

Erosion Thickness on Medial Axes of 3D Shapes

Yajie Yan* Kyle Sykes Erin Chambers David Letscher Tao Ju
 Washington Univ. in St. Louis St. Louis Univ. St. Louis Univ. St. Louis Univ. Washington Univ. in St. Louis

Abstract

While playing a fundamental role in shape understanding, the medial axis is known to be sensitive to small boundary perturbations. Methods for pruning the medial axis are usually guided by some measure of significance. The majority of significance measures over the medial axes of 3D shapes are locally defined and hence unable to capture the scale of features. We introduce a global significance measure that generalizes in 3D the classical Erosion Thickness (ET) measure over the medial axes of 2D shapes. We give precise definition of ET in 3D, analyze its properties, and present an efficient approximation algorithm with bounded error on a piecewise linear medial axis. Experiments showed that ET outperforms local measures in differentiating small boundary noise from prominent shape features, and it is significantly faster to compute than existing global measures. We demonstrate the utility of ET in extracting clean, shape-revealing and topology-preserving skeletons of 3D shapes.

Keywords: Medial axis, skeletons, shape analysis

Concepts: •Computing methodologies → Shape analysis;
 •Theory of computation → Computational geometry;

1 Introduction

Since its introduction by Blum [1967], the *medial axis* has become the basis for many shape descriptors. Simply defined as the set of points with more than one nearest neighbors on the shape boundary, the medial axis has many desirable properties for shape description: it has a lower dimension than the shape, captures components and protrusions, and preserves the homotopy of the shape [Lieutier 2003].

A notable issue of medial axis that prevents its wider adoption is its sensitivity to boundary perturbations: a bump on the shape boundary, no matter how tiny, is captured by some branch of the medial axis (Figure 1). Extensive research has been conducted on taming such sensitivity (see Section 2). The majority of these methods works by pruning branches that arise from boundary noise. These methods are usually guided by some form of *significance measure* that rates the importance of the shape feature represented by each medial axis point.

A variety of significance measures have been formulated on the graph-like medial axis of a 2D shape [Shaked and Bruckstein 1998], and they generally fall into two categories. While *local* measures are based on the immediate geometry around a medial axis point,

*e-mail:yajieyan@wustl.edu

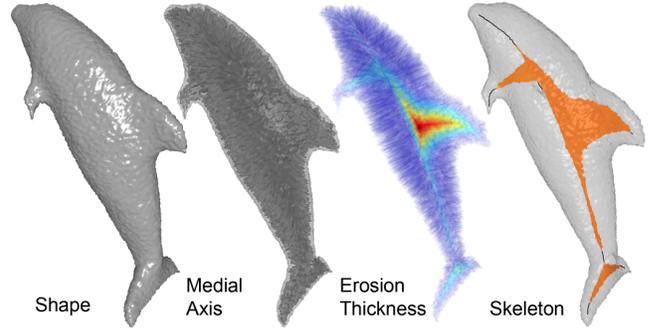


Figure 1: The medial axis of a bumpy 3D shape contains numerous spurious sheets. Erosion Thickness (ET) highlights parts of the medial axis that represent significant shape features. We use ET to create a family of clean, topology-preserving skeletons made up of both 2D sheets and 1D curves.

global measures consider shape in a larger neighborhood. The extended context allows global measures to outperform local measures in differentiating boundary noise from major shape features. A popular global measure is the Erosion Thickness (ET), which captures the amount of shape loss due to the erosion of the medial axis [Ho and Dyer 1986; Brandt and Algazi 1992; Niblack et al. 1992; Arcelli and di Baja 1993; Attali et al. 1995] (see a precise definition in Section 3).

The complexity of the medial axis of a 3D shape makes it much more challenging to formulate global significance measures than in the 2D case. While many pruning heuristics are based on principles similar to ET, they do not lead to any continuously defined measures over the medial axis. The only well-defined global measure that we are aware of is the Medial Geodesic Function (MGF) [Dey and Sun 2006], which measures the geodesic distance between the two nearest neighbors of a medial axis point on the shape boundary. However, computing MGF is expensive due to the need to query geodesic distances between many pairs of surface points.

In this paper, we define, analyze and compute Erosion Thickness on the medial axis of a 3D shape. ET in 2D can be defined as the difference between the arrival time of a fire front propagating over the medial axis and the radii of maximal balls [Liu et al. 2011]. We follow the same idea and introduce a burning process over the medial axis in 3D. We show that several key properties of ET in 2D, such as finiteness, continuity, and lack of local minima, extend nicely to our definition of ET in 3D. Guided by the burning process, we develop an efficient approximation algorithm with bounded error on a piecewise linear medial axis. Our experiments showed that ET in 3D, like its counterpart in 2D, can effectively differentiate small boundary noise from prominent shape features (Figure 1). On the other hand, it is orders of magnitude faster to compute than MGF.

We demonstrate the utility of ET by developing a skeletonization algorithm guided by ET. The skeletons are made up of both 2-dimensional sheets of the medial axis with high ET values and 1-dimensional curves that follow the “ridges” of ET. Our method allows a user to control the pruning of the sheets and curves independently using two intuitive parameters. The skeletons preserve the topology of the 3D shape and are observed to be stable under boundary perturbations (Figure 1).

Contributions: To our best knowledge, this work introduces the first well-defined and efficient-to-compute global significance measure over the medial axis of a 3D shape. It makes the following theoretic and algorithmic contributions:

1. Formalizing Erosion Thickness over the medial axis of a 3D shape (Section 4) and analyzing its theoretical properties (Section 5).
2. Developing an efficient algorithm for approximating ET over a piece-wise linear medial axis and proving its error bound (Section 6).
3. Developing a skeletonization method guided by ET that produces stable, topology-preserving, and shape-depicting skeletons (Section 7).

2 Related works

We briefly review works on computing and simplifying medial axes. Please refer to the excellent book by Siddiqi and Pizer [2008] for more extensive and in-depth discussions.

Computing medial axes Algorithms for computing exact medial axes in 3D are available for only limited classes of shapes such as small polyhedra [Culver et al. 2004] and unions of balls [Amenta and Kolluri 2001]. Approximation strategies include thinning on voxelized models [Bertrand 1995], locating singularities of distance fields [Siddiqi et al. 2002], and computing a subset of the Voronoi diagram of boundary samples [Amenta et al. 2001; Dey and Zhao 2004]. Many of these methods produce a piece-wise linear representation (i.e., a mesh), on which our algorithms can be directly applied.

Simplifying medial axes There are three general approaches to medial axes simplification. The first approach reduces the complexity of the medial axis by *smoothing* the shape boundary before extracting the medial axis [Dill et al. 1987; Pizer et al. 1987; Giesen et al. 2009; Miklos et al. 2010]. However, the medial axis of the smoothed boundary may exhibit a different structure, and even topology, from the original shape [Shaked and Bruckstein 1998]. The second approach, pioneered by Golland et al. [2000] and Pizer et al. [2003a; 2003b], is to deform a noise-free, *template* medial axis into the shape. The deformation approach was found particularly useful in analyzing biological shapes [Fletcher et al. 2004], but it is limited to shapes with known structures or a group of similar shapes. The third approach, which we adopt, is directly *pruning* medial axis components arising from boundary noise. Guided by a significance measure, pruning methods often employ a contraction scheme [Siddiqi et al. 2002; Tam and Heidrich 2003; Sud et al. 2005; Liu et al. 2010; Li et al. 2015] to ensure that the pruned subset retains the topology of the medial axis.

Significance measures Most commonly, the significance of a medial axis point is measured by the *local* configuration of a medial axis point and its nearest neighbors on the shape’s boundary. For example, *object angle* measures the angle spanned by the vectors from a medial axis point to its two nearest neighbors on the boundary [Attali and Montanvert 1996; Amenta et al. 2001; Dey and Zhao 2004; Foskey et al. 2003; Sud et al. 2005]. Measures related to object angle include the propagation velocity of [Blum 1973], the outward flux of [Siddiqi et al. 2002], and the stability ratio of [Li et al. 2015]. Another measure considers the *circumradius* of the nearest neighbors on the boundary [Chazal and Lieutier 2004; Chaussard et al. 2009]. The subset of the medial axis where the circumradius is above a constant λ (known as the λ -medial axis) is provably stable under certain class of boundary perturbations.

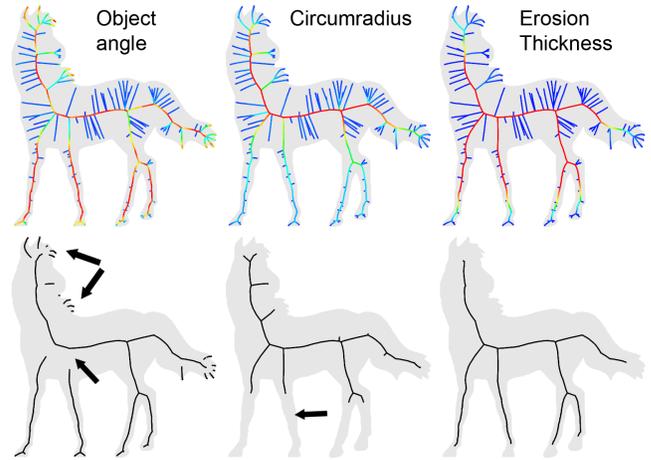


Figure 2: Top: Medial axes of the same 2D shape colored (blue for low and red for high) by the object angle measure, the circumradius measure, and Erosion Thickness. Bottom: Subsets of the medial axes where the respective measure is higher than some threshold. Local measures, such as object angle and circumradius, fail to differentiate major shape features (e.g., the horse legs) from boundary noise (e.g., hair).

The main drawback of local measures is that they cannot capture the size of features (Figure 2). Being scale-invariant, object angle is high both within prominent features and inside small boundary bumps (e.g., horse hair). Object angle is also low in the transition area between shape parts (e.g., between horse leg and body). On the other hand, the circumradius only captures thick parts of the shape and easily misses long but thin features (e.g., horse leg).

Various measures of *global* significance were proposed on medial axes of 2D shapes [Shaked and Bruckstein 1998]. A common scheme is measuring the amount of shape loss as a result of pruning the medial axis. The Erosion Thickness (ET), used by various authors [Ho and Dyer 1986; Brandt and Algazi 1992; Niblack et al. 1992; Arcelli and di Baja 1993; Attali et al. 1995; Shaked and Bruckstein 1998; Liu et al. 2011], measures the distance by which a protrusion shrinks after a medial axis branch is shortened (a precise definition will be reviewed in Section 3). Since ET captures the length of a protrusion, it outperforms local measures in differentiating boundary noise from long but thin features (Figure 2 top-right). ET also has several desirable properties. In particular, ET has no local minima on the medial axis. As a result, thresholding ET never disconnects the medial axis (Figure 2 bottom-right). Other global measures in 2D that capture shape loss include the area of erosion [Shaked and Bruckstein 1998; Attali et al. 1995] and the Potential Residue [Ogniewicz and Ilg 1992; Ogniewicz and Kübler 1995].

Global significance measures on the medial axes of 3D shapes are rare. Several researchers followed the general idea of capturing shape loss and developed heuristics for simplifying medial axes [Styner et al. 2003; Tam and Heidrich 2003; Liu et al. 2010; Li et al. 2015], although they do not lead to well-defined measures over the medial axis. The only well-defined global measure that we know of is the Medial Geodesic Function (MGF) proposed by Dey and Sun [2006] ((a variant was considered by Reniers et al. [2008])). Given a point x on the medial axis, MGF is defined as the geodesic distance between the two generators of x on the shape’s boundary surface. MGF extends the Potential Residue measure in 2D [Ogniewicz and Ilg 1992]. While enjoying a simple and elegant definition, MGF is expensive to compute: evaluating MGF at each medial axis point involves a geodesic distance calculation on the shape boundary. In

addition, MGF can be sensitive to perturbations that significantly affect the geodesic distances.

We introduce a new well-defined global measure in 3D by extending ET. Compared with MGF, ET’s definition is more involved, but its computation is orders of magnitude faster. In addition, our experiments show that ET is stable under geodesic perturbations.

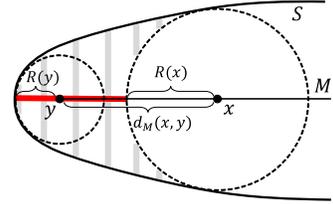
Medial curves Applications such as character animation and path finding make use of 1-dimensional curves centered within the shape, or *medial curves*. Current methods for computing medial curves mostly resort to heuristics, such as mesh contraction and projection (see surveys by Cornea et al. [2007] and more recently by Tagliasacchi et al. [2016]). These methods usually have many parameters to tune, and the computed curves lack a clear mathematical formulation. Despite many efforts, a definition of the medial curve that inherits the essential properties of the medial axis (e.g., being thin and topology-preserving) is still at large. The best attempts at this goal are perhaps the methods of Dey and Sun [2006] and Reniers et al. [2008], who compute medial curves that approximately capture the singularities of MGF. Similarly, our medial curves follow the singularities of MGF. We then combine the pruned medial curves with the pruned medial axis to produce a family of curve and surface skeletons.

3 Defining Erosion Thickness in 2D

We first review the definition of Erosion Thickness (ET) over the medial axis of a 2D shape. Consider an open set S of \mathbb{R}^2 . Its medial axis M , in general, is a graph made up of smooth curve segments meeting at *junctions*. Each point x on M is the center of a maximal disk in S with radius $R(x)$.

As mentioned earlier, ET captures how much the shape has shrunk after eroding the medial axis from its end points. Here, the shrunken shape refers to the union of the maximal balls centered on the eroded medial axis.

As illustrated in this picture, eroding away a medial axis segment yx where y is an end point of M results in the loss of the shaded portion of S . ET at x measures the “depth” or “thickness” of this lost portion as the length of the red curve,



$R(y) + d_M(x, y) - R(x)$, where d_M is the geodesic distance on M ¹. Note that ET is small if y lies within a small surface bump, in which case the lost portion would be thin.

To evaluate ET at an arbitrary location on the medial axis, earlier methods resort to a pruning heuristic that alternates between evaluating ET on peripheral branches and pruning the least significant branch [Shaked and Bruckstein 1998]. A precise definition was given by Liu et al. [2011] as the *the difference between the burn time of a fire moving along the medial axis and the maximal ball radii*. Imagine that the curves of M are made up of a uniform combustible material, and fire is ignited from each end point y at time $R(y)$. The fire moves at constant, unit speed along M , burning away points as the fronts pass them. When a fire front comes to a junction point, it *dies out* unless all but one branch at the junction have been burned away. In the latter case, the front continues burning along that remaining branch.

¹The form is similar to the Hyperbolic distance between two balls [Choi and Seidel 2001], where the Euclidean distance between x, y is used instead of the geodesic distance $d_M(x, y)$.

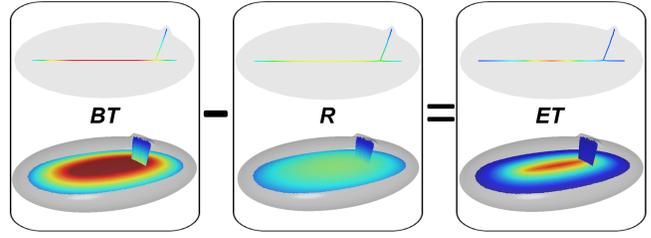


Figure 3: The burn time (BT , left), maximal ball radii (R , middle), and Erosion thickness (ET , right) in a 2D ellipse (top) and a 3D ellipsoid (bottom) with small bumps.

The time at which a point $x \in M$ is burned away (which we call the *burn time function*, or BT) can be obtained by

$$BT(x) = \sup_{yz} \min(R(y) + d_M(x, y), R(z) + d_M(x, z)). \quad (1)$$

where the supremum is taken over all curve segments yz of M that contain x . ET is then defined by

$$ET(x) = BT(x) - R(x). \quad (2)$$

The three functions, BT , R , and ET , for an ellipse shape with a small bump are shown in Figure 3 (top). Note that while both BT and R are non-trivial on the medial axis branch that reaches out to the bump, their difference, ET , is small.

The definition reveals an intuitive connection between ET and *tubularity* of shape. While the burn time measures the lateral distance to the shape boundary along the medial axis, R measures the vertical distance to the shape boundary. As a result, high ET values appear inside long and thin “tubes”.

Many properties of ET can be explained using the burning analogy. Since burning is continuous, so is its burn time and, in turn, ET. As fire never “pierces” through the middle of a segment (due to dying-out at junctions), the burn time is free of local minima. This claim can be extended to ET by noting that the gradient magnitude of R on M never exceeds that of burn time (which is constant 1). Burning also lends naturally to an iterative algorithm for computing the burn time and ET on the medial axis graph [Liu et al. 2011].

4 Defining Erosion Thickness in 3D

As in 2D, Erosion Thickness in 3D captures the shrinkage of the 3D shape when the medial axis is eroded. Our definition closely mimics that of Liu et al. [2011] in 2D: *ET in 3D is defined as the difference between the burn time of a fire over the medial axis and the maximal ball radii* (see an illustration in Figure 3 bottom). While ET in 2D reflects the *tubularity* of shape, ET in 3D captures the *plate-likeness*. High ET values arise at the center of wide and thin “plates”, where the difference between the lateral distance (i.e., burn time) and the vertical distance (i.e., maximal ball radius) to the shape boundary is large.

To define ET in 3D, the key challenge is therefore to formulate burning on the complex structure of the medial axis of a 3D shape. We first review the structure of a 3D medial axis, and then introduce the burning process before giving a precise definition of burn time (and in turn ET). As we will show in the next section, our definition of ET in 3D is able to retain key properties of ET in 2D, such as continuity and lack of local minima (see Section 5).

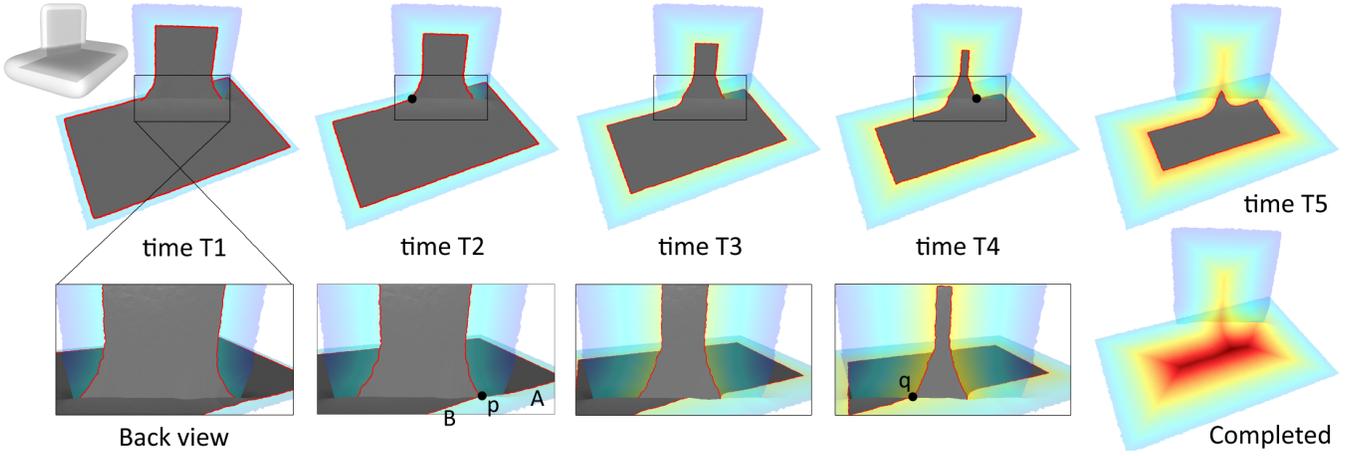


Figure 4: Burning on the medial axis of a 3D shape (shown in the top-left) at five time points and at completion. Fire fronts are indicated by red curves, unburnt portion of the medial axis is colored gray, and the burned portion is colored by their burn time and rendered in transparency.

4.1 Preliminaries

We consider a shape S as a bounded open subset of \mathbb{R}^3 , and denote ∂S as its boundary. Define $M = \{x \in S \mid \|B(x)\| > 1\}$, where $B(x)$ is set of nearest neighbors of x on ∂S . We consider the closure of M , $M = \overline{M}$, as the medial axis. The distance from $x \in M$ to ∂S is denoted by $R(x)$.

A point x on M is *manifold* if its local neighborhood on M is homeomorphic to a disk or half-disk, otherwise it is *singular*. x is a *boundary point* if there is no disk on M that contains x , otherwise it is an *interior point*. Denote the set of all manifold points, singular points, and boundary points as $M_2, M_S, \partial M$. These points are illustrated in Figure 5. In the picture on the left, points x_1, x_2 are manifold while others are singular. Points x_2, x_6 are boundary points while others are interior points. Note that a boundary point can be either manifold (e.g., x_2) or singular (e.g., x_6).

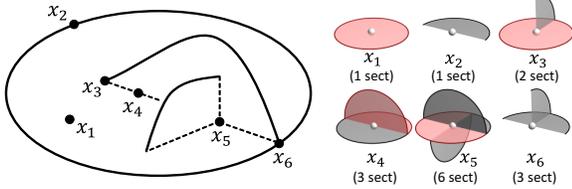


Figure 5: Left: illustration of a medial axis with three sheets. Right: local topology of 6 locations on the medial axis, each showing the number of sectors and a possible exposing set (red).

A connected component of the manifold set M_2 is called a *sheet*. The singular set M_S forms a graph whose nodes are called *junctions* and edges are called *seams*. Intuitively, a seam is where multiple sheets meet, and a junction is either an end of a seam or where multiple seams meet.

We will assume that M is finite and generic. The finiteness can be ensured if S is defined by a finite set of analytic equations and inequalities [Chazal and Soufflet 2003]. By [Giblin and Kimia 2004], any point on a generic medial axis has one of the five local topologies depicted by x_1, \dots, x_5 in Figure 5. The local topology of x_6 only arises during burning, as we shall see below (e.g., point p in Figure 4).

4.2 Burning the medial axis

Imagine that the sheets of M are made up of a uniform combustible material. Fire is ignited at every point $y \in \partial M$ at time $R(y)$ and propagates geodesically on M at constant, unit speed. Recall that, in 2D, the fire front dies out at a junction point unless there is only one remaining branch at the junction. To generalize this rule to 3D, we ask that the fire burns away a point x only if x is on the boundary of the unburned portion of M . For example, the fire front would die out as it reaches a seam where two or more sheets have not yet been burned.

The burning process is illustrated in Figure 4. As shown in the top-left, S consists of a “fin” on top of a thicker and wider “board”. M is made up of two sheets, called respectively the “fin sheet” and “board sheet”, that meet along a seam. Burning starts earlier along the boundary of the fin sheet than the board sheet, because the former has smaller maximal balls. At the beginning, fire on the fin sheet dies out as it reaches the seam, because the board sheet has not been burned away (see time T1). An interesting event takes place at time T2, when the fire front on the board sheet “catches up” with that on the fin sheet. Note that the local topology of the point where the fronts meet (p) in the unburned portion of M is the same as x_6 in Figure 5. Thereafter, the front splits into two segments, one (“A”) merged with the fire front on the fin sheet and the other (“B”) moving behind the fin sheet and dying out on the seam (see time T3). Later on, at time T4, the latter segment (“B”) merges with the other fire front on the fin sheet (at point q).

4.3 Defining burn time

To define burn time, we draw inspiration from geodesic distances. Note that burn time agrees exactly with the geodesic distance to the boundary if M has no singular points and R is zero everywhere. Recall that the geodesic distance is defined by the length of the shortest path. Similarly, we define burn time as the length of the shortest *tree* of paths that branch at the singular points.

Motivation To motivate our definition, we assume for now that R is zero everywhere on ∂M . That is, burning starts simultaneously along the boundary. The burn time of a point $x \in M$ is thus the length of the shortest path along which burning can proceed continuously (i.e., without dying out) from the boundary. If such path, say γ , passes a point z on a seam, the other sheet at z must have been burned before the fire front on γ reaches z . This implies that there is

some path γ' on that other sheet connecting z to the boundary that is shorter than the segment of γ between z and the boundary (see illustration in Figure 6 left). Note that if z has a more complicated local topology (e.g., x_5 in Figure 5), there should be enough number of paths like γ' at z such that removing these paths, together with the segment of γ between z and the boundary, from M would “expose” z as a boundary point.

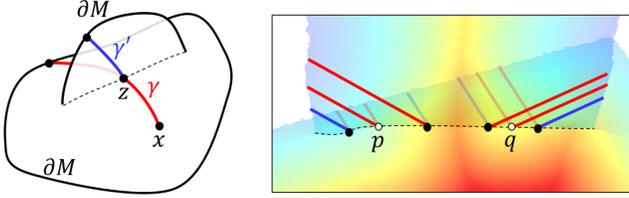


Figure 6: Left: An illustrative exposing tree at x . Right: Exposing trees that realize burn time at different locations on the seam of the medial axis of Figure 4. The primary path in each tree is colored red, and the other tree edges are colored blue.

The path γ' , in turn, may spawn more paths to the boundary as it passes other singular points. If we recursively collect all paths spawned by γ and its spawned paths, we get a tree of paths to the boundary that is rooted at x and branches at singular points. We call this tree an *exposing tree*, because each vertex of the tree is exposed as a boundary point if all its child branches are removed from M . Note that γ is the longest root-to-leaf path in the tree. To define burn time, we explore all exposing trees rooted at x and seek the one whose longest rooted path is shortest.

Formal definition We start by formalizing the notion of “exposure”, which is key to defining exposing trees. Given $x \in M$, we consider the regular neighborhood of x on M , denoted as $N(x)$. We call each component of $N(x) \cap M_2$ a *sector*. Intuitively, a sector corresponds to a (local) sheet that contains x . While a manifold point has a single sector, a singular point has multiple sectors. Figure 5 (right) plots the regular neighborhood and note the number of sectors for each sample point.

We say x is exposed by sectors $\{s_1, \dots, s_k\}$ if removing them from $N(x)$ leaves no complete disks in $N(x)$ (and hence x becomes a boundary point). More precisely,

Definition 4.1 $x \in M$ is exposed by sectors $\{s_1, \dots, s_k\}$ if there is no 2-dimensional disk $D \subseteq N(x)$ such that $x \in D$, $\partial D \subseteq \partial N(x)$, and $s_i \cap D = \emptyset, \forall i = 1, \dots, k$. The set $\{s_1, \dots, s_k\}$ is called an exposing set of x .

Examples of exposing sets are highlighted in Figure 5 (right). Note that a singular point can have more than one exposing sets. Consider point x_4 : any two of x_4 ’s three sectors form an exposing set, and so is the set of all three sectors. An empty set can also be an exposing set if the point is already on ∂M .

Alternatively, exposing sets can be defined by considering the boundary of the regular neighborhood, $\partial N(x)$. This boundary is a graph where each edge corresponds to a sector of x . If we remove from this graph those edges corresponding to an exposing set, the remaining subgraph is free of cycles.

Next we introduce exposing trees, their length, and burn time.

Definition 4.2 An image of a finite tree on M is called an exposing tree of $x \in M$ if it is rooted at x , each edge lies on M_2 , and each vertex is exposed by the sectors that contain the child edges of that vertex.

This definition naturally implies that each leaf vertex of Γ must lie on ∂M , the only place where points are exposed by empty sets. An exposing tree can consist of a single point (if x is on ∂M), a single path, or a collection of paths.

Definition 4.3 Given an exposing tree Γ at $x \in M$, a rooted path from x to a leaf vertex $y \in \partial \Gamma$ is called the primary path if

$$y = \arg \max_{y' \in \partial \Gamma} (R(y') + d_\Gamma(x, y')), \quad (3)$$

where d_Γ measures path distance on Γ . The value $R(y) + d_\Gamma(x, y)$ is called the length of Γ , denoted as $L(\Gamma)$.

Definition 4.4 Let \mathcal{T}_x be the set of all exposing trees at $x \in M$. The burn time function, $BT : M \rightarrow \mathbb{R} \cup \{\infty\}$ is

$$BT(x) = \inf_{\Gamma \in \mathcal{T}_x} L(\Gamma) \quad (4)$$

We use *infimum* instead of minimum, because burn time may only exist in the limit. An example is given in the Supplementary Materials. Finally, Erosion Thickness (ET) is defined as the difference between BT and the radius function R , as in Equation 2.

Examples As examples, Figure 6 (right) plots the exposing trees at various locations on the seam in the medial axis of Figure 4. In this example, the burn time at each location is exactly realized by the length of the exposing tree. Each exposing tree consists of two edges that lie respectively on the board sheet and fin sheet. Recall that, during burning, the merging of fire fronts happens at two locations on the seam, first at p and later at q (Figure 4). For a point located away from the seam segment between p and q , the fin sheet is burned away first, and hence the primary path (red) of the exposing tree lies on the board sheet. For a point located within the segment between p and q , the board sheet is burned away first (from behind the fin sheet), and hence the primary path lies on the fin sheet. The fire fronts from the fin sheet and board sheet reach p (resp. q) at the same time, hence the exposing tree of p (resp. q) has two primary paths of equal lengths.

5 Properties of Erosion Thickness

Guided by the same burning analogy, our definition of ET in 3D naturally inherits the properties of ET in 2D [Liu et al. 2011]. We give intuitive explanation of these properties here and refer readers to the Supplementary Materials (Propositions 1.1, 1.2, 1.4) for precise statements and formal proofs of these properties.

Finiteness ET in 2D is finite everywhere except on closed cycles of the medial axis. Similarly, ET in 3D is finite anywhere away from the closed surfaces in M . This property is evident from the burning analogy, since burning terminates either when M is completely burned or the unburned residue has no boundary. Infinite ET happens rarely: in general, M contains a closed surface only when the shape S is bounded by multiple surfaces (e.g., a ball with interior cavities that are disconnected from the outside).

Continuity ET in 2D is continuous everywhere over the medial axis. Even at a junction point, ET is continuous over a segment formed by two branches (see Figure 2). Such continuity extends nicely to 3D: at any point $x \in M$ (manifold or singular), ET is continuous on some disk neighborhood of x . This disk is the last remaining neighborhood of x right before x is burned away. As an example, at a point x on the seam in Figure 6 (right), ET is continuous either on the board sheet (if x is not on the seam segment between p, q) or on a disk formed by the fin sheet and board sheet. Note that, in comparison, MGF is generally discontinuous at singular points.

Local minima Like ET in 2D, ET in 3D is free of local minima on the medial axis. Intuitively, since points are burned away only when they lie on the boundary of the unburned subset, the burn time lacks local minima. As R never varies faster than the burn time, ET is free of local minima as well.

Unlike the 2D case, the lack of local minima in 3D does not warrant a topology-preserving subset after thresholding. This is because ET may contain saddle points (critical points of a 2-dimensional function that are neither locally maximal or minimal), and the medial axis can become disconnected using thresholds lower than the values at the saddles. In our skeletonization algorithm, we shall use a contraction method to ensure that the resulting skeleton preserves both the significant parts and the topology of M .

Consistency with 2D definition Finally, we remark that our definition of burn time in 3D, with minor modifications, reduces to Liu’s definition in 2D (Equation 1). Consider now the medial axis M of a 2D shape S . In this case, $x \in M$ is *exposed* if at least all but one of x ’s adjacent segments are removed. If M is free of cycles, x has only two exposing trees, one on each side of x , that correspond to the two components of the residue $M \setminus \{x\}$. It is easy to verify that the greater length of the two trees evaluates to be BT as defined in Equation 1.

6 Computing Erosion Thickness in 3D

As mentioned earlier, exact or approximate algorithms for computing medial axes in 3D often produce a piece-wise linear structure, or a triangular mesh. Given such a mesh M and a well-defined radius function R (e.g., piece-wise linear interpolation of radii at vertices), our goal is to compute a good approximation of ET. While we leave exact computation of ET as future work, we also note that seeking exact values may not be necessary when M itself is an approximation of the true medial axis.

Our algorithm is inspired by Lanthier’s method [1997] for approximating geodesic distances over triangulated manifolds. His idea is to construct a graph G , whose vertices are sampled from triangle edges of M , and compute the shortest path distances on G . We use a similar graph construction, based on which we formulate and compute a version of ET that is restricted to the graph. The advantage of this approximation strategy is that we can bound the error of approximating the true ET by the sampling density.

6.1 Graph construction

We build an undirected graph $G = \{N, A\}$ with nodes N and arcs A . The nodes include the triangle vertices and extra sample points along triangle edges. As in [Lanthier et al. 1997], we allow the user to control the sampling density by specifying the maximal distance ω between two nodes on the same triangle edge. Additionally, our proof of error bound requires at least one sample point on each triangle edge. As a result, we add $\max(1, \lfloor l/\omega \rfloor)$ sample points with uniform spacing on an edge with length l .

The nodes on the edges of each triangle are connected by arcs in one of two ways, as illustrated in Figure 7 (left). First, any two nodes that do not lie on the same edge are connected by an arc (called a *triangle-arc*). Next, any two nodes are the adjacent on a triangle edge are connected by an arc (called an *edge-arc*). A notable distinction from [Lanthier et al. 1997] is that our graph maintains multiple edge-arcs between two nodes, one for each triangle that shares the edge. Since the triangle edge can be part of a seam where multiple medial axis sheets meet, the multiple edge-arcs capture the paths on different sheets that lie infinitesimally close to the seam. The length of an arc is the Euclidean distance between its ends.

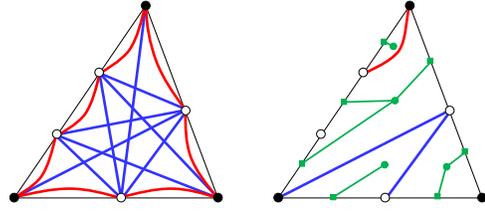


Figure 7: Left: A graph on a triangle, showing graph nodes (fill dots for vertices and hollow dots for edge samples) and arcs (blue for triangle-arcs and red for edge-arcs). Right: Arcs in burn trees (blue and red lines) and the dual medial curve (green).

6.2 Graph-restricted ET

An exposing tree is said to be *restricted to G* if vertices and edges coincide with some subset of N and A . We denote the set of all restricted exposing trees at a node v by $\mathcal{T}_{G,v}$. This allows us to define a restricted version of burn time and ET on the graph.

Definition 6.1 Given a graph $G = \{N, A\}$, the graph-restricted burn time function, $BT_G : N \rightarrow \mathbb{R} \cup \{\infty\}$ is defined for each $v \in N$ as

$$BT_G(v) = \min_{\Gamma \in \mathcal{T}_{G,v}} L(\Gamma) \quad (5)$$

The graph-restricted Erosion Thickness, $ET_G : N \rightarrow \mathbb{R} \cup \{\infty\}$ is

$$ET_G(v) = BT_G(v) - R(v) \quad (6)$$

It is not difficult to verify that the graph-restricted ET retains key properties of the continuous definition, including being finite anywhere away from the closed surfaces of M and free of local minima (the latter property additional requires that $\|R(u) - R(v)\|$ is no greater than the distance between u, v for any node pair $\{u, v\}$ that shares an edge).

Obviously, $ET_G(v)$ cannot be smaller than $ET(v)$, since the former considers a smaller set of exposing trees (i.e., those restricted to G). On the other hand, we show that $ET_G(v)$ cannot be too much longer than $ET(v)$ either (see proof in Supplementary Materials):

Proposition 6.2 Let $|M|$ count the number of triangles in M , g be the maximal gradient magnitude of R on any triangle edge on ∂M , and ω be the maximal distance between adjacent nodes in G on a triangle edge. For any node v in G ,

$$ET(v) \leq ET_G(v) \leq ET(v) + (2|M| + g)\omega \quad (7)$$

This bound is similar to that of approximating geodesic distances [Lanthier et al. 1997], with the difference of adding the variation of R . Although theoretically g should not exceed 1, we allow arbitrary g that might arise when M is an approximation of the true medial axis. The bound implies that the graph-restricted ET can be arbitrarily close to the true ET when edge sampling is sufficiently dense. In practice, we observed that a moderate sampling rate already achieves a good approximation, and denser sampling only brings minor improvement in accuracy (see Section 8).

6.3 Algorithm

To compute the graph-restricted burn time, our algorithm simulates burning on G . Like Dijkstra’s algorithm, we adopt a greedy expansion: burn times are estimated at nodes that have not yet been burned, and a node is burned if its estimate is minimal among all un-burned nodes *and* if its neighborhood is exposed.

Our algorithms uses the following data structures. We store at each node v its radius ($v.R$), an estimated burn time ($v.time$), a state flag ($v.burned$), and a list of its sectors ($v.sectors$). Each sector s keeps its own estimated burn time ($s.time$), its own state flag ($s.burned$), and a list of arcs that lie on that sector ($s.arcs$). During the algorithm, $s.time$ maintains the earliest time of arrival of the fire front from neighboring, burned nodes. Specifically, let $a.len$ be the length of an arc a , $s.time$ is the minimum of ($a.len + u.time$) over those arcs a in $s.arcs$ whose other end node, u , is already burned. For back-tracking purposes, the arc that realizes the minimum, called the *primary arc*, is stored in $s.primeArc$. The minimal $s.time$ among all unburned sectors s at node v is stored in $v.time$. The sector that realizes the minimum, called the *primary sector*, is stored in $v.primeSector$.

The pseudo-code of the algorithm is provided in Figure 8. The algorithm keeps all unburned nodes in a list Q . Initially, all nodes and sectors are unburned, and they all have infinite burn time except for the boundary nodes whose burn times are their radii. The main loop of the algorithm extracts the unburned node v with smallest estimated burn time. Its primary sector is considered burned, as well as any exposed sectors. Here, an unburned sector s is said to be exposed if it cannot form a 2-dimensional disk with any other unburned sectors at v . We set an exposed sector's primary arc to *null* because it is burned by an arc on another sector. If all sectors at v are burned, the node v is considered burned, and the estimated burn time at neighboring nodes and sectors are updated. Otherwise, v is re-inserted into Q after updating its estimated burn time using the remaining unburned sectors.

The algorithm exactly computes the graph-restricted burn time (see proof in Supplementary Materials):

Proposition 6.3 *At the termination of algorithm **Burn** in Figure 8, $v.time = BT_G(v)$ for every node v .*

The complexity of the algorithm is dominated by operations on Q and the updates of burn time between neighboring nodes. Since the number of times that a node v is inserted into Q is no more than the number of its sectors, and since the latter is less than the number of arcs at v , both operations $Q.extractMin$, $Q.addWithKey$ are called no more than $O(|A|)$ times where $|A|$ is the total number of arcs. Similarly, calling times for both the burn time updates and $Q.updateKey$ are no more than $O(|A|)$. Using a priority queue for Q , the complexity of the algorithm is $O(|A| \log |N|)$, where $|N|$ counts the total number of nodes. This is the same complexity as Dijkstra's algorithm.

Besides burn time, the exposing trees whose lengths realize burn time are also readily available from the algorithm. This tree, called a *burn tree*, can be constructed at each node v by tracing $s.primeArc$ over all sectors s of v back to the boundary. Figure 9 (top-middle) plots all burn trees computed by our algorithm on a synthetic medial axis. In the next section, we will use these trees to compute the medial curves. Here, we make the important observation that, like shortest paths on a surface, no two burn trees intersect (see proof in Supplementary Materials).

7 Computing skeletons

As a significance measure, ET in 3D excels in differentiating boundary noise from prominent features. However, naive pruning of the medial axis by thresholding ET may not produce a desirable skeleton. First, since ET captures plate-likeness, the pruned medial axis may lose tubular shape parts (e.g., fingers in Figure 10 top-left). Secondly, as mentioned earlier, thresholding ET may change the topology of the medial axis in 3D due to the presence of saddle points.

```

// G: graph (∂G denotes boundary nodes)
Burn (G)
  Q.init()           // Initialize
  for each node v of G
    for each s ∈ v.sectors
      s.time ← ∞
      s.burned ← False, s.primeArc ← null
  if v ∈ ∂G
    v.time ← v.R
  else
    v.time ← ∞
    v.burned ← False, v.primeSector ← null
    Q.addWithKey(v, v.time)
  while Q is not empty // Main loop
    v ← Q.extractMin()
    if v.primeSector ≠ null
      v.primeSector.burned ← True
    for each exposed sector s ∈ v.sectors
      s.burned ← True
      s.time ← v.time, s.primeArc ← null
    U ← all unburned sectors in v.sectors
    if U = ∅ // Burn v and update its neighbors
      v.burned ← True
      for each s ∈ v.sectors and each a ∈ s.arcs
        u ← other end node of a
        t ← sector of u that contains a
        if u.burned=False and t.burned=False
          h ← a.len + v.time
          if h < t.time
            t.time ← h, t.primeArc ← a
          if h < u.time
            u.time ← h, u.primeSector ← t
            Q.updateKey(u, u.time)
    else // Update burn time of v
      v.primeSector ← arg mins∈U s.time
      v.time ← v.primeSector.time
      Q.addWithKey(v, v.time)

```

Figure 8: Computing graph-restricted burn time.

To recover tubular features and retain topology, we augment the thresholded medial axis with 1-dimensional *medial curves*. These curves roughly follow the “ridges” of ET, which we observe to lie centered with respect to lateral expansion of the shape. Our method proceeds in three steps on a triangulated medial axis. First, using the same graph structure for computing ET, we extract a discrete set of ridge curves of ET. Second, we introduce a significance measure over the medial curves as a variant of ET over the medial axis of a 2D shape. Finally, significant subsets of the medial axis and medial curves are combined to form topology-preserving skeletons.

7.1 Computing the medial curves

While it would be ideal to precisely define the medial curves, we have found that it is difficult to come up with a definition that has all the desirable properties, such as being 1-dimensional and topologically equivalent to the 3D shape. Instead, we achieve these properties using a discrete algorithm that approximately traces the singularities of ET.

Our algorithm is again motivated by the burning analogy. Note that the process of burning retains the topology of the medial axis until fire fronts from different boundary parts collide and quench. The quench sites become the ridges of *BT*, which coincide with the ridges of ET (since the variation of R is never faster than that of

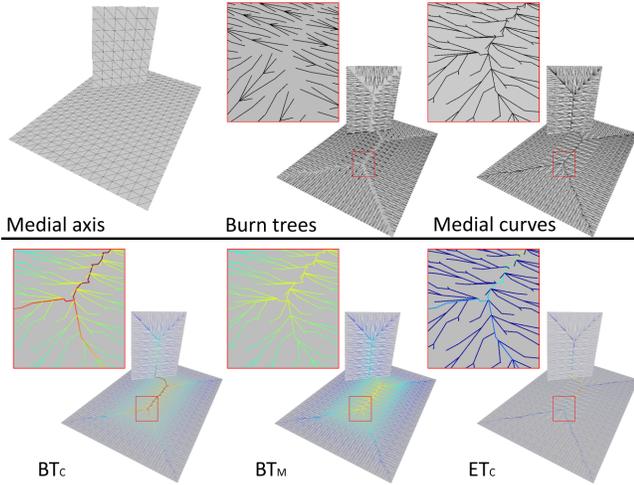


Figure 9: Top: A synthetic medial axis mesh, burn trees computed by our algorithm *Burn*, and the medial curves obtained by dualizing the space between the burn trees. Bottom: burn time over the medial curves (BT_C), burn time over the medial axis (BT_M), and their difference, the Erosion Thickness over the medial curves (ET_C). Close-up views are shown in the inserts.

BT). On the other hand, the *burn trees* computed by our algorithm of ET are exactly the paths of burning restricted to the graph G . Fire fronts moving along these paths would quench in the space between the paths. This motivates our strategy for computing medial curves: extracting the center lines of the areas of medial axis *between* paths of burn trees.

We use a simple dualization technique to extract these center lines. Consider a subdivision of each triangle on the medial axis M by all arcs used in burn trees (Figure 7 right). This is possible since burn trees do not intersect each other, as we have mentioned earlier. Note that a degenerate face with zero area can be formed by a triangle edge segment and a triangle-arc of G (e.g., the top-left face in the picture). We create a dual vertex on each edge segment (green square) and within each face (green dot). They are located at the average locations of the ends of the segment or corners of the face weighted by their burn time. Finally, the dual vertex of each face is connected to all dual vertices on the face’s boundary (green lines).

We show the result of our dualization method on a synthetic medial axis in Figure 9 (top-right). Note that the medial curves form a connected, tree-like structure. In fact, we can prove that these curves form a deformation retract of M and hence is homotopy-equivalent to M (see Supplementary Materials).

7.2 Erosion Thickness on the medial curves

The computed medial curves may contain many spurious branches, due in part to the noisiness of the medial axis and in part to the discrete nature of our algorithm. Our goal is to identify a subset of the medial curves that captures major tubular features.

We define a significance measure over the medial curve by making a small change to the definition of ET on the medial axis of a 2D shape. Specifically, the radius function is replaced by the burn time over M . To avoid confusion, let us denote $BT_M(x)$ as the burn time over M at a point $x \in M$. Given the medial curves C , we define the Erosion Thickness over C at a point $x \in C$, $ET_C(x)$, in a form similar to Equations 2,1:

$$ET_C(x) = BT_C(x) - BT_M(x). \quad (8)$$

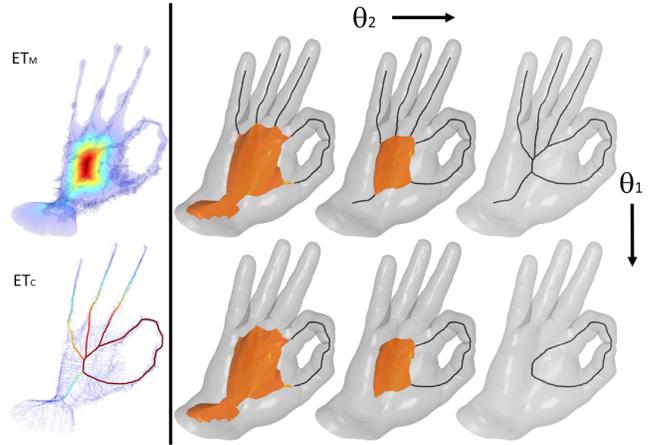


Figure 10: Left: ET on the medial axis (ET_M) and medial curve (ET_C) of a genus-1 hand. Right: skeletons created with different settings of significance thresholds. Note that the hand topology is preserved in every skeleton.

where $BT_C(x)$ is the *burn time function* over the medial curve C , defined as

$$BT_C(x) = \sup \min_{yz} (BT_M(y) + d_C(x, y), BT_M(z) + d_C(x, z)). \quad (9)$$

where the supremum is taken over all segments yz on C between two end points y, z that contain x . Both BT_C and ET_C can be computed using the same reduction process for computing ET in 2D [Liu et al. 2011]. The functions, BT_C , BT_M and ET_C , are shown on a synthetic medial axis in Figure 9 (bottom).

Conceptually, ET_C measures the shrinking of a tubular protrusion as a result of pruning the medial curve. We have observed that ET over the medial curves can effectively distinguish spurious branches from center curves of prominent tubular features, such as the fingers of the hand (Figure 10 bottom-left). ET_C is finite anywhere except over cycles of C (i.e., within a topological handle, as in this example) and closed surfaces of M (i.e., within a shell bounded by multiple surface components).

7.3 Generating skeletons

Finally, we combine significant parts of the medial axis and medial curves (as deemed by their ET) into a skeleton that also preserves the topology of the shape. This is done using a homotopy-preserving contraction operator on general cell complexes [Liu et al. 2010].

We start by creating a refined subdivision of each triangle using both the arcs of the burn trees and the segments of the medial curves (i.e., the union of red, blue and green lines in Figure 7 right). The medial axis M is then represented as a cell complex, whose 2-, 1-, and 0- cells are the faces of the subdivision, their bounding segments, and their corners. Each 2-cell is assigned a priority as the average ET_M at the corners of the cell, and each 1-cell is assigned a priority as the average ET_C at the ends of the cell. A 2-cell (resp. 1-cell) whose priority is greater than a user-specified threshold θ_2 (resp. θ_1) is marked as *protected*. For cleaner skeletons, protected 2-cells (resp. 1-cells) that form connected components of size smaller than a constant ϕ_2 (resp. ϕ_1) become un-protected. Contraction of M , ordered by the priority of cells, yields a minimal sub-complex that contains all protected cells while preserving the topology of M [Liu et al. 2010].

Our skeletonization method allows flexible, independent control over 1-dimensional and 2-dimensional components of the skeleton (Figure 10 right). By the definition of ET, the threshold θ_2 (resp. θ_1) roughly corresponds to how far the outer rim (resp. end point) of a plate-like (resp. tubular) part of the shape shrinks inward after the medial axis is pruned. Visually, a lower θ_2 (resp. θ_1) retains more sheets (resp. curves) and captures more plate-like (resp. tubular) feature of the shape. Setting $\theta_2 = \infty$ produces an all-curve skeleton (except for shapes bounded by multiple surfaces), and setting $\theta_1 = \infty$ produces an surface-dominant skeleton where curves are present only for preserving topology. Setting both thresholds to ∞ yield a minimal structure that only preserves the topology, such as a point (for a genus-0 shape), a cycle set (for a non-genus-0 shape, as in Figure 10 bottom-right) or a set of closed surfaces (for a shape bounded by multiple surfaces).

8 Results

We implemented the algorithms for computing Erosion Thickness and the skeletons in C++. We use the software provided by [Miklos et al. 2010] to generate the triangulated medial axes. To be able to use consistent parameter settings, we scale each shape so that the maximal dimension of the bounding box is 1. All tests are performed on a quad-core PC with 3.2GHz CPU and 12 GB RAM.

Quality and speed To study the approximation error of our algorithm **Burn**, we consider the analytical medial axis of a simple rectangular box (made up of 13 planar polygons), on which the exact burn time can be analytically determined based on the burning process (Figure 11 (a)). To compare the exact burn time with that computed by **Burn** on the approximate triangulated medial axis, we select ten vertices on the latter (dots in Figure 11 (a)) and plot the result of **Burn** divided by the exact burn time at each vertex under different values of ω (Figure 11 (b)). Note that our discrete algorithm is able to approximate the true burn time with high accuracy, and increasing the sampling rate (lowering ω) only brings minor improvement.

On the same medial axis, we plot the number of graph arcs and running time of **Burn** as functions of ω (Figure 11 (c,d)). Both quantities grow roughly quadratically with the sampling rate, which is consistent with our complexity analysis of **Burn**.

In all our experiments, we set $\omega = 0.004$ to balance accuracy and efficiency. At this setting, our algorithm completes under just a few seconds on a triangulated medial axis with 50K vertices (e.g., for Dolphin in Figure 12). As a comparison, even using the most efficient approximation method available for geodesic distances [Crane et al. 2013], evaluating MGF at all vertices of the same medial axis, assuming a typical complexity of 30K vertices on the shape boundary, requires at least 15 minutes.

Comparison of significance measures We have performed extensive tests to compare ET with object angle, circumradius and MGF on smooth and perturbed shapes (see Supplementary Materials). Two examples (Dolphin and Elephant) are shown in Figure 12. Each significance measure is shown first by color and transparency over the medial axis and next as a sub-set of the medial axis triangles with significance values higher than a threshold (same thresholds are used before and after boundary perturbation).

Consistent with our observations in 2D (Figure 2), ET in 3D excels over local measures in differentiating noise from features. While highlighting thin features (e.g., the tail of Dolphin and ear of Elephant), object angle is also high near boundary noise (note the large number of islands in the thresholded medial axis of the perturbed shapes). While circumradius is much less sensitive to boundary noise, the values are overly low on thin features. In contrast, ET

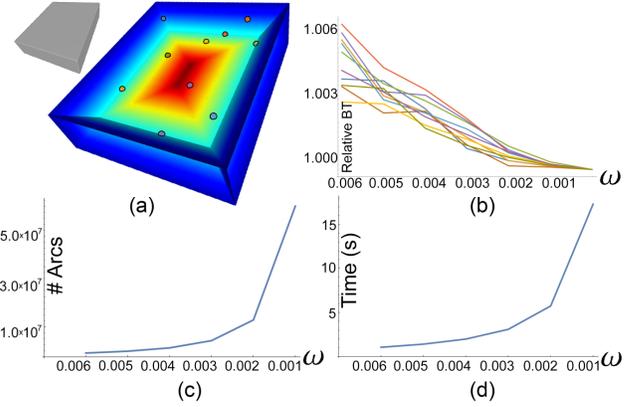


Figure 11: The exact medial axis of a box colored by exact burn time (a) and the plots under decreasing ω (increasing sampling) for approximation error of our burn time algorithm (b), number of graph arcs (c), and running time (d) on a triangulated medial axis.

is sensitive to thin, plate-like features while being insensitive to boundary noise.

Compared with MGF, ET is often better at highlighting plate-like features (e.g., Dolphin’s tail and Elephant’s ears). We also found ET to be more stable under boundary perturbations that significantly alter geodesic distances. As shown in Figure 13, after ripples are added to one side of a box, the distribution and range of MGF values notably change while ET stays roughly the same.

Skeletons Guided by ET, our skeletonization algorithm produces skeletons that are stable under boundary perturbations. Figure 14 compares the skeletons of Dolphin and Elephant before and after boundary perturbation (see more examples in the Supplementary Materials). The thresholds θ_2 and θ_1 used in our examples typically range between $[0.01, 0.03]$ and $[0.02, 0.06]$, respectively, except for all-curve skeletons where $\theta_2 = \infty$. For each shape, we keep the same thresholds before and after perturbation. In all of our examples, we use the component size limits $\phi_2 = 0.0005$ and $\phi_1 = 0$.

We compare with skeletons produced by a state-of-the-art pruning method of Li et al. [2015], called Q-MAT, in Figure 15. Besides pruning spurious branches, Q-MAT also simplifies the mesh structure in the un-pruned region, and both goals are achieved within a unified edge-contraction framework. The method offers a family of skeletons with varying levels of details, from a complex surface mesh to a coarse polygonal curve. In comparison, by allowing independent control of the 2D and 1D components (via θ_2 or θ_1), our method creates a greater variety of skeletons. For example, we can increase the pruning on the skeleton surface without losing the skeleton curves within fine features (e.g., Dinosaur’s teeth). Furthermore, since edge-contraction in Q-MAT is guided by a combination of a local significance measure (known as stability ratio) and a heuristically defined distance metric, it can be difficult for the user to interpret their pruning parameter. On the other hand, the pruning parameters (θ_2 or θ_1) in our method have more intuitive meanings (as the amount of shrinking of plate-like or tubular shape features from their extremities after medial axis pruning).

We also compare with the discrete Scale Axis Transform (SAT) of Miklos et al. [2010] in Figure 16. The SAT method inflates the shape, by a scaling parameter, before extracting the medial axis. For noisy inputs, a small amount of inflation is insufficient to create a clean skeleton (Figure 16 left). However, greater inflation can cause nearby shape features to merge, thus creating a skeleton protruding

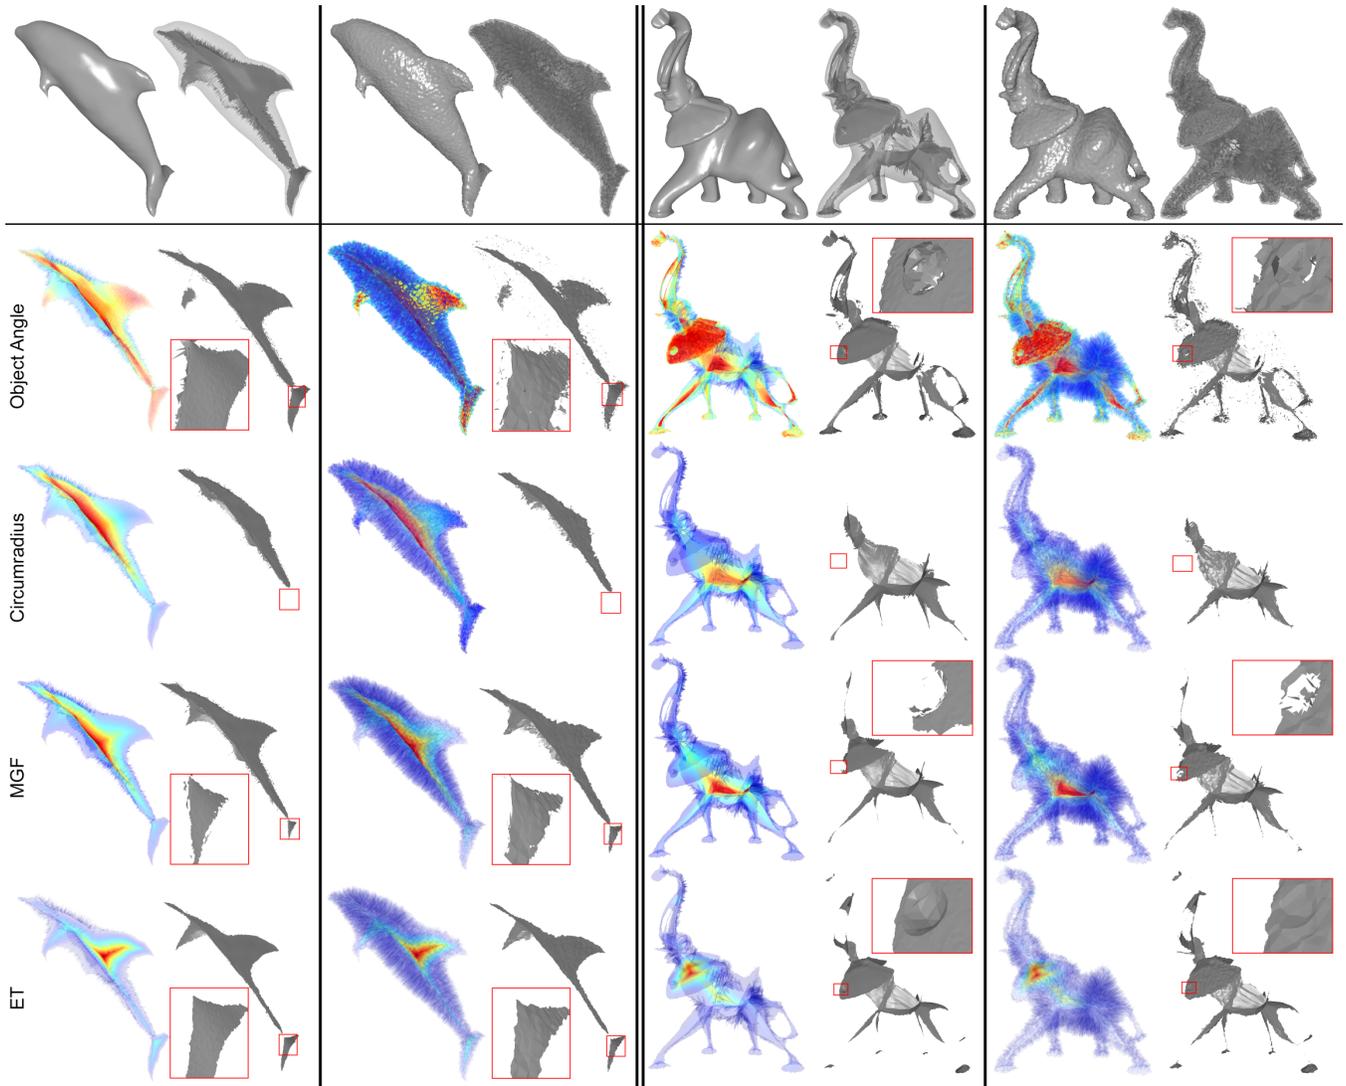


Figure 12: Comparison of significance measures on Dolphin (left), perturbed Dolphin (middle left), Elephant (middle right) and perturbed Elephant (right). First column for each shape visualizes the significance measure by heat color (higher values use hotter color) and transparency (higher values are more opaque). Second column shows the portion of the medial axis whose significance values are higher than a threshold. Note that ET is sensitive to thin, plate-like features (see inserts) while being robust to perturbation. See more examples in Supplementary Materials.

outside the original shape and even with a wrong topology (e.g., between the fingers in Figure 16 middle). In contrast, our skeletons are guaranteed to stay inside the shape and preserve its topology.

9 Conclusion and discussion

We introduce a novel global significance measure on the medial axes of 3D shapes that extend the Erosion Thickness measure in 2D. We give a precise definition, prove its key theoretical properties, and present an efficient discrete approximation with bounded error. The measure is used to guide the creation of clean, topology- and shape- preserving skeletons.

Future work We would like to investigate a formal definition of medial curves. Our medial curve (as ridge of Erosion Thickness) can also be understood as the quench site of burning over the medial axis. This interpretation is in fact consistent with the medial axis itself, which is the quench site of Blum’s grassfire burning from the

boundary of the 3D shape. Naturally, one may conjecture the existence of a hierarchy of structures (3D shape, 2D medial axis, 1D medial curve, 0D medial point), such that each lower-dimensional structure is the quench site of burning on the higher-dimensional structure. It would also be interesting to see if exact geodesic algorithms on meshes, which are known to have a low polynomial complexity, can be extended to compute exact ET on a triangulated medial axis with a similar complexity.

Applications Skeletons have been widely used in applications such as animation, shape analysis, object retrieval, and medical image analysis [Tagliasacchi et al. 2016]. As 3D ET measures the plate-likeness of shape parts, we are particularly interested in applying ET and the skeletons in part-based shape segmentation. This can be achieved, for example, by defining signature functions over the shape boundary that capture plate-likeness and tubularity.

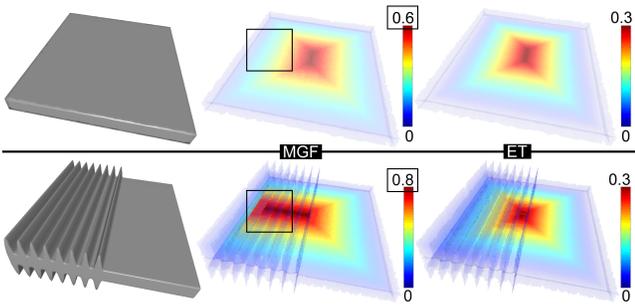


Figure 13: Comparing MGF (middle) and ET (right) on the medial axis of a box before (top) and after (bottom) adding ripples to one side of the box. Note that MGF has a significant shift in its distribution (see the outlined areas) and a big increase in its range of values, while ET remains largely unchanged.

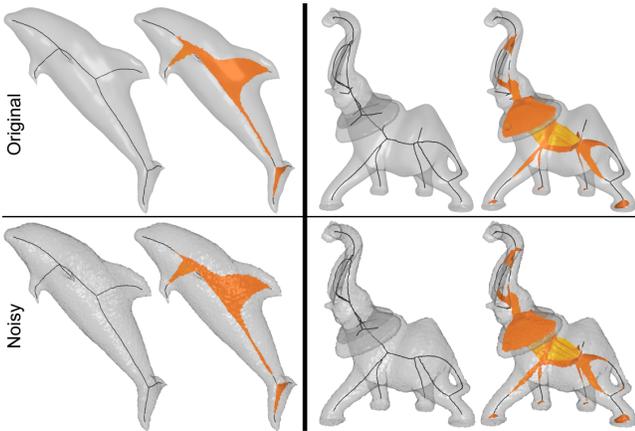


Figure 14: Skeletons generated by our method before and after boundary perturbation.

Acknowledgements We thank the Princeton Shape Benchmark and Aim@Shape for providing the models used in this paper. The work is supported in part by NSF grants IIS-1319573, DBI-1356388, and CCF-1054779.

References

AMENTA, N., AND KOLLURI, R. K. 2001. The medial axis of a union of balls. *Comput. Geom.* 20, 1-2, 25–37.

AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust. In *SMA '01: Proceedings of the sixth ACM symposium on Solid modeling and applications*, 249–266.

ARCELLI, C., AND DI BAJA, G. S. 1993. Euclidean skeleton via centre-of-maximal-disc extraction. *Image Vision Comput.* 11, 3, 163–173.

ATTALI, D., AND MONTANVERT, A. 1996. Modeling noise for a better simplification of skeletons. In *Proceedings 1996 International Conference on Image Processing, Lausanne, Switzerland, September 16-19, 1996*, 13–16.

ATTALI, D., DI BAJA, G. S., AND THIEL, E. 1995. Pruning discrete and semicontinuous skeletons. In *Image Analysis and Processing, 8th International Conference, ICIAP '95, San Remo, Italy, September 13-15, 1995, Proceedings*, 488–493.

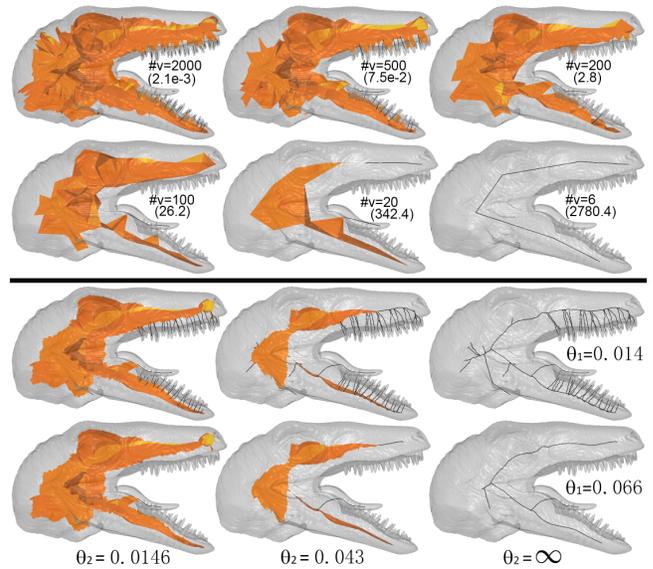


Figure 15: Comparing the skeletons of Q-MAT at different simplification levels, showing both remaining number of vertices and the accumulated error (top), and ours with different θ_2, θ_1 . Note that our method allows independent control over the pruning of the surfaces and curves on the skeletons.

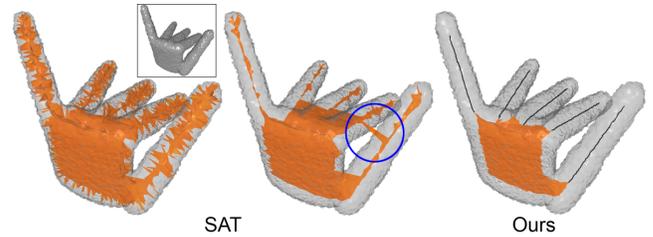


Figure 16: Comparing the skeletons of SAT at two scaling parameters and ours on a bumpy shape. Note the noisiness of SAT skeleton at a low scale (left) and the incorrect topology of SAT skeleton at a high scale (middle).

BERTRAND, G. 1995. A parallel thinning algorithm for medial surfaces. *Pattern Recogn. Lett.* 16, 9.

BLUM, H. 1967. A transformation for extracting new descriptors of form. *Models for the Perception of Speech and Visual Form*, 362–80.

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

BRANDT, J., AND ALGAZI, V. 1992. Continuous skeleton computation by voronoi diagram. *CVGIP: Image Understanding* 55, 3, 329 – 338.

CHAUSSARD, J., COUPRIE, M., AND TALBOT, H. 2009. A discrete lambda-medial axis. In *DGCI*, vol. 5810, 421–433.

CHAZAL, F., AND LIEUTIER, A. 2004. Stability and homotopy of a subset of the medial axis. In *Proceedings of the Ninth ACM Symposium on Solid Modeling and Applications, SM '04*, 243–248.

CHAZAL, F., AND SOUFFLET, R. 2003. Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems* 10, 2, 149–170.

- CHOI, S. W., AND SEIDEL, H.-P. 2001. Hyperbolic hausdorff distance for medial axis transform. *Graph. Models* 63, 5 (Sept.), 369–384.
- CORNEA, N. D., AND MIN, P. 2007. Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics* 13, 3, 530–548. Member-Silver, Deborah.
- CRANE, K., WEISCHEDL, C., AND WARDETZKY, M. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph.* 32, 5 (Oct.), 152:1–152:11.
- CULVER, T., KEYSER, J., AND MANOCHA, D. 2004. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design* 21, 1, 65–98.
- DEY, T. K., AND SUN, J. 2006. Defining and computing curve-skeletons with medial geodesic function. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 143–152.
- DEY, T. K., AND ZHAO, W. 2004. Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design* 36, 2, 195–202.
- DILL, A. R., LEVINE, M. D., AND NOBLE, P. B. 1987. Multiple resolution skeletons. *IEEE Trans. Pattern Anal. Mach. Intell.* 9, 4, 495–504.
- FLETCHER, P. T., LU, C., PIZER, S. M., AND JOSHI, S. C. 2004. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imaging* 23, 8, 995–1005.
- FOSKEY, M., LIN, M. C., AND MANOCHA, D. 2003. Efficient computation of A simplified medial axis. *J. Comput. Inf. Sci. Eng.* 3, 4, 274–284.
- GIBLIN, P. J., AND KIMIA, B. B. 2004. A formal classification of 3d medial axis points and their local geometry. *IEEE Trans. Pattern Anal. Mach. Intell.* 26, 2, 238–251.
- GIESEN, J., MIKLOS, B., PAULY, M., AND WORMSER, C. 2009. The scale axis transform. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, SCG '09, 106–115.
- GOLLAND, P., ERIC, W., AND GRIMSON, L. 2000. Fixed topology skeletons. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, vol. 1, 10–17 vol.1.
- HO, S.-B., AND DYER, C. R. 1986. Shape smoothing using medial axis properties. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8, 4, 512–520.
- LANTHIER, M., MAHESHWARI, A., AND SACK, J.-R. 1997. Approximating weighted shortest paths on polyhedral surfaces. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, SCG '97, 485–486.
- LI, P., WANG, B., SUN, F., GUO, X., ZHANG, C., AND WANG, W. 2015. Q-mat: Computing medial axis transform by quadratic error minimization. *ACM Trans. Graph.* 35, 1 (Dec.), 8:1–8:16.
- LIEUTIER, A. 2003. 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*, SM '03, 65–75.
- LIU, L., CHAMBERS, E. W., LETSCHER, D., AND JU, T. 2010. A simple and robust thinning algorithm on cell complexes. *Comput. Graph. Forum* 29, 7, 2253–2260.
- LIU, L., CHAMBERS, E. W., LETSCHER, D., AND JU, T. 2011. Extended grassfire transform on medial axes of 2d shapes. *Computer-Aided Design* 43, 11, 1496–1505.
- MIKLOS, B., GIESEN, J., AND PAULY, M. 2010. Discrete scale axis representations for 3d geometry. *ACM Trans. Graph.* 29, 4 (July), 101:1–101:10.
- NIBLACK, W., GIBBONS, P., AND CAPSON, D. 1992. Generating connected skeletons for exact and approximate reconstruction. In *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR '92., 1992 IEEE Computer Society Conference on*, 826–828.
- OGNIEWICZ, R., AND ILG, M. 1992. Voronoi skeletons: theory and applications. In *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR '92., 1992 IEEE Computer Society Conference on*, 63–69.
- OGNIEWICZ, R. L., AND KÜBLER, O. 1995. Hierarchic voronoi skeletons. *Pattern Recognition* 28, 3, 343–359.
- PIZER, S. M., OLIVER, W. R., AND BLOOMBERG, S. H. 1987. Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE Trans. Pattern Anal. Mach. Intell.* 9, 4, 505–511.
- 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., AND CHANEY, E. L. 2003. Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision* 55, 2-3, 85–106.
- PIZER, S. M., GERIG, G., JOSHI, S. C., AND AYLWARD, S. R. 2003. Multiscale medial shape-based analysis of image objects. *Proceedings of the IEEE* 91, 10, 1670–1679.
- RENIERS, D., VAN WIJK, J., AND TELEA, A. 2008. Computing multiscale curve and surface skeletons of genus 0 shapes using a global importance measure. *IEEE Transactions on Visualization and Computer Graphics* 14, 2, 355–368.
- SHAKED, D., AND BRUCKSTEIN, A. M. 1998. Pruning medial axes. *Comput. Vis. Image Underst.* 69, 2, 156–169.
- SIDDIQI, K., AND PIZER, S. M. 2008. *Medial Representations*. Springer.
- SIDDIQI, K., BOUIX, S., TANNENBAUM, A., AND ZUCKER, S. W. 2002. Hamilton-jacobi skeletons. *International Journal of Computer Vision* 48, 3 (July), 215–231.
- STYNER, M., GERIG, G., JOSHI, S. C., AND PIZER, S. M. 2003. Automatic and robust computation of 3d medial models incorporating object variability. *International Journal of Computer Vision* 55, 2-3, 107–122.
- SUD, A., FOSKEY, M., AND MANOCHA, D. 2005. Homotopy-preserving medial axis simplification. In *SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling*, ACM, New York, NY, USA, 39–50.
- TAGLIASACCHI, A., DELAME, T., SPAGNUOLO, M., AMENTA, N., AND TELEA, A. 2016. 3d skeletons: A state-of-the-art report. In *Eurographics*.
- TAM, R. C., AND HEIDRICH, W. 2003. Shape simplification based on the medial axis transform. In *14th IEEE Visualization 2003 Conference (VIS 2003), 19-24 October 2003, Seattle, WA, USA*, 481–488.