Generalizing Medial Axes with Homology Switches

by

Elizabeth Stephenson

May, 2023

A thesis submitted to the
Graduate School
of the
Institute of Science and Technology Austria
in partial fulfillment of the requirements
for the degree of
Master of Science

Committee in charge:

Herbert Edelsbrunner
Chris Wojtan

ISTA | **Institute of Science and Technology Austria**

Abstract

We introduce the notion of a Faustian interchange in a 1-parameter family of smooth functions to generalize the medial axis to critical points of index larger than 0. We construct and implement a general purpose algorithm for approximating such generalized medial axes.

Acknowledgements

Elizabeth would like to thank her supervisor Herbert for his mathematical wisdom. She would also like to thank her friends and family for their support, especially Martin and Ondra for helping with code queries, Teresa for helping with homology problems, Mathijs for medial axis insight, her parents for their unwavering assurance that a doughnut is, in fact, a coffee cup, and her cat Velcro for attempting to eat the thesis multiple times.

Contents

Chapter 1

Introduction

Given an embedded curve in the Euclidean plane, the *medial axis* consists of the points $t \in \mathbb{R}^2$ that have at least two distinct closest points on the curve [2]. Equivalently, the smallest circle centered at t that has a non-empty intersection with the curve touches the curve in at least two points; see Figure 1.1. More precisely, the medial axis is the *closure* of these points, which effectively adds the centers of osculating circles that meet the curve in only one point. A related concept is the *symmetry set*, which is the closure of the set of centers of circles that touch the curve in at least two distinct points [3]. This condition is weaker, which implies that the symmetry set contains the medial axis. By requiring that the points at which the circle touches the curve minimize the distance from the center, the medial axis selects a more manageable subset that captures important aspects of the shape. Indeed, the medial axis enjoys applications in many fields, while the symmetry set is too unwieldy to be of much practical relevance.
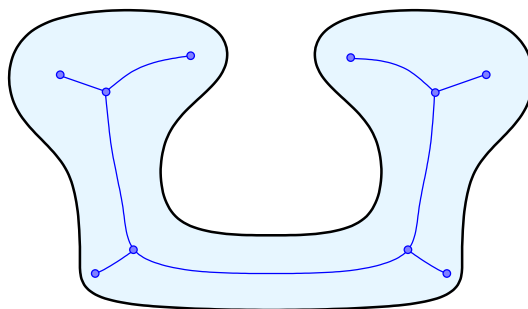
Figure 1.1: The medial axis of a closed curve. We only show the component inside the region enclosed by the curve.

In both definitions, we can interpret the required double-tangency as an interchange of critical points in the parametrized family of squared distance functions

on the curve. Each function in this family corresponds to a point $t \in \mathbb{R}^2$ and maps every point of the curve to the Euclidean distance from t. In the case of the medial axis, this interchange is between the two minima with smallest value, while in the case of the symmetry set, these can be any two critical points, as long as they share the same value.

We suggest that the closest two minima owe their significance to the topological roles they play in the evolution of the sublevel set when we increase the threshold: the first minimum pairs up with the empty set to end the reign of nothingness, while the second minimum starts the first gap between two sprouting empires. By focusing on the simultaneous occurrence of the first two minima, the medial axis traces the fine line along which the competition between who ends and who starts is yet undecided. In the reduced homology interpretation of the sublevel set filtration, this balancing act is between the death of the (-1)-dimensional class and the birth of the first 0-dimensional class.

The main contribution of this thesis is the generalization of the topologically motivated selection criterion to classes of dimension higher than -1 and 0.

Prior work. In the late 1960s, Blum [2] revolutionized the automated classification of shapes arising from biology by suggesting the medial axis as a tool fit to describe amorphous blobs, as opposed to the rectilinear methods used at the time. Since then it has been used as a tool in shape classification, animation, computer graphics, and other fields that benefit from skeletonization of an unwieldy object of study [1, 16, 17]. The distance from a submanifold in Euclidean space is studied in the smooth case by Mather [12]. It has also proved useful in the theoretical setting, often paired up with notions of curvature and local feature size. In fact, the medial axis was introduced in this setting by Federer [7, Def. 4.1 on page 432] almost a decade prior to Blum's paper, but due to the dense nature of Federer's writing and the unassuming name he gave it, Federer's contribution is frequently overlooked, and Blum given the sole credit for its introduction.

Outline. Chapter 2 provides the necessary technical background. Chapter 3 generalizes the medial axis. Chapter 4 explains how the generalized medial axis can be computed. Chapter 5 exhibits examples of the generalized medial axis. Chapter 6 concludes the paper.

Chapter 2

Background

We need background from Morse theory, Cerf theory, homology, and persistent homology. As sources for more extensive background, we recommend [13] for Morse theory, [9, 14, 15] for singularity theory including Cerf theory, [10] for algebraic topology including homology, and [6] for persistent homology.

2.1 Morse Functions

Let $g\colon M \to \mathbb{R}$ be a smooth function on a k-dimensional manifold, so derivatives of all orders exist, although we will need only those of first and second order. A *critical point* of g is a point $a \in M$ at which the first derivative vanishes. All other points of M are *regular* or *non-critical*. A *critical value* of g is the value of a critical point. All other values are *regular* or *non-critical*, but note that it is quite possible that a non-critical point has a critical value, which it necessarily shares with a critical point. A critical point, a, is *non-degenerate* if the Hessian of g at a is invertible. Equivalently, a has a parametrized neighborhood such that

$$g(x) = g(a) - x_1^2 - \ldots - x_p^2 + x_{p+1}^2 + \ldots + x_k^2, \tag{2.1}$$

for every $x = (x_1, x_2, \ldots, x_k)$ in this neighborhood. The number of minus signs is called the *index* of the critical point, denoted $\text{index}(a) = p$. The critical point is a *minimum* if $p = 0$, a *maximum* if $p = k$, and a *saddle* if $0 < p < k$. Non-degenerate critical points are necessarily isolated. If M is compact, this implies that g has only finitely many non-degenerate critical points.

Definition 1 (Morse function): *A smooth function on a manifold, $g\colon M \to \mathbb{R}$, is*

Morse *if all its critical points are non-degenerate, and no two of them have the same value.*

2.2 Interchanges and Cancellations

We are interested in parametrized families of smooth functions, and it is too restrictive to require that all their functions be Morse. Letting $G\colon M \times \mathbb{R} \to \mathbb{R}$ be a smooth map, we call the $g_t\colon M \to \mathbb{R}$ defined by $g_t(x) = G(x,t)$ a *1-parameter family* of smooth functions.

Definition 2 (Cerf family of smooth functions): *The family of functions g_t, with $t \in \mathbb{R}$, is Cerf if g_t is Morse except for a finite number of values of t, and for every such exceptional value, $s \in \mathbb{R}$, there is only one violation of g_s being Morse, namely*

I *two critical points of g_s share the same value, or*

II *one critical point, $a \in \mathbb{R}$, of g_s is degenerate, and (a, s) has a neighborhood such that*

$$g_t(x) = g_s(a) + [x_1^3 \pm t x_1] - x_2^2 - \ldots - x_p^2 + x_{p+1}^2 + \ldots + x_k^2, \qquad (2.2)$$

for every $(x = (x_1, x_2, \ldots, x_k), t)$ in this neighborhood.

We call a violation of type I an *interchange*, and we call a violation of type II a *cancellation* if the sign of the linear term in (2.2) is positive, and an *anti-cancellation* if the sign is negative. In a cancellation, two critical points get destroyed when t passes from the negative to the positive side of the real line, and with the notation in (2.2), the indices of these critical points are $p-1$ and p. In an anti-cancellation, two critical points of index $p-1$ and p get created as t passes from the negative to the positive side of real line. By the *essential property* in Cerf theory, any 1-parameter family that starts with a Morse function and ends with a Morse function, can be approximated by a Cerf family of smooth functions.

2.3 Stratified Control Spaces

We borrow the term control space from catastrophe theory [8, 14, 15], where the control space is used to identify sudden changes caused by smooth alterations to a system. An example is the phase transition diagram between solid, liquid, and gaseous states of water. Next, we take some definitions taken from [14]. We consider a *smooth family* of functions $f_u(x)$ (also denoted $F(x, u)$), where $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^a$.

We require that F is smooth in both u and x (so F is smooth in $n + a$ variables). Then we can define the *catastrophe set* of F to be

$$C_F = \left\{ (x, u) \in \mathbb{R}^n \times \mathbb{R}^a \mid \frac{\partial f_u}{\partial x_1}(x) = \cdots = \frac{\partial f_u}{\partial x_n}(x) = 0 \right\}.$$

Intuitively, the catastrophe set is the set of points $(x, u) \in \mathbb{R}^{n+a}$ such that f_u has a critical point at x. An important subset of the catastrophe set is the *singular set*, denoted Σ_F, which is the set of points in C_F where the critical point is *degenerate*, meaning the determinant of the Hessian at x is zero. Next, we define the *bifurcation set* Δ_F. If we let π_F be the map projecting C_F onto the parameter space $\mathbb{R}^a$, then

$$\Delta_F = \pi_F(\Sigma_F) = \{ u \in \mathbb{R}^a \mid \exists x \in \mathbb{R}^n, (x, u) \in \Sigma_F \}.$$

In this paper, we will focus on a closely related but slightly larger set: the set of all control points at which the function fails to be Morse.

Definition 3 (Stratification): *A stratification of the control space is a partition* $\mathbb{R}^d = X_0 \sqcup X_1 \sqcup \ldots \sqcup X_d$ *such that*

(i) X_q *is a smoothly embedded q-manifold in $\mathbb{R}^d$, called the q-stratum, for $0 \le q \le d$;*

(ii) $X_0 \sqcup X_1 \sqcup \ldots \sqcup X_q$ *is closed, for $0 \le q \le d$, and equal to $\mathbb{R}^d$, for $q = d$.*

We note that individual strata are not necessarily connected. We require that the family of smooth functions admits a stratification of the control space such that the drawing of interest is the union of the strata of dimension at most $d - 1$. Hence g_t is Morse whenever t is a point of the d-stratum. If this is the case, we call the g_t a *d-parameter Cerf family* of smooth functions.

A stratification that we will make use of in this paper is the symmetry set [3]. For a smooth manifold M in $\mathbb{R}^n$, the *symmetry set* $S(M)$ is defined to be the closure of the set of points $u \in \mathbb{R}^n$ which are centers of spheres tangent to M at two or more distinct points. The connection with singularity theory and control spaces is that the symmetry set $S(M)$ can be described alternatively as the levels bifurcation set of the family of distance-squared functions on M. In our paper, we will look at a subset of the symmetry set. As [3] notes, the canonical stratification of the bifurcation set in its entirety is a "hopelessly complicated object".

2.4   Homology

We use the homology of the sublevel sets to turn a continuous map into an algebraic object, which can be studied with discrete methods. Assuming a smooth function $g\colon M \to \mathbb{R}$, where M is a manifold in $\mathbb{R}^n$, we write $M_r = g^{-1}(-\infty, r]$ for the *sublevel set* at $r \in \mathbb{R}$. Homology is an algebraic formalism that captures the p-dimensional connectivity of M_r in terms of an abelian group, referred to as the *p-th homology group*, denoted $H_p(M_r)$. There is some variability depending on the coefficients used to add classes in this group, and we keep things simple by adopting modulo-2 arithmetic, for which $H_p(M_r)$ is a vector space. The rank of that vector space, denoted $\beta_p(M_r) = \operatorname{rank} H_p(M_r)$ and referred to as the *p-th Betti number*, is interpreted as the number of *independent* and *non-bounding* p-cycles in M_r. The algebraic details that explain what exactly these cycles are can be found in standard topology texts, including [10].

There are different theories that all lead to the same numbers, two of which are used in this paper. In *singular homology*, the p-cycles are sums of maps of the standard p-simplex to M_r in such a way that their boundaries cancel. This is mathematically convenient since it does not require any additional structure on top of the sublevel set. When we compute Betti numbers, we use *simplicial homology*, which first triangulates M_r and second expresses the vector spaces in terms of the boundary matrices that connect the simplices in adjacent dimensions. More about this in Section 4, where we talk about algorithms.

As a substitute for the algebraic details, we illustrate the Betti numbers with an example. Let M be the torus with a nose in Figure 2.1, and let $g\colon M \to \mathbb{R}$ be the height function. It has six critical points, which are all non-degenerate: two minima, u_1 and u_2, three saddles, u_3, u_4, and u_5, and one maximum, u_6. Choose interleaved non-critical values, $r_0 < g(u_1) < r_1 < \ldots < g(u_6) < r_6$, and simplify notation by writing $M_i = g^{-1}(-\infty, r_i]$. Since M is 2-dimensional, all Betti numbers in dimension beyond 2 vanish. We use reduced homology, so 0-cycles are gaps between components rather than components themselves, and the empty set is a (-1)-cycle. In sequence, M_0 is empty, M_1 is a disk, M_2 is two disks, M_3 is a cylinder and a disk, M_4 is a cylinder, M_5 is a torus with a cap removed, and M_6 is a torus. The non-zero Betti numbers are therefore, $\beta_{-1}(M_0) = 1$, $\beta_0(M_2) = 1$, $\beta_0(M_3) = 1$, $\beta_1(M_3) = 1$, $\beta_1(M_4) = 1$, $\beta_1(M_5) = 2$, $\beta_1(M_6) = 2$, $\beta_2(M_6) = 1$. All other Betti numbers vanish, and in particular all Betti numbers of M_1 vanish since we work with reduced homology.
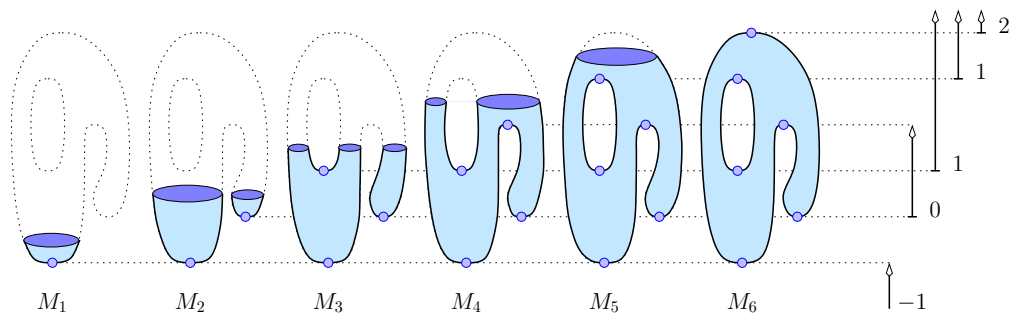
Figure 2.1: The six non-empty sublevel sets at regular values interleaved between the critical values of the height function on the torus-with-a-nose. On the *right*, we see the barcode that shows for which value a gap, loop, or closed surface exists in the sublevel set.

2.5   Persistent Homology

Classical homology is useful for studying topological features in data, but has the disadvantage of lacking the discernment to pick the correct scale to extract features from. Persistent homology remedies this shortcoming by studying nested sequences of complexes, instead of a fixed one, thus extracting features across many scales.

We return to the example of the torus, M, in Figure 2.1. The height function, $g \colon M \to \mathbb{R}$, yields inclusions $M_i \subseteq M_j$ for $i \leq j$, $i, j \in \mathbb{N}$, some of which are represented in the figure. This sequence of sublevel sets is called a *filtration*, and it naturally gives rise to a sequence of homology groups and linear maps induced by the inclusions between the sublevel sets. Rather than considering homology groups for each dimension individually, we write $H(M_i)$ for the direct sum over all dimensions p. This allows us to write the *persistence module* of the filtration compactly in a single line:

$$\ldots \to H(M_{i-1}) \to H(M_i) \to \ldots \to H(M_{j-1}) \to H(M_j) \to \ldots. \qquad (2.3)$$

Composing the maps between consecutive groups, we get a map between any two groups in the module. We say a homology class $\alpha \in H(M_i)$ is *born* at M_i if it is not in the image of the map from $H(M_{i-1})$ to $H(M_i)$. If α is born at M_i, it *dies entering* M_j if the image of the map from $H(M_{i-1})$ to $H(M_{j-1})$ does not contain the image of α, but the image of the map from $H(M_{i-1})$ to $H(M_j)$ does. The *persistence* of α is the difference between the function values at its birth and its death. If g is a Morse function on a manifold, then precisely one Betti number changes when the threshold passes a crtical value. If the index of the corresponding

critical point is p, then either a p-dimensional class is born, so β_p increases by one, or a $(p-1)$-dimensional class dies, so β_{p-1} decreases by one.

Referring to the right half of Figure 2.1, we use a *persistence barcode* to encode the birth-death information of all of the p-dimensional homology classes of M arising from the sublevel set filtration induced by the height function. In this particular barcode, we see a (-1)-dimensional class which lasts from the beginning of time until the first component appears, a 0-dimensional class which lasts from the arrival of the second compoment to its merging with the first, two 1-dimenional classes, which correspond to the two loops of the torus, and a 2-dimensional class representing the closed surface. Note that in this particular example, the 1- and 2-dimensional classes never die.

2.6   Parametrized Persistence

The concept of parametrized functions and, correspondingly, of parametrized persistence modules, plays an important role in this paper. Take for example the torus-with-a-nose in Figure 2.1, and instead of the height function in the vertical direction, consider the height functions in all possible directions. More concretely, let M be the torus in an arbitrary but fixed position in $\mathbb{R}^3$, let $u \in \mathbb{S}^2$ be a direction, and let $g_u \colon M \to \mathbb{R}$ defined by $g_u(x) = \langle x, u \rangle$ be the height function in the direction u. We thus have a 2-parameter family of functions, and similarly a 2-parameter family of persistence modules.

A motivating factor in the construction of parametrized persistence modules is the Stability Theorem of persistent homology originally proved in [4]. It asserts that similar functions have similar barcodes. More precisely, the *bottleneck distance* between the barcodes of functions $g_u, g_v \colon M \to \mathbb{R}$ is bounded from above by the infinity norm of their difference:

$$W_\infty(\mathrm{Bar}(g_u), \mathrm{Bar}(g_v)) \leq \|g_u - g_v\|_\infty, \tag{2.4}$$

in which $\mathrm{Bar}(g)$ is the barcode of g, a matching between the two barcodes is quantified by the supremum of the max-distances between matched bars, and W_∞ is the infimum over all possible matchings, in which we allow the introduction of zero-length bars, which we preferably match with short bars in the barcodes. Intuitively, the bottleneck distance describes the worst disparity between the best matching of points in persistence diagrams: the worst disparity is the "bottleneck" preventing a

smaller distance. See [6, Chapter VIII] for details.

For example, the difference between the height functions along directions $u, v \in \mathbb{S}^2$ has infinity norm at most $\|M\|(1 - \langle u, v \rangle)$, in which $\|M\|$ is the maximum Euclidean distance of point in M from the origin. It follows that various aspects of the height function, as measured by persistence, vary continuously with the direction. One example is the length of the longest bar in the barcode, or really the length of any bar in the barcode.

This gives rise to the concept of *vineyard*, introduced in [5]. It formalizes the idea that a feature of g_u is still recognizable in g_v, provided the two functions u and v are not too far apart. Features are bars in the barcode, and the association is a matching between the bars of g_u and of g_v. As it turns out, the situation is sensitive to numerical error, so that computing the matching from the two barcodes is not practical. Instead, [5] proposes an algorithm that traces the features (bars) while continuously deforming g_u into g_v. This is what we call the *vineyard algorithm*. We will use it in Chapter 4 to trace the features we use to define the generalized medial axis. Further details can be found there.

Chapter 3

Generalized Medial Axes

In this section, we introduce the generalization of the classic medial axis, which we refer to as the 0-*th medial axis*. For technical convenience, we limit ourselves to smooth functions on submanifolds of Euclidean space. The ideas extend beyond this setting, but we prefer to avoid the technical complications that come with more general settings and distract from the main construction.

3.1 Pointed Squared Distance Functions

In this paper, a *submanifold* of $\mathbb{R}^d$ is a smoothly embedded manifold of dimension $k < d$ in d-dimensional Euclidean space. We refer to it as a *curve*, if $k = 1$, and a *surface*, if $k = 2$.

Definition 4 (Squared Distance Functions): *Given a submanifold $M \subseteq \mathbb{R}^d$ and a point $t \in \mathbb{R}^d$, the* (pointed) distance function *at t is $f_t \colon M \to \mathbb{R}$ defined by $f_t(x) = \|x - t\|^2$. We write $F \colon M \times \mathbb{R}^d \to \mathbb{R}$ defined by $F(x, t) = f_t(x)$ for the corresponding d-parameter family.*

When we think of F as a family, we take $t \in \mathbb{R}^d$ as the parameter that distinguishes different members. Alternatively, we can take $x \in M$ as the parameter, but the function x defines records the distance to a single point, namely x, which is not very interesting.

Lemma 5 (Smoothness of Squared Distance Functions): *Let M be a submanifold of $\mathbb{R}^d$. Then $F \colon M \times \mathbb{R}^d \to \mathbb{R}$ is smooth. It follows that for every $t \in \mathbb{R}^d \setminus M$, the pointed squared distance function, $f_t \colon M \to \mathbb{R}$, is smooth.*

The derivative of F decomposes into the derivative of the embedding of M at x and the derivative of $\|x - t\|^2$ at t. To see that both parts are smooth, it is helpful to consider an example (pictured in Figure 3.1). Let M be a closed curve in

10

$\mathbb{R}^2$, and let x be a point on M and t a point in $\mathbb{R}^2 \backslash M$. Note that x has only one degree of freedom, while t has two. The pointed squared distance function from t to x is a paraboloid, which is indeed a smooth surface. Moving t around in $\mathbb{R}^2 \backslash M$ just explores the surface of the paraboloid; in particular the line between t and x is projected to a parabola. The distance relationship is symmetric $(d(x,t) = d(t,x))$, so the reverse direction is also smooth. Note that, had we chosen to use only the pointed distance function instead of squaring it, we would have a hyperboloid instead of a paraboloid, and in particular it would be pointy at $x = t$, and therefore not everywhere smooth. We avoid this annoyance by squaring the distance. Stepping back from our example, we have the following proof.

Proof. We let M be a submanifold of $\mathbb{R}^d$ and observe that the pointed squared distance function for $x \in M$ and t in $\mathbb{R}^d \backslash M$ is still symmetric with respect to t and still a paraboloid centered at x, and is therefore smooth. The other part of the decomposition, the embedding of M at x, is also differentiable because $\mathbb{R}^d$ is a smooth manifold and M a submanifold of it. $\qquad\square$
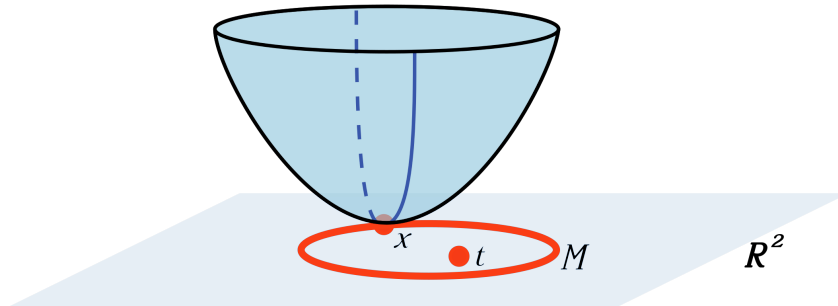
Figure 3.1: An illustration of the pointed squared distance function. As in the above example, we have M a closed loop in $\mathbb{R}^2$, $x \in M$, and t a point in $\mathbb{R}^2 \backslash M$. The squared distance function from t to x forms a paraboloid at x.
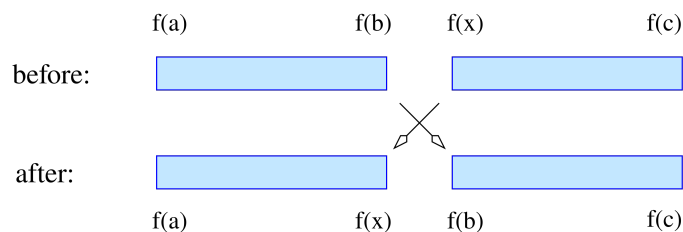
3.2   Faustian Interchanges

We are interested in a special type of interchange, namely one that alters the pairing in the persistence module. Such an interchange affects two critical points, b and x, that satisfy the following four conditions:

1. $\mathrm{index}(b) = \mathrm{index}(x)$, and we write p for this common index;
2. $f(b) = f(x) < f(c)$ for all other index-p critical points c of f;
3. before the interchange, b gives death and x gives birth;

4. after the interchange, x gives death and b gives birth.

Of course, strictly before and strictly after the interchange, we are talking about different functions in the 1-parameter family, and different critical points altogether. We can however make the change at the time of the interchange, preparing the function for the time after the interchange. The time-lag between before and after the interchange is therefore zero, so it makes sense to talk about the same critical point before and after. We call an interchange that satisfied Conditions 1, 2, 3, 4 an *index-p Faustian interchange*. Drawing the function values from left to right, we can visualize the effect of the Faustian interchange on the pairing in the persistence module as illustrated in Figure 3.2.

Figure 3.2: Effect of a Faustian interchange on the pairing of critical points in the persistence module. The deal between b and x is that b may give birth in exchange for allowing x the pleasure to give death.

Definition 6 (Faustian Switch): *We call $s \in \mathbb{R}^d$ a Faustian p-switch if there is a smooth curve, $\gamma \colon \mathbb{R} \to \mathbb{R}^d$, with $s = \gamma(0)$, such that the family of functions $f_t \colon M \to \mathbb{R}$, with $t = \gamma(u)$, is Cerf and the violation of f_s to being Morse is an index-p Faustian interchange.*

Recall that we use reduced homology to construct the persistence module of f_t. At a points $s \in \mathbb{R}^d$ that has two closest points, $b, x \in M$, we can find a smooth curve $\gamma \colon \mathbb{R} \to \mathbb{R}^d$, with $\gamma(0) = s$, that defines a Cerf family of smooth functions such that b and x define an index-0 Faustian interchange at s. Hence, such a point s is a Faustian 0-switch.

3.3 Main Definition

We are ready to present the main new concept in this paper. Recall that a Faustian interchange in a 1-parameter Cerf family of smooth functions is an interchange that has a particular effect on the associated persistence module, and that a Faustian switch is the point in $\mathbb{R}^d$ at which the Faustian interchange occurs.

Definition 7 (Generalized Medial Axis): *Let M be a submanifold of $\mathbb{R}^d$, and let the $f_t\colon M \to \mathbb{R}$, with $t \in \mathbb{R}^d$, be its family of squared distance functions. The (p)-th medial axis of M is the closure of the Faustian p-switches in $\mathbb{R}^d$, denoted $\mathcal{A}_p(M)$.* Assuming $k < d$ is the dimension of M, the p-th medial axis is defined for $0 \leq p \leq k$. For $p = 0$, we get $\mathcal{A}_0(M)$, which agrees with the medial axis as introduced in [2]. Indeed, we have already observed that the points with two closest points in M are Faustian 0-switches. The 0-th medial axis is the closure of these points.

3.4   Geometric Interpretation

Since the generalized medial axes are defined in terms of the squared Euclidean distance from a submanifold, $M \subseteq \mathbb{R}^d$, there is a direct geometric interpretation of the concept. Write $S = S(r, t)$ for the $(d-1)$-sphere with center $t \in \mathbb{R}^d$ and radius $r > 0$. We say this sphere *touches* M in $x \in M$ if $f_t(x) = \|x - t\|^2 = r^2$ and x is a critical point of f_t. Clearly $x \in S \cap M$, and if x is a minimum, then there is a neighborhood of x in M such that $\|y - t\|^2 > r^2$ for all points $y \neq x$ in this neighborhood.

Letting $B(x, \varepsilon)$ be the closed ball with center $x \in \mathbb{R}^d$ and sufficiently small radius $\varepsilon > 0$, $N(x, \varepsilon) = M \cap B(x, \varepsilon)$ is such a neighborhood. Write $\partial N = \partial N(x, \varepsilon)$ for the boundary of this neighborhood, and define

$$\partial N_r = \{y \in \partial N \mid \|y - t\|^2 \leq r^2\}. \tag{3.1}$$

If x is a minimum, then $\partial N_r = \emptyset$, but for critical points of index $p > 0$, the situation is more interesting, namely ∂N_r has the homotopy type of $\mathbb{S}^{p-1}$. If two index-p critical points perform a Faustian interchange, then this local picture is complemented by a global, topological picture.

To describe the global picture and corresponding geometric interpretation, we first let M be a submanifold of $\mathbb{R}^d$. Then $t \in \mathbb{R}^d$ is a Faustian p-switch iff there is a closed ball, $B = B(r, t)$, such that

- its bounding sphere, ∂B, touches M in (at least) two points, b and x;

- $M \cap B$ contains a family of non-bounding p-cycles, all of which pass through b and x;

- r is the smallest radius for which a ball centered at t contains a non-bounding p-cycle.

We walk the reader through these observations. The first follows from the definition of a Faustian p-switch. The second and third are better illustrated with the example pictured in Figure 3.3.
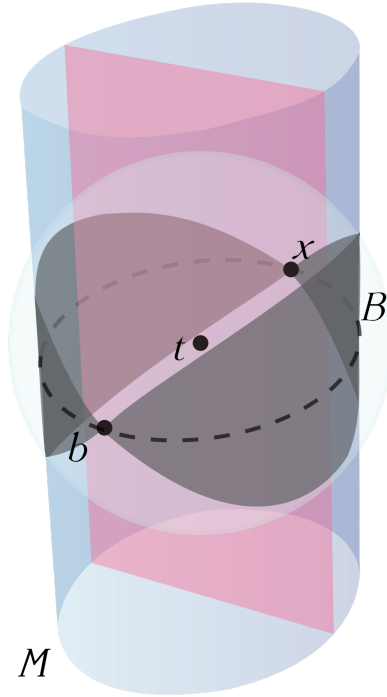
Figure 3.3: Here we illustrate all three points in the global situation described above. M in this case is the blue squished cylinder, and t is a point in the pink bisecting plane of the cylinder. The pink plane is not meant to be part of M, but is drawn to make identification of t's location easier. B is a white ball centered at t and b and x are the points whose inclusion in $M \cap B$ (drawn in grey) finally allows for a 1-cycle to form in $M \cap B$ which wraps around the cylinder (in other words, the radius of B is the smallest radius which admits a non-bounding 1-cycle). We illustrate this cycle as an ellipse with endpoints b and x, and observe that there is room in $M \cap B$ for a whole family of non-bounding 1-cycles passing through b and x.

Some comments. If M is a one-dimensional manifold, such as a closed curve, we can already begin to see that ∂N_r has the homotopy type of $\mathbb{S}^{p-1}$. Take again the example of an ellipse. If x is outside of the ellipse, we first encounter M with $B(x, \epsilon)$ at an index-0 critical point, namely, a minimum. As we continue to grow r, B increases in size, and so does its intersection N with M. Eventually, N is so large that it reaches the top of the ellipse, closing the growing curve and forming a 1-cycle. This corresponds with a 2-index critical point, the maximum. These examples are clear, but not very interesting. The more interesting examples occur

in higher dimensions, where we can have more types of critical points, namely, more interesting saddles. There, the generalized medial axis really shines.

Chapter 4

Computation

The construction of the medial axis is numerically demanding and generally unstable, and so is the construction of the generalized medial axis. We therefore aim for an approximation algorithm sufficient to showcase the concept while leaving room for improvement. It combines elements of image processing with the vineyard algorithm reviewed in the background section. The goal for implementation is to generate points of the p-th medial axis by comparing the changes in the persistence pairing that occur when the object of study is filtered by the squared radius from a point as the point moves through space. The collection of points we add to approximate the p-th medial axis are those points at which Faustian p-switches occur. More concretely, we take a simplicial complex as input, sample a grid across it, find the filtration of the complex arising from the squared distance to the complex from each grid point, and compare changes in the persistence pairing that arise when travelling from a grid point to its neighbor. If a Faustian p-switch occurs when travelling from a grid point to its neighbor, we might think to include their midpoint as a naïve point on the p-th medial axis. This is going in the right direction, but is not sufficient, as this does not guarantee us to observe all of the switches in the ordering one by one. We are quite likely to miss some important transitions, unless we get extremely lucky with our grid sampling. Therefore we invoke the power of the vineyard algorithm, using it to traverse the distance between grid points, guaranteeing that we capture every change in the ordering between grid neighbors.

4.1 The Setting

As before, let M be a k-dimensional submanifold of $\mathbb{R}^d$. The computations are done for a finite sample, $A \subseteq M$, and a simplicial complex, K, with vertex set A that

triangulates M. Recall that $f_t\colon M \to \mathbb{R}$ maps every point $x \in M$ to its squared Euclidean distance from $t \in \mathbb{R}^d$. Let $\bar{f}_t\colon |K| \to \mathbb{R}$ be the piecewise linear extension of f_t evaluated at the points of A:

$$\bar{f}_t(x) = \lambda_0 f_t a_0) + \lambda_1 f_t(a_1) + \ldots + \lambda_p f_t(a_p), \tag{4.1}$$

in which the a_i are the vertices of a p-simplex in K, x is a point of this p-simplex, and the λ_i are the non-negative parameters that satisfy $\lambda_0 + \lambda_1 + \ldots + \lambda_p = 1$ and $\lambda_0 a_0 + \lambda_1 a_1 + \ldots + \lambda_p a_p = x$.

Suppose A and K are such that drawing straight edges and higher-dimensional simplices connecting the points in A produces no spurious intersections. This is not necessary for our purposes but reasonable to assume since spurious intersections between simplices are indications of insufficiently dense sampling. Note however that even in this reasonable case, $\bar{f}_t$ is not the pointed squared distance function of $|K|$ at $t \in \mathbb{R}^d$. Nevertheless, $\bar{f}_t$ is a reasonable approximation of f_t provided A is a sufficiently dense sample of M.

4.2   1-parameter Vineyard

Suppose $t_0 \neq t_1$ are points in $\mathbb{R}^d$. Write $\bar{f}_0, \bar{f}_1\colon |K| \to \mathbb{R}$ for the piecewise linear functions that approximate the pointed squared distance functions at t_0 and t_1. The straight-line homotopy from $\bar{f}_0$ to $\bar{f}_1$ consists of the maps $\bar{f}_\lambda\colon |K| \to \mathbb{R}$ defined by

$$\bar{f}_\lambda(x) = (1 - \lambda)\bar{f}_0(x) + \lambda \bar{f}_1(x), \tag{4.2}$$

for $0 \leq \lambda \leq 1$. We review the vineyard algorithm for the straight-line homotopy and refer to [5] for details. It begins with the lower star filtration of simplicial complex K defined by $\bar{f}_0$. Specifically, we map each simplex, $\sigma \in K$, to the maximum value of its vertices, and we list the simplices in the order of non-decreasing value, while ordering lower- before higher-dimensional simplices if they share the same value. Write $\sigma_1, \sigma_2, \ldots, \sigma_n$ for the resulting sequence, and let $K_i = \{\sigma_1, \sigma_2, \ldots, \sigma_i\}$ be the i-th complex in this filtration. The *ordered boundary matrix* is

$$\partial[i, j] = \begin{cases} 1 & \text{if } \sigma_i \subseteq \sigma_j \text{ and } \dim \sigma_i = \dim \sigma_j - 1, \\ 0 & \text{otherwise.} \end{cases} \tag{4.3}$$

Initialize $R = \partial$ and U to the n-by-n identity matrix, so that $RU = \partial$. To get the algorithm started, we use the classic left-to-right column operations to reduce R. We apply each operation also to U in order to maintain the relation $RU = \partial$.

The difference between $\bar{f}_0$ and $\bar{f}_1$ is in the sequence of simplices, which determines the ordering of rows and columns in ∂. We transform the sequence defined by $\bar{f}_0$ one transposition at a time to the sequence defined by $\bar{f}_1$. The ordering of the transformations follows the straight-line homotopy, which prescribes the value of every vertex at any moment of time, from $\lambda = 0$ to $\lambda = 1$. Each transposition amounts to a sequence of matrix operations, which are detailed in [5]. After performing all transpositions, we arrive at the reduced matrix relation, $RU = \partial$, but now for the ordering of the simplices defined by $\bar{f}_1$.

4.3   Detecting Interchanges and (Anti-)Cancellations

An elementary step is the transposition of two consecutive simplices in the ordering. We follow the description in [5] but focus on the actions rather than their correctness. To transpose simplices σ_i and σ_{i+1}, we exchange the corresponding two rows and columns, both in R and in U, which we write by multiplying with the permutation matrix:

$$RU = \partial \quad \longrightarrow \quad (PRP)(PUP) = P\partial P, \tag{4.4}$$

because PP is the identity matrix. The trouble is that in some cases, PRP is not reduced and PUP is not upper triangular. The vineyard algorithm fixes this with a constant number of row and column operations. We write $\mathrm{low}(j)$ for the largest row index so that $R[\mathrm{low}(j), j] = 1$.

Case 1: σ_i and σ_{i+1} both give birth. Let $U[i, i+1] = 0$, just in case.

    Case 1.1: there are columns k and ℓ with $\mathrm{low}(k) = i$, $\mathrm{low}(\ell) = i+1$, and $R[i, \ell] = 1$.

        Case 1.1.1: $k < \ell$. Add column k of PRP to column ℓ; add row ℓ of PUP to row k.

        Case 1.1.2: $\ell < k$. Add column ℓ of PRP to column k; add row k of PUP to row ℓ. We witness a change in the pairing but not of Faustian type.

    Case 1.2: such columns k and ℓ do not exist. Done.

Case 2: σ_i and σ_{i+1} both give death.

> **Case 2.1:** $U[i, i+1] = 1$. Add row $i+1$ of U to row i; add column i of R to column $i+1$.
>
> > **Case 2.1.1:** $\mathrm{low}(i) < \mathrm{low}(i+1)$. Done.
> >
> > **Case 2.1.2:** $\mathrm{low}(i+1) < \mathrm{low}(i)$. Add column i of PRP to column $i+1$; add row $i+1$ of PUP to row i. We witness a non-Faustian type change of the pairing.
>
> **Case 2.2:** $U[i, i+1] = 0$. Done.

Case 3: σ_i gives death and σ_{i+1} gives birth.

> **Case 3.1:** $U[i, i+1] = 1$. Add row $i+1$ of U to row i; add column i of R to column $i+1$. Furthermore, add column i of PRP to column $i+1$; add row $i+1$ of PUP to row i. We witness a Faustian type change in the pairing.
>
> **Case 3.2:** $U[i, i+1] = 0$.

Case 4: σ_i gives birth and σ_{i+1} gives death. Set $U[i, i+1] = 0$, just in case.

According to the reduced matrix, every simplex gives either birth or death. In other words, no simplex has a neutral effect, which implies that cancellations and anti-cancellations cannot happen. This is of course not possible, and the resolution to the apparent dilemma is that consecutive simplices may enter the complex at the same moment in time. To recognize this case, we store with each vertex the distance to t, and with each simplex of dimension 2 or higher a pointer to its vertex with maximum distance from t.

4.4   Staircase Approximations

Using the vineyard algorithm to go from t_0 to t_1, we construct a staircase approximation of the generalized medial axis. As mentioned above, we assume a finite sampling, $A \subseteq \mathbb{R}^d$, of a submanifold, M of $\mathbb{R}^d$, and a simplicial complex, K, with vertex set A that triangulates M. Letting $k < d$ be the dimension of M, the generalized medial axis, $\mathcal{A}_p(M)$, is defined for $0 \leq p \leq k$. Generically, the dimension of $\mathcal{A}_p(M)$ is $d - 1$ and thus independent of p.

To approximate $\mathcal{A}_p(M)$, we fix a *resolution*, $\varrho > 0$, and write $\varrho\mathbb{Z}^d$ for the scaled integer lattice in $\mathbb{R}^d$. Two points in $\varrho\mathbb{Z}^d$ are *adjacent* if the Euclidean distance

between them is ϱ. In this case, there are integers $i_1, i_2, \ldots, i_d$ and $1 \le j \le d$ such that one of the points is $t_0 = (i_1, i_2, \ldots, i_d)\varrho$ and the other is $t_1 = (i_1, \ldots, i_j + 1, \ldots, i_d)\varrho$. The $(d-1)$-dimensional *surface element* that separates the two points is

$$A(t_0, t_1) = [i_1 - \tfrac{1}{2}, i_1 + \tfrac{1}{2}]\varrho \times \ldots \times (i_j + \tfrac{1}{2})\varrho \times \ldots \times [i_d - \tfrac{1}{2}, i_d + \tfrac{1}{2}]\varrho. \qquad (4.5)$$

Write $\bar{\mathcal{A}}_p(M)$ for the approximation of the p-th medial axis of M, which we define as a collection of surface elements. To decide whether or not $A(t_0, t_1)$ is part of this approximation, we run the vineyard algorithm from t_0 to t_1. Let $\mathrm{I}_p(t_0, t_1)$ be the index-p Faustian interchanges that occur in the linear homotopy from t_0 to t_1. Then

$$A(t_0, t_1) \begin{cases} \in \bar{\mathcal{A}}_p(M) & \text{if } \#\mathrm{I}_p(t_0, t_1) \text{ odd,} \\ \notin \bar{\mathcal{A}}_p(M) & \text{if } \#\mathrm{I}_p(t_0, t_1) \text{ even,} \end{cases} \qquad (4.6)$$

see Figure 4.1 for an example. With a straightforward implementation of the vine-
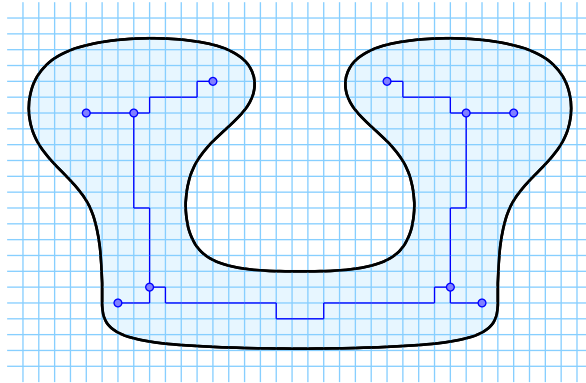
Figure 4.1: Staircase approximation of the 1-st medial axis shown in Figure 1.1. To simplify the picture, we show only the part inside the region bounded by the closed curve.

yard algorithm, we can decide the membership of any surface element in $O(n^3)$ time, in which n is the number of simplices in K. There are numerous practical considerations and possible improvements, of which we list a few.

- Instead of the entire lattice, $\varrho\mathbb{Z}^d$, we explore only a bounded piece, eg. the points inside $[-1, 1]^d$. The running time is $O(n^3/\varrho^d)$.

- Along any line, two vertices change order only once, so we can compute the membership of the surface elements along a coordinate line in $O(n^3)$ time in total. The overall improvement is to $O(n^3/\varrho^{d-1})$ time.

- The Euclidean distance is 1-Lipschitz, so distances change only gradually and we may be able to discard large regions of $[-1, 1]^d$ without running the vineyar algorithm.

Chapter 5

Experiments

In this chapter, we showcase the results of implementing the algorithm in the previous chapter. We verify that our method finds the standard medial axis (which we now call the 0-th medial axis), and we present pictures of the new 1-st medial axis. We note that for simplicity, we use the naïve method mentioned in the introduction to the previous chapter, namely we sample a grid over our input complex, and test for Faustian p-switches between neighboring grid cell centers. If we find such a switch, we draw the edge that the two grid cells share. We realize that we incur some noise, miss some detail, and have a slower running time by not making full use of the vineyards algorithm to find precise locations of switches and ensure we don't miss any. Refining with the vineyards algorithm is an obvious next step for future implementation. A second obvious next step would be to extend our computations to 3D.

5.1   Some Pretty Pictures

Using the (naïve version of the) algorithm proposed in the previous chapter, we can already see an advantage over the standard Voronoi-based algorithm for computing the medial axis. In Figure 5.1, we compute an approximation of the 0-th medial axis of a flower shape. We purposely choose a very sparse point sample (only 20 points) to illustrate that whereas the Voronoi method cannot improve itself without adding more points to the input sample, our method can improve its result by tuning the grid, even with the same bad 20-point input. However, a downfall of our method is that one must balance a system of parameters to achieve success. We see that a much denser input sample of 200 points doesn't improve our result if we don't calibrate the grid as well.
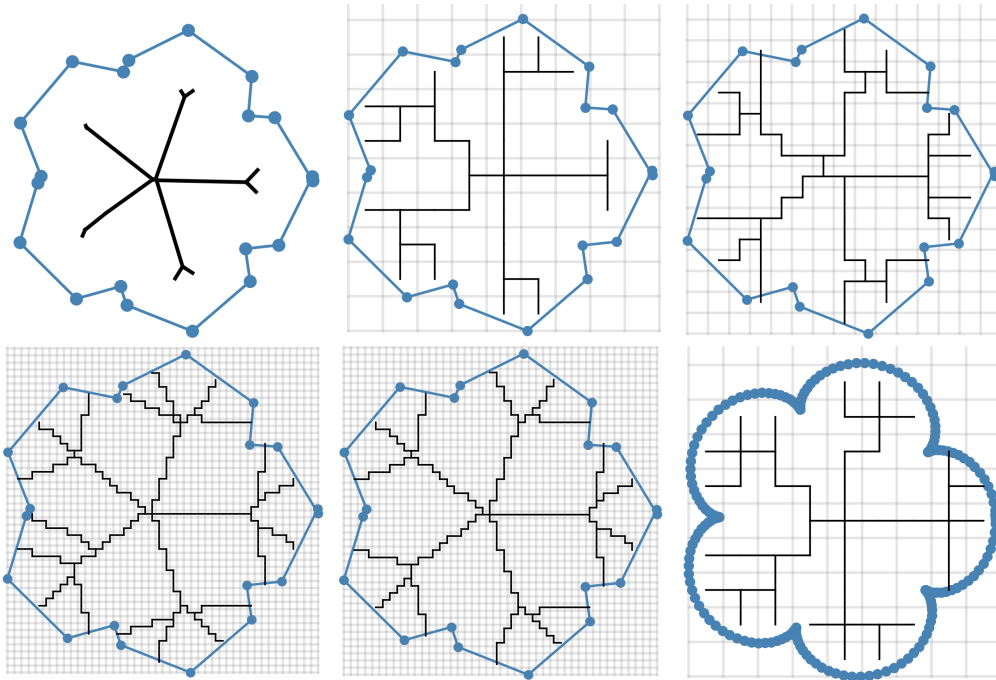
Figure 5.1: Top left: the best the standard Voronoi method can do with 20 input points. Continuing: we refine the grid and get a better approximation with our method, still with 20 input points. Last: our method performs roughly the same with 200 points as it does with 20 if we do not refine the grid with the sampling (the grid density in the bottom right image is the same as that in the top center image).

We show in Figure 5.2 the staircase approximation of the medial axis for several different inputs. The strength of the regular grid is also its downfall: if the grid edges happen to fall where the medial axis edges should be, it can cause confusion for the algorithm. See Figure 5.3 for an example of a misaligned medial axis that is caused by a grid placement problem. Additionally, Figure 5.3 showcases the 1-st medial axis.

As an alternative that is less impacted by poor placement (but looks more wobbly and tends to have more spurious edges), we run the naïve algorithm from the previous chapter, but instead of testing for an index-p Faustian switch between neighboring grid cells, we do the following:

1. Run a Poisson point process in a rectangle that covers the input shape, with intensity i

2. Compute the Voronoi tesselation of the Poisson points

3. Look at the dual Delaunay edges of the Voronoi edges whose vertices are both contained in the interior of the input shape. If, by running along the Delaunay edge, one encounters an index-p Faustian switch, draw the corresponding Voronoi
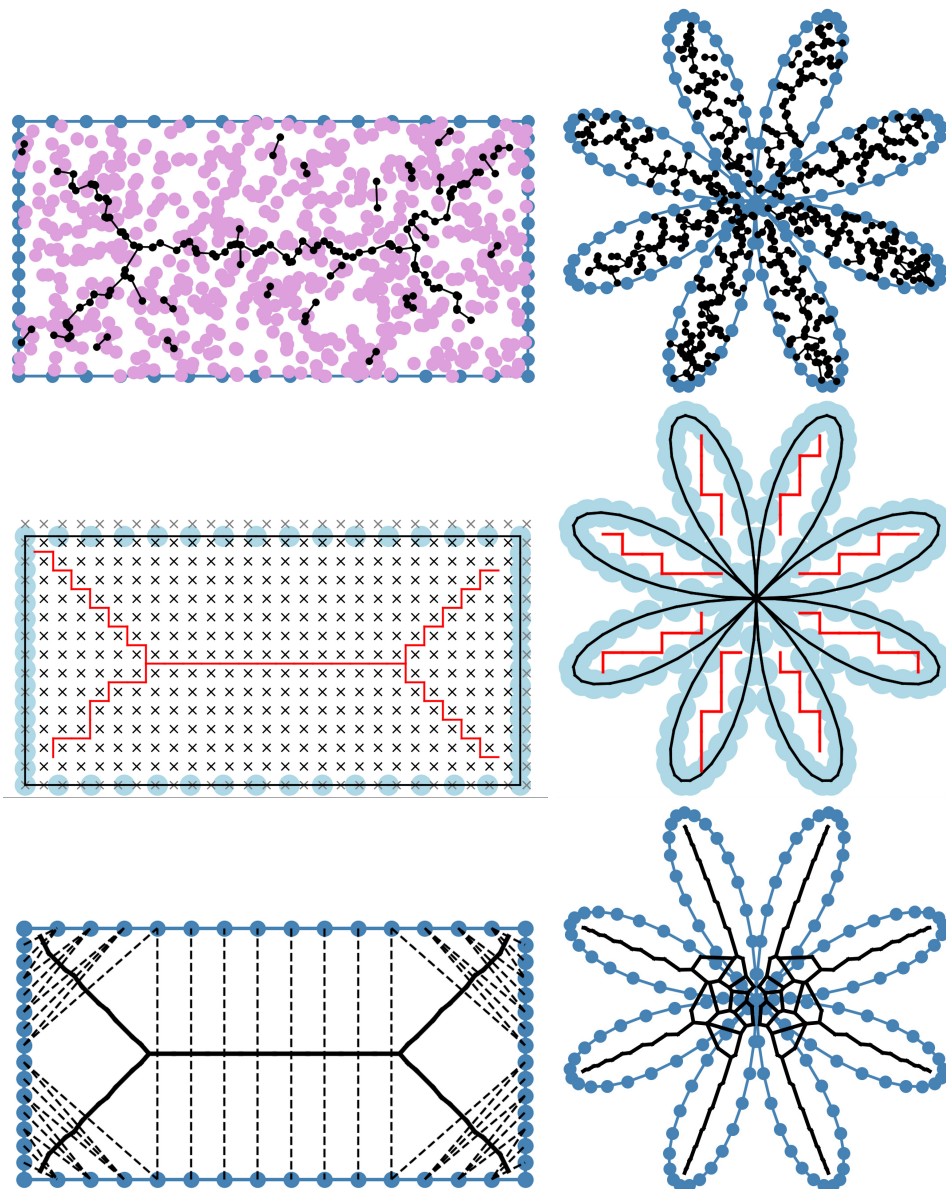
Figure 5.2: Our Poisson and staircase, and the standard Voronoi medial axis approximations of the rectangle and flower. See data in Table 5.3.

edge as part of the medial axis approximation. One can also exchange the Delaunay and Voronoi diagrams in these steps, drawing Delaunay edges if their corresponding Voronoi edges encounter an index-p Faustian switch.

Figure 5.4 illustrates the process, and Figure 5.5 shows how the approximation changes when increasing the Poisson intensity i and the pruning parameter ϵ. See also Figure 5.2.

For both the grid and Poisson methods, we can tune various parameters to affect the accuracy of the medial axis approximation. In the grid setting, we can refine the
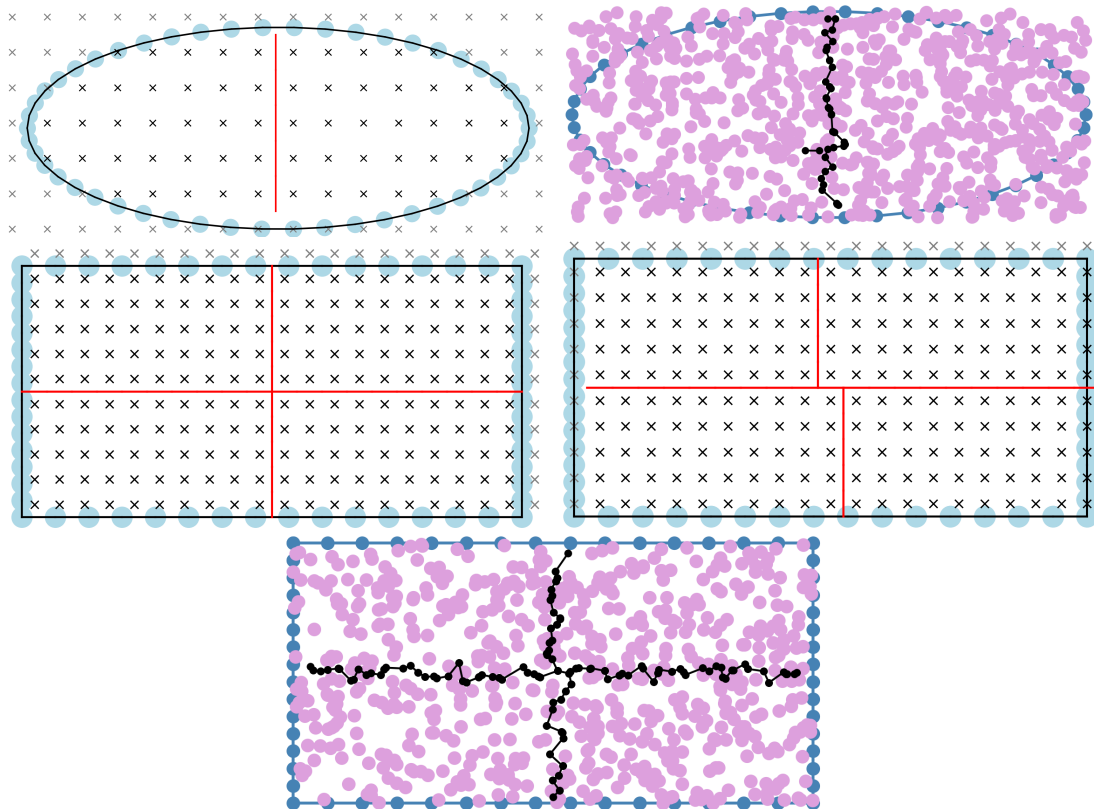
Figure 5.3: Here we demonstrate the first medial axis of the ellipse and the rectangle, with both grid and Poisson approaches. We cannot show the traditional Voronoi approach, as it is only defined for the traditional, or 0-th medial axis. We also show that grid alignment matters. See Table 5.4 for parameters.

grid density with parameter g, which is the distance between neighboring cells. We can also increase n, the number of points we sample along the input shape. For the Poisson method, we can increase the intensity i, sampling more points. Increasing any of these parameters results in an approximated medial axis which covers more of the real medial axis, but also results in more spurious edges. To prune some of these, we compare the distance $|p_1 - p_0|$ between the test points (the grid cell centers or the Delaunay vertices) to the distance $|x_1 - x_0|$ between the simplices of the input shape that are involved in the index-p Faustian interchange. If $|x_1 - x_0| > |p_1 - p_0|\epsilon$, we keep the edge $\vec{p}$. Increasing ϵ thus prunes away more edges.

In Figure 5.6, we demonstrate the infamous instability of the medial axis by adding a bump to our input shape, and confirm that our method is just as unstable as the traditional one.
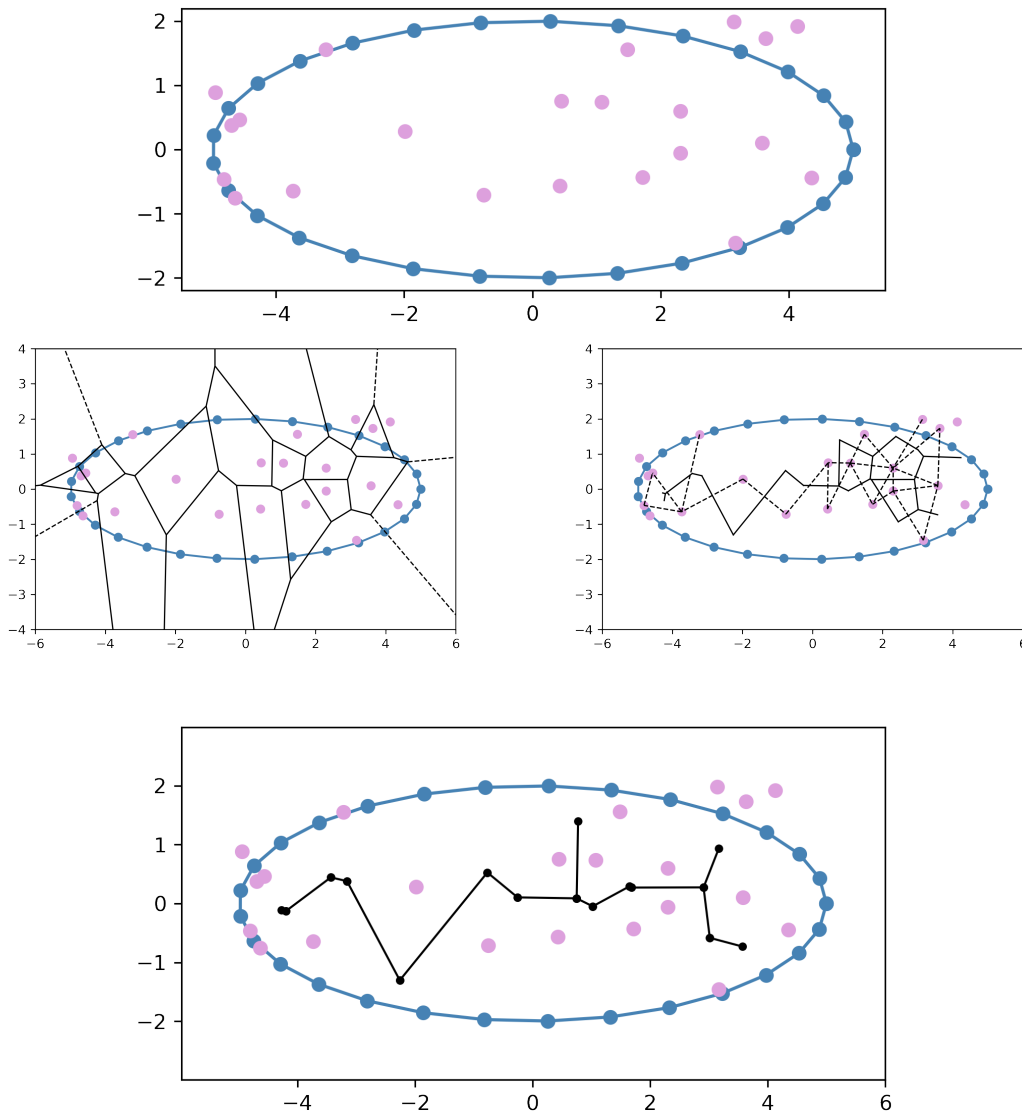
Figure 5.4: Top: we sample points around an ellipse and run a Poisson point process to generate the light purple points. Middle left: we take the Voronoi tesselation of the Poisson points. Dashed lines go to infinity. Middle right: we show the restriction of the Voronoi tesselation to the interior of the ellipse, and their corresponding dual Delaunay edges (dashed edges). Bottom: we further restrict the Voronoi edges to those whose dual Delaunay edges traverse an index-p Faustian interchange. This is an approximation of the ellipse's medial axis.
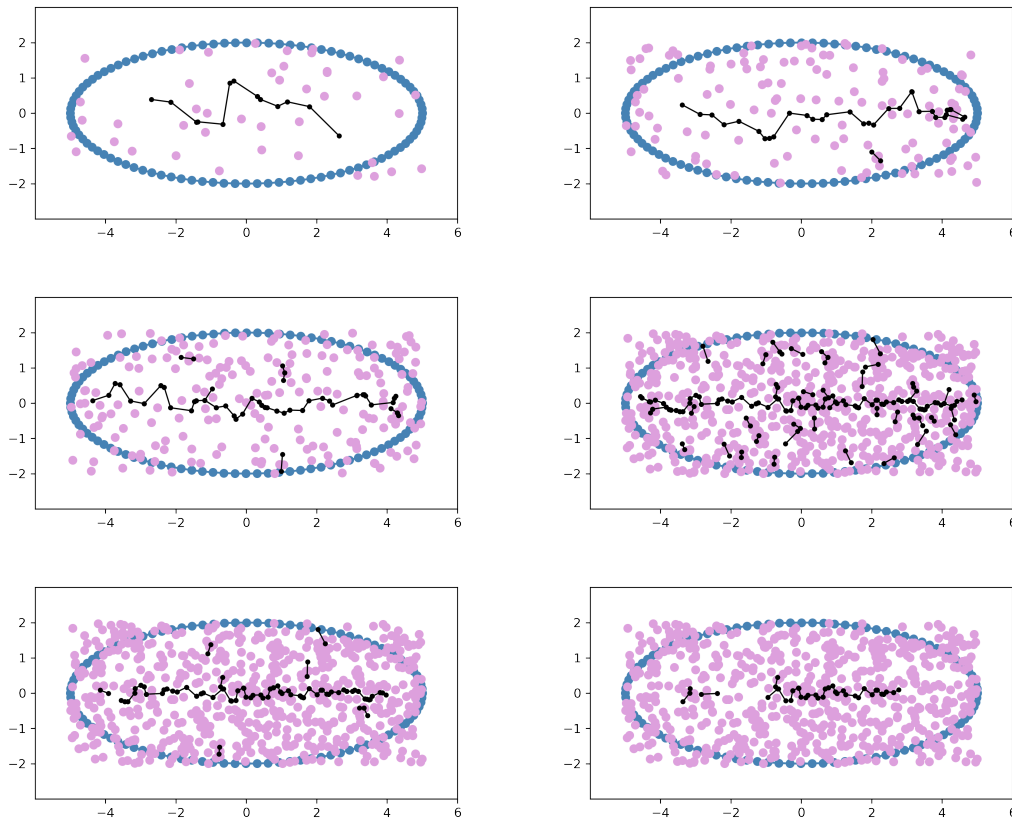
Figure 5.5: Increasing the Poisson intensity improves the medial axis approximation. However, spurious edges enter the scene. We can prune them (bottom two images), but this can also go too far and remove parts of the medial axis we want to keep (last image). For parameters, see Table 5.1.

One problem with the Poisson-based Voronoi/Delaunay approach is that the sampled Poisson points are not evenly distributed, and this leads to noise in the approximated medial axes that is difficult to eliminate. We can combat this by adding "blue noise" to the Poisson point process: we add Poisson points one by one, and check that each new point is outside a predetermined radius of each other point. See Figure 5.7. This prevents Poisson points from being too close together. Now the grid edges are more regular, and pruning works better. One problem with the blue noise, however, lies in a future avenue of exploration: there have been some nice results [11] approximating intrinsic volumes with Poisson-Delaunay mosaics, and we think we could use these results to determine bounds on the distortion factor a Poisson-Voronoi/Delaunay medial axis would have with respect to the length of the theoretical medial axis, but by adding the blue noise, we could no longer rest on the

Table 5.1: Data for Fig 5.5. n is the number of points used to sample the ellipse, ϵ is the pruning parameter, and i is the Poisson point process intensity. Fig number is ordered top down, left to right.

fig	n	ϵ	i
1	100	2	1
2	100	2	3
3	100	2	5
4	100	2	15
5	100	4	15
6	100	9	15

Table 5.2: Data for Fig 5.6. n is the number of points used to sample the ellipse, g is the grid density, ϵ is the pruning parameter, and i is the Poisson point process intensity.

| Method	n	g	i	ϵ
Poisson	100	–	15	3
Grid-based	100	0.3	–	4
Voronoi	100	–	–	–

laurels of these previous results.
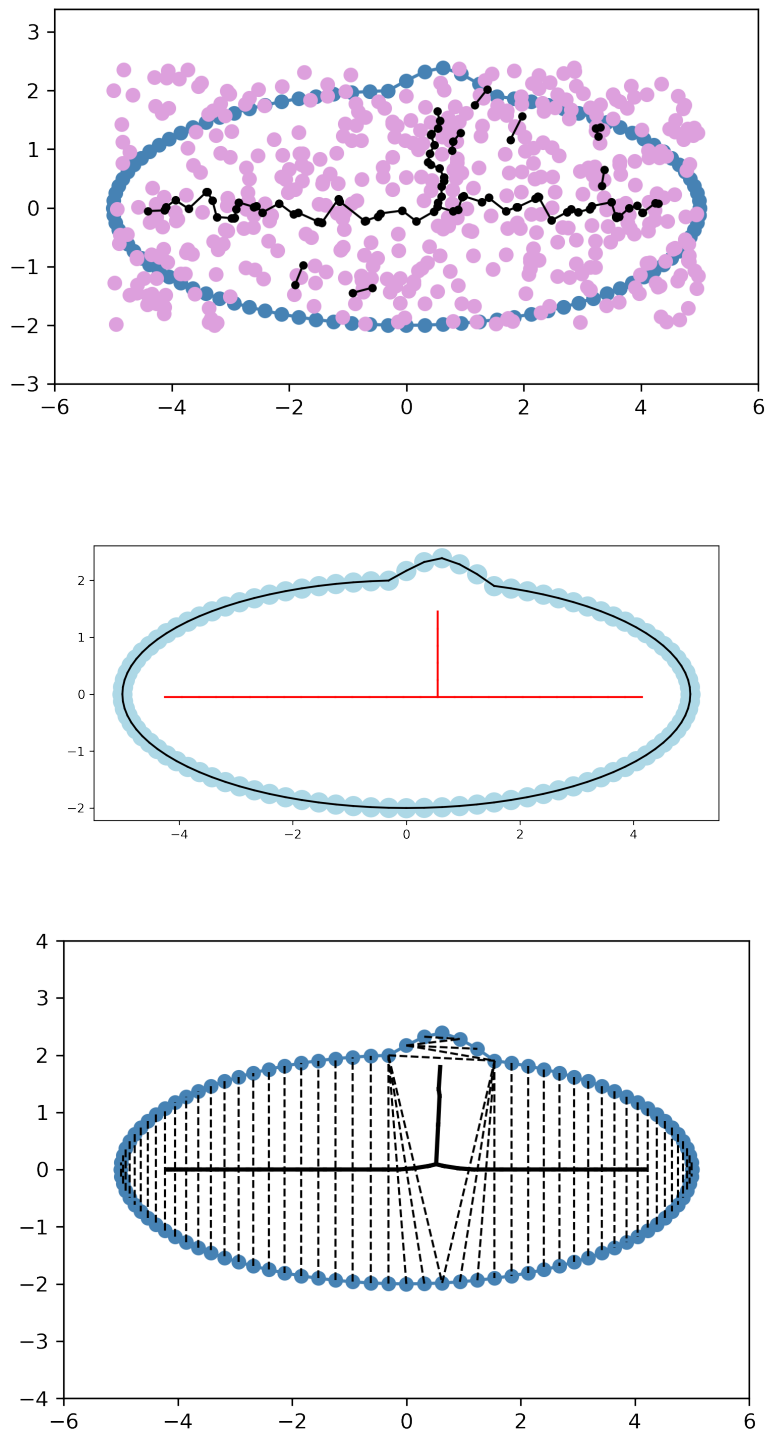
Figure 5.6: Instability of our Poisson and grid-based approaches, as well as the traditional Voronoi approach. see Table 5.2 for parameters.

Table 5.3: Data for Fig 5.2. Fig number is ordered top down, left to right.

fig	n	g	ϵ	i
1	60	-	4	20
2	150	-	2	500
3	160	0.3	2	-
4	150	0.1	2	-
5	60	-	-	-
6	150	-	-	-

Table 5.4: Data for Fig 5.3. Fig number is ordered top down, left to right.

fig	n	g	ϵ	i
1	50	0.7	1	-
2	50	-	5	19
3	60	0.4	2	-
4	60	0.4	2	-
5	60	-	4	19
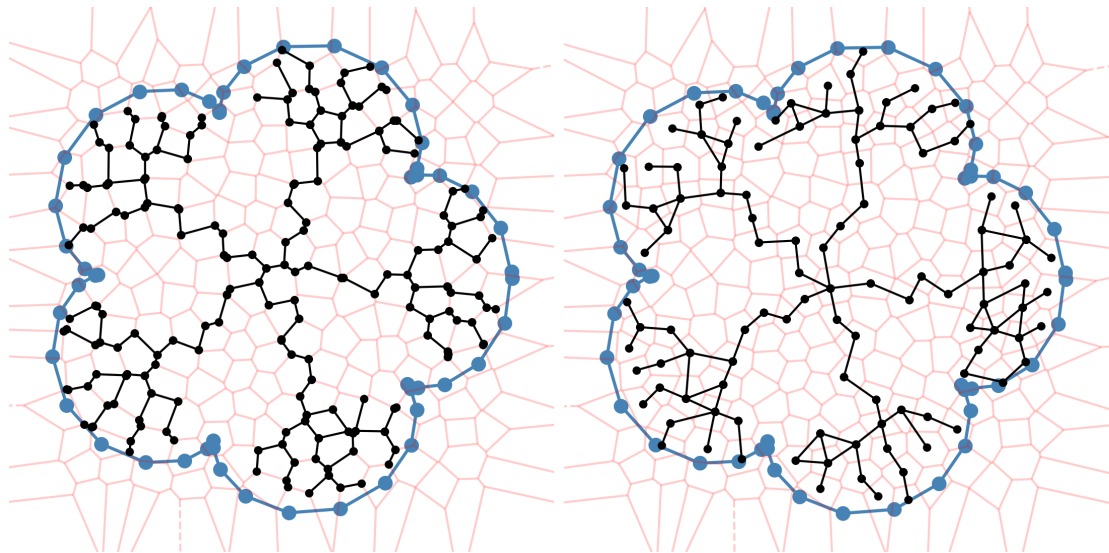
Figure 5.7: The Poisson-based grid method with blue noise added to make the "grid" more regular. On the left, the edges of the medial axis are Voronoi edges, and on the right, they are Delaunay edges. Note that there are cycles in the medial axis that should not be there; this is most likely an artifact of using the naïve method and skipping the sanity check the vineyards algorithm provides.

Chapter 6

Discussion

The main contribution of this paper is a generalization of the medial axis using the persistent homology of pointed squared distance functions, and an algorithm for approximating it. The algorithm is used to illustrate the new concept but has not been optimized for speed and appearance.

Open mathematical questions.

- Is $\mathcal{A}_p(M)$ homotopy equivalent to the solid body bounded by M (if it exists)?

- Can the λ-medial axis be extended to a stable version of the generalized medial axis?

- Is there a scale-dependent version of the generalized medial axis based on the persistent homology of the pointed squared distance functions?

Open computational questions.

- Do we need adaptive sampling as we change the point t for the pointed squared distance function?

- Can we use the derivative of the persistence diagram to trace out $\mathcal{A}_p(M)$?

- Is there an analog of the fast distance transform that can be used to compute an approximation of the generalized medial axis?

- Is there a hierarchical approach to computing the generalized medial axis that avoids spending time in portions of space that are empty of the medial axis?

Bibliography

[1] E. Bittar, N. Tsingos, and M. Gascuel. Automatic reconstruction of unstructured 3d data: Combining a medial axis and implicit surfaces. *Computer Graphics Forum*, 1995.

[2] H. Blum. A transformation for extracting new descriptors of shape. *Models for the Perception of Speech and Visual Form*, pages 362–380, 1967.

[3] J.W. Bruce, P.J. Giblin, and C.G Gibson. Symmetry sets. *Proceedings of the Royal Society of Edinburgh Section A: Mathematics*, 101:163–186, 1985.

[4] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37:103–120, 2007.

[5] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. Vines and vineyards by updating persistence in linear time. *Proc. 22nd Sympos. Comput. Geom.*, pages 119–126, 2006.

[6] H. Edelsbrunner and J. Harer. *Computational Topology. An Introduction.* Amer. Math. Soc., Providence, Rhode Island, 2010.

[7] H. Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93:418–491, 1959.

[8] R. Gilmore. *Catastrophe Theory for Scientists and Engineers*. Wiley, New York, New York, 1981.

[9] M. Golubitsky and V. Guillemin. *Stable Mappings and their Singularities*. Springer-Verlag, New York, New York, 1973.

[10] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, Cambridge, England, 2002.

[11] P. Koehl, A. Akopyan, and H. Edelsbrunner. Computing the volume, surface area, mean, and gaussian curvatures of molecules and their derivatives. *J. Chem. Inf. Model.*, 63(3), 2023.

[12] J.N. Mather. Distance from a submanifold in Euclidean space. *Singularities, Proceedings of Symposia in Pure Mathematics*, 40, Part 2:199–216, 1983.

[13] J. Milnor. *Morse Theory*. Princeton Univ. Press, Princeton, New Jersey, 1963.

[14] J. Montaldi. *Singularities, Bifurcations and Catastrophes*. Cambridge Univ. Press, Cambridge, England, 2021.

[15] T. Poston and I. Stewart. *Catastrophe Theory and its Applications*. Dover Publications, Mineola, New York, 1978.

[16] R. Tam and W. Heidrich. Shape simplification based on the medial axis transform. *IEEE Visualization*, 2003.

[17] B. Yang and J. Yao. Dmat: Deformable medial axis transform for animated mesh approximation. *Computer Graphics Forum*, 2018.

Signature Page

The Master's thesis of Elizabeth Stephenson, titled Generalizing Medial Axes with Homology Switches, is approved by:

Supervisor: Herbert Edelsbrunner, IST Austria, Klosterneuburg, Austria

Signature: _____

Committee Member: Chris Wojtan, IST Austria, Klosterneuburg, Austria

Signature: _____

Signed page is on file

Copyright