



Stability and Computation of Medial Axes: a State-of-the-Art Report

Dominique Attali, Jean-Daniel Boissonnat, Herbert Edelsbrunner

► To cite this version:

Dominique Attali, Jean-Daniel Boissonnat, Herbert Edelsbrunner. Stability and Computation of Medial Axes: a State-of-the-Art Report. T. Möller and B. Hamann and R. Russell. Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration, Springer-Verlag, pp.109-125, 2009, Mathematics and Visualization. hal-00468690

HAL Id: hal-00468690

<https://hal.science/hal-00468690>

Submitted on 31 Mar 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stability and Computation of Medial Axes — a State-of-the-Art Report

Dominique Attali¹, Jean-Daniel Boissonnat², and Herbert Edelsbrunner³

¹ LIS, ENSIEG, Domaine Universitaire, BP 46, 38402 Saint Martin d'Hères, France. Dominique.Attali@lis.inpg.fr

² INRIA, 2004 Route des Lucioles, BP 93, 06904 Sophia-Antipolis, France. Jean-Daniel.Boissonnat@sophia.inria.fr

³ Department of Computer Science, Duke University, Durham, and Raindrop Geomagic, Research Triangle Park, North Carolina, USA. edels@cs.duke.edu

Summary. The medial axis of a geometric shape captures its connectivity. In spite of its inherent instability, it has found applications in a number of areas that deal with shapes. In this survey paper, we focus on results that shed light on this instability and use the new insights to generate simplified and stable modifications of the medial axis.

1 Introduction

In this paper, we survey what is known about the medial axis of a geometric shape. To get an intuitive feeling for this concept, consider starting a grass fire along a curve in the plane, like the outer closed curve in Figure 1. The




Fig. 1: Medial axis of shape whose boundary is the outer closed curve. The distance-to-boundary function has three critical points, one saddle and two maxima. One of the maxima coincides with a branch point.

fire starts at the same moment, everywhere along the curve, and it grows at constant speed in every direction. The medial axis is the set of locations where the front of the fire meets itself. In mathematical language: it is the set of points that have at least two closest points on the curve. If we start the fire along the boundary of a geometric shape in $\mathbb{R}^k$ we generically get a medial axis of dimension $k - 1$, one less than the dimension of the space.

In the plane, the medial axis is a (one-dimensional) graph whose branches correspond to regions of the shape it represents. Its structure has found applications in image analysis for shape recognition [52] and in robotics for motion planning [42]. The distance-to-boundary recorded at points of the medial axis provides information about local thickness, which can be used to segment the shape, separating it into large regions with relatively narrow connections [21, 25, 45]. In reverse engineering, the medial axis appears naturally as a tool to characterize the sampling density needed to reconstruct a curve in the plane and a surface in space [1, 2]. Other applications include domain decomposition in mesh generation [46, 50], feature extraction in geometric design [38, 39], and tool-path creation in computer-aided manufacturing [36].

In this paper, we make no attempt to cover the large amount of work on medial axes in digital image processing and instead refer to texts in the area [35, 40, 41]. Whenever possible, we state the definitions and results for $\mathbb{R}^k$, where k is an arbitrary but fixed constant, but sometimes we need to limit ourselves to $\mathbb{R}^2$ and $\mathbb{R}^3$. Since most applications are in two and three dimensions, this limitation implies only a minor loss of relevance. The style adopted in this paper is not that of a typical survey paper. Rather than aiming at a broad coverage of the literature, we focus on a small number of results that we deem important. Those are centered around questions of stability and computation of the medial axes. We encapsulate the various topics in a relatively large number of small and by-and-large independent sections. Starting with fundamental properties, we slowly progress towards more advanced results. In Section 2, we define medial axes and skeletons. Sections 3 and 4 state properties of the medial axis that concern its finiteness and its homotopy type. Sections 5, 6, and 12 discuss the stability of the medial axis under various notions of distance. Section 7 recalls that computing the medial axis exactly runs into obstacles except for certain classes of shapes, and Section 8 introduces the approximation paradigm designed to circumvent these obstacles. Sections 9, 10, 11 and 13 describe steps and aspects of this paradigm. We find that topics of stability and computation are related, which is the reason for interleaving the sections as enumerated.

2 Medial axis and skeleton

There is no generally agreed upon definition for either the medial axis or the skeleton of a shape; the precise meaning of these terms changes from one author to another. The medial axis has been introduced by Blum [12] as a

tool in image analysis. In this paper, we adopt the definitions given in [43]. Let X be a bounded open subset of the k -dimensional Euclidean space, $\mathbb{R}^k$. The *medial axis*, $\mathcal{M}[X]$, is the set of points that have at least two closest points in the complement of X . We call an open ball $B \subseteq X$ *maximal* if every ball that contains B and is contained in X equals B . The *skeleton* is the set of centers of maximal balls. The two notions are closely related but not the same. Specifically, the medial axis is a subset of the skeleton which is, in turn, a subset of the closure of the medial axis [44, chapter 11]. In the most general case, the skeleton is not necessarily closed and the last inequality is strict. Examples of shapes in $\mathbb{R}^2$ whose skeletons are composed of an infinite number of curves and which are not closed can be found in [16, 22]. A simple example of a skeleton in $\mathbb{R}^3$ that is not closed is given in [20]. Even though medial axes are not necessarily the same as skeletons, the two concepts are too similar to warrant a balanced treatment of both. The rest of this paper will therefore focus exclusively on the medial axis.

If we weight each point x of the medial axis with the radius $\rho(x)$ of the maximal ball centered at x , then we have enough information to reconstruct the shape. In other words, the medial axis together with the map ρ provides a reversible coding of shapes. This coding is not necessarily minimal and some shapes, such as finite unions of balls, can be reconstructed from proper subsets of their weighted medial axes.

3 Finiteness properties

There are cases in which the medial axis has infinitely many branches, even if the shape is bounded and its boundary is C^∞ -smooth [22]. To construct such an example, let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a C^∞ -smooth function defined by

$$f(x_1) = \begin{cases} 1 - e^{-\frac{1}{x_1^4}} \left(\sin \frac{\pi}{x_1} \right)^2 & \text{if } |x_1| < 1, \\ 1 & \text{if } |x_1| \geq 1. \end{cases}$$

Consider the set of points $X \subseteq \mathbb{R}^2$ above its graph, as shown in Figure 2. The medial axis of X consists of infinitely many branches, one for each oscillation of f . To obtain a bounded shape, we apply inversion, mapping every point $x \in X$ to $\iota(x) = x/\|x\|^2$. The inversion preserves circles and incidences between curves. It follows that the medial axis of $\iota(X)$ has the same structure and number of branches as the medial axis of X . More specifically, if we compactify $\mathbb{R}^2$ and join all branches of $\mathcal{M}[X]$ at the added point at infinity, then we have a homeomorphism between $\mathcal{M}[X]$ and $\mathcal{M}[\iota(X)]$. The point at infinity maps to the center of the circle that is the image of the line $x_2 = 1$. We note, however, that the homeomorphism between the two medial axes is different from ι , which does not preserve centers of circles. The above construction can be extended to produce medial axes with infinitely many branch points and similar pathological examples in higher dimension.




Fig. 2: The upper ‘half-plane’ X bounded by a smooth curve and its image under inversion. Both shapes have medial axes with infinitely many branches.

In the plane, Choi, Choi and Moon [22] establish that the medial axis of a bounded shape whose boundary is piecewise real analytic is a finite graph. Chazal and Soufflet [20] extend this result to *semianalytic* bounded open sets, which are bounded and open subsets $X \subseteq \mathbb{R}^k$ for which every point of $\mathbb{R}^k$ has a neighbourhood U with $X \cap U$ defined by a finite system of analytic equations and inequalities. They prove that such sets have medial axes which admit stratifications and satisfy finiteness properties. Specifically, the medial axis can be decomposed into a finite number of strata, each a connected i -manifold with boundary, for $i < k$. Furthermore, the medial axis has finite j -dimensional volume, where j denotes the largest dimension of any stratum. In addition, in $\mathbb{R}^3$ the boundary of the medial axis consists of a finite union of points and curves of finite length. For shapes in $\mathbb{R}^3$, a classification of points in the closure of the medial axis can be found in [4, 33].

4 Homotopy equivalence

In [43], Lieutier proves that any bounded open subset $X \subseteq \mathbb{R}^k$ is homotopy equivalent to its medial axis. Unlike earlier works [53, 51], he assumes no regularity condition on the boundary of X . Intuitively, this result implies that the medial axis and the shape are connected the same way, ignoring local dimensionality. To be formal, we say that two maps f and g from X to Y are *homotopic* if there exists a continuous map $H : X \times [0, 1] \rightarrow Y$ with $H(x, 0) = f(x)$ and $H(x, 1) = g(x)$. Using this definition, two spaces X and Y are *homotopy equivalent* if one can find two continuous maps $f : X \rightarrow Y$ and $g : Y \rightarrow X$ such that $g \circ f$ is homotopic to the identity on X and $f \circ g$ is homotopic to the identity on Y . To establish the homotopy equivalence between a shape X and its medial axis, Lieutier considers the distance-to-boundary function, ρ , which associates to each point $x \in X$ its distance to the complement of X and defines the vector field $\nabla \rho : X \rightarrow \mathbb{R}^k$ by

$$\nabla\rho(x) = \frac{x - c(x)}{\rho(x)}, \quad (1)$$

where $c(x)$ is the center of the smallest ball enclosing the set of points in the complement of X closest to x . Used before for the purpose of surface reconstruction in [29, 34], this vector field extends the gradient of ρ to points on the medial axis, where the gradient is not defined. The extended vector field is not continuous but induces flow lines used in the proof to map the shape to its medial axis. Specifically, each point x is mapped to the point $f(x)$ it occupies after flowing along the vector field $\nabla\rho$ for a sufficiently long but constant amount of time. The proof uses $f : X \rightarrow \mathcal{M}[X]$ and the inclusion map $g : \mathcal{M}[X] \rightarrow X$ to establish the homotopy equivalence.

5 Instability and semi-continuity

We think of $\mathcal{M}$ as a transform that maps the shape X to its medial axis, $\mathcal{M}[X]$. As emphasized in [43], geometric shapes are usually not known exactly and represented by approximations of one kind or another. For example, the boundary of a shape may be approximated by a triangulation obtained by software for surface reconstruction or segmentation. Under these circumstances, it would be important that the transform be continuous. In other words, one should be able to compute an arbitrarily accurate approximation of the output for a sufficiently accurate approximation of the input. Most commonly, one would use the Hausdorff distance to quantify the difference between two inputs and two outputs and this way define what it means for the transform to be continuous. Unfortunately, the medial axis transform is not continuous under this notion of distance: small modifications of the input shape can induce large modifications of its medial axis. This effect is illustrated in Figure 4, where we compare the medial axis of an oval on the left with the medial axis of a set whose Hausdorff distance to the oval is bounded from above by $\epsilon > 0$. The difficulty of approximating the medial axis due to its instability with respect to the Hausdorff distance is a well-known but until recently not well-understood problem.

One can observe experimentally that small modifications of a shape do not affect the entire medial axis. Typical effects for shapes in $\mathbb{R}^2$ are fluctuating branches that leave the rest of the medial axis unchanged. Similarly, for shapes in $\mathbb{R}^3$ we notice fluctuating spikes, added to or removed from the otherwise stable structure. This observation is consistent with the fact that the medial axis is semicontinuous with respect to the Hausdorff distance [44, chapter 11]. To explain this concept, we let A and B be subsets of $\mathbb{R}^k$ and write $d_H(A|B) = \sup_{x \in A} d(x, B)$ for the *one-sided Hausdorff distance* of A from B , where $d(x, B)$ is the infimum of the Euclidean distances between x and points y in B . Observe that $d_H(A|B) < \epsilon$ if and only if A is contained in the parallel body $B^{+\epsilon} = \{x \in \mathbb{R}^k \mid d(x, B) < \epsilon\}$. The *Hausdorff distance*

between A and B is $d_H(A, B) = \max\{d_H(A|B), d_H(B|A)\}$. We write A^c and B^c for the complements of A and B and note that the Hausdorff distance between A^c and B^c is generally different from that between A and B . Indeed, $d_H(A^c, B^c)$ is forgiving for small islands of A far away from B , while $d_H(A, B)$ is forgiving for small holes of A far away from B^c . With this notation, we are ready to define the concept of semicontinuity. Specifically, a transform $\mathcal{T}$ is *semicontinuous* if for every bounded open subset $X \subseteq \mathbb{R}^k$ and for every $\delta > 0$, there exists $\epsilon > 0$ such that for every open subset Y of $\mathbb{R}^k$,

$$d_H(X^c, Y^c) < \epsilon \implies d_H(\mathcal{T}[X] | \mathcal{T}[Y]) < \delta. \quad (2)$$

Note that ϵ depends on X . In words, small Hausdorff distance between the complements of X and Y implies that $\mathcal{T}[X]$ is contained in a tight parallel body of $\mathcal{T}[Y]$. As mentioned earlier, this condition is satisfied for $\mathcal{T} = \mathcal{M}$.

6 Stability under C^2 -perturbations

In [20], Chazal and Soufflet prove that the medial axis transform is continuous when C^2 -perturbations are applied to shapes in $\mathbb{R}^3$. To define what this means, we call two multilinear maps ϵ -close if the norm of their differences is less than ϵ . A map $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is an ϵ -small C^m -perturbation if

- (i) $f(x) = x$ outside some compact subset of $\mathbb{R}^3$,
- (ii) f is a C^m -diffeomorphism,
- (iii) the i -th derivatives of f and f^{-1} are ϵ -close to the i -th derivative of the identity map, for all points $x \in \mathbb{R}^3$ and all i from 0 to m .

Let X be an open subset of $\mathbb{R}^3$ whose boundary is a C^2 -smooth manifold [37]. Chazal and Soufflet [20] prove that a small C^2 -perturbation f implies a small Hausdorff distance between the medial axes of X and $f(X)$. Formally, for every $\delta > 0$, there exists $\epsilon > 0$ such that for every ϵ -small C^2 -perturbation f , $d_H(\mathcal{M}[X], \mathcal{M}[f(X)]) < \delta$. This result is optimal for m since the medial axis of a shape is already instable under C^1 -perturbations [20]. Therefore, for approximating the medial axis of X with the medial axis of Y , the boundary of Y must be close to the boundary of X both in position, normal direction and curvature. Unfortunately, effective implementations of exact algorithms for the medial axis are known only for restricted families of shapes, such as polyhedra, unions of balls and complements of discrete point sets, whose boundaries are generally not C^2 . In other words, it is unlikely that the positive approximation result for C^2 -perturbations can be turned into a practical algorithm.

7 Exact computation of medial axes

A fairly general class of shapes for which it is possible, in principle, to compute the medial axis exactly are the *semi-algebraic sets*, each the set of solutions

of a finite system of algebraic equations and inequalities. The medial axis of such a set is itself semi-algebraic and can be computed with tools from computer algebra. To describe this, let X be a shape in $\mathbb{R}^3$ whose boundary is a C^1 -smooth manifold. We introduce the *symmetry set* of X , consisting of the centers of spheres tangent to the boundary of X at two or more points. It contains all points of the medial axis but possibly additional points since the spheres are not constrained to bound balls contained in X . Suppose now the boundary of X is defined by the algebraic equation $f(x) = 0$ and 0 is a regular value of f . It follows that the gradient for all points of the boundary is non-zero, $\nabla f(x) \neq 0$. In this case, the symmetry set is the closure of the set of points z for which there exists points x and y that satisfy the following system of algebraic equations:

$$\begin{cases} f(x) = 0, \\ f(y) = 0, \\ (x - z) \times \nabla f(x) = 0, \\ (y - z) \times \nabla f(y) = 0, \\ \|x - z\|^2 = \|y - z\|^2, \\ t\|x - y\|^2 = 1 \end{cases}$$

In the last condition, t is an additional free variable that ensures that x and y are distinct. If 0 is not a regular value of f , we need to add $\nabla f(x) \nabla f(y) s = 1$ as yet another equation, with s a free variable. Finally, the medial axis is obtained by imposing the additional conditions that $\|u - z\|^2 \geq \|x - z\|^2$, for all points u on the boundary, and z be contained in X . Considering u a new free variable, we thus remove points from the solution, namely the points z for which $f(z) < 0$ or for which there exists u with $f(u) = 0$ and $\|u - z\|^2 < \|x - z\|^2$. This new set is still semi-algebraic since it is the difference between two semi-algebraic sets.

Although possible in principle, we are not aware of an implementation that effectively constructs the exact medial axes of general semi-algebraic sets. The most advanced effective implementations are limited to the planar case, to piecewise linear shapes, and to shapes constructed from finitely many balls. Even for shapes bounded by simple curves in the plane, the algebraic difficulties in computing medial axes are significant and satisfactory implementations are rare and far in between. Piecewise linear curves involve the comparison of expressions with two nested square roots [17] and efficient and fully robust implementations are few [36]. Ramamurthy and Farouki tackle the case of algebraic curve segments whose bisectors have rational parametrizations [47]. An exact algorithm for not-necessarily convex polyhedra in $\mathbb{R}^3$ can be found in [24]. For the complement of a union of balls in $\mathbb{R}^k$, the medial axis can be derived from the Apollonius diagram of the corresponding spheres or from convex hulls of finitely many points in $\mathbb{R}^{k+2}$ [11, 14]. Perhaps suprisingly, the medial axis of the union of finitely many balls is simpler than that of the complement. As first described in [5], it is piecewise linear and can be constructed from the Voronoi diagram of a finite set of points. As discussed in

more detail shortly, the cells of dimension less than k in this diagram may be interpreted as the medial axis of a punctured Euclidean space, a case that permits particularly simple exact algorithms.

8 Approximation paradigm for medial axes

The difficulty of computing medial axes exactly motivates a serious look at approximation algorithms. We describe a framework that captures a common line of attack to approximating the medial axis, as sketched in Figure 3. First, we find Y approximating X that belongs to a class of shapes for which the medial axis can be constructed exactly. Second, we construct the medial axis of Y . Third, we prune the medial axis of Y to get a subset $\mathcal{P}[\mathcal{M}[Y]] \subseteq \mathcal{M}[Y]$ that approximates the medial axis of X . The composition of the three steps provides the approximation of the medial axis of X . The most challenging




Fig. 3: An approximation $\mathcal{P}[\mathcal{M}[Y]]$ of the medial axis of a shape X can be found as part of the medial axis of a shape Y approximating X .

step in this paradigm is the extraction of a subset $\mathcal{P}[\mathcal{M}[Y]]$ of $\mathcal{M}[Y]$ that indeed approximates $\mathcal{M}[X]$. Recent mathematical results that rationalize this approach are discussed shortly.

The notion of approximation used in the first step varies between different implementations of the approximation paradigm. It either means that Y is the image of X under a small C^m -perturbation, or that the Hausdorff distance between the complements of X and Y is small, as in [19]. Other notions of approximation are conceivable.

9 Punctured Euclidean spaces

We start by identifying a class of shapes for which the medial axis can be constructed exactly and efficiently. We obtain shapes in this class by puncturing the k -dimension real space at a discrete set of locations. Equivalently, we consider the complement of a discrete set of points P in $\mathbb{R}^k$. The medial axis of this space is the *Voronoi graph* of P , which we define as the union of all cells in the Voronoi diagram of dimension $k - 1$ or less. Algorithms for constructing the Voronoi graph are well-studied in computational geometry

and implementations are available from the geometric software library CGAL [54]. For a set P of n points in $\mathbb{R}^k$, the graph can be constructed in time $O(n^{\lceil k/2 \rceil} + n \log n)$, which is optimal in the worst case because the graph can consist of a constant times $n^{\lceil k/2 \rceil}$ faces. In most practical applications, the number of faces, F , is much less and the output-sensitive algorithm in [18] constructs the graph in $\mathbb{R}^3$ in time $O((n + F) \log^2 F)$. Examples of point sets with provably small Voronoi graphs are so-called κ -light ϵ -samples of compact smooth generic surfaces in $\mathbb{R}^3$, with $F = O(n \log n)$ [8], and κ -light ϵ -samples of polyhedral surfaces in $\mathbb{R}^3$, with $F = O(n)$ [7]. Such samples will be studied in more detail shortly.

Consider a finite point set P whose Hausdorff distance to the boundary of a shape X is less than ϵ and write $\text{Vor}[P]$ for the Voronoi graph of P . Using the semicontinuity of the medial axis expressed in (2), we can show that the subset of $\text{Vor}[P]$ inside X contains an approximation of the medial axis of X . In the approximation paradigm for medial axes, this subset can be interpreted as part of the medial axis of a shape Y close to X . Following [19], we let Y be the parallel body $X^{+\epsilon}$ of X minus the points in P ; see Figure 4. Since the Hausdorff distance between P and the boundary of X is less than ϵ , the same is true for the complements of X and the thus constructed space: $d_H(X^c, Y^c) < \epsilon$. In summary, we have $\mathcal{M}[Y] \cap X = \text{Vor}[P] \cap X$.




Fig. 4: On the left, a shape X and its medial axis. In the middle, a finite set of points P whose Hausdorff distance to the boundary of X is less than ϵ and its Voronoi graph. On the right, $X^{+\epsilon} - P$ and its medial axis.

10 Voronoi graph and medial axis

We now consider results that focus on the detailed relationship between the Voronoi graph of a finite point set and the medial axis of the shape whose boundary the points sample. We need precise notions. A *sample* of the boundary of a shape X is a finite set of points (exactly and not just approximately) on that boundary. An ϵ -*sample* is a sample whose Hausdorff distance to the

boundary of X is less than ϵ . In other words, every point of the boundary is less than distance ϵ away from a point in the ϵ -sample. The ϵ -sample is κ -light if the number of sample points within distance ϵ is never more than κ . The ϵ -sample is *noisy* if points are not necessarily on the boundary but at Hausdorff distance less than ϵ to the boundary.

An early result on the connection between the Voronoi graph and the medial axis is due to Brandt [15]. Given a shape in $\mathbb{R}^2$, he takes an ϵ -sample on the boundary curve and considers the Voronoi edges and vertices that are completely contained in the shape; see Figure 5. He then proves that under some




Fig. 5: In $\mathbb{R}^2$, vertices and edges lying inside a shape and extracted from the Voronoi graph of an ϵ -sample of the boundary approximate the medial axis (courtesy of Attali and Montanvert [10]).

technical conditions on the boundary curve, the portion of the Voronoi graph defined by these edges and vertices approximates the medial axis. Amenta and Bern [1] point out that the direct extension of this result to shapes in $\mathbb{R}^3$ does not hold; see Figure 6. The validity of the extension is spoiled by the existence of slivers in three-dimensional Delaunay triangulations, which occur for ϵ -samples with arbitrarily small $\epsilon > 0$. Roughly, a *sliver* is a tetrahedron whose four vertices are almost cocircular. The location of the Voronoi vertex corresponding to the sliver depends on the four vertices but is generally unrelated to any feature of the surface and does not necessarily lie near the medial axis. As a first step to cope with slivers, Amenta and Bern eliminate all but a few Voronoi vertices they refer to as *poles*. Every sample point p generates a Voronoi polyhedron and the vertices furthest away from p on the two sides of the surface are the *poles* of p . Clearly, there are at most $2n$ poles

for a sample of n points. As proved in [3], for a shape whose boundary is a smooth C^1 -manifold, the poles tend to the medial axis of the shape and its complement as ϵ goes to zero.




Fig. 6: On the left we see a triangulation of the boundary of a shape in $\mathbb{R}^3$. Its vertices determine a Voronoi diagram whose vertices inside the shape are shown in the middle. The subset of poles inside the shape is shown on the right.

To extend the result of Brandt to $\mathbb{R}^3$, we need more than just points (the poles) near the medial axes, we also need to connect them to form a geometric structure approximating the medial axis. In [3], Amenta, Choi and Kolluri use simplices of the (weighted) Delaunay triangulation of the poles. To avoid the construction of this weighted Delaunay triangulation and connect the poles directly inside the Voronoi graph, we need to know about its local distance from the medial axis. Bounds on this distance can be found in [6, 13, 23]. Assuming the boundary of the shape is a smooth C^1 -manifold and using these bounds, among other things, Dey and Zhao [27] give an algorithm that identifies a subgraph of the Voronoi graph that approximates the medial axis for the Hausdorff distance. We note that the above results are limited to smooth surfaces and to samples of points that lie on that surface. The next two sections deal with more general data.

11 Pruning in the presence of noise

Assuming the medial axis of a shape Y approximating X has been constructed, we prune $\mathcal{M}[Y]$ to retrieve an approximation of $\mathcal{M}[X]$. In this paragraph, the terms shape, medial axis, and stable part refer to Y , $\mathcal{M}[Y]$, and $\mathcal{P}[\mathcal{M}[Y]]$, respectively. Pruning methods shorten peripheral branches of the medial axis, trying to capture its stable part. Typically, points on the border are successively removed until a stopping condition is satisfied. This condition may be a threshold on the difference between the initial shape and the shape reconstructed from the simplified medial axis [16, 23, 28, 48, 49], or it may be based on an estimate of the stability of portions of the medial

axis [9, 10, 26, 31, 32, 45]. We present experimental results due to Attali and Montanvert [10] that shed light on the latter approach. To each point $y \in Y$, we associate the distance to Y^c , the complement of Y , and the largest angle formed by points in Y^c closest to y :

$$\rho(y) = d(y, Y^c) \quad \text{and} \quad \theta(y) = \max_{a, b \in \Pi(y)} \angle ayb,$$

where $\Pi(y) = \{x \in Y^c \mid d(y, Y^c) = \|y - x\|\}$. We obtain the *parameter graph* by collecting, for all points y of the medial axis, the points $(\theta(y), \rho(y))$ in the two-dimensional parameter space [9, 10]. Points in this graph lie on




Fig. 7: Medial axis, parameter graph and simplified medial axis obtained by keeping points in the upper right quadrant of the parameter graph (courtesy of Attali and Montanvert [10]).

curves associated with branches of the medial axis, as illustrated in Figure 7. When noise is added to the boundary, new branches appear on its medial axis; see Figure 8. The corresponding effect on the parameter graph is the appearance of a hyperbola-like point cloud located near the coordinate axes. This experimental observation suggests a method for recognizing points on the medial axis that owe their existence to noise in the input data. As a first cut, we remove points y for which $\rho = \rho(y)$ is smaller than a first threshold or $\theta = \theta(y)$ is smaller than a second threshold. In order to refine the method, let us introduce $h(y) = \sup_{a, b \in \Pi(y)} \|a - b\|/2$. The new quantity, $h = h(y)$ is related to the previous ones by

$$\rho = \frac{h}{\sin(\theta/2)}.$$

For each fixed h we get ρ as a function of θ . By plotting a few of these functions, we can experimentally find a value of h for which the graph of this function approximates the hyperbola-like point cloud induced by the noise.




Fig. 8: A small amount of noise added to the boundary of a shape suffices to dramatically change its medial axis. The simplified medial axis is defined by points in the upper right quadrant of the parameter graph (courtesy of Attali and Montanvert [10]).

This suggests that the stable part of the medial axis corresponds to points y for which h exceeds a given threshold. The next section describes a theoretical justification of this experimental finding.

12 Stability of the λ -medial axis

Chazal and Lieutier [19] define the λ -medial axis of a bounded open subset X of $\mathbb{R}^k$ and prove its stability under the Hausdorff distance, for regular values of λ . Remember that this property is not shared by the medial axis transform. To describe their results, let $r(x)$ be the radius of the smallest ball enclosing $\Pi(x)$, the set of points in the complement of X with minimum distance to x . By definition, the λ -medial axis of X is

$$\mathcal{M}_\lambda[X] = \{x \in X \mid r(x) \geq \lambda\}.$$

For $\lambda > 0$, the λ -medial axis is a subset of the medial axis and the Hausdorff distance between the two tends to zero when λ goes to zero. We say that λ is a *regular* value of X if the function that maps $\mu \in \mathbb{R}$ to $\mathcal{M}_\mu[X]$ in $\mathbb{R}^k$ is continuous under the Hausdorff metric at $\mu = \lambda$. In other words, a small modification of a regular value λ implies a small modification of the λ -medial axis. Typical non-regular values are radii of locally largest maximal balls. We are now ready to give a precise statement of the result in [19]: if λ is a regular value of a shape X , the λ -medial axis transform is continuous at X for the Hausdorff distance. In other words, for every $\delta > 0$, there exists $\epsilon > 0$ such that for every open subset Y of $\mathbb{R}^k$,

$$d_H(X^c, Y^c) < \epsilon \implies d_H(\mathcal{M}_\lambda[X], \mathcal{M}_\lambda[Y]) < \delta. \quad (3)$$

Note the similarity with (2), which expresses the same property using one-sided instead of two-sided Hausdorff distance. As part of the approximation paradigm for medial axes, this result sheds new light on the pruning method described above, which is now seen as approximating the λ -medial axis. Furthermore, an approximation of the medial axis can be obtained by forcing λ to decrease as Y gets closer to X . This idea appears in [27] and can also be found in [19]. Specifically, we consider a sequence of shapes Y_ϵ whose Hausdorff distance to X is at most ϵ . Writing D for the diameter of X and introducing $g(\epsilon) = 10\sqrt{3}D^{3/4}\sqrt[4]{\epsilon}$, we get

$$\lim_{\epsilon \rightarrow 0} d_H(\mathcal{M}[X], \mathcal{M}_{g(\epsilon)}[Y_\epsilon]) = 0, \quad (4)$$

[private communication with André Lieutier]. Unlike the medial axis, the λ -medial axis is not necessarily homotopy equivalent to the shape. To shed light




Fig. 9: Left: a shape with weak feature size indicated by the arrow. Right: the λ -medial axis for a value of λ greater than the weak feature size. Each endpoint has two closest points on the boundary, whose distance from each other is 2λ .

on this phenomenon, Chazal and Lieutier [19] call a point $x \in X$ *critical* if the vector field $\nabla \rho$ defined in (1) vanishes at x ; see Figure 1. The *weak feature size* of X is the smallest distance between a critical point and the boundary of X . As proved in [19], the λ -medial axis is homotopy equivalent to X if λ is smaller than the weak feature size.

13 What now?

How do we best harness the power of the new insights, in particular the stability of the λ -medial axis? In this section, we speculate how this stability can be used to obtain improved implementations of the approximation paradigm for medial axes. We also mention some of the open issues that are still obstacles in our quest for a satisfactory solution in the absence of any knowledge on the shape X other than a possibly noisy finite sample of its boundary.

Given a finite point set P , we call the λ -medial axis of the complement the λ -Voronoi graph of P . The λ -complex consists of all simplices in the Delaunay triangulation that can be enclosed by a sphere of radius less than λ . The relation between the two is one of duality and complementarity: a Voronoi cell of dimension less than k belongs to the λ -Voronoi graph iff the dual Delaunay simplex does not belong to the λ -complex. To derive an alternative




Fig. 10: The λ -complex and the λ -Voronoi graph of a noisy sample of a simple closed curve in the plane.

description, let B be the set of open balls with radius λ whose centers are points in P . The λ -complex consists of all Delaunay simplices spanned by points in P whose balls have a non-empty intersection. This is similar to but slightly weaker than the condition for the simplex to belong to the α -shape [31], which requires that the balls and the (k -dimensional) Voronoi cells have a non-empty intersection. Indeed, it is not difficult to prove that for $\lambda = \alpha$, the λ -complex and the α -complex are homotopy equivalent. To construct the λ -Voronoi graph, we simply select the Delaunay simplices that belong to the λ -complex and collect the dual Voronoi cells of the remaining Delaunay simplices as the pieces of the λ -Voronoi graph. An example of this construction is shown in Figure 10. Assume now that P is a noisy sample of the boundary of an open set $X \subseteq \mathbb{R}^k$. If we know that P is an ϵ -sample of that boundary, we may set $\lambda = g(\epsilon)$ and compute the λ -Voronoi graph. The results presented in Sections 10 and 12 assure that as ϵ goes to zero, the λ -Voronoi graph restricted to X is an approximation of the medial axis of X . To recapitulate, we go through the following steps to obtain an approximation of the medial axis:

1. determine how small an ϵ is needed;
2. obtain an ϵ -sample of the boundary of X ;
3. construct the λ -Voronoi graph, with $\lambda = g(\epsilon)$;
4. select the part of the graph inside X .

In practice, we rarely have enough knowledge about X to know what $\epsilon > 0$ is sufficiently small, and even if we knew, we might not have the means to obtain an ϵ -sampling of the boundary. In exceptional cases, the boundary of X is defined mathematically, e.g. as the zero-set of an algebraic function $f : \mathbb{R}^k \rightarrow \mathbb{R}$, and we can determine sufficiently fine ϵ -samples and therefore λ -Voronoi graphs that approximate the medial axis, as in Figure 11. This




Fig. 11: Two λ -medial axes of the same shape, with λ increasing from left to right, constructed as a subset of the λ -Voronoi graph of a sample of the boundary.

approach to medial axes thus suffers from the same difficulties as the α -shape approach to surface reconstruction: it is usually not clear which value of λ (or α) is most appropriate, and in many cases there is no such most appropriate value. This suggests we re-trace some of the developments aimed at fixing this drawback for α -shapes, namely looking at the filtration (nested sequence) of λ -Voronoi graphs and use topological persistence [30] to select and combine pieces of λ -Voronoi graphs for different values of λ in different portions of X .

Acknowledgement. We thank Frédéric Chazal and André Lieutier for communications related to questions discussed in this paper. We thank David Cohen-Steiner for fruitful discussions. The first two authors acknowledge the support of the EU through the Network of Excellence AIM@SHAPE Contract IST 506766. The third author acknowledges the support of the NSF through grant CCR-00-86013.

References

1. N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.
2. N. Amenta, M. Bern, and D. Eppstein. The crust and the beta-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.

3. N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory Appl.*, 19:127–153, 2001.
4. E. V. Anoshkina, A. G. Belyaev, O. G. Okunev, and T. L. Kunii. Ridges and ravines: a singularity approach. *Internat. J. Shape Modeling*, 1:1–11, 1994.
5. D. Attali. *Squelettes et graphes de Voronoi 2-d et 3-d*. PhD thesis, Univ. Joseph Fourier, Grenoble, 1995.
6. D. Attali and J.-D. Boissonnat. Complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete Comput. Geom.*, 30:437–452, 2003.
7. D. Attali and J.-D. Boissonnat. A linear bound on the complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete Comput. Geom.*, 31:369–384, 2004.
8. D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces: the smooth case. In *Proc. 19th Ann. Sympos. Comput. Geom.*, pages 201–210, 2003.
9. D. Attali and J.-O. Lachaud. Delaunay conforming iso-surface. *Skeleton Extraction and Noise Removal*, 19:175–189, 2001.
10. D. Attali and A. Montanvert. Modeling noise for a better simplification of skeletons. In *Proc. Internat. Conf. Image Process.*, volume 3, pages 13–16, 1996.
11. F. Aurenhammer and H. Imai. Geometric relations among Voronoi diagrams. *Geom. Dedicata*, 27:65–75, 1988.
12. H. Blum. A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, editor, *Models for the Perception of Speech and Visual Form*, pages 362–380, Cambridge, MA, 1967. M.I.T. Press.
13. J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Comput. Geom. Theory Appl.*, 19:155–173, 2001.
14. J.-D. Boissonnat and M. Karavelas. On the combinatorial complexity of Euclidean Voronoi cells and convex hulls of d -dimensional spheres. In *Proc. 14th ACM-SIAM Sympos. Discrete Alg.*, pages 305–312, 2003.
15. J. W. Brandt. Convergence and continuity criteria for discrete approximations of the continuous planar skeletons. *CVGIP : Image Understanding*, 59:116–124, 1994.
16. J. W. Brandt and V. R. Algazi. Continuous skeleton computation by Voronoi diagram. *CVGIP : Image Understanding*, 55:329–337, 1992.
17. C. Burnikel. *Exact Computation of Voronoi Diagrams and Line Segment Intersections*. Ph.D thesis, Universität des Saarlandes, March 1996.
18. T. M. Chan, J. Snoeyink, and C. K. Yap. Primal dividing and dual pruning: Output-sensitive construction of 4-d polytopes and 3-d Voronoi diagrams. *Discrete Comput. Geom.*, 18:433–454, 1997.
19. F. Chazal and A. Lieutier. Stability and homotopy of a subset of the medial axis. In *Proc. 9th ACM Sympos. Solid Modeling Appl.*, 2004.
20. F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *J. Control Dyn. Syst.*, 10:149–170, 2004.
21. C. H. Chen, L. F. Pau, and P. S. Wang, editors. *Segmentation tools in Mathematical Morphology*. World Scientific, 1993.
22. H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific J. Math.*, 181:57–88, 1997.
23. S. W. Choi and H.-P. Seidel. Linear one-sided stability of MAT for weakly injective 3D domain. In *Proc. 7th ACM Sympos. Solid Modeling Appl.*, pages 344–355, 2002.

24. T. Culver. *Computing the medial axis of a polyhedron reliably and efficiently*. Depart. comput. sci., Univ. North Carolina, Chapel Hill, North Carolina, 2000.
25. T. K. Dey, J. Giesen, and S. Goswami. Shape segmentation and matching with flow discretization. In F. Dehne et al., editor, *Proc. Workshop Alg. Data Structures*, pages 25–36, 2003.
26. T. K. Dey and W. Zhao. Approximate medial axis as a Voronoi subcomplex. In *Proc. 7th ACM Sympos. Solid Modeling Appl.*, pages 356–366, 2002.
27. T. K. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38:179–200, 2004.
28. A. R. Dill, M. D. Levine, and P. B. Noble. Multiple resolution skeletons. *IEEE Trans. Pattern Anal. Mach. Intell.*, 9:495–504, 1987.
29. H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry — The Goodman-Pollack Festschrift*, pages 379–404. Springer-Verlag, Berlin, 2004.
30. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
31. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics*, 13:43–72, 1994.
32. M. Foskey, M. Lin, and D. Manocha. Efficient computation of a simplified medial axis. In *Proc. 8th ACM Sympos. Solid Modeling Appl.*, pages 96–107, 2003.
33. P. Giblin and B. B. Kimia. A formal classification of 3D medial axis points and their local geometry. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 26:238–251, 2004.
34. J. Giesen and M. John. Surface reconstruction based on a dynamical system. In *Proc. 23rd Ann. Conf. European Association for Computer Graphics (Eurographics), Computer Graphics Forum*, pages 363–371, 2002.
35. R. Haralick and L. Shapiro. *Computer and Robot Vision*, volume 1. Addison-Wesley Publishing Company, 1992.
36. M. Held. VRONI: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments. *Comput. Geom. Theory Appl.*, 18:95–123, 2001.
37. M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, 1988.
38. M. Hisada, A. G. Belyaev, and T. L. Kunii. A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. In *Computer Graphics Forum*, volume 21, pages 689–700, 2002.
39. C. Hoffmann. *Geometric and Solid Modeling*. Morgan-Kaufmann, San Mateo, CA, 1989.
40. B. Jähne. *Digital Image Processing*. Springer, 4th edition, 1997.
41. L. Lam, S.-W. Lee, and C. Y. Suen. Thinning methodologies - a comprehensive survey. *IEEE Trans. on PAMI*, 14(9):869–885, September 1992.
42. J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.
43. A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as its medial axis. In *Proc. 8th ACM Sympos. Solid Modeling Appl.*, pages 65–75. ACM Press, 2003.
44. G. Matheron. Examples of topological properties of skeletons. In J. Serra, editor, *Image Analysis and Mathematical Morphology, Volume 2: Theoretical Advances*, pages 217–238. Academic Press, London, 1988.

45. R. Ogniewicz. A multiscale MAT from Voronoi diagrams: the skeleton-space and its application to shape description and decomposition. In C. Arcelli et al., editors, *Aspects of Visual Form Processing*, pages 430–439. World Scientific, Singapore, 1994.
46. M. A. Price and C. G. Armstrong. Hexahedral mesh generation by medial surface subdivision: Part II. solids with flat and concave edges. *Internat. J. Numerical Methods Engineering*, 1997.
47. R. Ramamurthy and R. T. Farouki. Voronoi diagram and medial axis algorithm for planar domains with curved boundaries, I and II. *J. Comput. Appl. Math.*, 102:119–141 and 253–277, 1999.
48. G. Sanniti di Baja and E. Thiel. A multiresolution shape description algorithm. In D. Chetverikov et al., editor, *Lecture Notes in Computer Science*, volume 719, pages 208–215. Springer-Verlag, 1993.
49. D. Shaked and A. M. Bruckstein. Pruning medial axes. *Comput. Vis. Image Underst.*, 69:156–169, 1998.
50. A. Sheffer, M. Etzion, A. Rappoport, and M. Bercovier. Hexahedral mesh generation using the embedded Voronoi graph. *Engineering Comput.*, 15:248–262, 1999.
51. E. Sherbrooke, N. M. Patrikalakis, and F.-E. Wolter. Differential and topological properties of medial axis transforms. *Graphical Models and Image Processing*, 58:574–592, 1996.
52. M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. PWS Publishing, Pacific Grove, second edition, 1999.
53. F.E. Wolter. Cut locus & medial axis in global shape interrogation & representation. Technical Report Design Laboratory Memorandum 92-2, MIT, 1992.
54. <http://www.cgal.org/>.