_p$ where p is the nearest sample point to m .

PROOF. Let c be the center of the medial ball touching S at p . Since this ball does not contain any point of S inside, we must have c inside the Voronoi cell V_p . Also, $m \in V_p$ since p is the nearest sample point to m . First, both m and c have medial angles more than $\varepsilon^{1/3}$ by the property of M_ε . Applying Lemma 10 we get that m and c cannot be far apart, $\|m - c\| = O(\varepsilon^{2/3})\mu_p$ since both have a pole of p within $O(\varepsilon^{2/3})\mu_p$ distance.

By definition, $p \in S_\varepsilon$. Thus, according to Lemma 9 there exists a Delaunay edge pq with $\|p - q\| \geq \varepsilon^{1/4}\mu$ where μ is

the radius of the medial ball with the center m . The dual of this edge must belong to \mathcal{L}_ε by the Ratio condition. Lemma 8 implies that there is a point $w \in F = \text{Dual}pq$ so that $\|w - c\| = O(\varepsilon^{3/4})\mu_p$. This means

$$\|w - m\| \leq \|w - c\| + \|m - c\| = O(\varepsilon^{3/4})\mu_p.$$

□




Figure 12: Convergence of L_ε and M .

Theorem 2 $M \subseteq \lim_{\varepsilon \rightarrow 0} L_\varepsilon$.

PROOF. By definition of M_ε we have $M = \lim_{\varepsilon \rightarrow 0} M_\varepsilon$. Lemma 11 implies $\lim_{\varepsilon \rightarrow 0} M_\varepsilon \subseteq \lim_{\varepsilon \rightarrow 0} L_\varepsilon$. □

Theorem 1 and 2 together establish that, in the limit $\varepsilon \rightarrow 0$, L_ε converges to the medial axis M . A schematic diagram of the convergence of various structures is illustrated in Figure 12.

Remark: Note that it is essential to assume that no five sample points lie on the boundary of an empty sphere so that dual facets of all edges with an empty circumscribing sphere exist in the Voronoi diagram. This excludes, for example, sample points on a sphere.

Also, a careful examination of the proofs of Theorem 1 and Theorem 2 suggests that the only Ratio condition is sufficient to filter Delaunay edges. However, as we have indicated earlier, we need both conditions to deal with data sets appearing in practice. This is the reason why we designed MEDIAL with two conditions and prove their contributions towards convergence.

5 Experimental results

We implemented MEDIAL in C++ using the CGAL library [32] for the Voronoi diagram and the Delaunay triangulation. Figure 13 shows our result on some example data sets. In this picture we show only the ‘inner’ medial axis that is ‘enclosed’ by the surface. We computed the ‘inner’ medial axis as follows. A piecewise linear surface interpolating the sample points is computed with our TIGHT COCONE software

[31]. This surface is a subcomplex of the Delaunay triangulation and is guaranteed to be watertight. We output those Voronoi facets computed by MEDIAL that are dual to the Delaunay edges enclosed by the computed surface of TIGHT COCONE.

The approximation of the medial axis is better where the data is dense. Near high curvature regions, or non-smooth regions where undersampling happens, the approximation contains artifacts. For example, in the CLUB and ROCKER data, the medial axis is well approximated in most part where the density is high. But, near the handle of the CLUB and top of the ROCKER, the sampling density suffers due to high curvature and small features. As a result the approximation is poor in these regions.

object	# points	Delaunay time(sec.)	Filter time(sec.)
3HOLE	4000	2.37	0.82
KNOT	10001	8.36	2.66
ENGINE	11361	37.0	1.82
MANNEQUIN	12773	7.79	2.28
FANDISK	16475	13.4	3.0
CLUB	16865	18.81	3.06
FOOT	20021	13.24	3.71
HAND	25626	51.14	5.11
SCREWDRIVER	27152	51.87	5.36
DINOSAUR	28098	23.28	5.8
HEART	37912	32.49	6.58
ROCKER	40177	74.54	8.25
FEMALE	121723	292.85	24.95

Table 1: Time data.

SCREWDRIVER data is mostly very dense and almost uniform except near some small curvature regions. On the other hand, 3HOLE is relatively a sparse data set. In both cases, MEDIAL approximated the medial axis quite well. This shows that MEDIAL is tolerant to different levels of data densities.

The sample in ENGINE lie on several connected components. In this case we computed the entire medial axis since ‘inner’ and ‘outer’ distinction in this case does not have any meaning.

In general, the medial axis of a surface may contain elements of different dimensions. In MEDIAL we approximate all these elements with two dimensional Voronoi facets. The data set KNOT shows this interesting phenomenon. The one dimensional medial axis in this case has been approximated with very thin Voronoi facets.

The HEART data is extracted as an iso-surface from the intensity field of a volumetric image. The data contains some noise introduced by the iso-surface extraction procedure. MEDIAL could tolerate most of this noise as the output in Figure 13 suggests.

Four more examples HAND, MANNEQUIN, DINOSAUR and FEMALE show how MEDIAL computes clean medial



CLUB




3HOLE




ROCKER



SCREWDRIVER




ENGINE surface



medial axis of ENGINE




KNOT



HEART



HAND



MANNEQUIN



DINOSAUR



FEMALE

Figure 13: Medial axis with MEDIAL shown dark shaded.

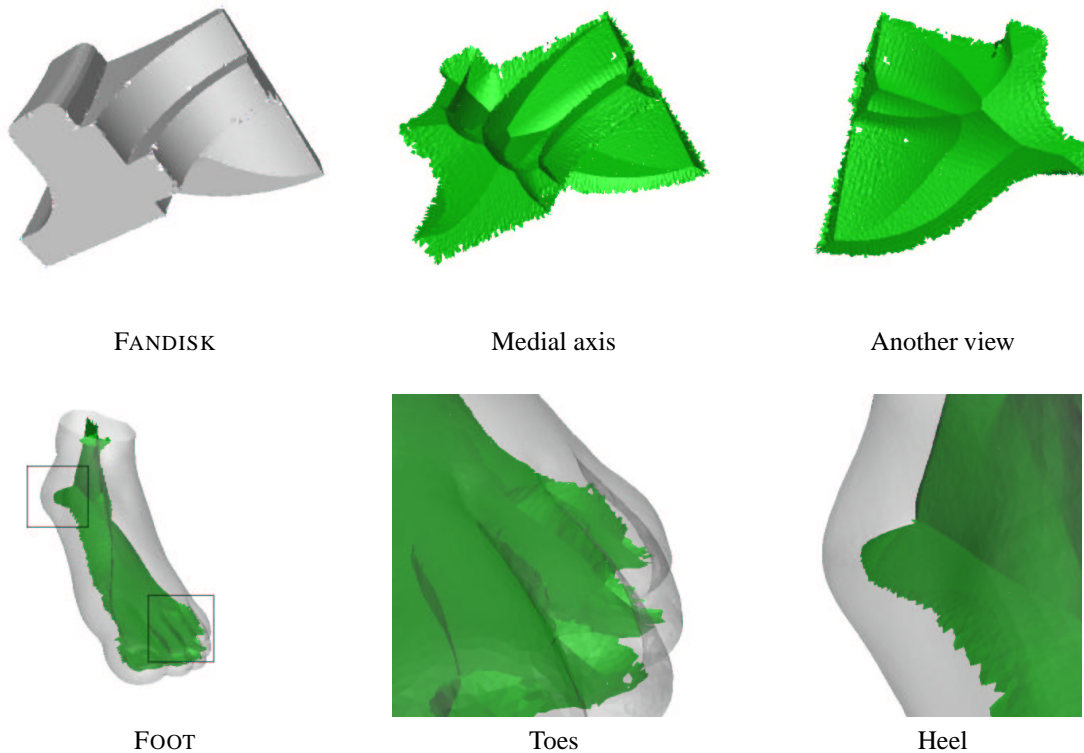


Figure 14: Computed medial axis of FANDISK and FOOT.

axes for complicated shapes.

In Figure 14, we show the two views of the medial axis for a CAD object FANDISK. The medial axis is well approximated except at the boundaries where it is ‘jagged’. The poor approximation at these places is caused by the sharp edges of the surface where it is nonsmooth and therefore the inherent problem of undersampling occurs, see [13]. In the FOOT data we zoom some places of the toes and the heel. The zoomed region in the toe has undersampling and the medial axis near the boundary has some roughness. However, the heel is well sampled, and the corresponding medial axis boundary is smoother. The FOOT data has abrupt density changes in some parts. As a result ‘only Ratio condition’ did not produce good result for any ratio. MEDIAL produced a good approximation to the medial axis. It shows that MEDIAL is impervious to nonuniformity in data.

5.1 Timing

We used the filtered floating point arithmetic provided by CGAL [32] for robust geometric computations. Experiments were conducted on a PC with 933 MHz CPU and 512MB memory. The code was compiled with CGAL2.3 library and g++ compiler with `o1` level of optimization. The time for the Delaunay triangulation and filtration are listed in Table 1.

6 Conclusions

In this paper we present an algorithm that approximates the medial axis from the Voronoi diagram of a set of sample points. Unlike previous approaches, this algorithm is scale and density independent. Experimental studies suggest that the algorithm computes clean medial axes without any fine tuning of the parameters.

Although this is the first algorithm that shows that medial axis can be computed as a subcomplex of the Voronoi diagram with guaranteed convergence, we could not prove that the output medial axis maintains the topology of the original one. From our experiment results, we observe that all significant branchings of the medial axis are approximated by MEDIAL. If the data set is dense, we obtain quite clean medial axis without much of sampling artifacts. Although proving a homeomorphism between the output of MEDIAL and the true medial axis will be difficult, if not impossible, we expect that it would be possible to show the homotopy equivalence between the two. We plan to investigate this aspect in future research.

In many applications it is useful to have a simplification of the medial axis. Obviously, larger parameter values simplify the medial axis. Is there a way to determine the values automatically that bring up the hierarchy of the stable parts of the medial axis? More investigations are necessary to answer

this question.

Acknowledgements. We thank Dominique Attali for sending the papers [5, 6] and the HEART data.

References

- [1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.* **22** (1999), 481–504.
- [2] N. Amenta, M. Bern and D. Eppstein. The crust and the β -skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, **60** (1998), 125–135.
- [3] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Applications*, **12** (2002), 125–121.
- [4] N. Amenta, S. Choi and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory and Applications* **19** (2001), 127–153.
- [5] D. Attali and J.-O. Lachaud. Delaunay conforming iso-surface, skeleton extraction and noise removal. *Comput. Geom. : Theory Appl.*, 2001, to appear.
- [6] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding* **67** (1997), 261–273.
- [7] J. D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbor interpolation of distance functions. *Proc. 16th. ACM Sympos. Comput. Geom.*, (2000), 223–232.
- [8] J. D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Comput. Geom. Theory Appl.* **19** (2001), 87–120.
- [9] S. Bouix and K. Siddiqi. Divergence-based medial surfaces. *Proc. European Conference on Computer Vision*, 2000.
- [10] J. W. Brandt. Convergence and continuity criteria for discrete approximation of the continuous planar skeletons. *CVGIP: Image Understanding* **59** (1994), 116–124.
- [11] J. W. Brandt and V. R. Algazi. Continuous skeleton computation by Voronoi diagram. *Comput. Vision, Graphics, Image Process.* **55** (1992), 329–338.
- [12] T. Culver, J. Keyser and D. Manocha. Accurate computation of the medial axis of a polyhedron. *Proc. ACM Sympos. Solid Modeling Applications.*, (1999), 179–190.
- [13] T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *Proc. 17th Ann. Sympos. Comput. Geom.* (2001), 257–263.
- [14] H. Edelsbrunner. Shape reconstruction with Delaunay complex. *LNCS 1380, LATIN'98: Theoretical Informatics*, (1998), 119–132.
- [15] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.
- [16] M. Etzion and A. Rappoport. Computing Voronoi skeletons of a 3D polyhedron by space subdivision. *Tech. Report*, Hebrew University, 1999.
- [17] P. J. Giblin and B. B. Kimia. A formal classification of 3D medial axis points and their local geometry. *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2000.
- [18] L. Guibas, R. Holleman and L. E. Kavradi. A probabilistic roadmap planner for flexible objects with a workspace medial axis based sampling approach. *Proc. IEEE/RSJ Intl. Conf. Intelligent Robots and Systems*, 1999.
- [19] H. N. Gursoy and N. M. Patrikalakis. Automated interrogation and adaptive subdivision of shape using medial axis transform. *Advances in Engineering Software* **13** (1991), 287–302.
- [20] C. Hoffman. How to construct the skeleton of CSG objects. *The Mathematics of Surfaces, IVA*, Bowyer and J. Davenport Eds., Oxford Univ. Press, 1990.
- [21] P. Hubbard. Approximating polyhedra with spheres for time critical collision detection. *ACM Trans. Graphics* **15** (1996), 179–210.
- [22] C. Niblak, P. Gibbons and D. Capson. Generating skeletons and centerlines from the distance transform. *CVGIP: Graphical Models and Image Processing* **54** (1992), 420–437.
- [23] R. L. Ogniewicz. Skeleton-space: A multiscale shape description combining region and boundary information. *Proc. Computer Vision and Pattern Recognition*, (1994), 746–751.
- [24] M. Schmitt. Some examples of algorithms analysis in computational geometry by means of mathematical morphological techniques. In J.-D. Boissonnat and J.-P. Laumond eds. *Proc. Workshop on Geometry and Robotics*, vol. 391, *Lecture Notes Comput. Sci.*, Springer-Verlag, 1989.
- [25] D. Sheehy, C. Armstrong and D. Robinson. Shape description by medial axis construction. *IEEE Trans. Visualization and Computer Graphics* **2** (1996), 62–72.

- [26] A. Sheffer, M. Etzion, A. Rappoport and M. Bercovier. Hexahedral mesh generation using the embedded Voronoi graph. *Engineering Comput.* **15** (1999), 248–262.
- [27] E. C. Sherbrooke, N. M. Patrikalakis and E. Brisson. An algorithm for the medial axis transform of 3D polyhedral solids. *IEEE Trans. Vis. Comput. Graphics* **2** (1996), 44–61.
- [28] D. Storti, G. Turkiyyah, M. Ganter, C. Lim and D. Stal. Skeleton-based modeling operations on solids. *ACM Sympos. Solid Modeling Applications.*, (1997), 141–154.
- [29] M. Teichman and S. Teller. Assisted articulation of closed polygonal models. *Proc. 9th Eurographics Workshop on Animation and Simulation*, 1998.
- [30] F.-E. Wolter. Cut locus & medial axis in global shape interrogation & representation. *MIT Design Laboratory Memorandum 92-2*, 1992.
- [31] <http://www.cis.ohio-state.edu/~tamaldey/cocone.html>.
- [32] <http://www.cs.ruu.nl/CGAL>.