

Washington University in St. Louis

Washington University Open Scholarship

All Theses and Dissertations (ETDs)

1-1-2011

Multi-Dimensional Medial Geometry: Formulation, Computation, and Applications

Lu Liu

Washington University in St. Louis

Follow this and additional works at: <https://openscholarship.wustl.edu/etd>

Recommended Citation

Liu, Lu, "Multi-Dimensional Medial Geometry: Formulation, Computation, and Applications" (2011). *All Theses and Dissertations (ETDs)*. 607.

<https://openscholarship.wustl.edu/etd/607>

This Dissertation is brought to you for free and open access by Washington University Open Scholarship. It has been accepted for inclusion in All Theses and Dissertations (ETDs) by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS
School of Engineering and Applied Science
Department of Computer Science and Engineering

Dissertation Examination Committee:

Tao Ju, Chair
Erin Wolf Chambers
Cindy Grimm
Robert Pless
Fred Prior
Kilian Weinberger

Multi-Dimensional Medial Geometry: Formulation, Computation, and Applications

by

Lu Liu

A dissertation presented to the Graduate School of Arts & Sciences
of Washington University in partial fulfillment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

December 2011
Saint Louis, Missouri

ABSTRACT OF THE THESIS

Multi-Dimensional Medial Geometry: Formulation, Computation, and Applications

by

Lu Liu

Doctor of Philosophy in Computer Science

Washington University in St. Louis, 2011

Research Advisor: Professor Tao Ju

Medial axis is a classical shape descriptor. It is a piece of geometry that lies in the middle of the original shape. Compared to the original shape representation, the medial axis is always one dimension lower and it carries many intrinsic shape properties explicitly. Therefore, it is widely used in a large amount of applications in various fields.

However, medial axis is unstable to the boundary noise, often referred to as its *instability*. A small amount of change on the object boundary can cause a dramatic change in the medial axis. To tackle this problem, a significance measure is often associated with the medial axis, so that medial points with small significance are removed and only the stable part remains. In addition to this problem, many applications prefer even lower dimensional medial forms, e.g., shape centers of 2D shapes, and medial curves of 3D shapes.

Unfortunately, good significance measures and good definitions of lower dimensional medial forms are still lacking. In this dissertation, we extended Blum's grassfire

burning to the medial axis in both 2D and 3D to define a significance measure as a distance function on the medial axis. We show that this distance function is well behaved and it has nice properties. In 2D, we also define a shape center based on this distance function. We then devise an iterative algorithm to compute the distance function and the shape center. We demonstrate usefulness of this distance function and shape center in various applications. Finally we point out the direction for future research based on this dissertation.

Acknowledgments

It is a great pleasure to thank all the people I owe great gratitude to.

I owe a debt of gratitude to my adviser Dr. Tao Ju. Without him, this dissertation can never be done. I would like to thank him for providing guidance and support through my Ph.D, for being so responsible to offer responsive help when needed, for being so inspirational through all projects, for setting up such a great role model who is always passionate, persistent and earnest about research.

I would like to thank my collaborators Dr. Erin Chamber and Dr. David Letscher from Saint Louis University for weekly discussions. Their broad mathematical knowledge offers me insights and inspiration into the research topic. Thanks also goes to my collaborators from Washington University in St. Louis, School of Medicine: Dr. Fred Prior, Dr. David R. Sinacore, Paul Commean, Dr. Charles Hildebolt, Kirk Smith, Dave Gutekunst. It was a pleasure to work with them all.

I thank Dr. Tao Ju, Dr. Cindy Grimm, Dr. Robert Pless, Dr. Kilian Weinberger, Dr. Fred Prior, and Dr. Erin Wolf Chambers for agreeing to be on my committee. It is a great honor for me to have you all.

I owe great thanks to all my friends, who stand by me through joy and sorrow.

Last but not least, my deepest gratitude goes to my beloved parents and sister, for being always encouraging and supportive. I would also like to express my deepest thanks to my fiance, Ondrej Stava, for simply being so great.

Lu Liu

*Washington University in Saint Louis
December 2011*

Dedicated to my parents.

Contents

Abstract	ii
Acknowledgments	iv
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Medial Axis	1
1.2 Instability of Medial Axis	3
1.3 Lower Dimensional Medial Forms	4
1.4 Our Contributions	5
1.5 Dissertation Overview	7
2 Related Works	8
2.1 Computation of the Medial Axis	8
2.1.1 Accurate Medial Axis Computation	9
2.1.2 Medial Axis Approximation	10
2.2 Significance Measures	14
2.2.1 Local Significance Measures	14
2.2.2 Global Significance Measures	17
2.3 Lower Dimensional Medial Forms	20
2.3.1 Center Points of 2D Shapes	20
2.3.2 Definition of Medial Curves	21
2.4 Summary	23
3 Extended Grassfire Transform on Medial Axes of 2D Shapes	24
3.1 Formulation	24
3.1.1 Motivation	24
3.1.2 Definitions	26
3.2 Properties	28
3.3 A Grassfire Analogy for Computation	30
3.3.1 Extended Grassfire Transform	30
3.3.2 Discrete Algorithm	31
3.4 Examples and Comparisons	32
3.4.1 Erosion Thickness	35

3.4.2	Potential Residue	36
3.5	Utility	37
3.5.1	Pruning Medial Axes	38
3.5.2	Shape Alignment	40
3.5.3	Shape Signature	40
3.6	Summary	42
4	Extended Grassfire Transform on Medial Axes of 3D Shapes . . .	44
4.1	A First Attempt: Largest-Plate Definition	44
4.1.1	Definition	45
4.1.2	Example and Issues	47
4.2	Burning-based Definition	50
4.2.1	Intuition	50
4.2.2	Notations	53
4.2.3	Burning Sets and Burning Times	55
4.3	Properties of Burning Times	58
4.4	Summary	60
5	Future Work	61
5.1	Theory	61
5.1.1	Burning Time and Shape Properties	61
5.1.2	Homotopy Preserving Lower Dimensional Medial Forms	62
5.1.3	Stability	63
5.2	Algorithms	64
5.3	Applications	65
Appendix A	Proofs of Propositions in Chapter 3	68
A.1	Proof of Proposition 1	68
A.2	Proof of Proposition 2	69
A.3	Proof of Proposition 3	73
Appendix B	Proofs of Lemmas and Propositions in Chapter 4 . .	74
B.1	Proof of Lemma 2	74
B.2	Proof of Proposition 4	75
B.3	Proof of Proposition 5	77
B.4	Proof of Proposition 6	79
References	80

List of Tables

List of Figures

1.1	Medial axis defined as the shock graph in the distance transform (a) (image courtesy of Attali et al. [7]), medial axis defined as the union of the centers of maximally bi-tangentially inscribed balls (b), and the spoke representation of medial axis.	1
1.2	Medial axis of a rectangle (a), and medial axis of a rectangle with a very small bump (b).	3
2.1	The Voronoi Diagram (a) and the medial axis (b) of a point set. The Voronoi Diagram partitions the space into cells, colored differently for each input point. The boundary of these cells is the medial axis of the input points since any point on the boundary has at least two closest boundary points.	9
2.2	The classified elements of the medial axis of a polyhedron (a), an input polyhedron (b) and its medial axis (c) (Image courtesy of Culver et al. [25]).	10
2.3	A shape made from a union of balls (a), its α -shape (in gray) with singular elements highlighted in red (b), and its medial axis in red and blue (c). The medial axis consists of the singular elements in the α -shape (in red) and the intersection (in blue) between the Voronoi Diagram (dashed lines) of the boundary singular points (black dots) and the α -shape (in gray) (image courtesy of Amenta and Kolluri [4]).	11
2.4	The separation angle $\theta(x)$ at the medial point x (a), the pruned medial axis with the threshold $\frac{\pi}{2}$ (b), and the disconnected pruned medial axis of another model with the threshold $\frac{\pi}{3}$ (c). The separation angle is defined by the two rays from the medial point x to its two closest boundary points y_1 and y_2 (a). Note that a naive pruning based solely on separation angle could easily break the topology (image courtesy of Sud et al. [68]).	15
2.5	An input model with its medial axis before and after pruning (a), the ρ and θ distribution of all medial points (b), and the pruned medial axis of a model with a snail on a tree branch (c). The medial points corresponding to the upper right quadrant are those preserved after pruning in the bottom figure of (a). Note that the prominent features, antennae of the snail, are accidentally pruned based on the combination measure of ρ and θ , because they have small thickness (ρ) (images (a)(b) courtesy of Attali and Montanvert [8]).	16

2.6	The medial axis (a) and its λ -medial axis when λ is slightly larger than its wfs (b). The difference between the local thickness $\rho(x)$ and the feature size $r(x)$ can be seen at the annotated medial point x (in black), where the red circle is the bi-tangentially inscribed medial circle and the blue circle is the minimal circle that encloses the two closest boundary points (in red). The three hollow points are the critical points and wfs is the smallest feature size of the three. Note that the topology of λ -medial axis is different from the original model when $\lambda > wfs$ for this model (image courtesy of Attali et al. [7]).	17
2.7	The erosion thickness (a), the three residues (b), and the erosion area (c). The erosion thickness at q is measured as the difference between the length of the branch and the thickness of the branch, i.e., the elongation of the branch. The length of the branch is measured as the summation of the geodesic distance $d(p, q)$ from the medial point q to the branch end point p and the local thickness $\rho(p)$ at p . The thickness of the branch is measured as the local thickness $\rho(q)$. $L(B)$, $L(A)$ and $L(C)$ are the lengths of the boundary curve, the arc, and the chord connecting the two closest boundary points respectively. The potential residue, the circularity residue and the chord residue at q are $L(B)$, $L(B) - L(A)$, and $L(B) - L(C)$. They all estimate the change in the boundary length when the branch attached to the medial point q is removed. The erosion area of q is colored in gray in (c), which estimates the change in the area of the reconstructed shape in the absence of the branch attached to q (image courtesy of Shaked and Bruckstein [61]).	18
2.8	Comparison of several center definitions. The last one is defined by our definition (more details in Chapter 3).	21
2.9	An input human model (a), its medial axis colored based on the medial geodesic function (MGF) (b), and the extracted medial curves (c) (image courtesy of Dey and Sun [29]).	22
3.1	Axes (blue) and their tubes (gray): (a) an axis with a single constrained end z , (b) an inscribed axis of x with a constrained end y , (c) an inscribed and maximal axis of x with two constrained ends y, z , (d) an inscribed and maximal axis of $x \in \tilde{M}$	25
3.2	An illustration of intermediate states in the extended grassfire burning on the medial axis. Yellow dots are the fire fronts, arrows indicate the burning direction, and the red, blue, and green squares are where fire fronts are ignited, annihilated, and quenched.	30

3.3	EDF and EMA in a simply connected shape (left) and a shape with an interior hole (right). The medial axis is colored by the EDF in (a) while the EMA is drawn in grey, (b) shows the boundary distance function in the background, and (c) plots both EDF and the boundary distance function as a 3D height map. The heat coloring scheme is used (blue is low and red is high).	33
3.4	EDF (a) and overlaid on the boundary distance function (b) for a shape with boundary noise.	34
3.5	Comparing the EDF (top row) and the Potential Residue (PR) measure (bottom row) on the medial axis under boundary perturbation: the EDF and EMA are stable under both uniform (a,b) and non-uniform (c,d) perturbations, while PR can have sudden jumps (e,f), and both PR and its local maximum can exhibit significant drift (g,h). Each picture is colored by the EDF or PR normalized by the maximum value of EDF or PR present on the medial axis. The inserts in (a,b,e,f) zoom in on the middle of the medial axis where there is a connectivity change after perturbation. The closest boundary points to the local maximum of PR are shown as crosses.	35
3.6	The Erosion Thickness (ET) measure (a) and Shape Tubularity (ST) measure (b) over a 2D shape, and pruning using a low ET threshold (c), a high ET threshold (d), and the combination of a low ET threshold with an ST threshold (e).	38
3.7	More pruning examples combining ET and ST. Examples in (a,b,c) include the original medial axes, the EDF and the boundary distance functions, and the pruned medial axes. Only pruned medial axes are shown in (d).	39
3.8	Top: several human shapes and their centroid (pink), local maximum of PR (blue), and EMA (red). Bottom: alignment using the centroids (b), local maximum of PR (c), and the EMAs (d).	41
3.9	Boundary signatures: local curvature (a), local feature size (b) showing also the medial axis, Boundary Eccentricity (BE) (c,d) showing also $E(x)$ over the medial axes, and a matching result using BE (e).	42
4.1	A 3D shape made from two rectangular cuboids with smooth boundary edges (a), the 3D shape with its medial axis embedded inside (b), a disk (in blue) on the medial axis surrounding medial point x (c), and the plate (in blue) corresponding to the disk (d). The LPDF at x is achieved by the red curve path in (c), where z is a boundary point of the blue disk and $R(z)$ is the radius of the medial ball (the red ball in (d)) at z . The plate is the union of all medial balls with centers from the blue disk. We highlighted a few balls inside the plate in blue and red.	45

4.2	The board and fin example (right) made from a long and wide “board” (left) and a long and tall “fin” (middle). The black lines are their boundary lines and the arrows represent infinite expansion. A non-manifold curve is formed at the intersection between the board and the fin.	47
4.3	The LPDF distribution on the board and fin example (top), and the paths (in white) that realize the LPDF of medial points from its three different zones (bottom). The value of the LPDF increases when the hue color changes from red to purple. The medial points with the same LPDF forms an iso-curve (in black in the top row), which is upper semi-continuous at the non-manifold point and only C^0 continuous in the manifold region. The LPDF of points in zone 1 and zone 3 are realized by their distance to the border of their zones, while the LPDF of points in zone 2 are realized by their distance to the corner point c	48
4.4	The burning time distribution on the board and fin example (top), and the burning front at three different time points (bottom). The grassfire starts from the borders, and burns forward at a geodesic uniform speed. At the beginning, the burning front on the top fin is ahead of that of the board (bottom left). This lasts until they reach the point k (called the “kink” point) (bottom middle). Afterwards, the burning fronts on the top fin merges with that of the front sheet, while the burning front of the back sheet advances faster (bottom right).	51
4.5	The decomposition of a medial axis (a), and the shortest path between points p_1, p_2 on $M^{(2)}$ realized by a non-crossing curve (in yellow) (b). The medial axis is decomposed into 2-dimensional ($M^{(2)}$, in grey), 1-dimensional ($M^{(1)}$, in blue), and 0-dimensional ($M^{(0)}$, in red) regions, and it is bounded by curves and points (∂M , in black). $M^{(2)}$ in general is noncompact (open at the dashed lines in b), so the shortest geodesic path may not always exist (e.g., p_1, p_2). To overcome this problem, we define <i>non-crossing</i> curves which are allowed to touch singular points but not allowed to cross singular curves, so that the shortest geodesic path between two points can always be realized by a non-crossing curve.	54
4.6	The regular neighborhood radius ($R_M(x)$) at a point x (a), and its disk types $D_M(x)$, including D_1, D_2 and D_3 (b).	56
4.7	The shortest geodesic path (in yellow) between points x, y on a medial axis made from a large board at the bottom and the side of a hollow cylinder on top. It is an example for demonstrating when the burning path can get arbitrarily close, but not equal to, the shortest geodesic path.	57

5.1	The LPDF distribution on the board and fin example (top), and the burning time distribution (bottom). The iso-curves are drawn at medial points with the same function value. The function distribution is also drawn only on the board (right column).	62
5.2	The medial curves defined from burning paths. Black arrows indicate the burning paths of the extended grassfire. The end points of these burning paths are medial curve points.	63
5.3	The medial axis (black line) of a very thin 2D shape and its EMA (black dot) (top), the medial axis (red curve) of its perturbed shape within ϵ -Hausdorff distance, whose EMA (red dot) is shifted.	64
5.4	A surface model of a protein (a) and the abstract illustration of its secondary structure (b) (α -helix in green and β -sheets in cyan).	66
5.5	Two bones from two different subjects (top and bottom). Each bone is illustrated by a surface model, the model with its medial axis embedded colored by an approximated burning time using our own discrete algorithm [46], and the model with the extracted medial curves embedded. Note that even size and shape of the same bone vary in the two subjects locally, the burning time and medial curves are rather similar. Therefore, they can be used for shape matching.	67
A.1	Notations used in the proofs.	70

Chapter 1

Introduction

1.1 Medial Axis

The medial axis was first introduced by Blum [14] via a grassfire analogy. A grassfire is lit at the boundary of an R^2 shape and burns inwards at a uniform speed. The medial axis is formed at the locations where the grassfires from two different directions meet and quenches (Fig. 1.1 (a)). Equivalently, the medial axis can be defined as the union of the centers of maximally inscribed discs which touch the shapes at least bitangentially [64]. This definition is the same as the union of the points that have at least two closest boundary points. Fig. 1.1 (b) shows a medial point that satisfies these conditions. From this definition, we can easily derive the spoke representation [64] at every medial point (Fig. 1.1 (c)). At least two equally long spokes start from any medial point and end at boundary points. The spoke representation also guarantees an accurate reconstruction of the original shape that indicates the equivalence of the medial axis to the original shape [64].




Figure 1.1: Medial axis defined as the shock graph in the distance transform (a) (image courtesy of Attali et al. [7]), medial axis defined as the union of the centers of maximally bi-tangentially inscribed balls (b), and the spoke representation of medial axis.

The medial axis explicitly encodes many important shape properties. The spoke representation makes it easy to see that shape boundaries are symmetric relative to the medial axis, and that the medial axis lies in the middle of the shape. The length of the spoke reflects the local thickness. Each branch corresponds to one part of the shape and the relationship amongst these branches suggests the relation amongst shape parts. In addition to these obvious properties, it has been proven that the medial axis always at least one dimension less than the original shape, i.e., the medial axis of an R^n shape is always at most R^{n-1} [14], and the medial axis has the same homotopy/connectivity as the original shape [45].

Due to this set of intrinsic properties, the medial axis is widely used in many applications and in many different fields. Medial axes are extracted from shapes as small as protein molecules [10, 9] and as large as galaxies [67]. Often, medial axis is used to visualize the structure and the connectivity of such complex shapes, since it faithfully captures their topology and it makes their visualization simpler due to its lower dimension. The centeredness of medial axis makes it possible to achieve a better coverage of wireless sensors [41]. Robots navigating along medial axis stay as far away from obstacles as possible [38]. The correspondence between a medial axis branch and a shape part leads to a natural algorithm of shape segmentation by decomposing shape parts at branch junctions [59, 60]. A local object coordinate system can be easily defined based on the spoke representation at every medial point, and any affine transformation defined on the medial point can be propagated to deform shape parts [76, 77]. The advantage of medial axis-based deformations is that they offer a more intuitive interaction for users to deform models. The homotopy equivalence between medial axis and its original shape has been used to analyze and modify the topology of the original shape [27, 42], or to identify the topological errors and to mend them [42]. The union of balls at all medial points results in the original shape and editing the medial axis could therefore be used to easily alter the appearance of the input shape. Some have used this power to simplify the original shape by simplifying its medial axis [70]. More sophisticated morphological transformations, such as careful enlarging and shrinking, have been used to smooth out small surface features like bumps and dents while retaining big features such as big sharp corners [49]. The hierarchical structure of medial axis can be easily described by a few simple rules that can be used as a grammar for a procedural modeling of trees [62]. The aforementioned applications are all intra-shape, but the medial axis is also used in

inter-shape applications such as shape matching and retrieval [69]. The similarity in the hierarchical structures of medial axes suggests the overall shape similarity between two shapes, and hence such techniques are less sensitive to the surface noise and isometric deformations when compared to techniques relying on surface features such as curvature.

1.2 Instability of Medial Axis

The damaging weakness of medial axis is its *instability*. Even a small amount of surface noise can create significant extraneous branches on medial axis (Fig. 1.2). These extraneous branches are often referred to as the unstable parts since a little perturbation on the surface can make them disappear from the medial axis. Medial axes with these unstable parts pruned are more desirable in many applications. In scientific visualization, biologists prefer to see only those parts of the medial axis that correspond to the structure of a protein, and they are not interested in the parts created by the bumpy surface features. The structure of medial axis is analyzed and decomposed into branches at junctions, so that the shape can be segmented accordingly. However, the extraneous branches do not correspond to any meaningful shape parts, and they should be ignored in the segmentation process. In the medial axis based shape matching and retrieval, if the noisy branches are not removed before comparing the medial axes, then two similar shapes like the two in Fig. 1.2 will be missed. Hence, identifying and pruning the unstable parts is important for many applications.




Figure 1.2: Medial axis of a rectangle (a), and medial axis of a rectangle with a very small bump (b).

To differentiate the stable parts from the unstable ones, it is necessary to define an significance metric that measures the stability of every medial point. We call this metric a significance measure and the larger the measure is, the more stable the medial point is. A lot of effort has been made to define different significance measures in order to effectively differentiate the stable and unstable parts. There exist two main categories of such measures: *local* and *global*. The majority of existing measures are local. Local measures can be computed efficiently based on the local information associated with the medial points, but they are scale unaware. In consequence, they may accidentally identify a long but thin tube as unimportant. In contrast, the global measures look at a large scale to compute the significance measure, and therefore they are more accurate. However, such measures are scarce and, in general, more expensive to compute. A good significance measure that can capture global shape properties in a large scale and yet can be computed efficiently is still needed.

1.3 Lower Dimensional Medial Forms

Lower dimensional medial forms, i.e., shape centers in R^2 and medial curves in R^3 , are more desirable than the medial axis in many applications. In R^2 , the shape centers can be used to translate two shapes to achieve a good initial alignment for shape correspondence establishment. When creating an animation, a motion path specified for the shape center is better than for an arbitrary point, especially if the shape is undergoing an isometric deformation following the motion path. Another application of the shape center is in map creation, as the shape centers are good locations to put text annotations. In R^3 , medial curves are mainly used in three fields: medical, computer graphics, and CAD. In the medical field, medial curves are used to describe thin structures such as the backbones of proteins [10], or the center lines of the tubular organs [11], such as bronchia [54] or intestines [40]. The center lines can be used to navigate a virtual camera in an endoscopy system, or they can be straightened to measure the length of the organ [66]. In computer graphics, medial curves are used for shape decomposition [58, 48] because one curve branch corresponds to one part of a shape. Another popular application of medial curves in computer graphics and computer vision is shape matching. It is especially effective to use medial curves in matching two shapes when their structures are similar but their postures are different

[22, 1]. Medial curves can also be used as handles to deform shapes by applying inverse kinematics (IK) [13]. In CAD, medial curves with an associated thickness can be used for collision detection [37], and many other applications.

Good definitions of the lower dimensional medial forms are still needed. In R^2 , there are different definitions of the shape centers, but none of these approaches can simultaneously guarantee that the centers are in the interior, are unique, and are stable. Note that these properties are very important for the aforementioned applications. In R^3 , many heuristics-based algorithms have been carefully devised to compute medial curves, but they are not based on a mathematical definition, and many of them are tuned for specific applications. As far as we know, there is only one definition of medial curve, given by Dey and Sun [29], but this definition has its own limitations (Section 2.3.2). Without a mathematical definition, it is not possible to study the properties of the computed lower dimensional medial forms, and the results can only be demonstrated pictorially on a small set of examples. In [23], Cornea et al. listed a set of criteria for judging the quality of medial curves. Medial curves should be centered, homotopic (having the same topology as the original shape), connected, invariant under isometric transform, robust (insensitive to boundary noise), and thin. Medial curves computed as the singularities in the distance field [57] or the potential field [24] cannot guarantee homotopy. The Voronoi Diagram based methods can guarantee homotopy but they are very sensitive to boundary noise [74]. Thinning on a discrete grid often results in medial curves that are not thin, and their shape changes when the model is rotated [53].

1.4 Our Contributions

In this dissertation, we address several problems that have been mentioned in the previous sections. Our contributions are:

- In R^2 , we propose a new unified definition for a global shape measure as well as the shape center. The proposed global shape measure, which we call the extended distance function (EDF), has several nice properties such as continuity, constant gradient, and many others. It captures the global shape elongation

property, and hence it is more stable to boundary noise than other existing measures. Our shape center, called the extended medial axis (EMA), is directly derived from the definition of the EDF. We show that our EMA is interior, unique and stable to boundary noise. We prove that the EDF is equivalent to the extended grassfire burning process on the medial axis, which leads us to devise an iterative discrete algorithm to compute the EDF and EMA. We demonstrate the usefulness of the EDF and EMA in various applications such as medial axis pruning, shape alignment, and boundary signature for shape matching.

- We demonstrate that a direct extension of EDF to 3D leads to a distance function that has several limitations. In contrast, we show that a distance function derived from the burning time of the extended grassfire is well behaved, even when used in 3D. We formalize the latter definition and we describe its properties.

This dissertation gives a good foundation to many problems for future research. Most of the open questions exist in R^3 , and they can be grouped into three categories: *theory*, *algorithms* and *applications*. Unlike in R^2 , the relation between the EDF and the shape properties is not obvious in R^3 , which makes it more difficult to use in practice. Another problem is how to define homotopy-preserving medial curves directly based on the EDF in R^3 , or more generally, whether we can define even lower dimensional medial form, such as shape centers, in R^3 . Even though empirically we have observed the stability of the EDF and the EMA under significant amount of boundary perturbation, a rigorous mathematical proof of this stability is still lacking in both R^2 and R^3 . Also, an accurate computational algorithm for EDF and EMA on the medial axis in R^3 is still needed. In case an accurate algorithm for general cases is infeasible, it would be interesting to explore some specific scenarios where it is possible. Another important step would be to derive an approximation algorithm for computing the EDF and EMA, with theoretical guarantees bounding the error. Last but not least, it is important to explore possible application of EDF and EMA in R^3 .

1.5 Dissertation Overview

In Chapter 2, we review the existing methods for computing the medial axis, the existing significance measures for identifying the stable parts of the medial axis, and the existing definitions of the lower dimensional medial forms and their computational algorithms. We also discuss the advantages and the drawbacks of these approaches.

In Chapter 3, we give a mathematical definition of our significance measure (EDF) and shape center (EMA) in R^2 . Based on the properties of our definition, we devise a discrete iterative algorithm to compute the EDF and the EMA. We compare our measure (EDF) and our shape center (EMA) with existing ones, and empirically show that our measure and shape center are more stable than existing approaches. Finally, we demonstrate the usefulness of the EDF and the EMA through applications.

In Chapter 4 we first show that a naive extension of EDF from 2D to 3D leads to a flawed measure that is difficult to compute. We then define a new measure that is directly derived from the extended grassfire burning on the medial axis in 3D and we show that this measure is well behaved.

In Chapter 5, we elaborate on the possible future directions outlined above.

Chapter 2

Related Works

In this chapter, we review the existing related works and discuss their advantages and drawbacks. These existing works include the computational algorithms of the medial axis, the significance measures to identify the stable parts of the medial axis, and the definitions of lower dimensional medial forms.

2.1 Computation of the Medial Axis

Even though the medial axis has simple definitions, nice properties and a close relation to the shape representation, an accurate computation of the medial axis is usually infeasible for generic shapes, because of its high algebraic degree. Therefore many algorithms focus on calculating an approximation to the medial axis instead. The exact medial axis can be computed only for a set of simple shapes, such as point sets, polyhedra, and union of balls. We review the exact computational algorithms for a set of simple shapes first, and then we give an overview of the approximation methods for generic shapes. Among the approximation methods, we provide a description of the state-of-the-art Voronoi diagram based methods, because this is the very method we use for approximating medial axis of generic shapes. Other approximation methods are only briefly reviewed and we refer interested readers to the recent book [64] for details, since approximation algorithms are not closely related to the topic of this dissertation.

2.1.1 Accurate Medial Axis Computation

There are three types of simple shapes, for which the accurate medial axis computation is feasible: Point sets, polyhedra, and a union of balls. The medial axes of these three shapes have at most quadratic algebraic degree, which makes them tractable.

Point sets The medial axis of a point set in R^n is equivalent to its Voronoi Diagram [12], so to compute its medial axis is to compute its Voronoi Diagram. Voronoi Diagram computation is a solved problem in the computational geometry community. The Voronoi Diagram decomposes the space into cells, one cell for each input point. The distance from any point in the interior of a cell to its corresponding input point is smaller than the distance to any other input points (Fig. 2.1 (a)). The elements of the Voronoi diagram that have dimension less than n are the medial axis (Fig. 2.1 (b)), since any point from these elements has at least two closest input points.




Figure 2.1: The Voronoi Diagram (a) and the medial axis (b) of a point set. The Voronoi Diagram partitions the space into cells, colored differently for each input point. The boundary of these cells is the medial axis of the input points since any point on the boundary has at least two closest boundary points.

Polyhedra Polyhedra are closed shapes bounded by piecewise linear elements, and medial axes of such shapes consist of at most quadratic elements. Many algorithms have been proposed to trace out the medial axis [25, 63, 50]. In particular, Sherbrooke et al. classified the elements of the medial axis into three types: Medial sheets, medial seams and medial junctions (see Fig. 2.2 (a)), and their relations were described in [63]. Based on these relations, they traced out the medial seams from the medial junction

points, and the medial sheets from the medial seams. Fig. 2.2 shows one example of the medial axis of a polyhedron. Culver et al. [25] improved this algorithm by representing each element as a quadratic equation to achieve an accurate computation, while [63] only represented the positions of the medial elements in floating precision arithmetic.




Figure 2.2: The classified elements of the medial axis of a polyhedron (a), an input polyhedron (b) and its medial axis (c) (Image courtesy of Culver et al. [25]).

Union of balls Attalia and Montanvert [6] first characterized the medial axis of a union of balls in R^n . They pointed out that the medial axis is piece-wise linear, and that it consists of only two types of points: (1) the singular faces in the α -shape of the union of balls [33], (2) the subset of the Voronoi diagram of the singular points on the boundary surface, which has closest boundary points only from these singular points. Singular faces in an α -shape are those that do not bound any element of dimension n , such as the highlighted edge in Fig. 2.3 (b). The singular points on the boundary of the union of balls are those points intersected by at least n balls in R^n . However, this characterization does not directly lead to a computational algorithm. In [4], Amenta and Kolluri characterized the medial axis into a simpler form and proposed an algorithm to compute it. They simplified the second type into the intersection between the Voronoi diagram of the singular boundary points and the α -shape. Fig. 2.3 shows a simple shape made from the union of balls, its α -shape and its medial axis.

2.1.2 Medial Axis Approximation

The high algebraic degree of the medial axis makes its accurate computation difficult, therefore the medial axis is in general approximated. The approximation algorithms




Figure 2.3: A shape made from a union of balls (a), its α -shape (in gray) with singular elements highlighted in red (b), and its medial axis in red and blue (c). The medial axis consists of the singular elements in the α -shape (in red) and the intersection (in blue) between the Voronoi Diagram (dashed lines) of the boundary singular points (black dots) and the α -shape (in gray) (image courtesy of Amenta and Kolluri [4]).

are categorized into four types based on the input shape representation and how the medial axis is approximated. These four categories are *Voronoi diagram based*, *distance filed based*, *thinning*, and *hybrid*.

Voronoi Diagram of sampled points on the surface of the shape Brant and Algazi [16] proved that, in R^2 , the inner Voronoi vertices and edges converge to the real medial axis when the sampling density goes to infinity, and hence the inner Voronoi vertices and edges can be a good approximation of the medial axis. In R^3 , Amenta et.al [2] showed that not all inner Voronoi vertices are close to the medial axis, but only a subset of the inner Voronoi vertices are, and this subset converges to the medial axis as the sampling density approaches infinity. Such Voronoi vertices are farthest away from their corresponding sampling points in the inner direction, and they are called the inner poles. In [2], the authors also proved that the union of the inner Voronoi balls and its power crust are both a good shape approximation with a bounded geometric error and with the same topology. An inner Voronoi ball is centered at an inner pole and has radius equal to the distance from this inner pole to its corresponding sample point. In [3], Amenta et al. gave an algorithm to compute a medial axis approximation, called the power shape. It consists of the inner poles, whose connectivity is established based on the connectivity of the power diagram of the inner Voronoi balls. Tam and Heidrich in [71] pointed out and fixed a few problems with the power shape. First, the power shape includes flat tetrahedra

which contradicts the fact that the medial axis is at most two dimensional in R^3 . Second, the computation of the power shape easily generates a degenerate point for duplicated times, which causes cracks in the result. To fix these problems, Tam and Heidrich [71] took the intersection between the power shape and the accurate medial axis of the union of inner Voronoi balls using the aforementioned algorithm [4], and they also detected and filtered the repeated degenerate points.

The medial axis approximated from the Voronoi diagram is homotopy equivalent to the original shape, and it has a convergence guarantee. This is the same technique we use to compute the medial axis in R^2 in Chapter 3. The drawback is that extra care needs to be taken to deal with the degenerate cases, especially in R^3 , and the sampling density is required to be very high near the surface features.

Distance field based medial axis approximation The distance field represents the distance from any interior point to its closest boundary point. The input shape could be a grid representation or a polygonal mesh. If the shape is a polygonal mesh, then its interior is discretized into a grid representation [35] so that the distance field can be represented by associating each interior point with a distance. Many algorithms have been proposed to approximate the distance field in various dimensions [26, 15, 72, 39, 65]. The medial points are detected based on the local Laplacian [75], Hessian matrix [75], or the average outward flux of the gradient field [32, 65]. The detected singular points may not be connected, and hence this technique is often followed by a homotopy-preserving thinning (see the following paragraph). The interior points with the same distance to the boundary form a level set, which can be evolved from the shape boundary. Some methods directly evolve the boundary surface inward based on the Eikonal equation, and during the evolution, the shock locations are detected and preserved at points where the normal or the curvature does not exist [55].

The distance based methods are, in essence, based on the definition of the medial axis. However, the detection of the singular points can be error prone. Extra singular points lead to extra branches after thinning. A good distance approximation on a grid representation is often very complex with a lot of issues to deal with.

Thinning on the grid representations of the input shapes Thinning methods are widely used in computer vision for pattern recognition since the data come naturally as grid images. The input image consists of the object and the background pixels, and thinning turns the object pixels into the background pixels layer by layer, mimicking the grassfire burning process, until there is only one thin layer left. The remaining thin layer is the approximated medial axis. The key part of these algorithms is a set of criteria specified on the local neighborhood ($2 * 2$ or larger in R^2) of a grid point to determine if this point is simple or if it is medial. A simple point can be removed without altering the topology, while a medial point should be preserved. Surveys [43, 78] offer a good in depth review of this category.

Thinning methods are very simple to implement and most of them are homotopy preserving. The major drawback is that they are biased by the direction of the grid and hence, they are not rotation invariant, as the underlying distance field generated by the thinning process is based on the city-block distance instead of the Euclidean distance. Also, the resulting medial axis may not lie in the middle. For example, a shape with an even number of layers can never have a layer that is right in the middle. The medial axis that is represented as pixels has a bumpy look, and a smoother medial axis requires a higher resolution of the original input, which in turn makes it more expensive to thin. Also, the discrete nature of the grid representation may introduce additional branches that do not correspond to any surface features.

Hybrid methods Different methods can be easily combined, so that their advantages are united. For example, [75, 51] computed the distance field and identified a set of medial points based on this distance field first, and then discretized the model into a grid representation that was further thinned with the detected medial points preserved. This combines the relatively more accurate medial points detection of the distance based methods and the homotopy preservation of the thinning methods.

In this section, we reviewed the existing algorithms for computing the medial axis accurately on a set of simple shapes and approximately on generic shapes. In particular, the Voronoi Diagram based methods are simpler than other approximation methods, and they have a theoretic convergence guarantee, which is the very reason we use them in this dissertation.

2.2 Significance Measures

The instability of the medial axis is a well-known problem. In most applications, it is more desirable to find the stable parts of a medial axis as they correspond to the global shape properties. Unstable extraneous branches correspond to small surface features and noise, and they should be removed from the medial axis. In order to detect the stable parts, medial points are often associated with a significance measure. Parts with significance measures smaller than a specified threshold are pruned, so that only the stable parts remain. Numerous strategies have been proposed in the literature to measure the significance. There exist two main categories, *local* and *global*, that we review in this section.

2.2.1 Local Significance Measures

The majority of pruning algorithms are based on local measures due to their simplicity, since they can be easily computed from the local information associated with the medial points. Below is a description of several most popular local measures.

The separation angle For any medial point, its separation angle is formed between the two rays starting from this medial point and ending at its two closest boundary points. If there are more than two closest boundary points, the separation angle is defined as the largest among all possible angles. A lot of pruning algorithms rely on this measure [8, 5, 30, 36, 68, 6]. Fig. 2.4 shows the definition of the separation angle (Fig. 2.4 (a)) and the pruned result based on this measure (Fig. 2.4 (b, c)). As indicated by Fig. 2.4 (c), a naive pruning based on the separation angle alone could lead to a disconnected result. In [68], Sud et al. proposed a homotopy preserving pruning strategy. The key idea is to prune from the border of the medial axis, so that an interior point is not removed when its measure is smaller than the threshold while being surrounded by medial points with measures larger than the threshold. Note that such a thinning strategy could be easily combined with any significance measure to preserve homotopy.




Figure 2.4: The separation angle $\theta(x)$ at the medial point x (a), the pruned medial axis with the threshold $\frac{\pi}{2}$ (b), and the disconnected pruned medial axis of another model with the threshold $\frac{\pi}{3}$ (c). The separation angle is defined by the two rays from the medial point x to its two closest boundary points y_1 and y_2 (a). Note that a naive pruning based solely on separation angle could easily break the topology (image courtesy of Sud et al. [68]).

The combination of the separation angle and the local thickness The key problem with the separation angle is that it is scale unaware. For example, separation angles of medial points corresponding to the bottom rectangle are the same as those of medial points corresponding to the small blunt bump in Fig. 2.4 (c). Attali and Montanvert [8, 5] combined the separation angle with the local thickness (the distance from the medial point to one of its closest boundary points). They observed that the surface noise can be identified by either a small separation angle or by a small local thickness (Fig. 2.5 (b)), therefore only points with both values above specified thresholds are preserved. For example, those medial points with their separation angles and their local thickness in the upper right corner in Fig. 2.5 (b) are preserved, resulting in the bottom one in Fig. 2.5 (a). However, the local thickness condition may eliminate some parts that are important to the shape composition, such as a part of a medial axis corresponding to a long, but thin, dangling tube (Fig. 2.5 (c)).

The λ -medial axis Chazal and Lieutier [19, 20] proposed a subset of the medial axis called the λ -medial axis and studied its properties. To describe it, they define a local measure called the feature size $r(x)$ at any medial point, which is the radius of the smallest ball that encloses all of the closest boundary points to this medial point. Note that this $r(x)$ is different from the local thickness. For example, the $r(x)$ of a medial point with two closest boundary points is the half length between the two boundary points, unlike the local thickness which could be larger (see Fig. 2.6 (a)).




Figure 2.5: An input model with its medial axis before and after pruning (a), the ρ and θ distribution of all medial points (b), and the pruned medial axis of a model with a snail on a tree branch (c). The medial points corresponding to the upper right quadrant are those preserved after pruning in the bottom figure of (a). Note that the prominent features, antennae of the snail, are accidentally pruned based on the combination measure of ρ and θ , because they have small thickness (ρ) (images (a)(b) courtesy of Attali and Montanvert [8]).

The λ -medial axis is defined as all the medial points whose $r(x)$ are no less than a given λ . Listed below are a few properties of the λ -medial axis that have been proven in [19]:

- If $\lambda \leq wfs$, then the λ -medial axis is guaranteed to have the same homotopy as the medial axis, where wfs is short for the weak feature size. The weak feature size is defined as the minimum of feature size $r(x)$ at all critical medial points. A medial point is critical if it is a critical point in the gradient field of the distance field.
- If the Hausdorff distance between two inputs is bounded, then the Hausdorff distance between their λ -medial axes is bounded too, when the λ satisfies certain conditions. This indicates that a λ -medial axis of an original shape could be approximated within a bounded geometric error as long as the shape approximation is within a bounded geometric error. Note that the approximated shape could be noisy.
- As the Hausdorff distance between the approximation and the shape approaches 0, then the λ -medial axis converges to the real medial axis (the λ should also be constrained by the Hausdorff distance).

Attali et al. [7] drew the connection between the λ -medial axis and their local measure which combines the separation angle and the local thickness. Note that $r(x) = \rho * \sin(\theta)$ holds at the majority of the medial points, where ρ is the local thickness and θ is half of the separation angle. To give a threshold on $r(x)$ is equivalent to draw a hyperbola in the plot of ρ and θ in Fig. 2.5 (b) which is similar to [7].




Figure 2.6: The medial axis (a) and its λ -medial axis when λ is slightly larger than its wfs (b). The difference between the local thickness $\rho(x)$ and the feature size $r(x)$ can be seen at the annotated medial point x (in black), where the red circle is the bi-tangentially inscribed medial circle and the blue circle is the minimal circle that encloses the two closest boundary points (in red). The three hollow points are the critical points and wfs is the smallest feature size of the three. Note that the topology of λ -medial axis is different from the original model when $\lambda > wfs$ for this model (image courtesy of Attali et al. [7]).

The major drawback is that the λ -medial axis is guaranteed to be homotopy preserving only when the λ is smaller than the wfs . If the λ is too big, then the λ -medial axis could have a different homotopy (see Fig. 2.6 for an example). Unfortunately, the wfs can be arbitrarily small for general shapes.

2.2.2 Global Significance Measures

Due to the limitations of local measures, many researchers have resorted to global measures. The existing global measures determine the significance of a sub-part or of a whole branch of the medial axis by the elongation (e.g. erosion thickness [61]) or by the change in the reconstructed shape caused by its absence (e.g., potential residue [52], circularity residue [52], chord residue [52], erosion area [61]).

Erosion thickness Roughly speaking, the erosion thickness is measured as the difference between the length of the branch and the thickness of the branch, therefore it effectively captures the elongation of the branch. Fig. 2.7 illustrates the accurate computation of the erosion thickness. At a medial point, the length of the branch is measured as the summation of the geodesic length along the medial axis from this medial point to the end point of the branch and the thickness at this end point, which is approximately the length from the medial point to a boundary point. The branch thickness is defined as the local thickness at the given medial point. The erosion thickness is scale dependent. For example, a large blunt feature can have a very large erosion thickness because of its size rather than because of the shape elongation. In order to remove the part of the medial axis corresponding to such a blunt feature, the threshold needs to be increased, which may accidentally remove a large part of the medial axis that corresponds to features at a small scale (see Fig. 3.6 in Chapter 3 for one example).




Figure 2.7: The erosion thickness (a), the three residues (b), and the erosion area (c). The erosion thickness at q is measured as the difference between the length of the branch and the thickness of the branch, i.e., the elongation of the branch. The length of the branch is measured as the summation of the geodesic distance $d(p, q)$ from the medial point q to the branch end point p and the local thickness $\rho(p)$ at p . The thickness of the branch is measured as the local thickness $\rho(q)$. $L(B)$, $L(A)$ and $L(C)$ are the lengths of the boundary curve, the arc, and the chord connecting the two closest boundary points respectively. The potential residue, the circularity residue and the chord residue at q are $L(B)$, $L(B) - L(A)$, and $L(B) - L(C)$. They all estimate the change in the boundary length when the branch attached to the medial point q is removed. The erosion area of q is colored in gray in (c), which estimates the change in the area of the reconstructed shape in the absence of the branch attached to q (image courtesy of Shaked and Bruckstein [61]).

Potential residue, circularity residue, chord residue These measures were first defined in R^2 [52], and later they were extended to R^3 [29]. In R^2 , *Potential residue* is defined as the length of the boundary curve at a medial point, which is measured as the geodesic distance between its two closest boundary points on the boundary curve. *Circularity residue* is measured as the difference between the length of the boundary curve and the length of the arc between the two closest boundary points, where the arc connects the two closest boundary points on the bi-tangentially inscribed medial ball at the medial point. *Chord residue* is measured as the difference between the length of the boundary curve and the length of the segment connecting the two closest boundary points. All three measures are similar because they all try to capture the change in the length of the reconstructed boundary if the branch of the medial axis is removed. The reconstructed boundary is represented as nothing, the arc, or the chord, for the three residues respectively. Fig. 2.7 shows them in different colors. The potential residue has been extended to R^3 as the geodesic distance between the two closest boundary points along the surface [29]. Note that the geodesic computation has to be carried out often, so the method can be very computationally expensive. There are two other very important limitations of potential residue. First, it is not continuous at junctions. Second, it is not robust with respect to boundary perturbations, which can significantly enlarge the boundary distance.

Erosion area In R^2 , the erosion area measures the change in area in the reconstructed shape by removing a segment on the medial axis (Fig. 2.7). Note that the change in the area is very expensive to compute. One can imagine extending this algorithm to R^3 by measuring the change in volume caused by removing a part of the medial axis, which is even more expensive to compute.

In this section, we reviewed several significance measures. In general, the local significance measures can be computed efficiently based on the local information at the medial points, but they can misclassify some important medial points as unimportant, or vice-versa. The global significance measures are computed at a large scale, but they are very expensive to compute.

2.3 Lower Dimensional Medial Forms

The lower dimensional medial forms (e.g., the shape centers in R^2 , the medial curves in R^3) are desirable in many applications (Chapter 1.3). There are many definitions of the shape centers in R^2 , but none of them are guaranteed to be in the interior, unique and stable simultaneously. Many heuristics-based algorithms have been proposed to compute the medial curves in R^3 , but only a few are based on a mathematical definition. As far as we know, there is only one mathematical definition of medial curves [29]. The lack of mathematical definition makes the computed medial curves hard to analyze and hence the results are often judged visually on a small set of shapes. In this section, we review the definitions of the shape centers and the medial curves, and we refer interested readers to the recent comprehensive survey of computational methods of the medial curves by Cornea and Silver [23].

2.3.1 Center Points of 2D Shapes

The most common way of defining a shape center is the “center of mass” or the centroid, which minimizes the sum of squared Euclidean distances to all points either on the boundary of the shape or over the entire shape. However, the centroid may lie outside the shape if it is non-convex, and can be unstable under large shape deformations (see Fig. 2.8).

In the computational geometry literature, there are a number of alternative definitions that utilize geodesic distances within the shape to prevent the center from going outside. The *geodesic center* [56] minimizes the maximum geodesic distance to any point in the shape. The *link center* [44] in a polygonal region minimizes the maximum number of straight line segments in the geodesic path to any point in the shape. The *geodesic median* [34] minimizes the average geodesic distance in the L_1 norm to any point in the shape. However, these center locations may lie on the boundary (e.g., a concave vertex in a non-convex shape), and hence they are not always strictly interior. Note that, among these definitions, the link center is not uniquely defined.

In geography, one way to define the center of a geographical region is the furthest point from the boundary (or center of the largest inscribed disk). Although strictly

interior, such a center is obviously not unique. This, and other centers, are compared with our extended medial axes (EMA) in a simple 2D shape in Fig. 2.8.

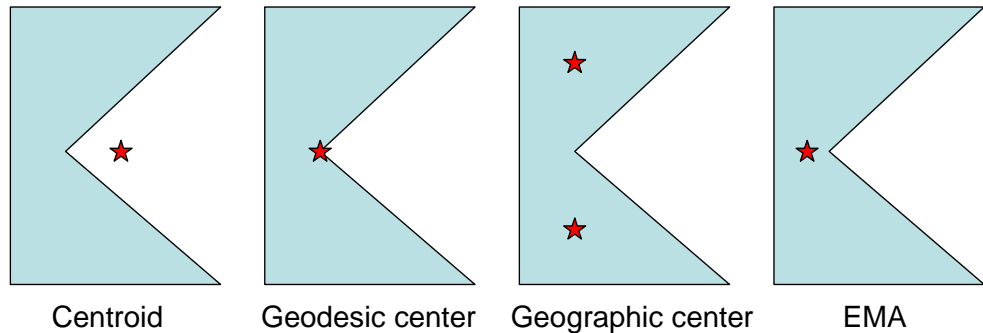


Figure 2.8: Comparison of several center definitions. The last one is defined by our definition (more details in Chapter 3).

Ogniewicz and Ilg [52] showed that the PR measure has a unique local maximum on the medial axis for a simply connected shape. This maximum is a unique, interior center point. However, as we demonstrate in Chapter 3, this local maximum can be sensitive to non-uniform boundary perturbations.

2.3.2 Definition of Medial Curves

As far as we know, there is only one mathematical definition of the medial curves, proposed by Dey and Sun [30]. Their definition is similar to the definition of the medial axis. They first defined a function called the medial geodesic function (MGF) on the medial axis, and then they defined lower dimensional medial points as the singular points of this function. The authors also proposed an approximation algorithm to compute the medial curves based on their definition.

MGF at a regular medial point is defined as the geodesic distance between the two closest boundary points along the boundary surface. A medial point is regular if it has only two closest boundary points. Fig. 2.9 (b) is a plot of MGF on the medial axis of the human model. Dey and Sun proved a few properties of MGF: (1) The singular points of MGF are thin, i.e., 1D or 0D; (2) MGF is continuous and differentiable in the manifold part of the medial axis. A point is manifold if its local neighborhood

can be flattened down to a plane without an overlap. For example, a point on a wall is manifold, but a point on the bottom line of a wall separating two rooms is not manifold. They conjectured two more properties and sketched the proofs: (1) The MGF does not have local minima; (2) The singular points of the MGF lie in the middle, or in other words, singular medial points have two boundary geodesic paths of equal lengths.




Figure 2.9: An input human model (a), its medial axis colored based on the medial geodesic function (MGF) (b), and the extracted medial curves (c) (image courtesy of Dey and Sun [29]).

Dey and Sun’s proposed method for approximating the medial curves consists of many steps. First, the medial axis is approximated and the MGF is computed based on its definition. Then, the singular points of the MGF are detected by thresholding the average outward flux, similar to the techniques used to detect the medial points in the distance field [65]. Finally, a homotopy-preserving erosion is run on the medial axis, while preserving the detected singular points. The erosion result is the medial curve.

Dey and Sun [30] gave the first definition of medial curves based on the MGF and proved the thinness of medial curves and the continuity of MGF. However, rigorous proofs for the nonexistence of local minima of the MGF and the centeredness of

medial curves are still lacking. In addition although they computed the curves in a homotopy preserving manner, which is guaranteed by the erosion algorithm, it is still unknown if the singular points of the MGF alone are homotopic equivalent to the input shape.

In this section, we reviewed the existing definitions for the shape centers and the medial curves, and we discussed their limitations. These limitations show that good definitions for the lower dimensional medial forms are still needed.

2.4 Summary

In this chapter, we reviewed the methods for computing and approximating the medial axis, the significance measures for pruning the instable parts of the medial axis, and the definitions for the lower dimensional medial forms. We pointed out that there is still a dire need for a significance measure that captures the global shape properties while being efficient to compute. Good definitions of the lower dimensional medial forms still need exploration. The main goal of this dissertation is to resolve these limitations.

Chapter 3

Extended Grassfire Transform on Medial Axes of 2D Shapes

In this chapter we present a uniform approach to define a global shape measure (called the *extended distance function*, or EDF) along the 2D medial axis as well as the center of a 2D shape (called the *extended medial axis*, or EMA), that we published in [47]. We reveal a number of properties of the EDF and EMA that resemble those of the boundary distance function and the medial axis, and show that EDF and EMA can be generated using a fire propagation process similar to Blum's grassfire analogy [14], which we call the extended grassfire transform. The EDF and EMA are demonstrated on many 2D examples, and are related to, and compared with, existing formulations. Finally, we demonstrate the utility of the EDF and EMA in pruning medial axes, aligning shapes, and shape description.

3.1 Formulation

3.1.1 Motivation

Our definitions of the extended distance function and the extended medial axis are motivated from those of the boundary distance function and the medial axis. In a 2D shape O , the distance from an interior point x to the boundary of O can be defined as the radius of the largest circle centered at x and inscribed in O . A point is on the medial axis if its largest inscribed circle touches the boundary of O at two or more

points. Due to the isotropic nature of the circle, the distance function at x captures the amount of uniform shape expansion around x , and the medial axis is where such expansion is constrained at two or more sites, and hence “maximal”.

To capture the elongation, or “side-ways” expansion, of a shape around a point x on the medial axis, our extended distance function essentially measures the half-length of the longest “tube” centered at x and inscribed in O . This tube extends longitudinally along the medial axis, rather than uniformly as in the case of the circle. The extended medial axis consist of those points on the medial axis whose longest inscribed tubes are confined at both ends, and hence are where the elongation is maximal.




Figure 3.1: Axes (blue) and their tubes (gray): (a) an axis with a single constrained end z , (b) an inscribed axis of x with a constrained end y , (c) an inscribed and maximal axis of x with two constrained ends y, z , (d) an inscribed and maximal axis of $x \in \tilde{M}$.

3.1.2 Definitions

We assume the 2D shape O is a closed set bounded by piece-wise C^2 smooth curves. The medial axis M of O are the closure of those points with two or more closest points on the boundary of O (or the “cut loci”) [73]. The regularity of the boundary implies a number of important properties of M , such as its homotopy equivalence to O [73], its low-dimensionality (1) [73], and its finite structure [21]. We denote the boundary distance function at $x \in O$ as $R(x)$.

We first introduce the notion of “axes” and “tubes”:

Definition 1. *Let $f : \mathbb{R}^1 \rightarrow \mathbb{R}^2$ be a local embedding of the real interval $D = [0, 1]$ onto M . The image $f(D) \subseteq M$ (noted simply as f hereafter) is called an axis.*

An axis is a path on M such that each interior point on the path has a manifold neighborhood. By local embedding (i.e., immersion), we allow an axis to be non-simple and hence contain loops. This relaxation is crucial to obtain some important properties later, such as the homotopy equivalence between the extended medial axis and the medial axis.

We call the union of all largest inscribed circles centered at points on an axis f the *tube* of f . Intuitively, the tube is formed by “rolling” a circle along f , while changing its radius according to the boundary distances on f . Note that when the axis is a non-simple path, the tube can “wrap around” and overlap itself.

Given some point x on an axis f , we are interested in the *radius* of f with respect to x , which is defined as the distance from x to the closer end of the tube of f :

Definition 2. *Given an axis f and a point $x \in f$,*

$$r_f(x) = \min_{y \in \partial f} (d_f(x, y) + R(y))$$

is called the radius of f with respect to x , where $d_f(x, y)$ is the geodesic length of segment $[x, y]$ on f . An end $y \in \partial f$ is called a constrained end with respect to x if it attains the minimum in this equation, and an unconstrained end otherwise.

Figure 3.1 (a,b,c) illustrate three different axes and their radii for a same point x on the medial axis. The first two axes each have one constrained end (z in (a) and y in (b)), where the radius of the axis is attained. The third axis has two constrained ends (y, z) due to symmetry of the shape, both attaining the radius of the axis.

The extended distance function at x is the largest radius of any axis containing x :

Definition 3. Given a point $x \in M$,

$$\tilde{R}(x) = \sup_{f \ni x} r_f(x)$$

is called the extended distance function (EDF) at x . The axis f that attains the supremum is called the inscribed axis at x .

Intuitively, the EDF captures the maximum amount of “side-ways” shape expansion on both sides of x . In Figure 3.1, the axes in (b,c) are inscribed axes of x and attain the EDF $\tilde{R}(x)$, which is the sum of the geodesic distance between x to the top-left end of the medial axis y and $R(y)$. The EDF can be understood as the half-length of the longest tube that can be centered at x .

The extended medial axis is the locus where the inscribed axis cannot be further expanded:

Definition 4. An axis f is called maximal if both of its ends are constrained. A point $x \in M$ lies on the extended medial axis (EMA) \tilde{M} if every inscribed axis of x is maximal.

Note that *all* inscribed axes of an EMA point need to be maximal. This requirement is important for distinguishing the center of the shape from the local symmetry centers. For example, even though the axis in Figure 3.1 (c) is an inscribed axis of x and is maximal, x has some other inscribed axis that is not maximal (e.g., (b)), and hence $x \notin \tilde{M}$. On the other hand, the point x in (d) is on \tilde{M} since all of its inscribed axes are maximal (one is shown in the picture). Observe that x in (d) is more “centered” with respect to the entire shape than the x in (c), the latter being centered only with respect to two symmetric shape parts.

We make a final note of the scenario in which the radius $r_f(x)$ is infinity, which happens when f is a non-simple path that travels infinitely on both sides of x . By the above definitions, both ends of f are constrained, hence f is maximal. On the other hand, $\tilde{R}(x) = \infty = r_f(x)$, hence f is an inscribed axis of x , and all inscribed axes of x have infinite radius. As a result, x is on the EMA \tilde{M} , since all of its inscribed, infinite-radius axes are maximal.

3.2 Properties

The EDF and EMA, defined on the medial axis, share several important properties with the boundary distance function and the medial axis, which we examine in this section (the proofs are provided in Appendix A).

We start by examining the range of values of the EDF, showing that it is lower bounded by the boundary distance function and is finite away from loops in the medial axis:

Proposition 1. *Let $x \in M$:*

1. $\tilde{R}(x) \geq R(x)$, and the equality holds only when $x \in \partial M$.
2. $\tilde{R}(x) = \infty$ iff there is some subset $S \subset M$ containing x such that $\partial S = \emptyset$.

The result in 1 aligns with the intuition that the “side-ways” expansion of the shape around x (captured by $\tilde{R}(x)$) is no smaller than the uniform expansion there (captured by $R(x)$). An immediate corollary of 2 is that \tilde{R} is finite everywhere on the medial axis M of a simply connected shape O (i.e., one without interior holes), since M is a tree and so is any of its subset. If O contains interior holes, the proposition implies that $\tilde{R}(x)$ is infinite over the largest subset of M that does not have open boundaries (similar to the 1-core of a graph), and finite everywhere else.

Next, we examine the local behavior of the EDF, and show that the EDF behaves like a geodesic function over the medial axis, similar to the boundary distance function

over the 2D shape. We consider the local behavior separately at the boundary of the medial axis, on the EMA, and on the rest of the medial axis:

Proposition 2. *Let $x \in M$:*

1. *If $x \in \partial M$, \tilde{R} has a directional gradient of 1 along M leaving x .*
2. *If $x \notin \partial M$ and $x \notin \tilde{M}$, $\tilde{R}(x)$ has a directional gradient of 1 on exactly one of the out-going branches of M at x , and a directional gradient of -1 on one or more of the out-going branches.*
3. *If $x \in \tilde{M}$ and $\tilde{R}(x) \neq \infty$, $\tilde{R}(x)$ has a directional gradient of -1 on two or more of the out-going branches at x .*
4. *If $x \in \tilde{M}$ and $\tilde{R}(x) = \infty$, \tilde{R} is infinite on two or more of the out-going branches at x .*
5. *In all three cases (2,3,4), \tilde{R} on the remaining out-going branches is bounded strictly below $\tilde{R}(x)$, and has constant gradient -1 .*

In short, any point on the medial axis that is not a boundary or part of the EMA has some neighborhood where the EDF is continuous and has a constant gradient of 1. As immediate corollaries, the function \tilde{R} is upper semi-continuous over all M , has no local minima except at the boundary ∂M , and is locally maximal at each point on \tilde{M} where \tilde{R} is finite. The last statement also implies that the part of \tilde{M} where \tilde{R} is finite consists of isolated points.

Finally, we show that the EMA preserves the topology of the medial axis, which in turn preserves the topology of the 2D shape:

Proposition 3. *\tilde{M} is homotopy equivalent to M .*

If the shape O is simply connected, the proposition implies that \tilde{M} has the homotopy of a point. Combined with the argument above that \tilde{M} consists of only isolated points, one can conclude that \tilde{M} is a unique point on the medial axis (e.g., a *center point*). If O contains interior holes, \tilde{M} has the homotopy of a system of loops. Hence \tilde{M} consists entirely of the part of M where $\tilde{R} = \infty$ without additional isolated points.




Figure 3.2: An illustration of intermediate states in the extended grassfire burning on the medial axis. Yellow dots are the fire fronts, arrows indicate the burning direction, and the red, blue, and green squares are where fire fronts are ignited, annihilated, and quenched.

3.3 A Grassfire Analogy for Computation

While explicitly defined, the EDF and EMA cannot be directly computed from their definitions, which involve exploring an infinite set of axes at each point. On the other hand, the properties of the EDF established by Proposition 2, particularly its lack of local minima and constant gradient, suggests that the function can be obtained by propagating values geodesically along the medial axis from its boundary points. In the following, we design a propagation method for computing EDF and EMA guided by their local properties. The propagation bears close resemblance with (and in some sense “continues”) Blum’s grassfire, and hence is called the *extended grassfire transform*.

3.3.1 Extended Grassfire Transform

Akin to Blum’s grassfire analogy, imagine the medial axis M is made up of a thin thread of grass. Each end $z \in \partial M$ is *ignited* at time $R(z)$, that is, when Blum’s grassfire reaches there. The fire propagates from those ends geodesically along M at a uniform speed. When a fire front comes to a junction, it continues onwards if there is exactly one remaining un-burned branch, and *annihilates* if there are two or more branches remaining. When multiple fire fronts meet at the same location, and if there

are no remaining un-burned branches, the fire fronts *quench* against each other. The process terminates when no fire front is active. A simple example illustrating this grassfire burning is shown in Figure 3.2. In this extended grassfire analogy, the EDF at a point on M is the time at which the point is burned by the fire, and infinity if the point is never burned. If the shape O is simply connected, the entire M will be burned out, and the EMA is the quench site of the fire fronts. Otherwise, if O contains interior holes, EMA is the remaining un-burned portion of M .

Note that our extended grassfire can be combined with Blum’s grassfire to a single fire-burning process, since the arrival time of Blum’s grassfire at a medial axis point ($R(x)$) is always earlier than the arrival time of the extended grassfire ($\tilde{R}(x)$). In this combined burning, the fire front starts from the boundary of the shape, quenches along the interior of the medial axis, and continues onto the medial axis from their boundaries.

3.3.2 Discrete Algorithm

The analogy gives rise to a simple, thinning-based algorithm that can compute $\tilde{R}(x)$ and \tilde{M} over a discretely represented medial axis M in time linear to the number of elements in M . The algorithm assumes that M is represented as a weighted graph that captures a piece-wise approximation of the medial axis. The graph nodes are vertices on the medial axis, and the weight of an arc between two nodes is the length of the line or curve segment connecting the two vertices. In addition, the distance to the shape boundary is given at each degree-1 node (the “end”) of M .

The algorithm iteratively reduces the graph M to compute \tilde{R}_i at each node i , which is initialized to be the boundary distance at each degree-1 node and infinity elsewhere. At each iteration, the degree-1 node i with the smallest \tilde{R}_i is removed with its incident arc. If the removal exposes a new degree-1 node, j , then \tilde{R}_j is updated as the sum of \tilde{R}_i and the weight of arc $\{i, j\}$. Iteration terminates when the graph is reduced to either a single node or a set of cycles, which are output as the EMA \tilde{M} .

In our experiments, we compute the discrete medial axis M as the subset of the Voronoi diagram of points sampled on O , consisting of the interior Voronoi vertices

and their connecting Voronoi arcs. This subset is a provenly good approximation of the 2D medial axis, and converges to the medial axis as the sampling density increases [17].

3.4 Examples and Comparisons

We start with two simple examples in Figure 3.3 computed using our discrete algorithm. One of the two shapes is simply connected, while the other contains an interior hole. We can observe the properties of the EDF and EMA discussed in Section 3.2 in this picture. In particular, the EDF is identical with the boundary distance at the ends of the medial axis, and increases at a constant gradient away from the ends (most notably in the 3D height map in (c)) while staying above the boundary distance function (as seen in the overlay in (b)). At each junction, the EDF is continuous along at least two branches. The EMA is the global maximum of the EDF in the simply connected shape (the gray dot), and a loop on the medial axis in the other shape (the gray loop) where the EDF is infinite.

Observe from Figure 3.3 that the medial-axis points with higher values of EDF lie in more elongated parts of the shape. The infinite EDF over a medial axis loop describes an infinite elongation there, since a tube can wrap around the loop for infinitely many times. More importantly, observe that the EDF is not sensitive to minor boundary perturbations in elongated shape parts. Intuitively, the EDF captures the half-length of a longest fitting tube, which is a global measure that does not change significantly by adding or removing small protrusions.

The noise-insensitivity of EDF is most notable in a complex example like the one in Figure 3.4, which contains a significant amount of boundary noise. Observe that the EDF along medial axis branches corresponding to small boundary noise are very close to the boundary distance function (which is most notable in the overlaying picture in (b)), whereas branches corresponding to prominent shape protrusions have a much higher EDF than the boundary distance function. Intuitively, EDF and boundary distance at a medial axis point capture respectively the “length” and “thickness” of the local shape, and hence their difference is a good measure of how protruded the




Figure 3.3: EDF and EMA in a simply connected shape (left) and a shape with an interior hole (right). The medial axis is colored by the EDF in (a) while the EMA is drawn in grey, (b) shows the boundary distance function in the background, and (c) plots both EDF and the boundary distance function as a 3D height map. The heat coloring scheme is used (blue is low and red is high).

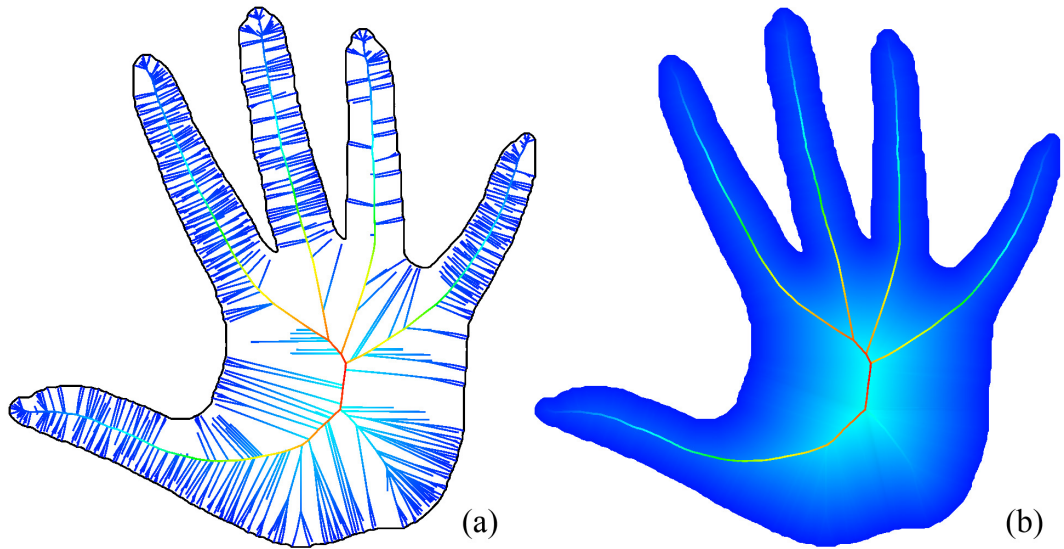


Figure 3.4: EDF (a) and overlaid on the boundary distance function (b) for a shape with boundary noise.

shape is. In fact, this measure has been widely used for identifying significant parts of the medial axes, although without an explicit formulation (see discussion below).

We further examine the stability of the EDF and the EMA under synthetic boundary perturbations in Figure 3.5 (top row). Here we perform a uniform perturbation of a square shape (a,b) and a non-uniform perturbation on one side of a key shape (c,d). Note that neither EDF nor EMA changes significantly, despite the change in the topology of the medial axis near the EMA (b) and the addition of a significant amount of medial axis branches (d).

We next compare EDF with two existing global measures on the medial axis, one based on heuristics (the Erosion Thickness) and the other formulated mathematically (the Potential Residue). In the first case, we show that the measure behaves similarly to EDF by giving an explicit characterization of the measure. In the second case, we demonstrate several advantages of EDF in terms of analytical properties and stability.

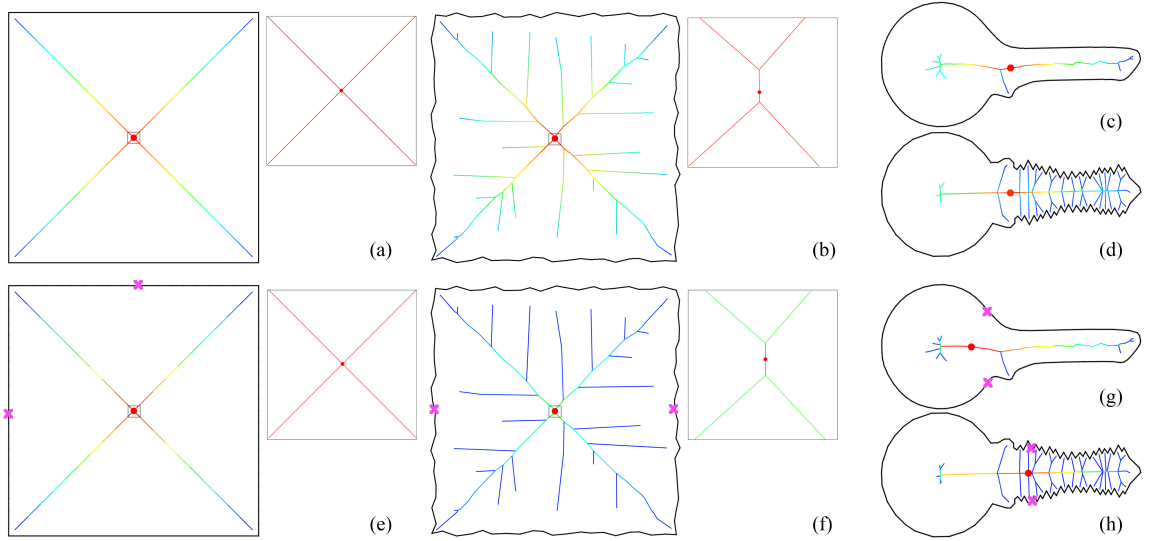


Figure 3.5: Comparing the EDF (top row) and the Potential Residue (PR) measure (bottom row) on the medial axis under boundary perturbation: the EDF and EMA are stable under both uniform (a,b) and non-uniform (c,d) perturbations, while PR can have sudden jumps (e,f), and both PR and its local maximum can exhibit significant drift (g,h). Each picture is colored by the EDF or PR normalized by the maximum value of EDF or PR present on the medial axis. The inserts in (a,b,e,f) zoom in on the middle of the medial axis where there is a connectivity change after perturbation. The closest boundary points to the local maximum of PR are shown as crosses.

3.4.1 Erosion Thickness

The Erosion Thickness (ET) measure captures the loss of the shape due to the pruning of a medial axis branch. For a point x located on a medial axis branch directly connected to an end of the medial axis y , ET is formulated as:

$$ET(x) = d(x, y) + R(y) - R(x)$$

Shaked and Bruckstein [61] proposed an extension of the measure to the entire medial axis using a *rate pruning paradigm*, where $ET(x)$ is determined as the time at which the pruning front reaches x . In this paradigm, the pruning front propagates similarly to the fire fronts in our extended grassfire analogy, with two differences. First, the pruning fronts start simultaneously from all medial axis ends at time 0. Second, the pruning front propagates at a non-uniform speed $1/(1 - R_\alpha(x))$, where $R_\alpha(x)$ is the gradient of R at x .

While the rate pruning paradigm can be implemented in a discrete algorithm (just like the extended grassfire), an explicit formulation of $ET(x)$ is not known. As a consequence, the meaning of $ET(x)$ beyond the medial axis branches directly connected to the ends is not clear. Here we give an explicit definition of $ET(x)$, which reveals its link to shape properties. In fact, ET is the residue of the EDF after subtracting the boundary distance function, that is,

$$ET(x) = \tilde{R}(x) - R(x)$$

The equivalence of this definition and the pruning time of the rate pruning paradigm can be verified by examining the property of the residue based on those of $\tilde{R}(x)$ (as discussed in Section 3.2) and $R(x)$ (i.e., continuous over M and has bounded derivative within $(-1, 1)$). As mentioned earlier, the difference of EDF and the boundary distance captures how much “longer” the shape is over its “thickness” around a medial axis point. As a result, medial axis points with higher ET values represent more prominent shape protrusions.

The definition of ET implies that it has very similar behavior to EDF. In particular, both functions share the same continuity and gradient direction over the entire medial axes. Also, both ET and EDF share the same set of local maxima, the EMA. From a practical point of view, the definition offers a simpler way of computing ET: rather than using the original rate pruning paradigm in [61] which requires a pruning rate that varies with a differential quantity (R_α), the ET can instead be obtained by a uniform-speed propagation (which gives $\tilde{R}(x)$) followed by subtraction of $R(x)$.

3.4.2 Potential Residue

The Potential Residue (PR) measures, at each medial axis point x of a simply connected shape, the shortest distance along the shape boundary between the two boundary points closest to x [52]. The intuition is that the closest boundary points for medial axis branches reaching to small boundary bumps are typically close together, hence the PR is small on those branches. Ogniewicz and Ilg showed that PR increases monotonically from the ends of the medial axis inward, and that there exists a unique local maximum of PR on the medial axis [52]. Note that the recently introduced definition

of the curve skeleton of a 3D shape by Dey and Sun [31] is in fact a 3D extension of the local maximum of the PR measure. The extended measure, called the Medial Geodesic Function (MGF), is the geodesic distance on the boundary surface between the two closest boundary points to a medial axis point.

We show several notable differences between the EDF and PR through analysis and experiments. First, although being monotonic like EDF, PR in general is not continuous at the junctions of the medial axis, which have three or more closest boundary points. In contrast, EDF is continuous along at least two branches at any junction.

Second, and more importantly, PR and its local maxima can change dramatically under boundary perturbations. We perform the same perturbation tests we had for EDF in Figure 3.5 for PR (bottom row). Observe that a slight change in the connectivity in the middle of the medial axis (see insert of (f)) causes a big increase in PR there after perturbation (note that the coloring of PR in both (e,f) is after normalization by the maximum PR over each medial axis). This is because the pair of closest boundary points to that part of the medial axis (shown as cross marks) changes dramatically after perturbation. Also, both PR and its local maxima are strongly affected by non-uniform boundary perturbations. In (g,h), perturbations on one part of the key shape cause PR and its local maximum to shift significantly towards that part, where the curve lengths increase greatly. Note that EDF and EMA are much more stable in both examples.

3.5 Utility

As demonstrated above, the EDF offers stable, global measures of shape elongations, and the EMA is a stable choice of shape center (for a simply connected shape). In this section, we show several ways in which these descriptors can be utilized for shape modeling.




Figure 3.6: The Erosion Thickness (ET) measure (a) and Shape Tubularity (ST) measure (b) over a 2D shape, and pruning using a low ET threshold (c), a high ET threshold (d), and the combination of a low ET threshold with an ST threshold (e).

3.5.1 Pruning Medial Axes

Since meaningful parts of the medial axis should capture elongated shape parts, we can use EDF to define significance measures for pruning the medial axis. As shown in Section 3.4.1, the difference between EDF and the boundary distance function is equivalent to the Erosion Thickness (ET) measure. Here we present another EDF-based measure called *Shape Tubularity* (ST), defined as:

$$ST(x) = \frac{\tilde{R}(x) - R(x)}{\tilde{R}(x)} = 1 - \frac{R(x)}{\tilde{R}(x)}$$

The measure is a scalar between $[0, 1]$, reaching 0 only at the boundary of the medial axis (where $\tilde{R}(x) = R(x)$) and 1 only at the EMA of a non-simply-connected shape (where $\tilde{R}(x) = \infty$).

Intuitively, ST captures the ratio (rather than the difference, as in ET) of the “thickness” over the “length” of the shape around a medial axis point. A high ratio implies a “sharp” protrusion (e.g., a needle), while a low ratio indicates a “blunt” one. Figure 3.6 compares ET (a) with ST (b) in a synthetic shape that contains both blunt (e.g., corners of the rectangles) and sharp (e.g., the smaller rectangle) protrusions. Note that ET treats the medial axis branches reaching diagonally to the corners of the bigger rectangle as important as those lying centered in the small rectangle. In contrast, ST along the diagonal branches is much lower than in the small rectangle, indicating that the later is a sharper protrusion.

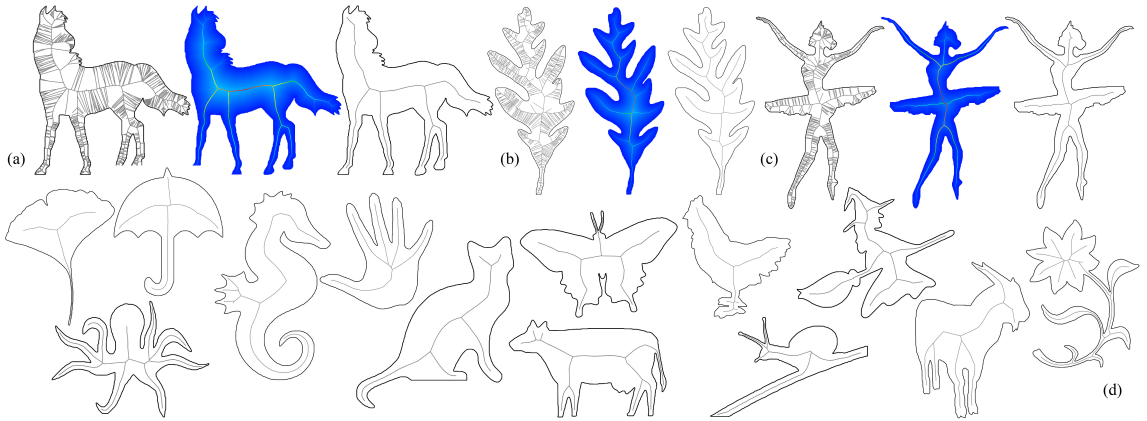


Figure 3.7: More pruning examples combining ET and ST. Examples in (a,b,c) include the original medial axes, the EDF and the boundary distance functions, and the pruned medial axes. Only pruned medial axes are shown in (d).

While capturing “sharpness” of protrusion, ST does not reflect the “size” of protrusion, and hence can be high around boundary noise (see the tip of the small bump in Figure 3.6 (b)). We therefore combine both ET (which captures the scale of protrusion) and ST for pruning, which removes parts of the medial axis representing shape features that are either small in size or weak in sharpness.

The advantage of using this combination over using ET alone is demonstrated in Figure 3.6 (c,d,e). Note that using ET alone with a low threshold (0.05 of the dimension of the shape in this example) is sufficient for removing branches caused by boundary noise (see (c)), but branches representing blunt features remain (e.g., the diagonal ones). While these branches can be removed using a high ET threshold (see (d)), branches representing sharp features are significantly shortened (e.g., the center axes of the small rectangle). The result in (e) is produced by using the same ET threshold as in (c) in combination with a suitable ST threshold ($1 - 1/\sqrt{2} \approx 0.3$) (that is, a medial axis point has to satisfy *both* ET and ST thresholds to be retained), which removes the diagonal branches but retain the length of the center axes in the small rectangle as in (c).

Note that the subset of the medial axis satisfying both ET and ST thresholds may not preserve the connectivity of the medial axis. For applications that require a topology-preserving curve skeleton, we may need to expand this subset to retain the original topology of the medial axis. This can be easily done by slightly modifying

the discrete thinning algorithm presented in Section 3.3.2. Instead of removing all degree-1 nodes, we shall preserve those with ST and ET values greater than or equal to the given thresholds. Since removing degree-1 nodes preserve the homotopy of the initial medial axis, the remainder after thinning is a topology-preserving subset. We show more examples of combined pruning with topology preservation in Figure 3.7, all computed using the same sets of thresholds as in Figure 3.6 (e).

3.5.2 Shape Alignment

Aligning two shapes is an important step for matching and recognition. Oftentimes, the first step of alignment is translating two shapes so that their centroid coincide. However, when two shapes that differ by a large variance such as isometric deformations (body movement in a human figure), the centroid can often change drastically (e.g., pink dots in Figure 3.8), and aligning the shape by their centroid would lead to unsatisfactory results (e.g., Figure 3.8 (b)).

In contrast, EMA, defined by the structure of the medial axes, is stable under a range of deformations including similarity transforms and isometric deformations. Hence EMA offers a good alternative to “centroid” for translational alignment of shapes undergoing these deformations. In the examples on the top of Figure 3.8, the EMAs are drawn as red dots. Note that they all lie roughly at the waist location of the human body. Alignment using EMA therefore achieves much better overlap between shapes, as shown in (d). We also compare with the local maxima of the Potential Residue (PR) measure, which are drawn as blue dots in the pictures. Note that these local maxima can sometimes drift significantly (see the last green figure), a phenomena that we already observed in the previous section. Hence the alignment using the PR local maxima (shown in (c)) does not look as good as the one produced by the EMA.

3.5.3 Shape Signature

While EDF offers a global shape metric over the medial axes, many applications such as shape matching require a descriptor (or *signature*) over the boundary of the 2D

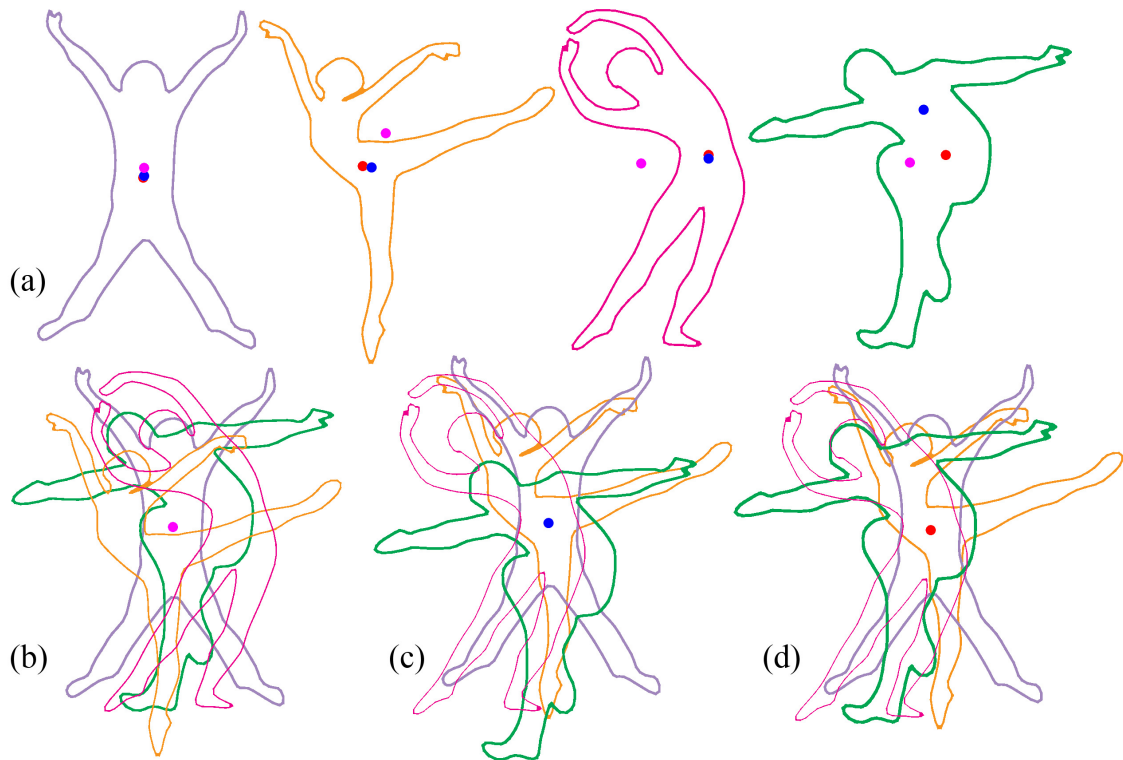


Figure 3.8: Top: several human shapes and their centroid (pink), local maximum of PR (blue), and EMA (red). Bottom: alignment using the centroids (b), local maximum of PR (c), and the EMAs (d).

shape. An ideal boundary signature should not only highlight local geometry, such as concavity or convexity, but also global shape properties, such as shape parts and extremities.

Common boundary signatures include curvature maps and local feature size (LFS), both providing only local shape information. As illustrated in Figure 3.9, the local curvature is homogenous on most parts of the boundary other than a few places where the curve bends strongly (see (a)). The LFS offers more information on the “thickness” of the local shape, as it measures the distance to the closest medial axis point. But LFS cannot differentiate parts with a common thickness. Also, note that the local nature of these descriptors make them very sensitive to boundary noise.

We introduce a new boundary signature for a simply connected shape, called *Boundary Eccentricity* (BE), which captures how far a boundary point is away from the EMA along the medial axis. For a point x on the medial axis M , let $E(x)$ be the




Figure 3.9: Boundary signatures: local curvature (a), local feature size (b) showing also the medial axis, Boundary Eccentricity (BE) (c,d) showing also $E(x)$ over the medial axes, and a matching result using BE (e).

geodesic distance from x to the EMA. For any boundary point p , let $X_p \in M$ be the set of points on the medial axis whose closest boundary point is p . BE is defined as:

$$BE(p) = \min_{x \in X_p} (E(x) + R(x))$$

Both BE and the function $E(x)$ over M are plotted in Figure 3.9 (c). Observe that, despite the noisiness of the shape, BE is a smoothly varying function that highlights shape parts and extremities that are away from the shape center. As shown in a different shape in (d), BE is also stable under isometric deformations. With these properties, BE is a good descriptor for matching 2D shapes that may be noisy and undergoing large deformations. We demonstrate this by matching the two curves in (c,d) guided by the BE values using a simple dynamic algorithm. The resulting correspondence for several critical points of BE on the first shape (c) is shown in (e).

3.6 Summary

In this chapter, we defined a global shape measure (EDF) over the medial axis that captured shape elongation, a shape center (EMA) where the elongation is maximal, and rigorously studied their properties. Both the EDF and EMA can be obtained using an extension of Blum's grassfire analogy onto the medial axis. The EDF and

EMA are demonstrated by examples, compared to related formulations, and used in several shape modeling applications.

Chapter 4

Extended Grassfire Transform on Medial Axes of 3D Shapes

In this chapter, we define a function on the medial axis of a 3D shape that captures the side-ways elongation that is similar to our *EDF* function in 2D. There are two equivalent ways of defining the EDF in 2D, one based on the length of the longest centered tubes, and the other based on an extended grassfire burning. This chapter starts by showing that a direct extension of the first definition from 2D to 3D does not yield a burning-like function on the 3D medial axis. Therefore, we define our function by extending the grassfire burning analogy and we show that the burning time of the grassfire over the 3D medial axis shares many properties with the 2D EDF function.

4.1 A First Attempt: Largest-Plate Definition

The elongation of the medial axis in 3D can be defined in a similar fashion as in 2D. In 2D, the elongation at a medial point x is measured as the half length of the longest “tube” centered at x and inscribed in the input shape O . This tube extends longitudinally along the medial axis. The counterpart elongation in 3D at a medial point x can be measured as the radius of the widest “plate” centered at x and inscribed in O . This plate extends sideways along the medial axis, therefore, it can be considered as the side-ways elongation of the 3D shape. We first give the definition using the plate radius, then we explain several issues with the definition through examples, which motivate our alternate definition in the next section.

4.1.1 Definition

In 2D, we first define any possible 1-dimensional local neighborhood of any medial point x along the medial axis as an axis f . The union of all of the medial disks with centers on f form a tubular shape, then the radius of f at x is defined as the shorter half length of the tubular shape. The radius is measured at all possible axes enclosing medial point x , and the largest among them is the EDF at x . This intuitively corresponds to the half length of the longest tube we can embed along the medial axis at the medial point x .




Figure 4.1: A 3D shape made from two rectangular cuboids with smooth boundary edges (a), the 3D shape with its medial axis embedded inside (b), a disk (in blue) on the medial axis surrounding medial point x (c), and the plate (in blue) corresponding to the disk (d). The LPDF at x is achieved by the red curve path in (c), where z is a boundary point of the blue disk and $R(z)$ is the radius of the medial ball (the red ball in (d)) at z . The plate is the union of all medial balls with centers from the blue disk. We highlighted a few balls inside the plate in blue and red.

These definitions in 2D can be directly extended to 3D as follows, which is illustrated through a simple example in Fig. 4.1. We start with the definition of a 2D neighborhood of any medial point x along the medial axis as disk f . The blue disk in Fig. 4.1 (c) is an example of such a disk neighborhood f . Note that by a disk, we mean only

a 2D area that is topologically equivalent to a disk, and it does not necessarily have to have a circular shape. The union of medial balls centered on medial points from f forms a “plate” (the counterpart of a 2D “tube” in 3D). The blue plate in Fig. 4.1 (d) corresponds to the disk in Fig. 4.1 (c), with a few interior medial balls highlighted in blue and red. The radius at the medial point x with respect to the disk is then defined as the minimum of $d_f(x, z) + R(z)$, where $d_f(x, z)$ is the geodesic distance between the two points on the medial axis, z is any point on the boundary of the disk f , and $R(z)$ is the radius of the medial ball at z . Note that the path x, z together with the radius connects point x to a boundary point on the 3D plate. Therefore, the radius essentially captures how wide the 3D plate is. In the example given in Fig. 4.1 (c), the radius of x with respect to the blue disk is the length of the curve marked in red. This radius computation is carried out for all possible plates that contain the medial point x , and our distance function at x (which we call the “largest-plate” distance) is defined as the maximum of them all. In this particular case, the distance at x is precisely $d_f(x, z) + R(z)$, since no other plates can have a larger radius with respect to x .

To be more precise, we define:

Definition 5. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be a local embedding of the disk $D = \{p | p \in \mathbb{R}^2, |p| \leq 1\}$ onto M . The image $f(D) \subseteq M$ is called a disk.

Given some point x on a disk f , we are interested in the *radius* of f with respect to x , which is defined as the distance from x to the shortest boundary point of the plate of f (which is the union of medial balls centered on f):

Definition 6. Given a disk f and a point $x \in f$,

$$r_f(x) = \min_{z \in \partial f} (d_f(x, z) + R(z))$$

is called the *radius of f with respect to x* , where $d_f(x, z)$ is the geodesic length of segment $[x, z]$ on f (the red path in Fig. 4.1 (c)).

The largest-plate distance function at x is the largest radius of any disk containing x :

Definition 7. Given a point $x \in M$,

$$\tilde{R}(x) = \sup_{f \ni x} r_f(x)$$

is the largest-plate distance function (LPDF) at x .

4.1.2 Example and Issues

We examine the behavior of the LPDF using a simple non-manifold medial axis shown in Fig. 4.2. The medial axis consists of two planar sheets joined at a non-manifold junction, one sheet parallel to the ground (called the “board”) and one sheet shooting straight-up (called the “fin”). For simplicity, we assume that the board is bounded on one side (denoted B) and extends infinitely away from that side (indicated by the arrows), and that the fin is bounded on two sides (denoted F and N and meeting at point c) and extends infinitely away from F and N . The two sheets meet at the non-manifold junction N . We assume that the angle between F and N is $\frac{1}{2}\pi$, and the angle between N and B is $\frac{3}{4}\pi$. In addition, we assume that the radius of the medial ball at each medial point is arbitrarily small. With this assumption, the LPDF at a medial point x reduces to the radius of the largest *circular* disk centered at x that lies on the medial axis.




Figure 4.2: The board and fin example (right) made from a long and wide “board” (left) and a long and tall “fin” (middle). The black lines are their boundary lines and the arrows represent infinite expansion. A non-manifold curve is formed at the intersection between the board and the fin.

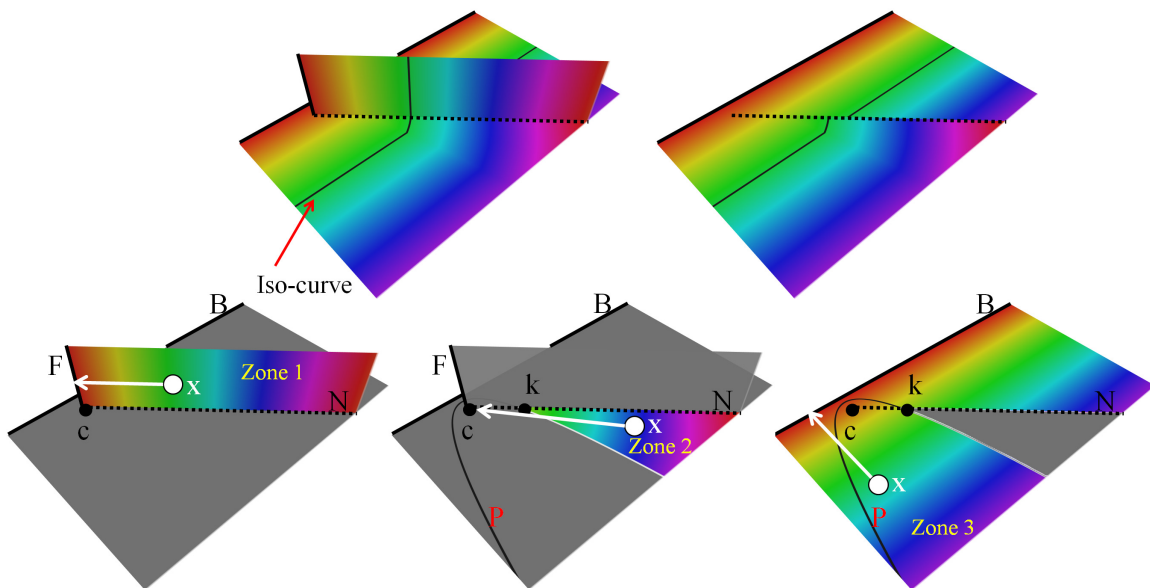


Figure 4.3: The LPDF distribution on the board and fin example (top), and the paths (in white) that realize the LPDF of medial points from its three different zones (bottom). The value of the LPDF increases when the hue color changes from red to purple. The medial points with the same LPDF forms an iso-curve (in black in the top row), which is upper semi-continuous at the non-manifold point and only C^0 continuous in the manifold region. The LPDF of points in zone 1 and zone 3 are realized by their distance to the border of their zones, while the LPDF of points in zone 2 are realized by their distance to the corner point c .

The LPDF on this medial axis is displayed using hue color in Fig. 4.3 (top). To explain this picture, we divide the medial axis into three zones and analyze the behavior of the function within each zone (see Fig. 4.3 bottom; each zone is hue-colored by distance):

- Zone 1 consists of the entire fin. A circular disk centered at any point x on the fin lies either completely on the fin or extends partly onto the board. The largest disk is confined by the fin's boarder F . Hence the LPDF in this zone increases linearly away from F .
- Zone 2 consists of a “wedge” region on the board, which is bounded on one side by the non-manifold junction N and on the other side by the parabola P consisting of points that are equally distant to B and c . The intersection point of P and N is denoted k . Among all possible circular disks centered at some point x in this wedge, the largest one lies partly on the fin and partly on the

board and touches the point c (note that the distance from x to c is greater than the distance from x to the board's border B). Hence the LPDF in this zone is a circular distance function centered at c .

- Zone 3 consists of the remainder of the board that is not in zone 2. The largest circular disk centered at a point x outside of zone 2 is one that lies completely on the board and is confined by the board's border B . Hence the LPDF in this zone increases linearly away from B .

Observe that the LPDF in each zone is a continuous function, and that the LPDF in two abutting zones meets continuously along the zone border. When three zones meet along a common border (e.g., the portion of N that extends infinitely away from point k), the LPDF in two of the zones meets continuously (in this case, zone 1 and zone 2) with values higher than those from the third zone. In other words, the LPDF is continuous everywhere on the medial axis away from the non-manifold junctions, and upper-semi continuous along the junctions. Such continuity closely mimics the continuity of our EDF on a 2D medial axis.

One of the key properties of our EDF definition in 2D is that it behaves like the burning time of a grassfire. While the LPDF has a similar behavior as burning time it increases away from the border of the medial axis, a closer look in the pictures in Fig. 4.3 (top) reveals a subtle but important difference. Along the parabola P that separates zone 2 and 3, the LPDF is only C^0 continuous, that is, it is not differentiable (this can be concluded from our prior analysis of the LPDF in each zone). Intuitively, the iso-curve of LPDF for values higher than the value at point k has a sharp and concave corner at P (as seen in Fig. 4.3 (top)). On the other hand, the firefront of a grassfire on the board should be everywhere smooth except at convex corners (where different fronts meet and quench). Hence the LPDF cannot be interpreted as the burning time of a grassfire propagation.

The practical implication of this difference is that the LPDF cannot be computed as easily as the EDF on 2D medial axis. The burning formulation of EDF allows the value at a medial point to be inferred locally from the values at its neighboring points, and hence a linear pass over the medial axis is sufficient to compute the entire function. However, the LPDF at a medial point depends on more global information (e.g., which border of the medial axis is used to compute the distance function),

which makes the design of simple, local computations difficult, if not impossible. From a theoretical point of view, the non-differentiability of LPDF also makes it a non-desirable choice.

4.2 Burning-based Definition

As discussed above, direct extension of the longest-tube definition of our 2D EDF onto the 3D medial axis loses the burning-time property of our 2D definition. To remedy this loss, we next attempt to derive a distance function directly from the grassfire burning process over the medial sheets of a 3D shape. We first describe the burning process in an informal language with illustrating examples. We then detail the necessary mathematical notations and the formal definitions.

4.2.1 Intuition

Recall that our extended grassfire transform over the medial axis of a 2D shape proceeds as follows:

- The fire is ignited at the ends of the medial axis at the time equal to the radius of the medial balls.
- Fire propagates geodesically along the medial axis at a uniform speed.
- When the fire front of one branch reaches a junction, it dies out if there are at least two unburned branches there.
- Fire fronts quench when they meet.

The following rules generalize these 2D rules to 3D, with specific extensions to the non-manifold structure of the 3D medial axis:

- The fire is ignited at the border of the medial axis. For a medial point x at the border, the ignition time is the radius of the medial ball at x .

- Fire propagates geodesically along the medial sheets at a uniform speed.
- When the fire front from one sheet comes to a non-manifold junction at point x , it dies out if there are at least two other unburned sheets at x .
- Fire fronts quench when they meet.

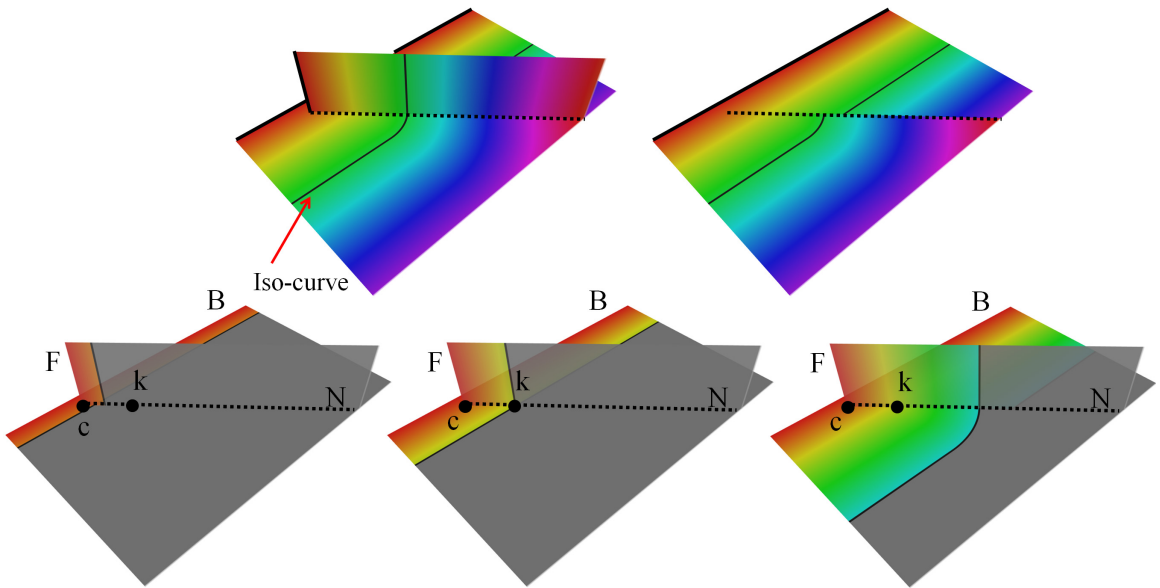


Figure 4.4: The burning time distribution on the board and fin example (top), and the burning front at three different time points (bottom). The grassfire starts from the borders, and burns forward at a geodesic uniform speed. At the beginning, the burning front on the top fin is ahead of that of the board (bottom left). This lasts until they reach the point k (called the “kink” point) (bottom middle). Afterwards, the burning fronts on the top fin merges with that of the front sheet, while the burning front of the back sheet advances faster (bottom right).

We illustrate this burning process using the same board-and-fin example as in Fig. 4.2. With the assumption that maximal balls are arbitrarily small at all medial points, the fire is lit at both borders B and F at time 0. Fig. 4.4 (bottom) shows the fire front (i.e., iso-curves of the burning time) at three different stages, which we explain as follows:

- At the beginning of the fire propagation, the fronts advance at uniform speed on the board and the fin from their respective borders (B or F). Note that the

fire on the fin stops along the non-manifold junction N , since the board is not yet burned (Figure 4.4 bottom-left).

- At time equal to the distance between c and k (which is the point on N that is equal-distant to c and B), the fronts on the board and the fin become connected at k (Figure 4.4 bottom-middle).
- Afterwards, because the portion of the board behind the fin forms a smaller angle $\frac{1}{4}\pi$ with the non-manifold junction N than the front side of the board and the fin, the fire front on the back side of the board will reach the non-manifold junction earlier than the fronts from the front side of the board and the fin. As a result, the fire front on the back side of the board will stop at the junction, while the remaining fronts propagate along the fin and the front side of the board. Intuitively, the back side of the board becomes the new “fin” while the remainder of the original fin and the front side of the board together become the new “board” (Figure 4.4 bottom-right).

The burning time is plotted as hue color in Fig. 4.4 (top). Observe that the function is quite similar to the LPDF function (Fig. 4.3), as it increases away from the border of the medial axis. Such similarity implies that the burning time, too, characterizes the expansion of shape. Like LPDF, the burning time is continuous everywhere away from the non-manifold junction. At a non-manifold medial point, the burning time is continuous over two sheets coming to that point, and is greater than the burning time over the third sheet. However, unlike LPDF, the burning time is differentiable at places where the LPDF is only C^0 continuous. Also, the erosion formulation implies that the burning time at a medial point only depends on the burning time of its neighboring points, hence making it possible to design efficient and local algorithms that propagate values over the medial axis from its borders.

It is also interesting to note that the pair of sheets with continuous burning time may vary along a non-manifold junction. In this example, for a medial point between the segment $\{c, k\}$, the burning time is continuous along the back and front sheets of the board. For a medial point after k on the same junction, the burning time is continuous along the fin and the front sheet of the board. We hence call k the “kink point” as it is where the behavior of the function changes. Note that kink points have no counterparts in 2D, as the only non-manifold entity in the medial axis of 2D shape

is a point, and hence there is no “change” in the behavior of the function along the non-manifold junction.

4.2.2 Notations

We now introduce some notions about medial axes in preparation for formally defining the burning time. Here, we assume the input medial axis M of a 3D shape is a piecewise smooth cell complex equipped with the distance function $f : \partial M \rightarrow \mathbb{R}$. f maps a point on the boundary of medial axis to the radius of the medial ball centered at this point. Note that f is a 1-Lipschitz function. A function is k -Lipschitz if $|f(x) - f(y)| \leq k d(x, y)$ and that Lipschitz functions are automatically continuous.

M naturally decomposes into 2-dimensional, 1-dimensional, and 0-dimensional regions as follows. The *manifold regions*, $M^{(2)}$, are any points that have a neighborhood in M which is homeomorphic to a disk where the point lies in either the interior or on the boundary of the disk. In the second case, the point is also on the boundary of M . *Singular curves*, denoted $M^{(1)}$, consists of all points $x \in (M - M^{(2)})$ with neighborhoods homeomorphic to a 1-dimensional disk. These are edges of the cell complex where three or more 2-cells meet. *Singular points*, denoted $M^{(0)}$, are the intersections between two (or more) singular curves, or $M \setminus (M^{(2)} \cup M^{(1)})$ (see Fig. 4.5 (a) for an annotated example). If a small neighborhood of a singular point was cut up along the singular curves, the remaining components are referred to as *sectors*. We refer to the union of $M^{(1)}$ and $M^{(0)}$ as the *singular set* of M and we use the notation $M^{(s)}$.

The *boundary* of M , denoted ∂M can be defined in terms of the manifold regions and singular curves:

$$\partial M = (\partial M^{(2)} \cap (M - M^{(s)})) \cup (\partial M^{(1)} \cap (M^{(1)} - M^{(0)}))$$

Notice that this is precisely the points in the portion of the boundary of the manifold region that are disjoint from the singular set and the endpoints of the singular curves that are not singular points.

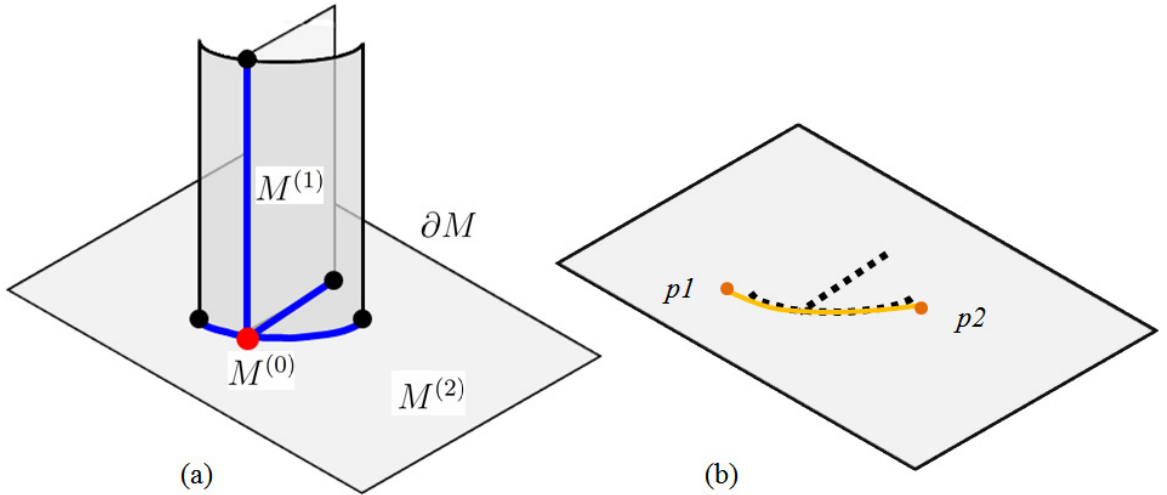


Figure 4.5: The decomposition of a medial axis (a), and the shortest path between points p_1, p_2 on $M^{(2)}$ realized by a non-crossing curve (in yellow) (b). The medial axis is decomposed into 2-dimensional ($M^{(2)}$, in grey), 1-dimensional ($M^{(1)}$, in blue), and 0-dimensional ($M^{(0)}$, in red) regions, and it is bounded by curves and points (∂M , in black). $M^{(2)}$ in general is noncompact (open at the dashed lines in b), so the shortest geodesic path may not always exist (e.g., p_1, p_2). To overcome this problem, we define *non-crossing* curves which are allowed to touch singular points but not allowed to cross singular curves, so that the shortest geodesic path between two points can always be realized by a non-crossing curve.

The extended grassfire locally burns along a geodesic path. We wish to constrain the curve on $M^{(2)}$ only, but $M^{(2)}$ in its nature is noncompact since it is not closed, which makes it impossible to assume that we have shortest geodesic paths between two arbitrary points and other useful properties. For example, the shortest path between p_1 and p_2 on $M^{(2)}$ can get arbitrary close to the length of the shortest path on M in Fig. 4.5 (b), but not equal, because the path touches the singular curve which does not belong to $M^{(2)}$. Therefore, we first define the shortest distance between two points on $M^{(2)}$ as $d_{M^{(2)}}(p_1, p_2) = \inf\{|P_{M^{(2)}}(p_1, p_2)| \mid P \text{ is a path between } p_1, p_2\}$. We then also define a *non-crossing* curve as a curve that is constrained in $M^{(2)}$ so that it touches the singular sets but never crosses them (e.g. the yellow path between p_1 and p_2 in Fig. 4.5 (b)). Otherwise, a curve that crosses a singular set is a *crossing* curve. A *non-crossing* curve can be infinitesimally perturbed to be completely inside of $M^{(2)}$. It turns out there always exists a non-crossing curve that realizes the shortest distance between two points on $M^{(2)}$, as the following lemma says:

Lemma 1. *Given $x, y \in M^{(2)}$ there exists a non-crossing curve $\gamma : I \rightarrow M$ with $\gamma(0) = x, \gamma(1) = y$ and $|\gamma| = d_{M^{(2)}}(x, y)$.*

4.2.3 Burning Sets and Burning Times

We formally define a burning process for M described in Section 4.2.1. The burning starts at the boundary at times specified by $f : \partial M \rightarrow \mathbb{R}$. We characterize the part of the medial axis M that is burned away at time t as the burning set $B_M(t)$. This characterization leads directly to an iterative definition. Intuitively, we know that the set $B_M(t)$ monotonically increases since more parts of the medial axis are burned away as time goes by, and finally it stays unchanged when burning ends. The monotonicity indicates that if a medial point p shows up in the burning set at time t , it will stay in the burning set thereafter, therefore the earliest time when $p \in B_M(t)$ is the burning time of p . We formalize these definitions in this section.

In the course of our burning, the grassfire reaches a medial point from all possible disk neighborhoods. For example, a medial point in the manifold area has a single disk neighborhood, and the extended grassfire burns in from one direction and exits from this point in another direction inside the disk. However, if the point has a non-manifold neighborhood, the disk neighborhood becomes more complex, so we instead consider the regular neighborhood radius and the set of disk types at any point:

Definition 8. *The regular neighborhood radius at a point x , $R_M(x)$, is the $\sup\{r \mid \forall r_1, r_2 \in (0, r), B(x, r_1) \cap M \cong B(x, r_2) \cap M\}$. The set of disk types at a point x , $\mathcal{D}_M(x)$, is the set of all possible combinatorial types of disks centered at x with radius equal to the regular neighborhood radius, $\{D \mid x \in D \subset M, d_M(x, y) = R_M(x) \forall y \in \partial D\}$.*

The grassfire could reach a medial point x on any of the combinatorial disks. Fig. 4.6 shows one example of $R_M(x)$ and $\mathcal{D}_M(x)$ at a medial point x . Note that x is burned away only when the fire in all disks $\mathcal{D}_M(x)$ have touched x . For example, point x in Fig. 4.6 is not burned away when the fire front from the lower-left border of the board first reaches there, as the disk $D_3(x)$ is not yet burned. We now formally define the burning set $B_M(t)$ as follows:




Figure 4.6: The regular neighborhood radius ($R_M(x)$) at a point x (a), and its disk types $D_M(x)$, including D_1 , D_2 and D_3 (b).

Definition 9. *The burning set of M at time t is defined as:*

$$\begin{aligned}
 B_M(t) = & \{x \in \partial M \mid f(x) \leq t\} \cup \\
 & \{x \in M \mid \forall D \in \mathcal{D}_M(x) \exists \gamma : I \rightarrow M \text{ such that } \gamma(0) = x, \gamma(1) \in B_M(t - |\gamma|), \\
 & \quad \gamma \text{ does not cross the singular set}, \exists \delta > 0 \text{ such that } \gamma([0, \delta]) \subset D\}
 \end{aligned}$$

In other words, a point x is in the burn set at time t either if it is on the boundary and has $f(x) \leq t$ (so that it already has “burned” by time t), or if for every possible disk type, there is a non-crossing path γ going in some direction, where the other endpoint of γ is in the burn set at time $t - |\gamma|$. γ in each disk is the slot burned away by grassfire by time t .

The following lemma justifies how we can view burning sets as erosion process as time progresses more points are “burned away”.

Lemma 2. *$s \leq t$ implies that $B_M(s) \subseteq B_M(t)$*

Finally, we are ready to define the burning time of a point based on our burning sets. From Lemma 2, it is easy to see that if $x \in B_M(t)$, then $x \in B_M(s)$, for any $s \leq t$. The burning time should be the earliest time that x appears in the burning set:

Definition 10. *The burning time $BT_M(x)$ of a point x is:*

$$BT_M(x) = \begin{cases} \inf\{t \mid x \in B_M(t)\} & \text{if } x \in \bigcup_{t \in \mathbb{R}} B_M(t) \\ \infty & \text{otherwise} \end{cases}$$




Figure 4.7: The shortest geodesic path (in yellow) between points x , y on a medial axis made from a large board at the bottom and the side of a hollow cylinder on top. It is an example for demonstrating when the burning path can get arbitrarily close, but not equal to, the shortest geodesic path.

Note that we use the infimum of t instead of the minimum of t in the definition, because the minimum may not exist in some cases. Fig. 4.7 shows one such example. This example consists of a board at the bottom and the side face of a hollow and long cylinder on top. The board and the cylinder meet at a non-manifold circular curve. Let y be the geodesically closest point to a point x on the border of the board, and suppose the shortest path between x and y on the board uses a segment of the non-manifold curve. Note x is in burning set $B_M(d_M(x, y) + \epsilon)$, $\epsilon > 0$, but x is not in the burning set $B_M(d_M(x, y))$, therefore only the infimum exists but not the minimum. The reason is that the grassfire cannot burn along the shortest path between x, y , which touches a part of the singular curve that is not burned away until it is reached by the grassfire from the top border of the cylinder. The grassfire can burn along a path that is arbitrary close to the shortest path between x, y with a length that is strictly larger than $d_M(x, y)$ by an arbitrarily small positive value. Note that in the simplest case when M is a manifold, the burning time is equal to $\min\{d_M(x, y) + f(y) \mid y \in \partial M\}$.

The burn process can get arbitrarily close to a medial point at multiple times, or in other words, a medial point is touched multiple times by the grassfire. For example, the medial point x in Fig. 4.6 is first touched by the grassfire from the lower-left

border of the board, then burned away by the grassfire from the border of the fin. We call the set of times when the grassfire gets arbitrary close *char times*. Apparently burning time is always in char times.

Definition 11. *The set of char times $CT_M(x)$ of a point x is:*

$$CT_M(x) = \{t \mid \exists \{x_i\} \rightarrow x \text{ with } \{BT_M(x_i)\} \rightarrow t\}$$

4.3 Properties of Burning Times

In this section we prove that burning times are well behaved, nice functions that in many ways generalize our definition of the EDF in 2D.

We begin by examining how distances affect the burning time, eventually proving that the function $BT_M(x)$ is 1-Lipschitz when restricted to either the manifold regions or singular curves, therefore the burning time is continuous on manifold regions and on singular curves.

Proposition 4. *For $i \in \{1, 2\}$ and $x, y \in M^{(i)}$,*

1. $x \in B_M(t)$ implies $y \in B_M(t + d_{M^{(i)}}(x, y))$
2. $|BT_M(x) - BT_M(y)| \leq d_{M^{(i)}}(x, y)$

The sketch of the proof is as follows (full proof can be found in Appendix B). When x, y are both in $M^{(i)}$, there exists a non-crossing curve α between x and y with length $d_{M^{(i)}}(x, y)$ based on Lemma 1. Since both x, y are in $M^{(i)}$, this α can serve as γ in the burning set definition (Definition 9) at the point x , which directly proves that $x \in B_M(t)$ implies $y \in B_M(t + d_{M^{(i)}}(x, y))$. Part 2 can be proved from the definition of the burning time and the conclusion of part 1.

Recall that a function $f : X \rightarrow \mathbb{R} \cup \infty$ is *upper semi-continuous* if for every $x_0 \in X$ and $\epsilon > 0$ there exists a neighborhood U of x_0 such that $f(x) \leq f(x_0) + \epsilon$ for all $x \in U$. In particular, this implies for any sequence $x_n \rightarrow x_0$ with $\{f(x_n)\}$ converging we have $f(x_0) \geq \lim_{n \rightarrow \infty} f(x_n)$. Lower semi-continuity is similarly defined.

Proposition 5. 1. $BT_M(x)$ is upper semi-continuous. Furthermore, for some sequence $\{x_n\}$ converging to x $BT_M(x) = \lim_{n \rightarrow \infty} BT_M(x_n)$.

2. $\min CT_M(x)$ is lower semi-continuous. Furthermore, $\min CT_M(x) = \lim_{n \rightarrow \infty} BT_M(x_n)$ for some sequence $\{x_n\}$ converging to x .

From Proposition 4 we know that if x is in the manifold regions, Proposition 5 automatically holds, since continuous functions must be both upper and lower semi-continuous. When x is on the singular curve, this proposition matches our intuitive extended grassfire burning rules. The extended grassfire burns towards x in multiple disks. If grassfire touches x when there is still an intact disk, this grassfire simply dies, which suggests the limit of the sequence following this grassfire is no larger than the burning time of x . On the other hand, the minimum of char times at x can be considered as the time of the sequence following the grassfire that first reaches x , and hence it is always smaller than or equal to the minimal char times in its local neighborhood.

Burning times are not a continuous function; however at any point there is a disk on which burning is continuous at that point. Intuitively, when the grassfire burns on the last intact disk at a medial point x , it leaves a continuous burning trace through x .

Proposition 6. For every $x \in M - \partial M$ there exists a disk D with $x \in D \subset M$ such that $BT_M|_D$ is continuous at x .

Unlike the distance function f , there are points where burning times could be infinite. It turns out that these points are precisely the largest subset of M without boundary, similar to the closed loops which have infinite values of EDF in 2D.

Proposition 7. $\{x \in M \mid BT_M(x) = \infty\}$ is equal to the maximal closed subcomplex of M .

Intuitively, this set is never burned by the extended grassfire since it does not have a boundary.

In sum, the burning time is continuous at any point on the medial axis within a small disk-neighborhood around the point, and it is finite when the point is not in

the closed subsets. These properties are reminiscent of those of the EDF over the 2D medial axis.

4.4 Summary

In this chapter, we explored two different extensions of EDF from 2D medial axes to 3D. The first extension, based on the radius of the largest plate, results in a continuous function but one that does not have an erosion formulation. The second extension explicitly formulates burning over the 3D medial axes, and we present several properties such as continuity and finiteness. These properties lay a theoretical foundation for developing computational algorithms of the burning time, as well as for defining lower-dimensional medial geometry in 3D such as medial curves and center points.

Chapter 5

Future Work

In this dissertation, we extended Blum’s grassfire burning to medial axis in both 2D and 3D. This extended grassfire burning leads to a function on the medial axis that captures intrinsic shape properties which further can be used to define lower dimensional medial forms.

This dissertation gives a good foundation to many problems for future research. Most of the open questions exist in R^3 , and they can be grouped into three categories: *theory*, *algorithms* and *applications*, which we discuss in the following sections.

5.1 Theory

5.1.1 Burning Time and Shape Properties

Unlike in R^2 , the relation between the burning time and the shape properties is not obvious in R^3 , which makes the burning time more difficult to use in practice. In R^2 , we have shown that EDF is the half the length of the longest tube that could be embedded inside the 2D shape, and hence it reflects how elongated a 2D shape is. Unfortunately, the burning time in R^3 does not correspond to the radius of the widest plate (LPDF) that can be possibly embedded inside the 3D shape (Section 4.1). But still, the burning time is somehow related to how wide the medial axis is at a medial point, as when the medial axis is wider, the burning time is larger. This behavior is also indicated by the highly similar value distribution of LPDF and the burning time on the board and fin example (Fig. 5.1).




Figure 5.1: The LPDF distribution on the board and fin example (top), and the burning time distribution (bottom). The iso-curves are drawn at medial points with the same function value. The function distribution is also drawn only on the board (right column).

In the future, it would be interesting to characterize the relation between the burning time and the 3D shape properties formally, so that the burning time can be used to help the understanding of 3D shapes.

5.1.2 Homotopy Preserving Lower Dimensional Medial Forms

Another problem is how to define homotopy preserving medial curves directly based on the burning time in R^3 , or more generally, whether we can define an even lower dimensional medial form, such as the shape center, in R^3 , since they are widely used in many applications (Section 2.3).

Defining lower dimensional medial curves on the medial axis in R^3 is not easy because of the non-manifold structure of the medial axis. We think a potential way to define

the medial curves is based on the burning paths of the extended grassfire. The burning paths are the curves that the extended grassfire burns along. Note that in Blum's grassfire burning analogy, the burning path starts at a boundary point and ends at a medial point. Similarly, on the medial axis of a 3D shape, the burning paths end at the medial curve points. We demonstrate that the medial curves created based on this definition in Fig. 5.2. Note that the medial curves are still disconnected, which is likely to be solved by carefully adding a few extra burning paths to connect them.




Figure 5.2: The medial curves defined from burning paths. Black arrows indicate the burning paths of the extended grassfire. The end points of these burning paths are medial curve points.

A formal characterization of the points on medial curves is still needed. The extended grassfire can also be used to define a deformation retract to prove the homotopy equivalence. When this problem is solved, it would be interesting to know if the grassfire can be further extended to the medial curves to define even a lower dimensional medial form, i.e., the shape center.

5.1.3 Stability

Even though we have empirically observed the stability of EDF and EMA under significant amount of boundary perturbation, a rigorous mathematical proof of this

stability is still lacking in both R^2 and R^3 . The stability is a very important property for a significance measure to be useful in practice (Chapter 2).




Figure 5.3: The medial axis (black line) of a very thin 2D shape and its EMA (black dot) (top), the medial axis (red curve) of its perturbed shape within ϵ -Hausdorff distance, whose EMA (red dot) is shifted.

By stability, we mean that the EDF and EMA stay almost unchanged when the boundary is perturbed. Let us focus on the EMA in 2D for now. A popular way to measure the change in boundary shape or in EMA is Hausdorff distance. Unfortunately, our EMA is not stable relative to this measure. One example is shown in Fig. 5.3. The black curve is the medial curve of a very thin 2D shape O whose thickness is almost zero, and the center black point is the EMA. The red curve is the medial axis of a perturbed shape of O which is within its ϵ Hausdorff distance, and the red point is its EMA. Note that we can add arbitrarily many red teeth on the left side, and this will lead to a big shift in the EMA. This simple example implies that our EMA is not stable under Hausdorff distance, and we need to apply a harsher constraint on the boundary noise. A possible candidate is the Hausdorff distance or some other curve metrics, such as Isotopic Frechet distance [18], normalized by the local thickness of the shape as indicated by the radius of the maximum balls.

5.2 Algorithms

In order to use the burning time in practice, it is necessary to devise a new computational algorithm. An accurate computational algorithm is preferred, because all the properties of the burning time hold true when the result is accurate. Apparently, the computation of burning time can only be more difficult than the computation of medial axis, and we know that an accurate computation of medial axis is limited to a few classes of shapes. As a result, we will first study accurate computational algorithms for the burning time on these shapes. We will also consider approximate

algorithms on approximations of the medial axis, such as on a discretized grid, and study the approximation accuracy of such algorithms.

We think that the accurate computation of the burning time on a piecewise linear medial axis is feasible. By piecewise linear, we mean that the medial axis consists only of points, straight line segments, and flat triangles (flat polygons can always be split into triangles without altering the shape). Note that the medial axis of a union of balls or a convex polygon is precisely piecewise linear, and the Voronoi Diagram based medial axis approximation algorithm results in a piecewise linear medial axis with both topology and convergence guarantees (Chapter 2).

Similar to the sweeping algorithm computation of Voronoi Diagram [28], we could develop a sweeping approach to mimic the extended grassfire burning on the piecewise linear shape. The sweeping algorithm maintains the burning front that is advanced in the normal direction. Note that the burning time at the boundary of the medial axis is not 0, because the boundary of the medial axis starts burning at time f , which is the radius of the medial ball at that boundary point. We know f is piecewise quadratic on the boundary of the medial axis of a union of balls. Also, similar to [28], we need to identify all the events that could happen, when we advance the burning front. The events not only include sweeping over a point, an edge or a triangle, but they also include the merge of burning fronts from different manifold pieces, which is more complicated to implement but still manageable.

5.3 Applications

The burning time and medial curves carry intrinsic shape information and they can be used in shape analysis and shape matching.

Shape analysis The burning time is related to the side-ways expansion of 3D shapes. A large burning time indicates the corresponding local shape is planar, while a small burning time indicates the corresponding local shape is tubular. The parts of medial axis that correspond to planar shape parts should be represented by medial sheets, and the tubular shape parts should be represented by medial curves. This

mixed-dimensional medial structure is useful in certain applications, for example, the visualization and analysis of the secondary structure of a protein, which consists of tubular α -helixes and planar β sheets (see Fig. 5.4). Of this mixed dimensional medial axis, the area of the medial sheets and the length of the medial curves can be measured and used as a shape signature. The medial curves defined based on our burning time can be used in various applications discussed in Section 2.3.




Figure 5.4: A surface model of a protein (a) and the abstract illustration of its secondary structure (b) (α -helix in green and β -sheets in cyan).

Shape matching In 2D, we mapped the EDF to the boundary curve of the shape to define a boundary signature called *boundary eccentricity*. Compared to other boundary signatures, ours is smoother and more stable with respect to boundary noise. We then applied boundary eccentricity to identify prominent shape features and to establish the correspondence between two similar shapes. One can imagine such boundary eccentricity to be directly extended to 3D, so that feature points can be identified for shape matching. For example, Fig. 5.5 shows two bones from two different subjects. The burning time and medial curves are approximated by our previously developed

discrete algorithm [46]. Even though the size and shape of the same bone vary between the two subjects, the overall burning time distribution and structure of the medial curves are rather similar, which indicates that our burning time and medial curves can be used for shape matching.




Figure 5.5: Two bones from two different subjects (top and bottom). Each bone is illustrated by a surface model, the model with its medial axis embedded colored by an approximated burning time using our own discrete algorithm [46], and the model with the extracted medial curves embedded. Note that even size and shape of the same bone vary in the two subjects locally, the burning time and medial curves are rather similar. Therefore, they can be used for shape matching.

Appendix A

Proofs of Propositions in Chapter 3

A.1 Proof of Proposition 1

Proof:

1. By triangle inequality, for any axes f containing distinct points x, y , we have $R(x) < d_f(x, y) + R(y)$. Hence $r_f(x) \geq R(x)$ with the equality attained iff x is an end of f . If $x \in \partial M$, all axes containing x will have x as an end, and hence $\tilde{R}(x) = R(x)$. Otherwise, there is some axes that does not have x as an end, and so $\tilde{R}(x) > R(x)$.

2. First, suppose there a subset $S \subset M$ containing x such that $\partial S = \emptyset$. Then it is possible to obtain an axes f where $r_f(x) = \infty$ by extending a path from x in both directions infinitely without encountering a boundary. Hence $\tilde{R}(x) = \infty$.

Next, suppose $\tilde{R}(x) = \infty$, which implies $r_f(x) = \infty$ for some axes f . Since O is bounded, R is finite, and hence both the geodesic distances from x to both ends of f need to be infinite. Note that O is bounded by piece-wise analytic curves, hence M contains a finite set of analytic curve arcs [21], and so M does not contain an infinite simple path. As a result, both segments of the axes f on the two sides of x need to overlap with themselves. It is easy to see that the subset of M covered by the segments of f on each side of x up to the first overlapping event is one without boundary.

□

A.2 Proof of Proposition 2

We begin by showing several lemmas that lead to the proof.

Lemma 3. *Extending an axes f from its ends does not reduce its radius with respect to some fixed $x \in f$.*

Proof: Denote the ends of f as z_0, z_1 , and the ends of the extended axes f' as z'_0, z'_1 . For each $i \in \{0, 1\}$, we have:

$$\begin{aligned} d_f(x, z_i) + R(z_i) &= d_{f'}(x, z_i) + R(z_i) \\ &\leq d_{f'}(x, z_i) + d_{f'}(z_i, z'_i) + R(z'_i) \\ &= d_{f'}(x, z'_i) + R(z'_i) \end{aligned}$$

Hence $r_{f'}(x)$ is no smaller than $r_f(x)$. \square

Let f be an axes containing two points x, y , we say y is on the *constrained side* (or *unconstrained side*) of x if y lies on the segment of f between x and a constrained (or unconstrained) end of f with respect to x . We have:

Lemma 4. *Let f be a non-maximal inscribed axes of $x \notin \tilde{M}$, or a maximal axes of $x \in \tilde{M}$ and $\tilde{R}(x) \neq \infty$. The following holds for any $y \in f$ that lies on the constrained side of x ,*

$$\tilde{R}(y) = \tilde{R}(x) - d_f(x, y)$$

Proof: Since f is an axes containing y , and since y is on the constrained side of x , we have

$$r_f(y) = r_f(x) - d_f(x, y) = \tilde{R}(x) - d_f(x, y).$$

We next show that there exists no other axes f' such that $r_{f'}(y) > r_f(y)$, and hence $\tilde{R}(y) = r_f(y)$. Suppose such f' exists. Denote the two ends of f as z_0, z_1 , so that y lies on the segment $[z_0, x]$ on f . Denote the two ends of f' as z'_0, z'_1 , so that the segment $[y, z'_0]$ on f' does not share the same half-disk neighborhood of y as the segment $[y, x]$ on f (see Figure A.1 (a)). Consider a new axes f'' made up by segments $[z'_0, y]$ on f'

and $[y, z_1]$ on f . Note that $x \in f''$, and

$$\begin{aligned} d_{f''}(x, z'_0) + R(z'_0) &= d_f(x, y) + d_{f'}(y, z'_0) + R(z'_0) \\ &\geq d_f(x, y) + r_{f'}(y) \\ &> d_f(x, y) + r_f(y) = r_f(x) \end{aligned} \tag{A.1}$$

On the other hand,

$$d_{f''}(x, z_1) + R(z_1) = d_f(x, z_1) + R(z_1) \geq r_f(x) \tag{A.2}$$

If the last equality in Equation A.2 holds, f is a maximal axes of x whereas f'' is an inscribed axes of x (because of the strict inequality in Equation A.1), which contradicts to the assumption of the lemma. Otherwise, f'' has a greater radius than f with respect to x , which contradicts with the fact that f is inscribed. \square




Figure A.1: Notations used in the proofs.

Let us further denote a *junction* of a set of curves as a point whose local neighborhood on the set contains more than two 1-D half-disks. We have a similar result as the previous lemma but concerning the unconstrained side of an axes:

Lemma 5. *Let f be an inscribed, non-maximal axes of $x \notin \tilde{M}$. The following holds for any $y \in f$ that lies on the unconstrained side of x ,*

$$\tilde{R}(y) = \tilde{R}(x) + d_f(x, y),$$

if the half-open interval $(x, y]$ does not contain any junction on M , and if

$$d_f(x, y) < \frac{\|d_f(x, z_0) + R(z_0) - d_f(x, z_1) - R(z_1)\|}{2} \quad (\text{A.3})$$

where z_0, z_1 are the two ends of f .

Proof: Using Lemma 4, we only need to show that f is an inscribed axes of y , is not maximal, and x lies on the constrained side of y on f . The last two properties are assured by the inequality in Equation A.3, which also implies that $r_f(y) = r_f(x) + d_f(x, y)$.

To show inscribedness, suppose on the contrary there exists f' containing y such that $r_{f'}(y) > r_f(y)$. Since the segment $(x, y]$ of f is free of junctions on M , and since we can always extend an axes without reducing its radius by Lemma 3, we can always find an f' that shares the segment $(x, y]$ with f , and hence $x \in f'$. Denote the two ends of f as z_0, z_1 , so that y lies on the segment $[x, z_1]$ on f . Denote the two ends of f' as z'_0, z'_1 , so that the segment $[y, z'_0]$ on f' contains x (see Figure A.1 (b)). Consider a new axes f'' made up by segments $[z'_0, x]$ on f' and $[x, z_1]$ on f . Using a similar argument as in Lemma 4, and since f is not a maximal axes of x , one can conclude that f'' has a greater radius than f with respect to x , reaching a contradiction with the fact that f is inscribed. \square

Now we are ready to prove Proposition 2:

Proof: We consider each case as follows:

1. If $x \in \partial M$, any axes f with one end at x is an inscribed, non-maximal axes of x (due to Proposition 1(i)). By Lemma 5, and due to the finite structure of M [21], there is some finite segment $[x, y]$ on f where \tilde{R} increases with constant gradient 1.
2. If $x \notin \partial M$ and $x \notin \tilde{M}$, x has at least one inscribed, non-maximal axes. Note that the unconstrained side of x in all these axes share the same half-disk neighborhood of x , or otherwise a longer axes could be constructed by concatenating two unconstrained segments on two inscribed axes. By Lemmas 4 and 5, \tilde{R} increases with gradient 1 along the shared unconstrained segment, and decreases with gradient -1 along the constrained segment of each inscribed axes.
3. If $x \in \tilde{M}$ and $\tilde{R}(x)$ is finite, x has at least one inscribed axes and all such axes are maximal. By Lemma 4, \tilde{R} decreases with a gradient of -1 on both sides of each of its inscribed axes.
4. If $x \in \tilde{M}$ and $\tilde{R}(x) = \infty$, by Proposition 1, x lies in a subset $S \subset M$ such that $\partial S = \emptyset$. Hence all points on the neighborhood of x in S have infinite \tilde{R} .
5. Consider a branch at x that is not part of any inscribed axes of x , and take a point y on the branch so that the segment (x, y) is free of junctions on M . Consider an inscribed axes f of x and denote its two ends as z_0, z_1 where z_1 is constrained. Consider an inscribed axes f' of y and denote its two ends z'_0, z'_1 , so that the segment $[y, z'_1]$ on f' contains the segment (y, x) (see Figure A.1 (c)). Again, such an axes f' can always be found due to Lemma 3. It is easy to see that the new axes f'' by joining segment $[z'_0, x]$ on f' and $[x, z_1]$ on f is an inscribed, non-maximal axes of y , and that

$$d_{f''}(y, x) < \frac{d_{f''}(y, z_1) + R(z_1) - d_{f''}(y, z'_0) - R(z'_0)}{2}.$$

By Lemma 5, for any point w on the open interval (y, x) , $\tilde{R}(w)$ increases with a constant gradient 1 as w moves from y to x . Combining with the above

equation, we have:

$$\begin{aligned}
\tilde{R}(w) &= \tilde{R}(y) + d_{f''}(y, w) \\
&< d_{f''}(y, z'_0) + R(z'_0) + d_{f''}(y, x) \\
&< d_{f''}(y, z_1) + R(z_1) - d_{f''}(y, x) \\
&= d_f(x, z_1) + R(z_1) = \tilde{R}(x)
\end{aligned}$$

Hence the limit of $\tilde{R}(w)$ as $w \rightarrow x$ is bounded below $\tilde{R}(x)$.

□

A.3 Proof of Proposition 3

Proof: To show homotopy equivalence, we use the common technique of constructing a *deformation retract* from M to \tilde{M} . We will find a mapping $h(t, x)$ that is continuous in both $t \in [0, t_0]$ for some $t_0 > 0$ and $x \in M$, so that $h(0, M) = M$ and $h(t_0, M) = \tilde{M}$.

We do so by establishing a “direction” field over M which will be followed by h . At each point $x \in M$ that does not belong to \tilde{M} , Proposition 2 implies that there is a unique out-going branch at x where \tilde{R} increases with the gradient of 1. This out-going direction is said to be the *flow direction* at x , $v(x)$. For $x \in \tilde{M}$, its $v(x)$ is set to null. Note that the flow directions are continuous. By Proposition 2, $v(y)$ at a point y in the neighborhood of x points away from x only when $v(x)$ points towards y , and points towards x if $v(x)$ is either null or points away from y .

We define $h(t, x)$ as the point on M that has travelled t time away from x at the geodesic speed of 1 following the field v . By the continuity of v , $h(t, x)$ is continuous in both t, x . Let $T = \sup_{x \in M, x \notin \tilde{M}} \tilde{R}(x) + 1$. Since \tilde{R} increases at least with the gradient of 1 along v (with possible jumps at junctions), $h(T, x)$ for $x \notin \tilde{M}$ must be at \tilde{M} , otherwise $\tilde{R}(h(T, x))$ would be greater than $\sup_{x \in M, x \notin \tilde{M}} \tilde{R}(x)$. Since $h(t, \tilde{M}) = \tilde{M}$ for $t \in [0, T]$, h is a deformation retract, and \tilde{M} is homotopy equivalent to M . □

Appendix B

Proofs of Lemmas and Propositions in Chapter 4

B.1 Proof of Lemma 2

Proof. Suppose $x \in B_M(s)$, so we know that x is in one of the two sets from definition 9.

Suppose x is in the first set, so x is on the boundary of M and $f(x) \leq s$. Since $s \leq t$ by assumption, we have $f(x) \leq t$ also and so $x \in B_M(t)$.

If x is not on the boundary, then for each disk type D_i in the set of disk types $\mathcal{D}_m(x)$, we have a curve γ_i which is initially contained in D_i , does not cross the singular set, and has $\gamma_i(0) = x$ and $\gamma_i(1) \in B_M(s - |\gamma_i|)$.

We will use each γ_i to construct a new curve α_i which will show that x is also in $B_M(t)$ for $t \geq s$. Choose a point on γ_i which is outside of $\gamma_i([0, \delta])$ (where δ is the length of γ_i which stays inside D_i from definition 9) and add a curve of length $t - s$ that begins and ends at this point which does not cross a singular set. Then from the definition, this curve α_i witnesses that $x \in B_M(t)$, since α_i has length $|\gamma_i| + t - s$ and $\alpha_i(1) = \gamma_i(1)$ so that $\alpha_i(1) \in B_M(s - |\gamma_i|)$. \square

B.2 Proof of Proposition 4

Proof. Part 1: First suppose we have x and y in the same component of $M^{(2)}$. By Lemma 1, we know that there is a non-crossing path $\alpha : I \rightarrow M$ realizing $d_{M^{(2)}}(x, y)$ which goes from x to y .

Since x and y are in $M^{(2)}$, there is only 1 disk type for each, so $|\mathcal{D}_M(x)| = |\mathcal{D}_M(y)| = 1$. In order to show that $y \in B_M(t + d_{M^{(2)}}(x, y))$, we simply take our γ in definition 9 to be the reversal of α above. Since $x \in B_M(t)$, this gives that $y \in B_M(t + |\gamma|) = B_M(t + d_{M^{(2)}}(x, y))$.

For two points $x, y \in M^{(1)}$, we simply note that any disk in $\mathcal{D}_M(x)$ contains part of the shortest path from x to y in $M^{(1)}$, since a portion of $M^{(1)}$ to any side of x must be in any disk containing x . So same use of α as in the $M^{(2)}$ case will prove our lemma for $M^{(1)}$.

Part 2: Suppose $BT_M(y) = t$. From our definition of BT_M as the infimum of all the burning sets that y belongs to, this means that for any $\epsilon > 0$, $y \in B_M(BT_M(y) + \epsilon)$.

Now, we use part 1. Since $y \in B_M(BT_M(y) + \epsilon)$, we get that for any $\epsilon > 0$, $x \in B_M(BT_M(y) + \epsilon + d_{M^{(i)}}(x, y))$. Now recall that by our definition, $BT_M(x)$ is the infimum of all the burning sets that x belongs to, and we know that for any $\epsilon > 0$, $x \in B_M(BT_M(x) + \epsilon)$. We conclude that $BT_M(x) \leq BT_M(y) + d_{M^{(i)}}(x, y)$.

The proof that $BT_M(x) \leq BT_M(y) + D_{M^{(2)}}(x, y)$ is completely symmetric, and the statement of the proposition follows immediately. \square

We will now address continuity, or rather exactly when burning time is discontinuous. First, however, we need the a few technical lemmas which describe what kind of paths γ (from the definitions of burning sets) we may use.

For the next two lemmas, we will use $B_M^k(t)$ to denote these burning sets (as defined in Definition 9, where the curves γ are restricted to have length $\leq k$). In a similar fashion, we define $BT_M^k(x) = \inf\{t | x \in B_M^k(t)\}$.

Lemma 6. *For $i \in \{1, 2\}$ and $x, y \in M^{(i)}$,*

1. If $x \in B_M^k(t)$ and $d_{M^{(i)}}(x, y) \leq k$, then $y \in B_M^k(t + d_{M^{(i)}}(x, y))$.
2. If $d_{M^{(i)}}(x, y) \leq k$, then $|BT_M^k(x) - BT_M^k(y)| \leq d_{M^{(i)}}(x, y)$

Proof. This follows immediately from the same proof as part 1 of Proposition 4. In part 1, the only difference is that α has length $\leq k$, which does not change the rest of the argument. In part 2, we deal with BT_M^k and B_M^k , but otherwise the proof is unchanged. \square

Next, we make a more general statement about the sets $B_M^k(t)$, proving that they are in fact completely equivalent to our original definition of burn sets.

Lemma 7. *For any value $k > 0$, we may assume that $|\gamma| < k$ (where γ is the curve in definition 9) without changing the burning sets $B_M(t)$.*

Proof. Clearly, we know that $B_M^k(t) \subseteq B_M(t)$, since any curve with length less than or equal to k is in the set of all possible curves.

So we must show that $B_M(t) \subseteq B_M^k(t)$. Consider $x \in B_M(t)$ which uses a γ that has length longer than k , where $\gamma(0) = x$ and $\gamma(1) = y$. Note that by lemma 1, we can assume that γ is a shortest path in $\overline{M^{(2)}}$, so it is piecewise geodesic in M .

We break γ into subpaths, each of which is a geodesic of length $\leq k$. Let x_1, x_2, \dots, x_l be the endpoints of these subpaths, where x_l is the point closest to y . We know that $BT_M^k(x_l) \leq BT_M^k(y) + d(x_l, y)$ from Lemma 4. We also know that $BT_M^k(x_i) \leq BT_M^k(x_{i+1}) + d(x_i, x_{i+1})$ for each $i \in 1, \dots, l-1$, using the same lemma. If we combine these inequalities, we get that $BT_M^k(x) \leq BT_M^k(y) + d(x, y)$.

But then this means that $x \in B_M^k(t)$ also, since x is in the burn set of y 's burn time plus the distance from x to y . \square

B.3 Proof of Proposition 5

Proof. Part 1: We will actually show something slightly stronger than upper semicontinuity. Instead, we'll show that for any point x , $BT_M(x) \geq \limsup_{x_n \rightarrow x} BT_M(x_n)$, and that in fact there is a sequence of points $x_n \rightarrow x$ such that $BT_M(x) = \limsup BT_M(x_n)$.

First, note that this holds trivially for any $x \in M^{(2)}$, since BT_M is 1-Lipschitz and therefore continuous on $M^{(2)}$.

Let x be a point in the singular set, and assume for the purposes of contradiction that there is a sequence $x_n \rightarrow x$ with $\limsup_{n \rightarrow \infty} BT_M(x_n) < BT_M(x)$. Since x_n is an infinite sequence and we have a finite number of disk types at any point, we can find an infinite subsequence of points which converge to x and lie entirely on one sector of $M^{(2)}$ which is in a small neighborhood of x ; we will thus assume that all of the points x_n lie on a single sector of $M^{(2)}$ near x or else they are entirely contained in $M^{(1)}$. In addition, we may likewise assume that for every n , $BT_M(x_n) < BT_M(x)$, since the sequence x_n must have an infinite subsequence with this property.

Consider any one of these points x_n . We know that there is a non-crossing path from x_n to x , so from definition 9 and definition 10 (using this path as our γ), we know that $BT_M(x_n) < BT_M(x) + d(x, x_n)$. (Note that this is not implied by our 1-Lipschitz proof, since x is not in $M^{(2)}$). Rearranging, we have that for any n , $BT_M(x_n) - BT_M(x) < d(x, x_n)$.

Now since $M^{(2)}$ can be extended to a compact closed manifold, we know that $\limsup BT_M(x_n)$ exists. Since this limit is strictly less than $BT_M(x)$, we can thus find a point x_m with $d_M(x, x_m) < |\limsup BT_M(x_n) - BT_M(x)|$. This directly contradicts our previous statement that for any n , $d(x, x_n) > BT_M(x_n) - BT_M(x)$, and we therefore conclude that the function BT_M is upper semicontinuous.

It remains to show that there is a sequence x_n where $\limsup BT_M(x_n) = BT_M(x)$. Suppose this is not the case, so that we have x with $BT_M(x) > \sup_{\{x_n\} \rightarrow x} \limsup_{n \rightarrow \infty} BT_M(x_n)$.
(

Now we know that BT_M is continuous in the interior of each sector since it is 1-Lipschitz on these regions, and we also know that there is a unique way to continuously

extend each sector to its boundary. Since we have finitely many sectors, we may in fact say that $BT_M(x) > \max_{\text{sectors } s} \lim_{n \rightarrow \infty} BT_M(x_n)$ where $\{x_n\}$ is *any* sequence converging to x on sector s .

Since we are taking the maximum over a finite number of sectors, each of which has a unique extension to its closure, we can set $b = \max_{\text{sectors } s} \lim_{n \rightarrow \infty} BT_M(x_n)$.

Now, pick any $\epsilon > 0$ which is less than $(BT_M(x) - b)/3$; we know this is a positive value since we have assumed that there is a gap between $BT_M(x)$ and the value b (or else BT_M would be upper semicontinuous at x).

From the definition of $BT_M(x) = \inf\{t \mid x \in B_M(t)\}$, we know that $x \in B_M(BT_M(x) + \epsilon)$ for any sufficiently small ϵ . From our definition, this means that for any $D \in \mathcal{D}_M(x)$, there exists a path $\gamma_D : I \rightarrow M$ with $\gamma_D(0) = x$, γ_D a noncrossing path, γ_D contained in D for some initial length, and $\gamma_D(1) \in B_M(BT_M(x) + \epsilon - |\gamma_D|)$.

By lemma 7, we may also assume without loss of generality that $|\gamma_D| < (BT_M(x) - b)/3$.

Now, using lemma 2, since $b < BT_M(x)$, we also have that $\gamma_D(1) \in B_M(b + |\gamma_D| + \epsilon)$, since γ_D is the curve from our definition of B_T and therefore must get a burn time no worse than any other possible γ .

So then we have that $x \in B_M(b + |\gamma_d| + \epsilon + |\gamma_D|)$, again since γ_D is the curve that realizes $B_M(x)$. Now $x \in B_M(b + \epsilon + 2|\gamma_D|)$, which means that $BT_M(x) \leq b + 2|\gamma_d| + \epsilon$, which (by assumption, since there is a gap between b and $BT_M(x)$) is less than $BT_M(x)$, giving a contradiction.

Part 2: Since we know that BT_M is 1-Lipschitz on $M^{(2)}$ and $M^{(1)}$, we know that BT_M is also continuous on those sets. Since the function is continuous on an open set, there is a unique way to extend that function to the closure in a way that maintains continuity. We have assumed that any point in $M^{(1)}$ has a finite number of half disks adjacent to it, since $\mathcal{D}_M(x)$ is finite, so any point $x \in M^{(1)}$ has a finite number of values in $CT_M(x)$. Since we are taking the minimum of a finite set, it will be lower semi-continuous.

□

B.4 Proof of Proposition 6

Proof. We know that for any $x \in M^{(2)}$, since BT_M is 1-Lipschitz (and thus continuous), the lemma trivially holds.

Consider $x \in M^{(1)}$. The local neighborhood of x consists of k of half disks, where each half disk extends into a component of $M^{(2)}$ but is bounded by the singular curve which x belongs to. Since BT_M is 1-Lipschitz, we know that BT_M on each half disk can be uniquely extended to the closure in a continuous fashion.

Each of these continuous extensions gives a burning value at the point on the closure that corresponds to x . We know that one of these values is equal to $BT_M(x)$, since from the previous lemma, we know that some sequence of points x_n achieves $\limsup BT_M(x_n) = BT_M(x)$. If another value is equal to $BT_M(x)$, then we are done, since these two sheets can be glued together to give a disk with BT_M on this disk being continuous at x .

So suppose every other half disk's extension gives a lower value than $BT_M(x)$. Consider the set of disks in $\mathcal{D}_M(x)$, and again consider our definition of burning sets. For each possible disk type, we now have a disk that contains a path γ that can avoid the half disk realizing $BT_M(x)$. In other words, γ can always find a burn set that is smaller, meaning that the burning time at x cannot be $BT_M(x)$.

Finally, take $x \in M^{(0)}$. If we intersect any small neighborhood of x with $M^{(2)}$, we get a series of *sectors*, which are either half disks (as in the singular curve case) or "quarter disks", where we have a region of $M^{(2)}$ which is bounded by a portion of each of the two singular curves which x lie upon. We will call these regions *sectors*. Note that again, each of these sectors has a unique extension of BT_M to its closure which maintains continuity.

We will construct a disk from these sectors with the property that BT_M on the disk is continuous at x .

□

References

- [1] Alexander Agathos, Ioannis Pratikakis, Panagiotis Papadakis, Stavros Perantonis, Philip Azariadis, and Nickolas S Sapidis. 3d articulated object retrieval using a graph-based representation. *The Visual Computer*, 26(10):1301–1319, 2010.
- [2] N Amenta, S Choi, and Ravikrishna Kolluri. The power crust, unions of balls, and the medial axis transform. *International Journal of Computational Geometry and its Applications*, 19(2-3):127–153, 2001.
- [3] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. *Proceedings of the sixth ACM symposium on Solid modeling and applications SMA 01*, 19(2-3):249–266, 2001.
- [4] Nina Amenta and Ravi Krishna Kolluri. The medial axis of a union of balls. *Computational Geometry*, 20(1-2):25–37, 2001.
- [5] D Attali. Delaunay conforming iso-surface, skeleton extraction and noise removal. *Computational Geometry*, 19(2-3):175–189, 2001.
- [6] D Attali and Annick Montanvert. Computing and simplifying 2d and 3d continuous skeletons. *Computer Vision and Image Understanding*, 67(3):261–273, 1997.
- [7] Dominique Attali, Jean-Daniel Boissonnat, and Herbert Edelsbrunner. Stability and computation of medial axes a state-of-the-art report. *Computer*, pages 1–18, 2004.
- [8] Dominique Attali and Annick Montanvert. *Modeling noise for a better simplification of skeletons*, volume 3, pages 13–16. IEEE, 1996.
- [9] Chandrajit Bajaj, Zeyun Yu, Rajit Bajaj Zeyun Yu, and Manfred Auer. Volumetric feature extraction and visualization of tomographic molecular imaging. *Journal of Structural Biology*, 144:1–2, 2003.
- [10] Matthew L Baker, Tao Ju, and Wah Chiu. Identification of secondary structure elements in intermediate-resolution density maps. *Structure London England* 1993, 15(1):7–19, 2007.
- [11] D Bartz, W Strasser, M Skalej, and D Welte. *Interactive exploration of extra- and intercranial blood vessels*, pages 389–547. Ieee, 1999.

- [12] M De Berg, O Cheong, M Van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*, volume 85. Springer, 2008.
- [13] Jules Bloomenthal. Medial-based vertex deformation. *ACM SIGGRAPH 2002*, page 147, 2002.
- [14] H. Blum. A transformation for extracting new descriptors of form. *Models for the Perception of Speech and Visual Form*, pages 362–80, 1967.
- [15] Gunilla Borgefors. Distance transformations in arbitrary dimensions. *Computer Vision Graphics and Image Processing*, 27(2):321–345, 1984.
- [16] J W Brandt and V R Algazi. Continuous skeleton computation by voronoi diagram. *CVGIP Image Understanding*, 55:329–338, 1992.
- [17] Jonathan W. Brandt. Convergence and continuity criteria for discrete approximations of the continuous planar skeleton. *CVGIP: Image Underst.*, 59:116–124, January 1994.
- [18] Erin Chambers, Tao Ju, David Letscher, and Lu Liu. Isotopic frechet distance. *Proceedings of 23rd Canadian Conference on Computational Geometry*.
- [19] F Chazal and A Lieutier. Stability and homotopy of a subset of the medial axis. *Proceedings of the ninth ACM symposium on*, page 248, 2004.
- [20] Frdric Chazal and Andr Lieutier. The -medial axis. *Graphical Models*, 67(4):304–331, 2005.
- [21] Hyeong In Choi, Sung Woo Choi, and Hwan Pyo Moon. Mathematical theory of medial axis transform. *Pacific J. Math*, 1997.
- [22] N D Cornea, M F Demirci, D Silver, S J Dickinson, and P B Kantor. 3d object retrieval using many-to-many matching of curve skeletons. *International Conference on Shape Modeling and Applications 2005 SMI 05*, pages 366–371, 2005.
- [23] Nicu D. Cornea and Patrick Min. Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 13(3):530–548, 2007. Member-Silver,, Deborah.
- [24] Nicu D Cornea, D Silver, X Yuan, and Raman Balasubramanian. Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer*, 21(11):945–955, 2005.
- [25] Tim Culver, John Keyser, and Dinesh Manocha. Accurate computation of the medial axis of a polyhedron. In *SMA '99: Proceedings of the fifth ACM symposium on Solid modeling and applications*, pages 179–190, New York, NY, USA, 1999. ACM.

- [26] P Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14(3):227–248, 1980.
- [27] Luis David Lopez, Deepk Shantharaj, Lu Liu, Hash Bais, and Jinyi Yu. Robust image-based 3d modeling of root architecture. *Proceedings of Computer Graphics International*, 2011.
- [28] F Dehne and R Klein. “the big sweep”: On the power of the wavefront approach to voronoi diagrams. *Algorithmica*, 17(1):19–32, 1997.
- [29] T K Dey and Jian Sun. Defining and computing curve-skeletons with medial geodesic function. *Computing*, pp:143C152, 2006.
- [30] Tamal Dey and Wulue Zhao. Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design*, 36(2):195–202, 2004.
- [31] Tamal K. Dey and Jian Sun. Defining and computing curve-skeletons with medial geodesic function. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing*, pages 143–152, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association.
- [32] Pavel Dimitrov, C Phillips, and K Siddiqi. *Robust and Efficient Skeletal Graphs*, volume 1, pages 417–423. 2000.
- [33] Herbert Edelsbrunner. The union of balls and its dual shape. *SCG 93 Proceedings of the ninth annual symposium on Computational geometry*, pages 218–231, 1993.
- [34] Sándor P. Fekete, Joseph S. B. Mitchell, and Karin Beurer. On the continuous fermat-weber problem. *Oper. Res.*, 53:61–76, January 2005.
- [35] James D Foley, Andries Van Dam, Steven K Feiner, and John F Hughes. *Computer graphics: Principles and practice*. Addison-Wesley Professional, 1995.
- [36] Mark Foskey, Ming C. Lin, and Dinesh Manocha. Efficient computation of a simplified medial axis. *Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 96–107, 2003.
- [37] N Gagvani and Deborah Silver. Shape-based volumetric collision detection. *of the 2000 IEEE symposium on Volume*, 2000.
- [38] Santiago Garrido, Luis Moreno, Mohamed Abderrahim, and Fernando Martin. Path planning for mobile robot navigation using voronoi diagram and fast marching. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (2006)*, pages 2376–2381, 2006.

- [39] Kenneth E Hoff, John Keyser, Ming Lin, Dinesh Manocha, and Tim Culver. Fast computation of generalized voronoi diagrams using graphics hardware. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques SIGGRAPH 99*, 33(Annual Conference Series):277–286, 1999.
- [40] Lichan Hong, Shigeru Muraki, Arie Kaufman, Dirk Bartz, and Taosong He. Virtual voyage: interactive navigation in the human colon. *Comput Graph SIGGRAPH Proc*, 31(Annual Conference Series):27–34, 1997.
- [41] Hongbo Jiang, Wenping Liu, Dan Wang, Chen Tian, Xiang Bai, Xue Liu, Ying Wu, and Wenyu Liu. Connectivity-based skeleton extraction in wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems*, 21(5):710–721, 2010.
- [42] Tao Ju, Qian-Yi Zhou, and Shi-Min Hu. Editing the topology of 3d models by sketching. *Cell*, 26(3):42, 2007.
- [43] L Lam, S W Lee, and C Y Suen. Thinning methodologies-a comprehensive survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):869–885, 1992.
- [44] W. Lenhart, R. Pollack, J. Sack, R. Seidel, and M. Sharir. Computing the link center of a simple polygon. In *Proceedings of the third annual symposium on Computational geometry*, SCG '87, pages 1–10, 1987.
- [45] André Lieutier. 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, pages 65–75, New York, NY, USA, 2003. ACM.
- [46] L. Liu, E. Chambers, D. Letscher, and T. Ju. A simple and robust thinning algorithm on cell complexes. *Computer Graphics Forum*, 29(7):2253–2260, 2010.
- [47] Lu Liu, W Chambers, David Letscher, and Tao Ju. Extended grassfire transform on medial axes of 2d shapes. *Computer-Aided Design*, 43:1496–1505, 2011.
- [48] Christian Lovato, Umberto Castellani, and Andrea Giachetti. Automatic segmentation of scanned human body using curve skeleton analysis. *Science*, pages 1–12, 2009.
- [49] Balint Miklos, Joachim Giesen, and Mark Pauly. Discrete scale axis representations for 3d geometry. *ACM Transactions on Graphics*, 29(4):1, 2010.
- [50] V Milenkovic. Robust construction of the voronoi diagram of a polyhedron. *Proc 5th Canad Conf Comput Geom*, pages 473–478, 1993.

- [51] A. Nedzved, S.V. Ablameyko, and S. Uchida. Gray-scale thinning by using a pseudo-distance map. pages II: 239–242, 2006.
- [52] R. Ogniewicz and M. Ilg. Voronoi skeletons: Theory and applications. In *in Proc. Conf. on Computer Vision and Pattern Recognition*, pages 63–69, 1992.
- [53] K Palagyi and Attila Kuba. A parallel 3d 12-subiteration thinning algorithm. *Graphical Models and Image Processing*, 61(4):199–221, 1999.
- [54] Diane Perchet, Catalin Fetita, and Françoise Prêteux. Advanced navigation tools for virtual bronchoscopy. In *In Proceedings SPIE Conference on Image Processing: Algorithms and Systems III Symposium on Electronic Imaging, Science and Technology 04*, 2004.
- [55] Stephen M. Pizer, Kaleem Siddiqi, Gabor Székely, James N. Damon, and Steven W. Zucker. Multiscale medial loci and their properties. *Int. J. Comput. Vision*, 55(2-3):155–179, 2003.
- [56] R. Pollack, M. Sharir, and G. Rote. Computing the geodesic center of a simple polygon. *Discrete Comput. Geom.*, 4:611–626, September 1989.
- [57] C Pudney. Distance-ordered homotopic thinning: A skeletonization algorithm for 3d digital images. *Computer Vision and Image Understanding*, 72(3):404–413, 1998.
- [58] Dennie Reniers, Dennie Reniers, and Alexandru Telea. Hierarchical part-type segmentation using voxel-based curve skeletons. *Computing*, 24(6):383–395, 2008.
- [59] Dennie Reniers and Alexandru Telea. Skeleton-based hierarchical shape segmentation. *IEEE International Conference on Shape Modeling and Applications 2007 SMI 07*, pages 179–188, 2007.
- [60] Dennie Reniers and Alexandru Telea. Patch-type segmentation of voxel shapes using simplified surface skeletons. *Computing*, 27(7):1837–1844, 2008.
- [61] Doron Shaked and Alfred M. Bruckstein. Pruning medial axes. *Comput. Vis. Image Underst.*, 69(2):156–169, 1998.
- [62] Arthur Shek and Tom Thompson. Art-directing disney’s tangled procedural trees. *Growth Lakeland*, pages 4503–4503, 2010.
- [63] Evan C Sherbrooke, Nicholas M Patrikalakis, and Erik Brisson. *Computation of the Medial Axis Transform of 3-D polyhedra*, pages 187–200. 1995.
- [64] K Siddiqi and S M Pizer. *Medial representations: mathematics, algorithms and applications*. Springer Verlag, 2008.

- [65] Kaleem Siddiqi, Sylvain Bouix, Allen Tannenbaum, and Steven W. Zucker. Hamilton-jacobi skeletons. *International Journal of Computer Vision*, 48(3):215–231, July 2002.
- [66] Deborah Silver and Nikhil Gagvani. *Unwinding the Colon*, volume d. IOS Press, 2002.
- [67] M U SubbaRao, M A Aragn-Calvo, H W Chen, J M Quashnock, A S Szalay, and D G York. Visualization of large scale structure from the sloan digital sky survey. *New Journal of Physics*, 10(12):125015, 2008.
- [68] Avneesh Sud, Mark Foskey, and Dinesh Manocha. Homotopy-preserving medial axis simplification. In *SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling*, pages 39–50, New York, NY, USA, 2005. ACM.
- [69] H Sundar, Deborah Silver, Nikhil Gagvani, and Sven Dickinson. Skeleton based shape matching and retrieval. *2003 Shape Modeling International*, 2003:130–139.
- [70] R Tam and W Heidrich. Shape simplification based on the medial axis transform. *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, pages 481–488, 2002.
- [71] R Tam and W Heidrich. Shape simplification based on the medial axis transform. *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, pages 481–488, 2002.
- [72] J Vleugels and M Overmars. Approximating voronoi diagrams of convex sites in any dimension. *Internat J Comput Geom Appl*, 8:201–222, 1998.
- [73] F. Wolter. Cut locus and medial axis in global shape interrogation and representation. *Design Laboratory Memorandum 92-2, MIT*, 1993.
- [74] Fu-Che Wu, Wan-Chun Ma, Rung-Huei Liang, Bing-Yu Chen, and Ming Ouhyoung. Domain connected graph: the skeleton of a closed 3d shape for animation. *The Visual Computer*, 22(2):117–135, 2005.
- [75] Hao Xia and Paul G Tucker. Finite volume distance field and its application to medial axis transforms. *International Journal for Numerical Methods in Engineering*, di:114–134, 2009.
- [76] Shin Yoshizawa, Alexander Belyaev, and Hans-Peter Seidel. Skeleton-based variational mesh deformations. *Computer Graphics Forum*, 26(3):255–264, 2007.
- [77] Shin Yoshizawa, Alexander G Belyaev, and Hans-Peter Seidel. Free-form skeleton-driven mesh deformations. *Proceedings of the eighth ACM symposium on Solid modeling and applications SM 03*, page 247, 2003.

- [78] Y.Y. Zhang and P.S.P. Wang. Analytical comparison of thinning algorithms. 7:1227–1246, 1993.

Multi-Dimensional Medial Geometry, Liu, Ph.D. 2011