

Delaunay Triangulations Approximate Anchor Hulls *

Tamal K. Dey
Dept. of CSE
The Ohio State University
Columbus, Ohio 43210, USA

Joachim Giesen
Theoretische Informatik
ETH Zürich, Switzerland

Samrat Goswami
Dept. of CSE
The Ohio State University
Columbus, Ohio 43210, USA

Abstract

Recent results establish that a subset of the Voronoi diagram of a point set that is sampled from the smooth boundary of a shape approximates the medial axis. The corresponding question for the dual Delaunay triangulation is not addressed in the literature. We show that, for two dimensional shapes, the Delaunay triangulation approximates a specific structure which we call *anchor hulls*. As an application we demonstrate that our approximation result is useful for the problem of shape matching.

1 Introduction

Shape modeling from point samples has recently gained considerable attention because of its great flexibility [12]. Researchers have started studying methods to extract various geometric structures of shapes from point samples. They include reconstructing boundaries of shapes [11], approximating their medial axes [4, 8, 14], identifying features [13] and many others.

A remarkable connection between the shape and its point sample was revealed by Brandt and Algazzi [9] when they showed that the Voronoi vertices of a dense point sample approximate the medial axis of a shape in

*This research is supported by NSF CARGO grants DMS-0138456 and DMS-0310642 and an ARO grant DAAD19-02-0347 and by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

two dimensions. Although this result does not hold in three dimensions, Amenta and Bern showed that the elongated Voronoi cells approximate the normals on the shape boundaries [3]. Later, Amenta, Choi and Kolluri [4] and Boissonnat and Cazals [8] established that a subset of Voronoi vertices does indeed approximate the medial axis in three dimensions. Dey and Zhao [14] showed that, not only a set of discrete Voronoi vertices approximates the medial axis, a subset of the Voronoi facets approximates it in three dimensions. Given these approximation results by Voronoi diagrams, a natural question arises ‘what does the Delaunay triangulation of a point sample approximate?’ Some properties of the Delaunay triangulation of a point sample from a curve or a surface have been studied recently. For example, the size complexity of the Delaunay triangulation of a point sample of a surface was recently investigated [7, 15]. However, no result exists to answer what structure of a shape is approximated by the Delaunay triangulations of a point sample.

In this paper we address this question for two dimensional shapes. We establish that a subset of the Delaunay triangulation approximates the *anchor hulls* of the shape. An anchor hull in a shape is the convex hull of the points where the maximal empty balls centering a medial axis point touch the boundary of the shape (see Figure 2). The anchor hulls cover the shape. Figure 1(a) shows the anchor hulls (shaded white) in a shape and Figure 1(b) shows the Delaunay triangulation of a point sample on its boundary. It exhibits how the anchor hulls are approximated by the Delaunay triangles.

After establishing the anchor hull approximation result, we apply it to the well known problem of shape matching [1, 5, 17, 18]. Since similar shapes have similar anchor hulls, one can segment almost identical shapes similarly using approximate anchor hulls. We use this segmentation to match two shapes with a score that reflects the similarity between their anchor hulls. Experimental results confirm that anchor hulls provide an effective tool for shape matching.

2 Definitions

2.1 Anchor hulls and medial axis.

In this paper the shape Σ is a compact, connected subset of \mathbb{R}^2 with smooth boundary $\partial\Sigma$.

Let $A : \Sigma \rightarrow \mathcal{P}(\partial\Sigma)$ be the function from Σ to the power set $\mathcal{P}(\partial\Sigma)$ of $\partial\Sigma$ that assigns to every point of Σ its nearest neighbors on the boundary,

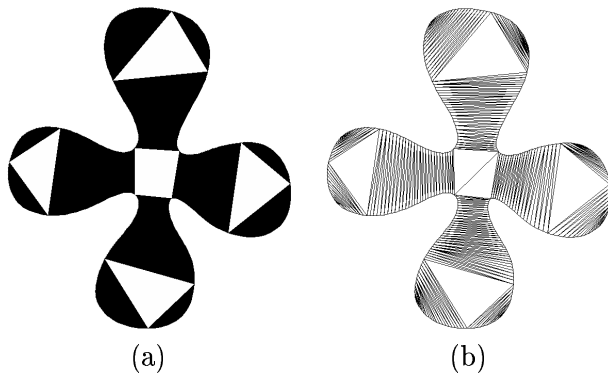


Figure 1: Anchor hull edges are approximated by thin Delaunay triangles while the anchor polygons are approximated by relatively fat Delaunay triangles.

i.e. $A(x) = \operatorname{argmin}_{y \in \partial\Sigma} \|x - y\|$. See Figure 2. The next observation follows immediately.

Observation 1 *For every $x \in \Sigma$ there exists a closed ball $B \subseteq \Sigma$ such that $A(x)$ is contained in the boundary of B .*

Anchor hull. For every $x \in \Sigma$, let $H(x)$ denote the convex hull of $A(x)$. We call $H(x)$ the *anchor hull* of x . See for example, Figure 2.

Observation 2 *For every $x \in \Sigma$, it holds that $H(x) \subseteq \Sigma$.*

PROOF. This follows from Observation 1 and the fact that balls are convex. \square

There is an interesting view of the anchor hulls in the *lifted* diagram of $\partial\Sigma$. Consider the map $h: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ where $h(x_1, x_2) = (x_1, x_2, x_1^2 + x_2^2)$. The curve $h(\partial\Sigma)$ is a lifting of $\partial\Sigma$ onto the paraboloid $x_3 = x_1^2 + x_2^2$. Let $\operatorname{Conv} h(\partial\Sigma)$ denote the convex hull of $h(\partial\Sigma)$. Each two dimensional anchor hull in Σ is the projection of a facet in $\operatorname{Conv} h(\partial\Sigma)$. What we are studying is the approximation of $\operatorname{Conv} h(\partial\Sigma)$ with the convex hull of a finite sample of $h(\partial\Sigma)$. Approximations of smooth convex bodies have been studied before [16]. However, these results cannot be called upon here straightforwardly since $\operatorname{Conv} h(\partial\Sigma)$ is not necessarily smooth.




Figure 2: Anchor sets and anchor hulls: $A(x) = \{b, c, d\}$ and $H(x)$ is the anchor triangle bcd , $A(y) = \{e, f\}$ and $H(y)$ is the anchor edge ef , $A(z) = \{a\}$ and $H(z)$ is the anchor point a .

Medial axis. The medial axis M of Σ is the closure of the set of center points of all maximal closed balls contained in Σ . The maximal balls are called *medial balls* which, by definition, are also tangent to $\partial\Sigma$.

The medial axis M , in general, is a geometric graph with branching points. It turns out that M can have infinitely many branchings if $\partial\Sigma$ is not well-behaved [10]. Although the examples are very pathological, we need to exclude these cases for our theoretical results. Also, to avoid complications in our proofs we need an assumption about the finiteness of the anchor sets.

Shape genericity assumptions. For the rest of the paper we assume that the shape Σ , other than being compact and bounded by a smooth curve $\partial\Sigma$, satisfies the following *generic* conditions.

- $A(x)$ is finite for all $x \in \Sigma$. That is, the anchor hulls are either points, line segments or polygons. We refer to them as *anchor points*, *anchor edges* and *anchor polygons*, respectively.
- The medial axis M has a finite graph structure; see Choi et al. [10] for the class of curves satisfying this condition.
- The medial balls with two points of tangency to $\partial\Sigma$ are non-degenerate. Consider any medial ball B with the center, say at m , where $A(m) = \{x, y\}$. When B is grown by moving its center along $m - x$, it cannot be tangent to another point, say z of $\partial\Sigma$, where z is arbitrarily close to y .

Let $N \subset M$ be the set of non-manifold points in M , i.e. the set of points that do not have an open neighborhood in M homeomorphic to an open

interval. Notice that N includes the boundary points of M as well as all branching points.

Observation 3 *For each point $m \in M \setminus N$ it holds that $|A(m)| = 2$.*

We define

$$\partial\Sigma^2 = \{x \in \partial\Sigma \mid \exists m \in M \setminus N \text{ with } x \in A(m)\}.$$

Pairing. We have a natural pairing μ on the points of $\partial\Sigma^2$, namely,

$$\mu : \partial\Sigma^2 \rightarrow \partial\Sigma^2, x \mapsto x' = \mu(x),$$

if there exists $m \in M \setminus N$ with $A(m) = \{x, x'\}$. See Figure 3 for an illustration.

Observation 4 *The pairing μ is a well defined function on $\partial\Sigma^2$.*

PROOF. By definition there exists a partner x' for every $x \in \partial\Sigma^2$. That is, all we have to show is that there exists only one such partner. Let x' be a partner of x and $m \in M \setminus N$ with $A(m) = \{x, x'\}$. Since we assumed that $A(m)$ is finite for all $m \in \Sigma$, there exists a unique ball $B \subset \Sigma$ with center m such that $B \cap \partial\Sigma = \{x, x'\}$. By construction the vector $m - x$ is orthogonal to the tangent of $\partial\Sigma$ at x and it points into the interior of the shape Σ . Since there can be only one such ball B , the partner x' of x has to be unique. \square

The continuity of μ implies the following nesting property of the pairing μ . To state this property we introduce the notation of a segment that turns out to be useful also later on. For $a, b \in \partial\Sigma$, we denote by $\sigma(a, b) \subset \partial\Sigma$ the shorter of the two open curve segments that connect a and b in $\partial\Sigma$. The closure of $\sigma(a, b)$ is denoted by $\overline{\sigma(a, b)}$.

Observation 5 *Let $\sigma(a, b) \subset \partial\Sigma^2$ be a curve segment and $\sigma(c, d) \subseteq \sigma(a, b)$. Then $\mu(\sigma(c, d)) \subseteq \mu(\sigma(a, b))$.*

We are also interested in $\partial\Sigma \setminus \partial\Sigma^2$. To classify these points we distinguish two types of points in N , namely boundary points m of M for which $|A(m)| = 1$ and non-boundary points m for which $|A(m)| \geq 3$. For $i > 0$, we define

$$\begin{aligned} \partial\Sigma^i &= \{x \in \partial\Sigma \mid \exists m \in M \text{ with } x \in A(m) \\ &\quad \text{and } |A(m)| = i\}. \end{aligned}$$

Note that, this definition subsumes the definition of $\partial\Sigma^2$. We have,

- (i) if $m \in M$ is a boundary-point, $A(m) \subseteq \partial\Sigma^1$,
- (ii) if $m \in M \setminus N$, $A(m) \subseteq \partial\Sigma^2$, and
- (iii) if $m \in N$ is non-boundary, $A(m) \subseteq \partial\Sigma^{>2}$.

See Figure 3 for an illustration.




Figure 3: Partition on $\partial\Sigma$: the points of some open curve segments are paired with points in its dashed partner. The endpoints of these curve segments are points in $\partial\Sigma^i$ for $i \neq 2$. For example, $a \in \partial\Sigma^1$ and $\{b, c, d\} \subset \partial\Sigma^3$.

2.2 Sampling.

In our setting a sampling of a shape Σ is a finite subset S of the boundary $\partial\Sigma$. Since we want a sampling to be feature adaptive, we adopt the notions of local feature size and ε -sample from Amenta, Bern and Eppstein [2].

Local feature size. The local feature size is a function f on $\partial\Sigma$ that assigns to every point $x \in \partial\Sigma$ its distance to the medial axis of Σ . Because of the second genericity conditions and compactness of $\partial\Sigma$, we have $\min_{x \in \partial\Sigma} f(x)$ lower bounded by a non-zero positive constant. It is a direct consequence of the triangle inequality that $f(\cdot)$ is 1-Lipschitz, i.e., $f(x) \leq f(y) + \|x - y\|$ for any two $x, y \in \partial\Sigma$.

ε -sample. A finite subset $S \subset \partial\Sigma$ is an ε -sample if every point $x \in \partial\Sigma$ has a point from S within a distance of $\varepsilon f(x)$.

Our main result concerns the approximation of anchor hulls by the Delaunay triangulation. Actually, we study the Delaunay triangulation of a

sample $S \subset \partial\Sigma$ restricted to Σ , denoted $\text{Del } S|_{\Sigma}$. To define this Delaunay triangulation we need the notion of Voronoi diagram. The Voronoi diagram of S is a cell decomposition of \mathbb{R}^2 into convex cells. Every *Voronoi cell* corresponds to exactly one sample point and contains all points of \mathbb{R}^2 that do not have a smaller distance to any other sample point.

Restricted Delaunay triangulation. The Delaunay diagram of a sample S restricted to Σ is a cell complex $\text{Del } S|_{\Sigma}$. The convex hull of three or more points in S defines a *Delaunay cell* in $\text{Del } S|_{\Sigma}$ if the intersection of the corresponding Voronoi cells is not empty in Σ and there exists no superset of points in S with the same property. Analogously, the convex hull of two points defines a *Delaunay edge* in $\text{Del } S|_{\Sigma}$ if the intersection of their corresponding Voronoi cells is not empty in Σ . Every point in S is a *Delaunay vertex* in $\text{Del } S|_{\Sigma}$. If the sample points in S are in general position, all Delaunay cells are triangles and we call the restricted Delaunay diagram $\text{Del } S|_{\Sigma}$ the restricted Delaunay triangulation of S .

We will show that all anchor edges are approximated by some Delaunay edges. Conversely, most Delaunay edges that are relatively long approximate some anchor edges. Some of the long Delaunay edges may not approximate any anchor edge. Their endpoints lie near the points in $\partial\Sigma^i$ for $i > 3$. They approximate a diagonal of the corresponding anchor polygon. In Figure 1, the edge dissecting the middle anchor polygon (a quadrilateral) is such an edge.

3 Approximation

3.1 Exclusions.

Our proofs proceed in two parts. First, we prove the results for all of $\partial\Sigma$ but some excluded regions. Then, we extend the results to these excluded regions.

Let $B(x, r)$ denote a ball of radius r with the center at x and $U_r(x) = \partial\Sigma \cap \text{interior}(B(x, r))$. For $\varepsilon_0 > 0$ define

$$\Gamma_{\varepsilon_0} = \text{closure} \left(\partial\Sigma \setminus \bigcup_{x \in A(N)} U_{\varepsilon_0 f(x)}(x) \right).$$

Figure 4 illustrates the exclusions for Γ_{ε_0} . We extend the exclusions a little

more to make them symmetric with respect to the pairing. This leads us to define,

$$\tilde{\Gamma}_{\varepsilon_0} = \text{closure}(\Gamma_{\varepsilon_0} \setminus \{x \in \Gamma_{\varepsilon_0} \mid \mu(x) \notin \Gamma_{\varepsilon_0}\}).$$

For $a \notin \partial\Sigma^2$ let $\xi(a)$ be the component in $\partial\Sigma \setminus \tilde{\Gamma}_{\varepsilon_0}$ that contains a , i.e., it is the open curve segment excluded from $\partial\Sigma$ by a . In what follows, by the endpoints of $\xi(a)$ we refer to the endpoints of the closure of $\xi(a)$.




Figure 4: Exclusions: The dark shaded curve segments are excluded from $\partial\Sigma$ to obtain Γ_{ε_0} (left). The excluded curve segment around a is extended and the curve segment $\sigma(s, t) = \xi(a)$ is excluded to obtain $\tilde{\Gamma}_{\varepsilon_0}$ (right).

Observation 6 For any $x \in \tilde{\Gamma}_{\varepsilon_0}$ we have $x' \in \tilde{\Gamma}_{\varepsilon_0}$.

PROOF. Follows immediately from the above definitions. \square

Observation 7 Let t be an endpoint of $\xi(a)$ for any $a \notin \partial\Sigma^2$. If $a \in \partial\Sigma^1$, the point t' is another endpoint of $\xi(a)$. Otherwise, t' is an endpoint of $\xi(b)$ where $\{a, b\} \subset A(n)$ for some non-boundary point $n \in N$.

PROOF. By construction $t \in \tilde{\Gamma}_{\varepsilon_0}$. Consider the open curve segment $\sigma(t, a)$. It is mapped to an open curve segment $\sigma(t', b)$ by μ where $\{a, b\} \subset A(n)$ for some $n \in N$. In case $a \in \partial\Sigma^1$, we have $a = b$.

Otherwise, we only need to show that t' is the endpoint of $\xi(b)$. If not, there are two possibilities: either $t' \in \xi(b)$, or there is an endpoint s of $\xi(b)$ so that $s \in \sigma(t', b)$. We handle both cases by exploiting the nesting property of the pairing μ , (Observation 5). In the first case we have $t' \notin \tilde{\Gamma}_{\varepsilon_0}$ which is a contradiction to $t' = \mu(t) \in \tilde{\Gamma}_{\varepsilon_0}$. In the second case we have on the one hand $\mu(s) \in \xi(a)$ and thus $\mu(s) \notin \tilde{\Gamma}_{\varepsilon_0}$. On the other hand we have $s \in \tilde{\Gamma}_{\varepsilon_0}$, because it is an endpoint of $\xi(b)$. This is a contradiction. \square

3.2 Sampling density.

In this subsection we detail the conditions on the sampling density ε that need to be satisfied for our results to hold.

CONDITION (i) For our proofs we need that $\varepsilon < \frac{1}{4}$.

CONDITION (ii) We need that $\xi(a) \cap \xi(b) = \emptyset$ for all $a, b \notin \partial\Sigma^2$. Actually we need that the excluded regions not only have empty intersection but are also well separated. We always can achieve such a well separation if we choose ε_0 for the exclusions sufficiently small. One should observe that this is the case, because we assume the finite graph structure of the medial axis. Since by the genericity assumption, the set N of non-manifold points are finite, they are isolated and so are the points in $A(N)$.

Let ε_1 be the smallest δ such that $\Gamma_\delta \subseteq \tilde{\Gamma}_{\varepsilon_0}$. This ε_1 will play an important role in our analysis.

Observation 8 For any two points $x, y \in \xi(a)$, it holds that $\|x - y\| \leq 2\varepsilon_1 f(a)$.

PROOF. This follows from the fact that $\xi(a)$ is included in $U_{\varepsilon_1 f(a)}(a)$. \square

For our proofs we will require that,

$$\varepsilon < \frac{\varepsilon_1}{4\Delta + 2\varepsilon_1},$$

where $\Delta = \sup_{x, y \in \Sigma} \frac{f(x)}{f(y)}$ is a shape dependent constant. Observe that $\Delta \geq 1$ implies that $\varepsilon < \varepsilon_1$.

CONDITION (iii) We need the following construction to explain this condition.

For $x \in \Gamma_{\varepsilon_1/2}$ let $B(m, \rho(x))$ be the corresponding medial ball, i.e. the ball with center m and radius $\rho(x) = \|x - m\|$ where $m \in M \setminus N$. Let $R(x)$ be the set that contains all $r > 0$ such the ball $B(m + r(m - x), (r + 1)\rho(x))$ intersects $\partial\Sigma$ in at least three connected components one of which is a curve segment $\overline{\sigma(a(r), b(r))}$ around x' . Notice that $R(x)$ can be empty. We set

$$\varepsilon_2 = \begin{cases} \infty & , \quad \forall x \in \Sigma, R(x) = \emptyset \\ \min_{x \in \Gamma_{\varepsilon_1/2}, R(x) \neq \emptyset, r \in R(x)} \|a(r) - b(r)\|/\rho(x), & \text{otherwise.} \end{cases}$$

We assume ε_2 to be strictly positive. It turns out that we do not need to use ε_2 in proofs when $R(x)$ is empty. That is why we set ε_2 arbitrarily to ∞ in the first case. For our proofs we will require that $\varepsilon < \frac{\varepsilon_2}{8}$.




Figure 5: Illustration for ε_2 : The medial ball for x and x' is grown to B when it intersects Γ in three connected components, x , c and $\sigma(a, b)$. By definition $\|a - b\| \geq \varepsilon_2 \rho(x)$.

3.3 Main result.

Our main result is the following approximation theorem.

Theorem 1 *Let Σ be a shape as described above and S an ε -sample of its boundary $\partial\Sigma$. Let $\varepsilon > 0$ fulfill the sampling conditions (i) - (iii). Then the following hold:*

- (1) *For any $x \in \partial\Sigma^2$ there is a Delaunay edge $pq \in \text{Del } S|_\Sigma$ where $\|x - p\| = O(\varepsilon_1)f(x)$ and $\|x' - q\| = O(\varepsilon_1)f(x')$.*
- (2) *Let $pq \in \text{Del } S|_\Sigma$ be a Delaunay edge where $\|p - q\| > \frac{2\varepsilon_1}{1-2\varepsilon_1}f(p) = \Omega(\varepsilon_1)f(p)$. Then, there exists an anchor hull $H(m)$ so that $\{x, x'\} \subseteq A(m)$ and $\|p - x\| = O(\varepsilon_1)f(x)$ and $\|q - x'\| = O(\varepsilon_1)f(x')$.*

Notice that (2) only refers to certain long Delaunay edges. This restriction cannot be avoided as some of the Delaunay edges approximate the edges that reconstruct $\partial\Sigma$ from S . These Delaunay edges do not approximate any anchor hull of a point on the medial axis. Also, if all anchor polygons are triangles, (2) implies that each long Delaunay edge approximates an anchor edge. As we mentioned earlier, in case of non-triangular anchor polygons, some of the long Delaunay edges may not approximate any anchor edge. Instead they approximate a diagonal of an anchor polygon $H(m)$ where $|A(m)| > 3$.

Also notice that the approximation guarantees are stated in terms of ε_1 . That is, at a first glance it seems that we do not get a better approximation for a denser sampling, i.e., smaller ε . However, the proper way to look at it, is as follows: Fix some constant ε_0 for the exclusion that works well. For any

ε that fulfills the conditions (i) – (iii) let ε'_0 be the infimum of all $\varepsilon_0 \geq \delta > 0$ such that if we replace ε_0 by δ the sampling density ε would still fulfill the conditions (i) – (iii). Obviously we have that ε'_0 goes to zero if ε does so. Let ε'_1 be defined via ε'_0 in the same way as ε_1 is defined via ε_0 . The continuity of μ on Σ^2 implies that ε'_1 goes to zero if ε'_0 does so. Thus in our guarantees we can replace ε_1 by ε'_1 . That shows that our guarantees actually imply that the approximations are getting better with higher sampling density ε .

4 Proofs

Lemma 1 *Let $B = B(x, r)$ be a ball centered at $x \in \partial\Sigma$ and $\overline{\sigma(a, b)}$ be the connected component in $B \cap \partial\Sigma$ containing x . If $\sigma(a, b)$ is empty of any sample point then*

$$\|x - a\| \leq \varepsilon f(x) \quad \text{and} \quad \|x - b\| \leq \varepsilon f(x).$$

PROOF. Suppose the contrary, i.e. $\|x - a\| > \varepsilon f(x)$ or $\|x - b\| > \varepsilon f(x)$. Since a and b are contained in the boundary of B , the radius r of B has to be larger than $\varepsilon f(x)$. By the sampling condition there exists a sample point $p \in S$ with $\|x - p\| \leq \varepsilon f(x)$. Thus, p is contained in B , but by our assumption it is not contained in $\sigma(a, b)$. Hence $B(x, r)$ intersects $\partial\Sigma$ in more than one component and thus has to contain a point of the medial axis, see [2]. This remains true if we shrink the radius r of B until B contains p in its boundary. That shows that $\|x - p\| \geq f(x)$ which contradicts $\|x - p\| \leq \varepsilon f(x)$ for $\varepsilon < 1$, a condition satisfied by our choice of ε , see condition (i) in the restriction on ε . \square

Lemma 2 *For any two points $x, y \in \partial\Sigma$ such that $\sigma(x, y) \cap S = \emptyset$, it holds that*

$$\|x - y\| \leq \frac{2\varepsilon}{1 - 2\varepsilon} f(x).$$

PROOF. Let s be the point where $\sigma(x, y)$ intersects the perpendicular bisector of xy . Apply Lemma 1 to the ball $B(s, \|x - s\|)$ to get $\|x - s\| \leq \varepsilon f(s)$ and $\|y - s\| \leq \varepsilon f(s)$. Applying the triangle inequality gives $\|x - y\| \leq 2\varepsilon f(s)$ and the Lipschitz property of f finally gives the desired result. \square

In the next lemma we show that all anchor edges connecting points outside the excluded regions are well approximated by Delaunay edges.

Lemma 3 *Let $x \in \tilde{\Gamma}_{\varepsilon_0}$. There is a Delaunay edge \overline{pq} in $\text{Del } S|_{\Sigma}$ with $\|x - p\| = O(\varepsilon)f(x)$ and $\|x' - q\| = O(\varepsilon)f(x')$.*

PROOF. Let $p \in S$ be the closest sample point to x , i.e., $\|x - p\| \leq \varepsilon f(x)$. If the segment $\sigma(x', p')$ does not contain any point from S then we have $\|x' - p'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(x')$ by Lemma 2. Otherwise, let $q \in \sigma(x', p') \cap S$ be the closest sample point in $\sigma(x', p')$ to x' . By Lemma 2, $\|x' - q\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(x')$. The point q' has to be contained in the segment $\sigma(x, p)$ by Observation 5 which implies again by Lemma 2 that $\|x - q'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(x)$. Both cases can essentially be treated the same way. Thus, we can assume without loss of generality that,

$$\|x - p\| \leq \varepsilon f(x) \text{ and } \|x' - p'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(x').$$

Since $p \in S$, it only remains to show that there exists $q \in S$ where pq is a Delaunay edge in $\text{Del } S|_{\Sigma}$ and $\|x' - q\| = O(\varepsilon)f(x')$.

From Observation 6 and $x \in \tilde{\Gamma}_{\varepsilon_0}$ we also have $x' \in \tilde{\Gamma}_{\varepsilon_0}$. Furthermore, p and p' are in $\Gamma_{\varepsilon_1/2}$ if $\varepsilon f(x) \leq \frac{\varepsilon_1}{2}f(a)$ and $\frac{2\varepsilon}{1-2\varepsilon}f(x') \leq \frac{\varepsilon_1}{2}f(a)$ for all $a \notin \partial\Sigma^2$. These constraints are satisfied by condition (ii) on ε_1 and ε .

Let $B = B(m, \|p - m\|)$ be the medial ball corresponding to p and p' , i.e. $A(m) = \{p, p'\}$. Grow B by moving its center m in the direction of $m - p$ while keeping p on its boundary. Stop the growth when B hits a sample point $q \in S$.

We claim that, when we stop, m is in Σ . If it were not, m would have crossed $\partial\Sigma$ at the point p' . At that moment B would have contained a curve segment $\sigma(z, p')$ empty of any sample point for some $z \in \partial\Sigma$ so that $\|z - p'\| \geq f(p')$. This violates Lemma 1. Thus, we have established $pq \in \text{Del } S|_{\Sigma}$.

Let $\overline{\sigma(a, b)}$ be the connected component of $B \cap \partial\Sigma$ containing p' . We want to establish $\|q - p'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(p')$. If $q \in \overline{\sigma(a, b)}$ we have $\|q - p'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(p')$ by Lemma 2. Otherwise, B meets $\partial\Sigma$ in at least three components and q is not contained in $\overline{\sigma(a, b)}$. Applying the definition of ε_2 to $p \in \Gamma_{\varepsilon_1/2}$ we get that

$$\varepsilon_2 \leq \frac{\|a - b\|}{\|p - m\|} = \frac{\|a - b\|}{\|p' - m\|} \leq \frac{\|a - b\|}{f(p')}.$$

Without loss of generality we can assume that $\|a - p'\| \geq \|b - p'\|$. That is, $\|a - p'\| \geq \frac{\|a - b\|}{2} \geq \frac{\varepsilon_2}{2}f(p')$. This contradicts Lemma 2 if $\frac{\varepsilon_2}{2} > \frac{2\varepsilon}{1-2\varepsilon}$. This again is satisfied by our choice of ε , see condition (iii) in the restriction on ε . Therefore, q must be contained in $\overline{\sigma(a, b)}$ and in fact be either a or b . So,

we conclude that $\|q - p'\| \leq \frac{2\varepsilon}{1-2\varepsilon}f(p')$. Using the triangle inequality

$$\begin{aligned} \|x' - q\| &\leq \|x' - p'\| + \|p' - q\| \\ &\leq \frac{2\varepsilon}{1-2\varepsilon}(f(x') + f(p')) \\ &\leq \frac{2\varepsilon(2-2\varepsilon)}{(1-2\varepsilon)^2}f(x'). \end{aligned}$$

This shows that $\|x - p\| = O(\varepsilon)f(x)$ and $\|x' - q\| = O(\varepsilon)f(x')$. \square

In the next lemma we eliminate the restriction over excluded regions from Lemma 3 and extend it to all of $\partial\Sigma$.

Lemma 4 *For any $x \in \partial\Sigma^2$ there is a Delaunay edge $pq \in \text{Del } S|_\Sigma$ where $\|x - p\| = O(\varepsilon_1)f(x)$ and $\|x' - q\| = O(\varepsilon_1)f(x')$.*

PROOF. If $x \in \tilde{\Gamma}_{\varepsilon_0}$ the claim is true by Lemma 3. So, assume otherwise. This means $x \in \xi(a)$ for some $a \notin \partial\Sigma^2$, i.e. x lies in some excluded region. There are two cases, either $a \in \partial\Sigma^1$ or $a \in \partial\Sigma^i$ for $i > 2$.

If $a \in \partial\Sigma^1$ then both x and x' belong to $\xi(a)$ which means $\|x - x'\| \leq 2\varepsilon_1f(a)$ by Observation 8. We know that if t is an endpoint of $\xi(a)$, the other endpoint is t' . The anchor edge tt' is approximated by a Delaunay edge $pq \in \text{Del } S|_\Sigma$ according to Lemma 3. This edge also approximates xx' with the stated bounds.




Figure 6: When $a \in \partial\Sigma^1$, both x and x' are in $\xi(a)$ shown by thickened curve segment. xx' is approximated by pq (left). The case when $x \in \xi(a)$ and $a \in \partial\Sigma^3$ is shown on the right.

Consider the remaining case where $a \in \Sigma^i$ for $i > 2$. Let n be the non-boundary point in N where $\{a, b\} \subset A(n)$ and $x' \in \xi(b)$. Let t be the endpoint of $\xi(a)$ so that x is contained in the segment $\sigma(t, a)$. By Observation 7, we have t' to be the endpoint of $\xi(b)$ where x' belongs to the

curve segment $\sigma(t', b)$. See Figure 6. By Lemma 3 there is a Delaunay edge $pq \in \text{Del } S|_\Sigma$ with $\|t - p\| = O(\varepsilon)f(t)$ and $\|t' - q\| = O(\varepsilon)f(t')$. By Observation 8 it also holds that $\|t - x\| = O(\varepsilon_1)f(t)$ and $\|t' - x'\| = O(\varepsilon_1)f(t')$. Using the triangle inequality, the Lipschitz property of f and $\varepsilon < \varepsilon_1$ we get $\|p - x\| = O(\varepsilon_1)f(x)$ and $\|q - x'\| = O(\varepsilon_1)f(x')$. \square

Next we proceed to prove part (2) of Theorem 1.

Lemma 5 *For each Delaunay edge $pq \in \text{Del } S|_\Sigma$ with $\|p - q\| > \frac{2\varepsilon}{1-2\varepsilon}f(p)$ and $p \in \tilde{\Gamma}_{\varepsilon_0}$ there exists an anchor hull $H(m)$ so that $A(m) = \{x, x'\}$ and $\|p - x\| = O(\varepsilon)f(x)$ and $\|q - x'\| = O(\varepsilon)f(x')$.*

PROOF. Let $B = B(c, \|p - q\|)$ be a Delaunay ball circumscribing the edge pq . First, B must intersect $\partial\Sigma$ in at least two disjoint segments, one containing p and another containing q . Otherwise, p and q belong to a single component of $\partial\Sigma \cap B$, the curve segment $\sigma(p, q)$ lies inside B and $\sigma(p, q)$ is empty of sample points. That creates a contradiction to Lemma 2 as $\|p - q\| > \frac{2\varepsilon}{1-2\varepsilon}f(p)$ by assumption.

Let σ_p and σ_q be the two components of $\partial\Sigma \cap B$ containing p and q , respectively. We claim that $\sigma_p \subset \Gamma_{\varepsilon_1/2}$. By Lemma 2 the Euclidean distance of p to any other point in σ_p cannot be more than $\frac{2\varepsilon}{1-2\varepsilon}f(p)$ since σ_p is empty of any sample point. Hence if $\frac{2\varepsilon}{1-2\varepsilon}f(p) \leq \frac{\varepsilon_1}{2}f(a)$ for all $a \notin \partial\Sigma^2$, we will have $\sigma_p \subset \Gamma_{\varepsilon_1/2}$. This condition is satisfied by the required relation between ε_1 and ε .

Now we shrink B radially till it meets one of σ_p or σ_q in a single point, say x . Shrink B further while keeping x on its boundary and moving its center c in the direction $x - c$ until it meets both σ_p and σ_q in a single point. One of these points is x and let the other be y . Also let $B' = B(c', r')$ be the new ball obtained by transforming B by the above motion.

First, we claim that $c' \in \Sigma$. Suppose it were not. Then, the center c of B during the above motion would have passed through a point, say z , in $\partial\Sigma$ since originally c lies in Σ as $pq \in \text{Del } S|_\Sigma$. At that moment B would have contained a curve segment of length larger than its radius while intersecting the curve in at least two connected components. This would mean that there is a curve segment $\sigma(w, z)$ empty of any sample point for some $w \in \partial\Sigma$ so that $\|w - z\| \geq f(z)$. This violates Lemma 1.

Next, we establish that $y = x'$. Suppose not. Then, B' meets $\partial\Sigma$ in at least three components and we can shrink B' further while keeping x on its boundary and moving its center c' in the direction $x - c'$ until B' meets

$\partial\Sigma$ in exactly two points. Let $B'' = B(c'', r'')$ be the new ball that we get from this transformation. We have that B'' is a medial ball, c'' a point on the medial axis with $A(c'') = \{x, x'\}$ and $y \neq x'$. Furthermore, reversing the transformation from B' to B'' shows that B' is a ball of the type that we used to define ε_2 . Let $\sigma(a, b)$ be the connected component of $B' \cap \partial\Sigma$ containing x' . The distance of x' to one of a and b is at least half the length of ab . Assume that

$$\begin{aligned} \|a - x'\| &\geq \frac{\|a - b\|}{2} \geq \frac{\varepsilon_2}{2} \|c'' - x\| \\ &= \frac{\varepsilon_2}{2} \|c'' - x'\| \geq \frac{\varepsilon_2}{2} f(x'). \end{aligned}$$

This contradicts Lemma 2 if $\frac{\varepsilon_2}{2} \geq \frac{2\varepsilon}{1-2\varepsilon}$ a condition satisfied by our choice of ε , see condition (iii) in our restriction on ε . Therefore, $y = x'$ and B' is a medial ball of Σ .

Without loss of generality assume $x \in \sigma_p$. Then, $\|x - p\| \leq \frac{2\varepsilon}{1-2\varepsilon} f(x)$ since otherwise we contradict Lemma 2 for $\sigma(x, p)$. The same holds for x' as it has to lie in σ_q which also is empty of any sample point. \square

In the next lemma we extend Lemma 5 by removing the restrictions on some of the excluded regions. Specifically, it extends Lemma 5 to all of $\partial\Sigma$ if Σ does not have any non-triangular anchor polygon.

Lemma 6 *Let $pq \in \text{Del } S|_\Sigma$ be a Delaunay edge where $\|p - q\| > \frac{2\varepsilon_1}{1-2\varepsilon_1} f(p)$. Then, there exists an anchor hull $H(m)$ so that $\{x, x'\} \subseteq A(m)$ and $\|p - x\| = O(\varepsilon_1) f(x)$ and $\|q - x'\| = O(\varepsilon_1) f(x')$.*

PROOF. If p or q belongs to $\tilde{\Gamma}_{\varepsilon_0}$, the claim follows from Lemma 5 and $\varepsilon < \varepsilon_1$. So, assume that neither of p or q belongs to $\tilde{\Gamma}_{\varepsilon_0}$. Let $p \in \xi(a)$ where $a \in A(n)$ for a non-manifold medial axis point $n \in N$.

First consider the case where n is a non-boundary point in N . We claim q has to lie in some $\xi(b)$ where $b \in A(n)$, see Figure 7. Each of the edges of the anchor polygon $H(n)$ has a Delaunay edge approximating it according to Lemma 4. The Delaunay edge pq cannot intersect any of these Delaunay edges except at their endpoints. Consider the space $X \subset \mathbb{R}^2$ delimited by these Delaunay edges and the shorter curve segments between their endpoints, see the shaded region in Figure 7. We claim that $q \in X$. Suppose not. Because of our choice of ε (condition (ii)) and Lemma 3, each of these curve segments contains a point c where $c \in A(n)$ and no other point from $\partial\Sigma^i$ for $i \geq 3$. Therefore, the edge pq has to intersect a Delaunay

edge as p lies inside X and q lies outside it. We reach a contradiction. Thus, $q \in X$. But then even $q \in \xi(b)$ for some $b \in A(n)$, because we assumed $q \notin \tilde{\Gamma}_{\varepsilon_0}$.

Next, we show $a \neq b$. Assume the contrary. Then both p and q have to be contained in $\xi(a)$. By Observation 8 it follows that $\|p - q\| \leq 2\varepsilon_1 f(a)$. Also from Observation 8 it holds $\|p - a\| \leq 2\varepsilon_1 f(a)$. The latter inequality together with the Lipschitz continuity of f imply that $\|p - a\| \leq \frac{2\varepsilon_1}{1-2\varepsilon_1} f(p)$. Using the Lipschitz continuity of f once more we get

$$\begin{aligned} f(a) &\leq f(p) + \|p - a\| \\ &\leq \left(1 + \frac{2\varepsilon_1}{1-2\varepsilon_1}\right) f(p) \\ &= \frac{1}{1-2\varepsilon_1} f(p). \end{aligned}$$

From our assumption on the length of pq it follows

$$\|p - q\| > \frac{2\varepsilon_1}{1-2\varepsilon_1} f(p) \geq 2\varepsilon_1 f(a).$$

This is a contradiction. Thus, both points p and q cannot lie simultaneously in $\xi(a)$. This means $a \neq b$ as the excluded regions are separated by condition (ii) on ε . Then, pq approximates the edge ab with the conditions as stated in the theorem.




Figure 7: The region X is shaded. Delaunay edge pq approximates the diagonal ab of the anchor hull with vertices $\{a, b, c, d\}$.

Next, consider the case where n is an endpoint of M , i.e. a boundary point in N , and $A(n) = a$. Let t and t' be the endpoints of $\xi(a)$. Using Observation 8 and the Lipschitz property of f we have

$$\|p - t\| \leq 2\varepsilon_1 f(a) \leq \frac{2\varepsilon_1}{1-2\varepsilon_1} f(t)$$

and

$$\|p - t'\| \leq 2\varepsilon_1 f(a) \leq \frac{2\varepsilon_1}{1 - 2\varepsilon_1} f(t').$$

By Observation 7, the segment tt' is an anchor edge. Also since t and t' belong to $\tilde{\Sigma}_{\varepsilon_0}$ by definition, we have, according to Lemma 3, a Delaunay edge, say uv , with $\|u - t\| = O(\varepsilon)f(t)$ and $\|v - t'\| = O(\varepsilon)f(t')$. By the construction used in the proof of Lemma 3 either u is the closest sample point to t in one direction along Σ or v is the closest sample point to t' in one direction along Σ . In both cases p has to be contained in $\sigma(u, v)$. We distinguish two cases either $p \in \{u, v\}$ or $p \notin \{u, v\}$.

In the first case we can assume without loss of generality that $p = u$. From Lemma 2 and $\|p - q\| \geq 2\varepsilon_1 f(p)$ we have that q can not be a neighbor of p along Σ . Thus, it has to be neighbor of v along Σ . Hence by Observation 8 it is $\|q - v\| \leq \frac{2\varepsilon}{1 - 2\varepsilon} f(v)$. Thus, we can conclude from Lemma 2 that

$$\|q - t'\| \leq \|q - v\| + \|v - t'\| \leq \frac{2\varepsilon}{1 - 2\varepsilon} f(v) + O(\varepsilon)f(t').$$

It follows from the Lipschitz continuity of f that $\|q - t'\| = O(\varepsilon)f(t') = O(\varepsilon_1)f(t')$ and pq approximates the anchor hull edge tt' with the properties stated in the theorem.

In the second case also q has to be contained in $\sigma(u, v)$. By Observation 8 we have $q \notin \sigma(t, t')$. We are left with two cases: if u is the closest sample point to t in one direction along Σ then $q \in \sigma(t', v)$ and if v is the closest sample point to t' in one direction along Σ then $q \in \sigma(u, t)$. Both cases can be dealt with in the same way. Here we consider the case that $q \in \sigma(t', v)$. By the construction used in the proof of Lemma 3 we have that v is the closest sample point to u' in one direction along Σ . That is, we either have $q \in \sigma(u, u')$ or $q = v$. In the latter case pq approximates the anchor hull edge tt' with the properties stated in the theorem. In the other case we even have $q \in \sigma(t', u')$ and $q' \in \sigma(t, u)$. Since by construction $\sigma(t, u)$ is empty of sample points also $\sigma(t, q')$ is empty sample points, i.e. $\|t - q'\| \leq \frac{2\varepsilon}{1 - 2\varepsilon} f(q')$ by Lemma 2. This implies,

$$\|p - q'\| \leq \|p - t\| + \|t - q'\| \leq \frac{2\varepsilon_1}{1 - 2\varepsilon_1} f(t) + \frac{2\varepsilon}{1 - 2\varepsilon} f(q').$$

Using the Lipschitz continuity of f we get $\|p - q'\| = O(\varepsilon_1)f(q')$. That is, pq approximates the anchor hull edge qq' with the properties stated in the theorem. \square

Lemma 4 and Lemma 6 prove Theorem 1.

5 An Application

We apply the anchor hull approximation to the problem of matching shapes. We follow the intuition that similar shapes often have similar anchor hulls. Formally, this intuition is justified by the fact that shapes under small C^2 -perturbations have small Hausdorff distance between their medial axes; see Attali et al. [6]. We devise an algorithm that first segments a shape based on the anchor hulls and then matches two shapes with respect to these segments.

Segmentation. Given a point sample S of the boundary $\partial\Sigma$ of a shape Σ , we identify the triangles in a triangulation of the anchor polygons from the Delaunay triangulation of S . Obviously, anchor polygons of Σ that have small edges compared to the local feature size are hard to identify. The Delaunay triangles approximating these anchor polygons become indistinguishable from other thin triangles that approximate anchor edges. Therefore, we look for approximation of the *prominent* anchor polygons. We say an anchor polygon is prominent in Σ if all of its edges have length more than 1.5 times the feature size at its vertices. It is known that, in two dimensions Voronoi vertices approximate the medial axis [9]. Therefore, for a point $p \in S$ we approximate $f(p)$ by $\hat{f}(p)$ which is the distance of p to the closest Voronoi vertex. We say a Delaunay triangle pqr in the Delaunay triangulation of S is *prominent* if all of its edges have lengths more than 1.5 times $\max\{\hat{f}(p), \hat{f}(q), \hat{f}(r)\}$.

Before proceeding to the segmentation, we reconstruct $\partial\Sigma$ from S so that we have $\text{Del } S|_{\Sigma}$. The underlying space $\hat{\Sigma}$ of $\text{Del } S|_{\Sigma}$ approximates Σ . Any of the curve reconstruction algorithms [11] can be chosen for that purpose. Let T be the set of prominent triangles in $\text{Del } S|_{\Sigma}$ and $|T|$ be the underlying space of their union. Then, each triangle in T and each connected component of $\hat{\Sigma} - |T|$ becomes a segment in our segmentation of $\hat{\Sigma}$. Figure 8 shows this segmentation of some shapes.

Signature. These segments of $\hat{\Sigma}$ are mapped to a set of weighted points called the *signature* of Σ . In order to measure the similarity of two shapes, we compare their signatures which boils down to matching two weighted point sets of small cardinality. This signature generation and a subsequent scoring process is very similar to a method described in [13]. We include the details for completeness.

Let $R_{S,\Sigma}$ denote the set of segments that are computed from a point sample of the shape boundary $\partial\Sigma$. To simplify notations we use R_{Σ} for




Figure 8: Segmentation using prominent Delaunay triangles

$R_{S,\Sigma}$. Let $r \in R_\Sigma$ be a segment where

$$r = \bigcup_{\sigma \in \text{Del}S|_\Sigma} \sigma.$$

Further let c_σ and v_σ denote the centroid and volume of σ respectively. The *representative point* r^* of a segment r and its weight \hat{r} are defined as follows.

If r is a prominent anchor triangle σ , then

$$\begin{aligned} \hat{r} &= \text{circumradius of } \sigma \text{ and} \\ r^* &= \text{circumcenter of } \sigma. \end{aligned}$$

Otherwise,

$$\begin{aligned} \hat{r} &= \sum_{\sigma \in r} v_\sigma \\ r^* &= \frac{\sum_{\sigma \in r} (v_\sigma \cdot c_\sigma)}{\hat{r}}. \end{aligned}$$

That is, when r is not a prominent triangle, the weight of r is its area and its representative point is the weighted average of the centroids of all $\sigma \in r$, weight being the area of each triangle.

A prominent anchor polygon $H(m)$ with more than three vertices are approximated by union of prominent triangles in the Delaunay triangulation. The circumcenters and circumradii of these triangles approximate the medial axis point m and the radius of the corresponding medial ball respectively. Thus, m and this radius and their approximations tend to be similar across similar shapes though the approximate triangulation of $H(m)$ may be different. This justifies our choice of r^* and \hat{r} in case r is a prominent triangle.

Given a segmentation R_Σ of a shape Σ , the signature $Sg(\Sigma)$ is defined as a set of weighted points, i.e.,

$$Sg(\Sigma) = \{(r^*, \hat{r}) \mid r \in R_\Sigma\}.$$

Scoring. The amount of similarity between two shapes is measured by scoring the similarity between their signatures. In order to score the similarity between two signatures $Sg(\Sigma_1)$ and $Sg(\Sigma_2)$, we need aligning them first.

Let r^*, s^* be the representative points in $Sg(\Sigma_1)$ and $Sg(\Sigma_2)$, respectively, with maximum weights. We first translate $Sg(\Sigma_2)$ so that r^*, s^* coincide. Then an alignment is obtained by rotating $Sg(\Sigma_2)$ so that a line segment between s^* and another point of $Sg(\Sigma_2)$ aligns with a line segment between r^* and another point in $Sg(\Sigma_1)$. Certainly, there are $\Theta(mn)$ alignments possible where $|Sg(\Sigma_1)| = m$ and $|Sg(\Sigma_2)| = n$. Since m, n are typically small (less than ten), checking all alignments is not prohibitive.

For each alignment we compute a score and the maximum of all the scores is taken to be the amount of similarity and corresponding transformations give the best alignment.

Before we compute the score, the weights of the segments are normalized so that each weight is between 0 and 1. Next, for each point $q^* \in Sg(\Sigma_2)$, we determine the Euclidean nearest neighbor, say p^* , in $Sg(\Sigma_1)$. If $\|p^* - q^*\|$ is less than a *threshold*, we compute a *similarity score* as

$$1 - \left| \frac{\hat{p} - \hat{q}}{\hat{p} + \hat{q}} \right|$$

where the *threshold* is a parameter that tells how much tolerance we can have for the proximity of two segments. The points in $Sg(\Sigma_1)$ and $Sg(\Sigma_2)$ that do not have nearest neighbors in the other set within threshold distance contribute to a *dissimilarity score* which is equal to the negative of their weights. Finally, we add both similarity and dissimilarity scores to obtain the score of matching between the two shapes Σ_1 and Σ_2 .

Experimental results. We implemented the above matching algorithm and experimented with it on a database of approximately 300 shapes. Figure 9 shows the result. The leftmost column contains the query shapes. Each row contains four shapes that matched with the query shape with the four highest scores. We also show the segmentation of each shape. It is evident that similar shapes are mostly segmented similarly as they have similar anchor hulls. We compared our results with the shape matching algorithm
















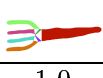
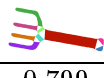
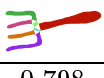
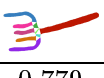
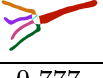
				
1.0	0.812	0.677	0.635	0.633
				
1.0	0.625	0.571	0.554	0.545
				
1.0	0.993	0.977	0.958	0.953
				
1.0	0.799	0.798	0.779	0.777

Figure 9: Matching results: Leftmost column shows the query shapes. Each row contains the shapes with the highest four scores.

of [13]. The results are very much comparable though the anchor hull based algorithm is simpler.

6 Conclusions

In this paper we established an approximation result between the Delaunay triangulation of a point sample from the boundary of a two dimensional shape and its anchor hulls. Obviously, the results hold for the complement of the shape as well. What about shapes in three dimensions? We believe that the results should extend to three dimensions as well. We plan to work out the details in future research.

Acknowledgements. We thank J.-D. Boissonnat, D. Attali and A. Lieutier for pointing out the need for the third genericity condition.

References

- [1] H. Alt and L. J. Guibas. Discrete geometric shapes: matching, interpolation, and approximation: a survey. Tech. report B 96-11, EVL-1996-142, Institute of Computer Science, Freie Universität Berlin, 1996.
- [2] N. Amenta, M. Bern and D. Eppstein. The crust and the β -skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, vol. 60 (1998), 125-135.
- [3] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.* **22** (1999), 481–504.
- [4] N. Amenta, S. Choi and R. Kolluri. The power crust, union of balls, and the medial axis transform. *Comput. Geom. Theory Appl.* **19** (2001), 127–153.
- [5] E. M. Arkin, L. P. Chew, D. P. Huttenlocher, K. Kedem and J. S. Mitchell. An efficiently computable metric for comparing polygonal shapes. *IEEE Trans. Pattern Analysis Machine Intelligence*, **13** (1991), 209–216.
- [6] D. Attali, J.-D. Boissonnat and H. Edelsbrunner. Stability and computation of medial axes – a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, T. Möller, B. Hamann and B. Rusell eds., Springer-Verlag (2005), to appear.
- [7] D. Attali, J.-D. Boissonnat and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces: the smooth case. *Proc. Annu. Sympos. Comput. Geom.* (2003), 201–210.
- [8] J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Comput. Geom. Theory Appl.* **19** (2001), 87–120.
- [9] J. W. Brandt and V. R. Algazi. Continuous skeleton computation by Voronoi diagram. *Comput. Vision, Graphics, Image Process.* **55** (1992), 329–338.
- [10] H. I. Choi, S. W. Choi and H. P. Moon. Mathematical theory of medial axis transform. *Pacific J. Mathematics* **181** (1997), 57–88.
- [11] T. K. Dey. Curve and surface reconstruction. Chapter in *Handbook on Discrete and Computational Geometry*, J. Goodman and J. O’Rourke eds. CRC press, 2nd edition, (2004), 677-692.
- [12] T. K. Dey. Sample based geometric modeling. *AMS/DIMACS volume on Computer-aided Design and Manufacturing*, D. Dutta, M. Smid and R. Jarnardan eds., to appear.
- [13] T. K. Dey, J. Giesen and S. Goswami. Shape segmentation and matching with flow discretization. *Proc. Workshop Algorithms Data Structures (WADS03)* (2003), LNCS 2748, F. Dehne, J.-R. Sack and M. Smid eds., 25–36.

- [14] T. K. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a Convergence Guarantee. *Algorithmica* **38** (2003), 179–200.
- [15] J. Erickson. Uniform samples of generic surfaces have nice Delaunay triangulations (2003), Manuscript.
- [16] P. M. Gruber. Aspects of approximation of convex bodies. *Handbook of Convex Geometry*, P. M. Gruber and J. M. Wills eds., North Holland, Amsterdam (2003), 319–345.
- [17] R. Osada, T. Funkhouser, B. Chazelle and D. Dobkin. Matching 3D models with shape distribution. *Proc. Shape Modeling Internat.*, 2001.
- [18] M. Tănase and R. C. Veltkamp. Polygonal decomposition based on the straight line skeleton. *Proc. 19th Annu. Sympos. Comput. Geom.* (2003), 58–67.