

Variational Pruning of Medial Axes of Planar Shapes

P. Rong and T. Ju 

Washington University in St. Louis, USA

Abstract

Medial axis (MA) is a classical shape descriptor in graphics and vision. The practical utility of MA, however, is hampered by its sensitivity to boundary noise. To prune unwanted branches from MA, many definitions of significance measures over MA have been proposed. However, pruning MA using these measures often comes at the cost of shrinking desirable MA branches and losing shape features at fine scales. We propose a novel significance measure that addresses these shortcomings. Our measure is derived from a variational pruning process, where the goal is to find a connected subset of MA that includes as many points that are as parallel to the shape boundary as possible. We formulate our measure both in the continuous and discrete settings, and present an efficient algorithm on a discrete MA. We demonstrate on many examples that our measure is not only resistant to boundary noise but also excels over existing measures in preventing MA shrinking and recovering features across scales.

CCS Concepts

• **Computing methodologies** → **Shape analysis**;

1. Introduction

The medial axis (MA) [Blu67] is a commonly used shape descriptor in computer graphics and computer vision. Defined as the loci of points with two or more closest points on the shape's boundary, MA enjoys several desirable properties including being thin, centered, capturing the shape's structure, and preserving its topology [Lie03]. As a result, MA has found utility in many applications including shape analysis, meshing, animations, to name a few [SP08].

A key limitation of MA is its sensitivity to noise. Small perturbations of the shape's boundary may result in the addition of numerous spurious branches on MA (e.g., see the Seahorse and its MA on the left of Figure 1). These branches need to be removed, or *pruned*, for MA to be useful in shape description. Existing pruning methods are often guided by some *significance measure* over MA. Such a measure assesses the importance of shape features represented by at each MA point [SB98]. Various significance measures have been proposed in the literature (see a brief review in Section 2). While *local* measures are defined based on the immediate neighborhood of an MA point (e.g., its closest points on the shape), *global* measures consider the overall shape and therefore can be more effective in distinguishing noise from salient shape features.

However, existing global measures share some common drawbacks. First, they are often biased towards the middle of MA and assume low values both near the ends of MA and along unwanted branches. Pruning therefore may come at the cost of shrinking the desirable branches of MA from its ends (see the boxed regions in Figure 1). Such shrinking may be detrimental, for example, if the MA is to be used for measuring the length of a shape, detecting

feature points on the shape boundary, or finding correspondences between shapes. Second, most global measures are scale-sensitive, in that they are generally higher in shape parts with larger sizes. As a result, pruning could result in disproportionate loss of features at different scales (see an example in Figure 2).

We propose a new significance measure for MA of 2D shapes that avoids these drawbacks while maintaining the key properties of MA (e.g., preserving structure and topology). Our measure is *variational* in that it is formulated as the solution to a constrained optimization problem. We introduce a score, called *Boundary-Skeleton Parallelism* (BSP), that evaluates how parallel an MA subset is to the shape boundary compared to a given angle α . The score encourages the inclusion of more (resp. less) MA branches that are more (resp. less) parallel to the boundary than α . We then seek, as α increases, a contracting sequence of MA subsets that maximize BSP while preserving the topology of MA. The significance at an MA point x , which we call the *Vanishing Angle* (VA) of x , is the highest α such that x remains in this sequence. We give formulations on both a continuous MA and a discrete approximation of MA, and develop an efficient, quadratic-time algorithm for the latter.

When evaluated on many planar shapes, we observe that VA, like other global measures, is robust against small boundary perturbations. Furthermore, pruning using VA leads to much less MA shrinkage than existing measures (e.g., Figure 1 last row), and its scale-invariance enables preservation of features across scales (e.g., Figure 2 last row).

The rest of the paper is organized as follows. After a brief review of the literature (Section 2), we introduce our variational formula-




Figure 1: Comparing our proposed significance measure, *Vanishing Angle* (VA), with two other global measures on the Seahorse (left: measures on the MA; right: pruned MAs at increasing thresholds; thresholds for VA are 30° , 45° , 60°). Our measure is as effective as other measures in removing unwanted branches but results in less shrinkage of MA (see boxed regions).

tion (Section 3), first on a continuous MA and then on a discrete approximation. Then we present our algorithm on a discrete MA (Section 4). Next we show the results (Section 5) and conclude with a discussion of limitations (Section 6).

2. Related works

We will briefly review works on improving the robustness of MA to boundary perturbations (also known as MA regularization), particularly in two dimensions. For in-depth discussions on the properties, computation, and applications of MA, please refer to the classical book by Siddiqi and Pizer [SP08] as well as the survey by Tagliasacchi et al. [TDS*16].

One approach to regularize MA is to extract MA after regularizing (i.e., smoothing) the shape boundary [DLN87, POB87, GMPW09, MGP10]. However, boundary smoothing may introduce significant changes to the shape, especially near highly convex or concave regions (see Figure 2 of [SB98] for an example). As a result, the MA of the regularized shape may not faithfully capture the structure and topology of the original shape, and it may also lie outside the original shape (see Figure 8). Another approach is to deform a clean MA template to fit a target shape




Figure 2: Comparing significance measures (left) and pruned MAs (right; VA threshold is 60°) on an ensemble of Seahorses at decreasing scales (top). Our measure, VA, is scale-invariant and hence preserves features of Seahorses at different scales equally well while removing unwanted MA branches.

[GEG00, PFJ*03, PGJA03]. This method is most useful for a family of similar shapes, such as characters and biological forms.

Yet another approach, which we take, is to directly prune MA to produce a subset that best represents the shape. Unlike boundary smoothing, MA pruning upholds two key properties of MA, namely being centered inside the shape and preserving its topology. To do so, one needs to distinguish MA branches that capture salient shape features as opposed to small boundary undulations. This is done by defining some *significance measure* over MA. Given such a measure, pruning can be done by either thresholding the measure, if all “level sets” of the measure preserve the MA topology, or using an erosion procedure guided by the measure to produce a topology-preserving subset [SBTZ02, SB98, SFM05].

Many existing measures consider the local neighborhood of an MA point x that consists of its closest boundary points $N(x)$. Two such local measures, which are shown in Figure 3, are the *Object Angle* (OA) [AM96, ACK01, DZ04, FLM03, SFM05], which is (half




Figure 3: Comparing various local (OA and Circumradius) and global (ET, CR and VR) significance measures (top) and their respective pruned MAs (bottom) on the Chicken (left). For local measures, an erosion heuristic is used [SBTZ02] to ensure topology preservation. Boxes highlight spurious branches (OA), missing branches (Circumradius), and shrunken ends (ET and CR) on the pruned MA. VA pruning threshold is 50° .

of) the angle between the vectors from x to $N(x)$, and the *Circumradius* [CL04, CCT09], which is the radius of smallest ball defined by $N(x)$. A higher OA means MA locally is more parallel to the boundary, and hence the local shape is more tubular, whereas a higher Circumradius implies a thicker shape. OA is also related to other measures such as the propagation velocity in [Blu73] and the outward flux in [SBTZ02, DPS00]. However, local measures are inherently limited in their ability to differentiate noise from salient features. For example, MA inside small bumps can have high OA (e.g., chicken fur in Figure 3), and MA inside thin but salient features has low Circumradius (e.g., chicken feet in Figure 3).

To better recognize noisy features, several global measures of significance have been proposed for MA of planar shapes. The underlying idea of these measures is to capture the amount of the shape information lost due to pruning. For example, *Erosion Thickness* (ET) [HD86, BA92, AdBT95, SB98, LCLJ11] approximates the distance from the end of the pruned shape to the original shape, and *Chord Residue* (CR) [OI92, OK95] considers the geodesic distance between $N(x)$ on the shape boundary subtracted by their Euclidean distance. As shown in Figure 3, both ET and CR correctly highlight MA branches lying in main shape parts (e.g., chicken body and feet) and assign low values to spurious branches. Both measures have been extended to MA of 3D shapes and found to be equally effective for pruning [RvWT08, DS06, YSC*16]. Other 2D global measures include the area of erosion [SB98, AdBT95] and the Delta MA [MLIM16]. However, as these measures all capture shape loss, which are scale-dependent, so are the measures themselves. Furthermore, since pruning to MA points located near the shape extremities would incur less loss than to MA points near the center of the shape, the latter are naturally considered to be more significant than the former. As a result, pruning spurious branches brings the side-effect of shrinking the remaining branches (e.g., Figure 1 and chicken feet in Figure 3).

Finally, a group of recent methods (mostly in 3D) aim to simultaneously simplify the discrete representation of an approximate MA while removing the spurious branches [FTB13, LWS*15, DLX*22]. These methods are typically guided by a combination of measures that concern both the saliency of shape features and the sampling density of MA.

3. Formulation

Our method builds on the local significance measure, *object angle* (OA) (see Section 2). Despite its sensitivity to noise (see Figure 3), OA outperforms global measures (e.g., ET or CR) in two aspects. First, while the global measures are biased towards the center of MA, which may result in significant erosion of extremities after pruning, OA is oblivious of the location of the point on MA due to its local definition. Hence OA can be equally high in the center of MA (e.g., chicken body in Figure 3) and at its ends (e.g., chicken feet in Figure 3). Second, OA is invariant under uniform scaling. However, the locality and scale-invariance also make OA sensitive to small boundary features.

To improve the robustness of OA to noise, while retaining its extremity-awareness and scale-invariance, we consider a more global picture beyond the MA point and its closest boundary points. The key observation is that segments of MA with high OA that are located in small boundary features (e.g., chicken fur in Figure 3) are often connected to the rest of MA via long segments with low OA. We thus formulate a variational problem that seeks, at a given pruning threshold, a *connected* subset of MA that contains as many (resp. few) points with high (resp. low) OA as possible. We additionally require that the sequence of such subsets monotonically contracts as the threshold increases, so that a significance measure can be derived as the highest threshold at which an MA point remains in the pruned set.




Figure 4: Local picture of a point x on a continuous MA (a) and an edge e on a discrete, Voronoi-based MA (b).

We first introduce our objective function based on OA (Section 3.1), which has an alternative geometric interpretation in terms of the length of MA and the boundary. We then define the variational pruning process and the significance measure (Section 3.2). Finally, we adapt our formulations to a discrete representation of MA (Section 3.3).

3.1. Boundary-Skeleton Parallelism

Consider a planar shape $S \in \mathbf{R}^2$ as an open set whose boundary is ∂S . The medial axis $\mathcal{M}$ of S is the set of points in S with two or more closest boundary points:

$$\mathcal{M} = \{x \mid x \in S, |N(x)| \geq 2\}.$$

Here, $N(x)$ denotes the points on ∂S that are closest to $x \in \mathbf{R}^2$. In general, $\mathcal{M}$ is a planar graph that consists of *regular* points where $|N(x)| = 2$ and *junction* points where $|N(x)| > 2$. The OA of a regular point $x \in \mathcal{M}$, denoted by $OA(x)$, is half the angle between the two vectors from x to the two points in $N(x)$ (see Figure 4 (a)). Intuitively, a higher OA indicates that MA at x is more parallel to the boundary.

Given some angle threshold $\alpha \in [0, \pi/2]$, we denote the difference between the OA of regular point $x \in \mathcal{M}$ and α as

$$f(x, \alpha) = \sin(OA(x)) - \sin(\alpha). \quad (1)$$

The use of sine will be explained in a moment. Intuitively, a positive (resp. negative) difference indicates that MA at x is more (resp. less) parallel to the boundary than the expected angle α . Integrating the difference over a subset $\mathcal{M}' \subseteq \mathcal{M}$ yields our objective function, called the *Boundary-Skeleton Parallelism* (or BSP),

$$\text{BSP}(\mathcal{M}', \alpha) = \int_{\mathcal{M}'_R} f(x, \alpha) dx, \quad (2)$$

where $\mathcal{M}'_R$ denotes the set of regular points of $\mathcal{M}'$. Note that maximizing BSP has the simultaneous effect of encouraging the inclusion of points with OA higher than α (which contribute positively to BSP) and discouraging the inclusion of points with OA lower than α (which contribute negatively to BSP).

BSP has a more direct interpretation in terms of the length of

MA and of the boundary, thanks to the use of sine in Equation 1. Specifically, substituting 1 into 2 yields:

$$\text{BSP}(\mathcal{M}', \alpha) = \int_{\mathcal{M}'_R} \sin(OA(x)) dx - \sin(\alpha) \int_{\mathcal{M}'_R} dx, \quad (3)$$

Observe from Figure 4 (a) that $\sin(OA(x))dx$ is the length of the infinitesimal segment of ∂S at one of the two closest points of x . Hence the first integral on the rhs of Equation 3 is half of the total length of all closest points of $\mathcal{M}'$ on ∂S . The second integral is simply the length of $\mathcal{M}'$ scaled by $\sin(\alpha)$. Let $N(\mathcal{M}')$ be the union of closest boundary points to $\mathcal{M}'$, Equation 3 can be re-written as:

$$\text{BSP}(\mathcal{M}', \alpha) = |N(\mathcal{M}')|/2 - \sin(\alpha)|\mathcal{M}'|. \quad (4)$$

In words, BSP of the subset $\mathcal{M}'$ is the difference between half length of the boundary curve represented by $\mathcal{M}'$ and the $\sin(\alpha)$ -scaled length of $\mathcal{M}'$.

3.2. Variational pruning

Given a threshold $\alpha \in [0, \pi/2]$, we prune $\mathcal{M}$ to a subset that maximizes BSP while satisfying two constraints. First, the subset is homotopy equivalent to $\mathcal{M}$, meaning that it is connected and retains all loops in $\mathcal{M}$. Second, as α increases, the pruned subsets form a contracting sequence. Contraction is required if we want to define a significance measure over $\mathcal{M}$ that can reproduce the pruned subsets as its “level sets”.

Formally, we call a function $\mathcal{M}(\alpha) : \mathbf{R} \rightarrow \mathcal{M}$ a *pruning function* if $\mathcal{M}(0) = \mathcal{M}$ and the following holds for all $\alpha \in (0, \pi/2]$:

$$\mathcal{M}(\alpha) \in \arg \max_{\mathcal{M}' \simeq \mathcal{M}; \mathcal{M}' \subseteq \mathcal{M}(\beta), \forall \beta \in [0, \alpha]} \text{BSP}(\mathcal{M}', \alpha). \quad (5)$$

Here, $\simeq$ indicates homotopy equivalence. Note that the rhs of Equation 5 may not be unique; that is, there may be multiple constraint-satisfying subsets $\mathcal{M}'$ with equal and maximal $\text{BSP}(\mathcal{M}', \alpha)$ for some α . As a result, there may exist multiple pruning functions $\mathcal{M}(\alpha)$.

Given a pruning function $\mathcal{M}(\alpha)$, we define the significance at a point $x \in \mathcal{M}$ as the smallest angle threshold α at which x is no longer in the pruned set. This measure, which we call the *Vanishing Angle* (VA), has the form:

$$\text{VA}(x) = \inf\{\alpha \mid \alpha \in [0, \pi/2], x \notin \mathcal{M}(\alpha)\}. \quad (6)$$

We use $\inf$ instead of $\min$ because the minimum may be only reached at the limit (e.g., if $x \in \mathcal{M}(\text{VA}(x))$ but $x \notin \mathcal{M}(\alpha)$ for any $\alpha > \text{VA}(x)$).

VA has several desirable properties as a significance measure. First, since $\mathcal{M}(\alpha)$ contracts as α increases, the part of $\mathcal{M}$ where VA is above a given α is exactly $\mathcal{M}(\alpha)$. As a result, and unlike many existing measures (e.g., OA, circumradius, CR, etc.), thresholding VA always results in a topology-preserving subset of $\mathcal{M}$.

Furthermore, just like OA, VA is invariant to uniform scaling. This is because a resizing of S by a factor of s transforms $\mathcal{M}$ (and any of its subsets) proportionally, and in turn multiplies s to the BSP of any subset $\mathcal{M}'$ (which is the difference between curve lengths). As a result, a pruning function $\mathcal{M}(\alpha)$, after resized by s , remains a pruning function on the resized $\mathcal{M}$, and hence the corresponding VA at each point of $\mathcal{M}$ remains the same after resizing.

3.3. Discrete formulation

To compute VA, we adapt our continuous formulations above to a discrete approximation to $\mathcal{M}$, represented as a Euclidean graph M . We assume that each edge e of M is equipped with an OA measure, denoted by $\text{OA}(e)$. The graph M can be obtained, for example, as the interior vertices and edges of the Voronoi Diagram of point samples along ∂S . This method results in a convergent approximation to the MA [BA92] and has been considered by many researchers [OI92, AM97, SM13]. In this case, each edge $e \in M$ is closest and equidistant to two boundary samples $N(e)$, whose Voronoi cells are incident to e , and $\text{OA}(e)$ is half of the angle spanned by the vectors from the midpoint of e to $N(e)$ (see Figure 4 (b)).

Given a threshold α , we first obtain the difference between the (sine of) $\text{OA}(e)$ of an edge e of M and α as,

$$f(e, \alpha) = \sin(\text{OA}(e)) - \sin(\alpha), \quad (7)$$

and we approximate the BSP of a subgraph $M' \subseteq M$ by a summation,

$$\text{BSP}(M', \alpha) = \sum_{e \in M'} |e| f(e, \alpha), \quad (8)$$

where $|e|$ is the length of edge e .

The pruning function $\mathcal{M}(\alpha)$, as α increases from 0 to $\pi/2$, takes the form of a finite sequence of contracting graphs, $M_\Sigma = \{M_0 = M \supset M_1 \supset \dots \supset M_k\}$, and a sequence of increasing thresholds, $\alpha_\Sigma = \{\alpha_0 = 0 < \alpha_1 < \dots < \alpha_k \leq \alpha_{k+1} = \pi/2\}$, that satisfy the following properties:

1. $M_i \simeq M$ for all $i = 0, \dots, k$.
2. For $i = 0, \dots, k$, $M_i \in \arg \max_{M' \simeq M, M' \subseteq M_i} \text{BSP}(M', \alpha)$ for any $\alpha \in [\alpha_i, \alpha_{i+1}]$.
3. For $i = 1, \dots, k$, $\text{BSP}(M_i, \alpha) > \text{BSP}(M_{i-1}, \alpha)$ for any $\alpha > \alpha_i$.

In words, each M_i remains, among all its topology-preserving subgraphs, the one with the highest BSP for α in the range $[\alpha_i, \alpha_{i+1}]$, and it is surpassed by the BSP of the subgraph M_{i+1} for $\alpha > \alpha_{i+1}$. We call the tuple $\{M_\Sigma, \alpha_\Sigma\}$ a *pruning sequence*. Just like the pruning function, the pruning sequence may not be unique, as there could be multiple subsets of M_i whose BSP are greater than that of M_i for $\alpha > \alpha_{i+1}$.

Given a pruning sequence, the VA of an edge $e \in M$ is the last threshold α_{i+1} after which e is pruned away:

$$\text{VA}(e) = \alpha_{1+\max\{i | e \in M_i\}} \quad (9)$$

Figure 5 shows an example pruning sequence of a simple shape (with a hole) and its VA. Observe that a noisy branch of MA (see the red arrow) is pruned at a low threshold ($\alpha_1 \approx 33^\circ$), even though the MA edge at the end of that branch has very high OA. This is because the great part of the branch has low OA, and hence the branch overall contributes negatively to BSP at $\alpha > \alpha_1$. On the other hand, the edge at the end of the main MA branch (see the blue arrow) remains in the pruned graph until a much higher threshold ($\alpha_2 \approx 65^\circ$), as it is connected to the rest of MA via mostly high-OA edges. As a result, VA does a better job than OA in differentiating noise from salient features, while still protecting the extremities of MA from over-erosion. Note that all pruned graphs $M_0, \dots, M_5$

are homotopy equivalent to the MA (e.g., they all include the loop around the hole).

4. Algorithm

We will describe how to compute a pruning sequence on a discrete MA, as formulated in Section 3.3. The definition of the pruning sequence has the flavor of a Prize-Collecting Steiner Tree (PCST) problem, as it seeks a connected subgraph that maximizes a summative metric (i.e., BSP). But our problem is more challenging, because BSP is not fixed and varies with the parameter α , and we need to find the value of the parameter at which the solution subgraph changes (from M_i to its subgraph M_{i+1}). Since PCST is already an NP-hard problem, our problem seems intractable at the first sight. However, as we will show, by constraining the solutions to only *topology-preserving* subgraphs, our problem is in fact much easier to solve than PCST.

We start by making an observation of topology-preserving subgraphs (Section 4.1), which we then leverage to design an incremental algorithm for constructing a pruning sequence and resulting VA (Section 4.2). We end with a discussion on implementation and complexity analysis (Section 4.3).

4.1. Topology-preserving subgraphs

As defined in Section 3.3, each graph M_i in a pruning sequence is a subgraph of M_{i-1} that maintain the latter's topology. We will show that such a subgraph can only be obtained by removing a specific type of graph components that we call *outer trees*.

Consider a graph G with vertices V and undirected edges E . Let $\mathcal{E}$ be the set of directed edges induced by E , so that each edge of E is included in $\mathcal{E}$ twice, one in each direction. We call a directed edge $\vec{e} = \{a, b\} \in \mathcal{E}$ *removable* if (1) removing the undirected edge e from E breaks the undirected graph G into two connected components, and (2) one of the two connected components that contains vertex a is acyclic. We call the acyclic component in (2), together with $\vec{e}$ (but not including vertex b), the *outer tree* of $\vec{e}$.

In the example of the insert, edge $\{a, b\}$ is removable, and the vertices and edges in its outer tree are highlighted in red. On the other hand, the directed edge in the reverse direction, $\{b, a\}$, is not removable, because the connected component in $G \setminus \{b, a\}$ containing b has a cycle. Note that neither $\{b, c\}$ nor $\{c, b\}$ is removable, because $G \setminus \{b, c\}$ remains connected.

Our key observation is that a topology-preserving subgraph of G can be only obtained by removing one or more outer trees from G . We say two outer trees are *disjoint* if they do not share a common vertex or edge. Note that two disjoint outer trees may be incident to the same vertex in G . For example, the outer trees of $\{a, b\}$ and $\{d, b\}$ in the insert above are disjoint, but they are both incident to vertex b . We will prove in Appendix A that,

Proposition 1 Let G' be a subgraph of G . Then $G' \simeq G$ if and only if $G \setminus G'$ consists of a set of mutually disjoint outer trees of G .




Figure 5: Left: OA on a discrete MA. Dotted gray lines connect the mid-point of each MA edge e to its nearest boundary samples $N(e)$. Middle: the complete pruning sequence including the thresholds α_i and pruned graphs M_i . Pictures on the top color each MA edge e by the sign of $f(e, \alpha)$ in Equation 7 (black/gray for positive/negative) for some α within each threshold range. Right: VA.

4.2. Incremental construction

Our algorithm constructs a pruning sequence and the resulting VA incrementally from the input graph M . Note that, if there exist multiple pruning sequences, the output of the algorithm is not necessarily the best in any sense of optimality. At each iteration, the algorithm finds the next graph M_{i+1} and threshold α_{i+1} from the previous graph M_i and threshold α_i .

Based on the characterization in Section 3.3, the next threshold, α_{i+1} , is the infimum of α at which the BSP of some topology-preserving proper subgraph of M_i exceeds that of M_i . We note the following identity for any α and subgraph M' of M_i ,

$$\text{BSP}(M_i, \alpha) - \text{BSP}(M', \alpha) = \text{BSP}(M_i \setminus M', \alpha) \quad (10)$$

That is, the difference in BSP between M_i and its subgraph M' is the BSP of their difference. Note that the rhs of Equation 10 strictly decreases with α if M' is a proper subgraph of M_i (hence $M_i \setminus M'$ has non-zero length). As a result, α_{i+1} is the smallest α such that $\text{BSP}(M_i \setminus M', \alpha) = 0$ for some topology-preserving proper subgraph M' . On the other hand, for such M' , Proposition 1 states that the difference $M_i \setminus M'$ is a set of outer trees of M_i . We conclude that α_{i+1} is the smallest α such that at least one outer tree of M_i has zero BSP.

To compute α_{i+1} , let $\alpha(M')$ be the value of α where $\text{BSP}(M', \alpha) = 0$ for a subgraph M' of M . We can derive $\alpha(M')$ from Equations 7 and 8 as

$$\alpha(M') = \arcsin\left(\frac{\sum_{e \in M'} |e| \sin(\text{OA}(e))}{\sum_{e \in M'} |e|}\right). \quad (11)$$

Let $\mathcal{T}(M_i)$ denote the set of all outer trees of M_i . Assuming the set is not empty, and following the argument above, we can obtain α_{i+1} as

$$\alpha_{i+1} = \min_{T \in \mathcal{T}(M_i)} \alpha(T). \quad (12)$$

Accordingly, M_{i+1} can be obtained by removing from M_i an outer tree $T \in \mathcal{T}(M_i)$ such that $\alpha(T) = \alpha_{i+1}$. We call such T a *vanishing outer tree*. Note that, if there are multiple vanishing outer trees, our construction results in a non-strictly increasing sequence of thresholds (e.g., $\alpha_{i+1} = \alpha_{i+2}$). While we could enforce the strictly increasing order of thresholds by removing as many (mutually disjoint) vanishing outer trees as possible in each iteration, the resulting VA will not be affected. In the case that $\mathcal{T}(M_i) = \emptyset$, which implies that M_i has no proper subgraph that preserves its topology, we complete the pruning sequence by setting $\alpha_{i+1} = \pi/2$ and $k = i$.

The pseudo-code of our algorithm is presented in Algorithm 1. The algorithm does not maintain the entire pruning sequence, but only the current threshold (as α) and pruned graph (as M'). The VA of graph edges are assigned (as the current threshold) when they are pruned.

4.3. Implementation and analysis

We make a remark on how to implement the algorithm efficiently. In particular, at each iteration of Algorithm 1, one needs to collect all removable edges R of the current graph M' and compute, for removable edge e , the quantity $\alpha(T_e)$ where T_e is the outer tree of e in M' . A brute-force computation would take quadratic time for these tasks. We next present a much faster (linear time) approach by observing that the removability and $\alpha(T_e)$ of an edge e can be inferred from those of its adjacent edges.

Specifically, consider a directed edge $e = \{a, b\}$ in M' , and let A_e be the set of remaining edges in M' incident to vertex a and oriented towards a . Denote the numerator and denominator inside the arcsin in the rhs of Equation 11 as $x(M')$ and $y(M')$ respectively (that is, $\alpha(M') = \arcsin(x(M')/y(M'))$). It is easy to show that:

- e is removable if either $A_e = \emptyset$, or every edge in A_e is removable.
- $x(T_e) = |e| \sin(\text{OA}(e)) + \sum_{e' \in A_e} x(T_{e'})$

Algorithm 1: Computing VA

Input: Voronoi-based discrete medial axis M
Output: VA of each edge $e \in M$
forall edge $e \in M$ **do**
 $VA(e) \leftarrow 0$
end
 $M' \leftarrow M$;
while True **do**
 $R \leftarrow$ all removable edges of M' ;
 if $R = \emptyset$ **then**
 forall edge $e \in M'$ **do**
 $VA(e) \leftarrow \pi/2$;
 end
 break
 end
 forall $e \in R$ **do**
 $\alpha_e \leftarrow \alpha(T_e)$; */
 end
 $\alpha \leftarrow \min_{e \in R} \alpha_e$;
 $e^* \leftarrow$ an edge in R such that $\alpha_{e^*} = \alpha$;
 forall edge $e \in T_{e^*}$ **do**
 $VA(e) \leftarrow \alpha$;
 end
 $M' \leftarrow M' \setminus T_{e^*}$;
end
return VA;

$$\bullet \ y(T_e) = |e| + \sum_{e' \in A_e} y(T_{e'})$$

Based on these relations, we can compute the removability and $\alpha(T_e)$ of all edges e in M' starting from edges incident to a degree-1 vertex (where $A_e = \emptyset$). The removability and quantities x, y are then propagated to other edges e once such information is available on all edges in A_e . Such a propagation takes time linear to the total number of edges of M' , as each (directed) edge is processed once.

We conclude with a complexity analysis of Algorithm 1. Consider an input graph M with n edges. Using the propagation approach described above, each iteration of the algorithm takes $O(n)$ time. Since each iteration (except for the last one) prunes at least one edge, the total number of iterations is capped by n . Hence the entire algorithm takes $O(n^2)$ time.

5. Results

We present more results of our method in this section. Unless otherwise stated, we use 50° as the pruning threshold for VA. All visualizations of VA in heat color use the same color range so that dark blue (resp. red) maps to VA value of 0 (resp. $\pi/2$).

To evaluate the sensitivity of VA to noise, we apply different types of synthetic distortions to the same Cat shape in Figure 6. Observe that, in each case, VA can distinguish noisy MA branches from those representing the major features of the Cat, producing qualitatively similar pruning results despite the variations in boundary distortions. Also note that the pruned MA exhibit little shrink- ing from their ends.




Figure 6: VA (middle column) and pruned MA at 50° (right column) of Cat after distorted by fine-scale noise (top row), coarse-scale noise (middle row), and pixelation (bottom row).

We demonstrate VA's robustness and scale-invariance in Figure 7. The input shape consists of the same bumpy starfish repeated at different scales and connected to each other. Existing global significance measures, such as ET and CR, are scale-dependent and have difficulty in differentiating between boundary bumps on the larger starfishes and the arms of the smaller starfishes. As a result, the smallest starfish loses its arms after pruning (see the green boxes) while some spurious branches still remain on the largest starfish (see the blue boxes). In contrast, VA highlights the arms of starfishes at different scales equally well and distinct from the bumps, resulting in a pruned MA that fully captures the arms of all starfishes.

We next compare with the alternative approach for regularizing MA, namely smoothing the shape's boundary. A notable method following this approach is the Scale Axis Transform (SAT) [GMPW09], which computes the MA of an inflated shape consisting of the maximal balls centered at the MA points of the original shape after their radii are multiplied by a scale parameter. Like other boundary-smoothing methods, the inflation in SAT may adversely alter the original shape so that the resulting MA no longer preserves the topology or structure. As demonstrated in Figure 8, as the scale parameter increases, inflation connects two nearby parts of Omega (top arrow), thereby creating a loop in the MA, and merges two hind legs of Giraffe into one (bottom arrow). In both cases, the pruned MA also has parts protruding outside the original shape. In contrast, the pruned MA produced by our method always lies inside the shape while preserving its topology and structure.

Finally, we show VA and the pruned MAs on a collection of pla-




Figure 7: Comparing ET, CR, VA, and the pruned MA using each measure on a cluster of “bumpy” starfishes. The pruning threshold for VA is 50° . The pruning thresholds for ET and CR are chosen to achieve the best balance between preserving features of smaller starfishes (green boxes) and removing noise on larger starfishes (blue boxes).

nar shapes in Figure 9. Note that the pruned MAs well capture the salient shape features while ignoring small boundary irregularities. In addition, the extremities of MA in the salient features are well protected from erosion (e.g., the branches of the tree and octopus).

6. Discussion

We present a new method for pruning the medial axes (MA) of planar shapes. Using a variational formulation, our new significance measure (Vanishing Angle) combines the extremity-awareness and scale-invariance of local measures (e.g., Object Angle) with the noise resistance of global measures (e.g., Chord Residue and Erosion Thickness).

While VA is particularly suited for tubular shape parts, where MA has high OA, it can have difficulty in capturing more rounded parts where OA is low. In the example of Figure 10, the round-shaped flower has much lower OA (and hence VA) than the tube-shaped stem, resulting in the loss of the entire flower in the pruned MA at the default threshold 50° . A work-around is to lower the pruning threshold, which comes at the risk of possibly including spurious branches representing noisy features. In the case of Figure 10, reducing the threshold by 5 degrees (to 45°) is sufficient to preserve the flower without introducing unwanted MA branches.

Another limitation of VA is that its maxima may not be located at the natural “center” of the shape. By construction, our pruning sequence is attracted to regions of MA with high-OA points. As a result, the maxima of VA often lies in off-centered, but highly tubular shape parts (e.g., a finger of the Hand or an arm of the Dancer in Figure 9). Also, as the pruning sequence for a given MA may not be unique, the choice of which tubular part that the VA maxima is found can be arbitrary if several similar tubular parts exist (e.g., arms of the Octopus in Figure 9).

Our formulations and algorithm can be directly applied to prun-

ing a curve skeleton of a 3D shape. The only requirement is that each skeleton edge must be associated with an object angle. It is also possible to extend our BSP definition to measure the degree of parallelism between a subset of the two-dimensional MA of a 3D shape and the boundary surface. However, computing the pruning sequence and VA measures on MA in 3D (which is a non-manifold network of sheets) will be more challenging due to its lack of a natural graph structure. Our work also opens up several theoretical questions about the VA measure that we will investigate, such as its convergence under increased boundary sampling density and its stability under boundary perturbations.

Acknowledgements

This work is supported in part by NSF grant EF-1921728. We would also like to thank the anonymous reviewers for their suggestions.

References

- [ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust. In *SMA '01: Proceedings of the sixth ACM symposium on Solid modeling and applications* (2001), pp. 249–266. 2
- [AdBT95] ATTALI D., DI BAJA G. S., THIEL E.: Pruning discrete and semicontinuous skeletons. In *Image Analysis and Processing, 8th International Conference, ICIAP '95, San Remo, Italy, September 13-15, 1995, Proceedings* (1995), pp. 488–493. 3
- [AM96] ATTALI D., MONTANVERT A.: Modeling noise for a better simplification of skeletons. In *Proceedings 1996 International Conference on Image Processing, Lausanne, Switzerland, September 16-19, 1996* (1996), pp. 13–16. 2
- [AM97] ATTALI D., MONTANVERT A.: Computing and simplifying 2d and 3d continuous skeletons. *Computer vision and image understanding* 67, 3 (1997), 261–273. 5
- [BA92] BRANDT J. W., ALGAZI V. R.: Continuous skeleton computation by voronoi diagram. *CVGIP: Image understanding* 55, 3 (1992), 329–338. 3, 5




Figure 8: Comparing Scale Axis Transform (SAT) [GMPW09] at various scaling parameters (s) with our method (showing VA and the pruned MA at different VA thresholds) on Omega (top) and Giraffe (bottom). The gray regions in SAT are the inflated shapes. Arrows point to where SAT alters the shape's topology (in Omega) or structure (in Giraffe, merging two legs into one).

[Blu67] BLUM H.: A transformation for extracting new descriptors of form. *Models for the Perception of Speech and Visual Form* (1967), 362–80. 1

[Blu73] BLUM H.: Biological shape and visual science (part i). *Journal of Theoretical Biology* 38, 2 (1973), 205–287. 3

[CCT09] CHAUSSARD J., COUPRIE M., TALBOT H.: A discrete lambda-medial axis. In *DGCI* (2009), vol. 5810, pp. 421–433. 3

[CL04] CHAZAL F., LIEUTIER A.: Stability and homotopy of a subset of the medial axis. In *Proceedings of the Ninth ACM Symposium on Solid Modeling and Applications* (2004), SM '04, pp. 243–248. 3

[DLN87] DILL A. R., LEVINE M. D., NOBLE P. B.: Multiple resolution skeletons. *IEEE Trans. Pattern Anal. Mach. Intell.* 9, 4 (1987), 495–504. 2

[DLX*22] DOU Z., LIN C., XU R., YANG L., XIN S., KOMURA T., WANG W.: Coverage axis: Inner point selection for 3d shape skeletonization. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 419–432. 3

[DPS00] DIMITROV P., PHILLIPS C., SIDDIQI K.: Robust and efficient skeletal graphs. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)* (2000), vol. 1, IEEE, pp. 417–423. 3

[DS06] DEY T. K., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *SGP '06: Proceedings of the*

fourth Eurographics symposium on Geometry processing (Aire-la-Ville, Switzerland, Switzerland, 2006), Eurographics Association, pp. 143–152. 3

[DZ04] DEY T. K., ZHAO W.: Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design* 36, 2 (2004), 195–202. 2

[FLM03] FOSKEY M., LIN M. C., MANOCHA D.: Efficient computation of A simplified medial axis. *J. Comput. Inf. Sci. Eng.* 3, 4 (2003), 274–284. 2

[FTB13] FARAJ N., THIERY J.-M., BOUBEKEUR T.: Progressive medial axis filtration. In *SIGGRAPH Asia 2013 Technical Briefs*. 2013, pp. 1–4. 3

[GEG00] GOLLAND P., ERIC W., GRIMSON L.: Fixed topology skeletons. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on* (2000), vol. 1, pp. 10–17 vol.1. 2

[GMPW09] GIESEN J., MIKLOS B., PAULY M., WORMSER C.: The scale axis transform. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry* (2009), SCG '09, pp. 106–115. 2, 7, 9

[Hat02] HATCHER A.: *Algebraic topology*. Cambridge University Press, 2002. 10

[HD86] HO S.-B., DYER C. R.: Shape smoothing using medial axis properties. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8, 4 (1986), 512–520. 3

[LCLJ11] LIU L., CHAMBERS E. W., LETSCHER D., JU T.: Extended grassfire transform on medial axes of 2d shapes. *Computer-Aided Design* 43, 11 (2011), 1496–1505. 3

[Lie03] LIEUTIER A.: Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. In *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications* (2003), SM '03, pp. 65–75. 1

[LWS*15] LI P., WANG B., SUN F., GUO X., ZHANG C., WANG W.: Q-mat: Computing medial axis transform by quadratic error minimization. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 1–16. 3

[MGP10] MIKLOS B., GIESEN J., PAULY M.: Discrete scale axis representations for 3d geometry. *ACM Trans. Graph.* 29, 4 (July 2010), 101:1–101:10. 2

[Mik10] MIKLÓS B.: *The Scale Axis Transform*. PhD thesis, ETH Zurich, 2010.

[MLIM16] MARIE R., LABBANI-IGBIDA O., MOUADDIB E. M.: The delta medial axis: a fast and robust algorithm for filtered skeleton extraction. *Pattern Recognition* 56 (2016), 26–39. 3

[OI92] OGNIWICZ R., ILG M.: Voronoi skeletons: theory and applications. In *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR '92., 1992 IEEE Computer Society Conference on* (1992), pp. 63–69. 3, 5

[OK95] OGNIWICZ R. L., KÜBLER O.: Hierarchic voronoi skeletons. *Pattern Recognition* 28, 3 (1995), 343–359. 3

[PFJ*03] PIZER S. M., FLETCHER P. T., JOSHI S. C., THALL A., CHEN J. Z., FRIDMAN Y., FRITSCH D. S., GASH A. G., GLOTZER J. M., JIROUTEK M. R., LU C., MULLER K. E., TRACTON G., YUSHKEVICH P. A., CHANEY E. L.: Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision* 55, 2-3 (2003), 85–106. 2

[PGJA03] PIZER S. M., GERIG G., JOSHI S. C., AYLWARD S. R.: Multiscale medial shape-based analysis of image objects. *Proceedings of the IEEE* 91, 10 (2003), 1670–1679. 2

[POB87] PIZER S. M., OLIVER W. R., BLOOMBERG S. H.: Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE Trans. Pattern Anal. Mach. Intell.* 9, 4 (1987), 505–511. 2

[RvWT08] RENIERS D., VAN WIJK J., TELEA A.: Computing multi-scale curve and surface skeletons of genus 0 shapes using a global importance measure. *IEEE Transactions on Visualization and Computer*




Figure 9: VA and pruned MA (at 50°) on a gallery of shapes.




Figure 10: VA can be low in rounded parts of the shape (e.g., the flower), and a lower threshold (45° in this case) is needed to preserve such features.

Graphics 14, 2 (2008), 355–368. doi:<http://dx.doi.org/10.1109/TVCG.2008.23>. 3

[SB98] SHAKED D., BRUCKSTEIN A. M.: Pruning medial axes. *Comput. Vis. Image Underst.* 69, 2 (1998), 156–169. 1, 2, 3

[SBTZ02] SIDDIQI K., BOUX S., TANNENBAUM A., ZUCKER S. W.: Hamilton-jacobi skeletons. *International Journal of Computer Vision* 48, 3 (July 2002), 215–231. 2, 3

[SFM05] SUD A., FOSKEY M., MANOCHA D.: Homotopy-preserving medial axis simplification. In *SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling* (New York, NY, USA, 2005), ACM, pp. 39–50. doi:<http://doi.acm.org/10.1145/1060244.1060250>. 2

[SM13] SCHMITT M., MATTIOLI J.: *Morphologie mathématique*. Presses des MINES, 2013. 5

[SP08] SIDDIQI K., PIZER S. M.: *Medial Representations*. Springer, 2008. 1, 2

[TDS*16] TAGLIASACCHI A., DELAME T., SPAGNUOLO M., AMENTA N., TELEA A.: 3d skeletons: A state-of-the-art report. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 573–597. 2

[YSC*16] YAN Y., SYKES K., CHAMBERS E., LETSCHER D., JU T.: Erosion thickness on medial axes of 3d shapes. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–12. 3

Appendix A: Proof of Proposition 1

Proof We shall refer to a set of mutually disjoint outer trees of G as an *outer forest* of G . Any homotopy equivalent subgraph of G can be obtained from G via a deformation retract [Hat02], which is a sequence of reductions that each removes a degree-1 vertex together with its only incident edge. We need to show that (i) the union of vertices and edges removed by any deformation retract is an outer forest, and (ii) any outer forest can be removed by some deformation retract.

We first show (i) by induction on the number of reductions. If G' is obtained by one reduction from G , the removed degree-1 vertex and its incident edge forms an outer tree. Suppose (i) holds for any deformation retract consisting of m reductions ($m \geq 1$). Consider a G' obtained from G via $m+1$ reductions. Let G^* be the intermediate result after the first m reductions. By induction hypothesis, $F = G \setminus G^*$ is an outer forest. Suppose the $(m+1)$ -th reduction removes a degree-1 vertex $a \in G^*$ and its only incident edge $\{a, b\}$. We now consider two cases. First, if no outer tree in F is incident to a , then a has degree 1 in G as well, and the vertex a and edge $\{a, b\}$ form an outer tree of G disjoint from all outer trees in F . Second, suppose one or more outer trees in F are incident to a , and denote them by F_a . Since the outer trees in F_a are mutually disjoint and also disjoint from outer trees in $F \setminus F_a$, the union of F_a remains acyclic

and disconnected from $G \setminus F_a$, which implies that edge $\{a, b\}$ is removable. The union of F_a with vertex a and edge $\{a, b\}$ forms a single outer tree that is disjoint from other outer trees in $F \setminus F_a$. In both cases, the union of F with vertex a and edge $\{a, b\}$ remains an outer forest.

To show (ii), observe that each outer tree can be removed by a sequence of reductions, each removing a leaf vertex with its parent edge. Since the outer trees are mutually disjoint in an outer forest, each tree in the forest can be removed independently, and hence the entire forest can be removed by a deformation retract. $\square$