

# Medial Spheres for Shape Representation

Svetlana Stolpner

Doctor of Philosophy

School of Computer Science

McGill University

Montréal, Québec

December, 2011

A thesis submitted to McGill University  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy

© Svetlana Stolpner, 2011



## ACKNOWLEDGEMENTS

I feel extremely lucky to have had two wonderful supervisors: Kaleem Siddiqi and Sue Whitesides. I want to thank Kaleem for introducing me to the fascinating world of ‘medial representations’, for sharing his faith in the utility of these representations, and for his expertise in this subject matter. I would like to thank Kaleem for supporting my endeavors, and for his intelligence in being able to follow my not always clear and rushed descriptions of results from computer graphics and computational geometry and to offer advice. I am also extremely grateful to Kaleem for his warmth, kindness, and attention to all aspects of life, not just academics, making me feel very much at home in Montreal. I would like to thank Sue for teaching me rigour in mathematical arguments and in scientific writing. Sue found the patience to bear with me as I unraveled the mess that became the contents of Chapters 3 and 4. Working with Sue, I never ceased to be amazed and inspired by her great intellect, her unrelenting energy, and her passion for mathematics and science.

I want to thank Sasha for giving me support and a listening ear. These past 5 years of PhD have been extremely trying at times, and there are few words to describe the thankfulness I feel that he stood by me through it all.

I am grateful to my mother for demonstrating her understanding of mathematics when I was a child. I would like to thank my father for sharing his enthusiasm for collecting and organizing knowledge, in his case, of historical and geographical facts. My deep gratitude is to my sister Nadia for taking an interest in the “shrunk pear” and offering unlimited encouragement.

My labmates over the years have made working in the lab a very desirable social experience. I want to thank Ryan Eckbo, Peter Savadjiev, Scott McCloskey and Nils Tilton for being fun to be around and for being my friends. Over time, I was rewarded with new friends: Parya Momayyez-Siahkal and Emmanuel Piuze-Phaneuf (who also helped edit Section 1.1), who made coming to McGill thoroughly rewarding, shared my hardships and happiness, and provided constant support.

I would like to thank Paul Kry for his insights, for listening to my concerns, and for offering encouragement in our collaborative work on the material in Part III, and for sharing the code that enabled the physically-based simulations. I am grateful to Nina Amenta for encouraging me to pursue the research direction that resulted in Part III of this thesis during our discussion in Barbados in 2009, when I was having doubts. I am very grateful to Patrice Koehl for sharing with me the *AlphaBall* software that computes the volume of a ball restricted to its power cell and thus opening up new research possibilities explored in this thesis.

I would like to thank the members of my examination committee for their efforts in reading the thesis and for providing valuable suggestions for improvement.

This research was supported by an FQRNT scholarship.



## ABSTRACT

This thesis presents a particular set of spheres as new representation of the shape of a 3D solid. The spheres considered are maximal inscribed spheres in the solid and their centres are chosen in such a way that at most one sphere centre lies in a cubic region of space.

The shape representation proposed is a discretization of the medial surface transform of a solid. Part I of this thesis presents algorithms for the computation of this representation given a boundary representation of a solid by approximating its medial surface transform. Properties of those medial spheres that are not detected by our algorithm in 3D are described and a complete characterization of those medial circles that are not detected by a 2D version of our algorithm is given.

In Part II, recent results from differential geometry are used to compute principal curvatures and principal curvature directions on the boundary of the smooth solid represented using the union of medial spheres. This computation is performed using only the medial sphere centres and a pair of points on each medial sphere that lies on the surface of the solid being modeled. It is shown how the union of medial spheres allows a part-based description of the solid, with a significance measure associated with each part.

In Part III, it is shown that our shape representation can offer a tight volumetric fit to a polyhedron, using a small number of spheres. The spheres used in our representation can be quickly updated as the solid undergoes a certain class of deformations. It is shown how our set of medial spheres allows efficient and accurate proximity queries between polyhedra.

## ABRÉGÉ

Cette thèse présente une nouvelle représentation de la forme de solides 3D, qui est un ensemble spécifique de sphères. Les sphères de cette représentation sont des sphères inscrites maximalement dans le solide et leur centre sont choisis de telle manière que tout au plus un seul centre de sphère occupe une région cubique de l'espace.

La représentation de forme proposée est une discrétisation de la surface médian d'un solide. La première partie de cette thèse présente des algorithmes pour le calcul de cette représentation à partir d'une description de la frontière d'un solide, en approximant sa surface médian. Certaines propriétés des sphères internes qui ne sont pas détectées par notre algorithme en 3D sont décrits et une caractérisation complète de ces cercles internes qui ne sont pas détectés par une version 2D de l'algorithme est donnée.

En deuxième partie, les résultats récents en géométrie différentielle sont utilisés pour calculer les courbures principales et les directions principales de courbure sur la surface d'un solide lisse représenté à l'aide de l'union des sphères médianes. Ce calcul est effectué en utilisant seulement le centre et une paire de points relatifs à chaque sphère médiane qui se trouve sur la surface du solide. Il est aussi démontré comment l'union des sphères médianes permet une description par partie du solide, tout en donnant une mesure de l'importance associée à chaque partie.

En troisième partie, il est démontré que notre représentation de forme peut servir à l'ajustement volumétrique serré d'un polyèdre, tout en utilisant un nombre restreint de sphères. Les sphères utilisées dans notre représentation peuvent être rapidement adaptées au solide lorsqu'il est sujet à une certaine classe de déformations. Notre ensemble de

sphères médianes permet d'effectuer des requêtes de proximité efficaces et précises entre des polyèdres.

## DECLARATION

This thesis presents original scholarship by the author. The results reported in Chapters 2-8 reflect work done largely by the author, with assistance offered by Dr. Kaleem Siddiqi, Dr. Sue Whitesides, and Dr. Paul Kry, as is reflected in the co-authorship of publications arising from this thesis described in Section 1.3.3. The primary contributor to the development of algorithms, the analysis of their correctness, and their implementations, as well as the technical descriptions, was the author of this thesis.

## TABLE OF CONTENTS

ACKNOWLEDGEMENTS . . . . .	ii
ABSTRACT . . . . .	iv
ABRÉGÉ . . . . .	v
LIST OF TABLES . . . . .	xii
LIST OF FIGURES . . . . .	xiii
1 Introduction . . . . .	1
1.1 Shape Representation . . . . .	1
1.1.1 Parametric/Boundary Representations . . . . .	4
1.1.2 Implicit/Interior Representations . . . . .	5
1.2 Preliminaries . . . . .	8
1.2.1 Medial Surface Transform . . . . .	8
1.2.2 Space Partitions . . . . .	14
1.3 Thesis Overview . . . . .	18
1.3.1 Objective and Outline . . . . .	18
1.3.2 List of Novel Contributions . . . . .	19
1.3.3 Publications Arising from This Thesis . . . . .	21
PUBLICATIONS . . . . .	22

<b>I</b>	<b>Medial Spheres – Construction</b>	<b>23</b>
2	Medial Surface Approximation . . . . .	24
2.1	Background . . . . .	24
2.2	Previous Work . . . . .	27
2.2.1	Voronoi Methods . . . . .	28
2.2.2	Tracing Methods . . . . .	32
2.2.3	Spatial Subdivision Methods . . . . .	32
2.3	Voxel-Based Approximation . . . . .	34
2.3.1	Foundations . . . . .	35
2.3.2	Approximating the Medial Surface with Voxels . . . . .	38
2.4	Point-Based Approximation . . . . .	40
2.4.1	Precise Medial Points . . . . .	41
2.4.2	Well-Spaced Salient Medial Points . . . . .	47
2.4.3	Handling Narrow Features . . . . .	50
2.4.4	Degeneracy Issues . . . . .	52
2.5	Grouping Medial Points into Medial Sheets . . . . .	54
2.6	Experimental Results and Discussion . . . . .	56
2.7	Chapter Summary . . . . .	61
3	Theoretical Analysis: 3D . . . . .	63
3.1	Basic Algorithm . . . . .	63
3.2	Deep Samples . . . . .	64
3.3	Building the Foam . . . . .	67
3.4	Measuring Missed Medial Point Quality . . . . .	73
3.5	Chapter Summary and Future Work . . . . .	76
4	Theoretical Analysis: 2D . . . . .	79
4.1	Tools for Identification of Medial Points . . . . .	80
4.2	Algorithm . . . . .	84
4.3	Missing a medial point . . . . .	85
4.3.1	Case 1 . . . . .	87
4.3.2	Case 2 . . . . .	92
4.3.3	Error Bound . . . . .	95
4.4	Medial Axis Approximation . . . . .	96
4.5	Chapter Summary and Discussion . . . . .	97

## **II Quality of Shape Description 99**

5	Boundary Differential Geometry from Medial Geometry . . . . .	100
5.1	Differential Geometric Shape Operator . . . . .	101
5.2	Discrete Differential Geometry . . . . .	102
5.3	Boundary Geometry from Medial Geometry . . . . .	103
5.4	Boundary Geometry from Sampled Medial Geometry . . . . .	106
5.4.1	Imposing Smoothness on the Boundary . . . . .	108
5.4.2	Derivatives on Medial Sheets . . . . .	110
5.4.3	Estimating Surface Curvature . . . . .	111
5.4.4	Sources of Error . . . . .	113
5.4.5	Experimental Results and Analysis . . . . .	115
5.4.6	Discussion . . . . .	119
5.5	Chapter Summary and Future Work . . . . .	121
6	A Significance Measure for Medial Sheets . . . . .	123
6.1	Related Work . . . . .	123
6.2	Significance Measure . . . . .	124
6.3	Experimental Results and Discussion . . . . .	126
6.4	Chapter Summary and Future Work . . . . .	127
6.4.1	Homotopy-preserving Medial Surface Simplification . . . . .	128
6.4.2	Shape Indexing and Retrieval using Medial Surfaces . . . . .	129

## **III Fast and Tight Shape Approximation 130**

7	Fast and Tight Spheres . . . . .	131
7.1	Generating Spheres: Previous Work . . . . .	132
7.2	Our Sphere-Based Shape Approximation . . . . .	133
7.3	Volumetric Error for Unions of Spheres . . . . .	136
7.3.1	Volumetric Error: Exact and Lower Bound . . . . .	136
7.3.2	Experimental Results and Discussion . . . . .	137
7.4	Fast Updates under Deformation . . . . .	142
7.4.1	Update Method . . . . .	142
7.4.2	Volumetric Error: Upper Bound . . . . .	144
7.4.3	Experimental Results and Discussion . . . . .	145
7.5	Chapter Summary . . . . .	147

8	Application to Proximity Queries . . . . .	149
8.1	Background . . . . .	149
8.2	Improving Boundary Coverage by Conservative Dilation . . . . .	151
8.3	$(\sigma, \theta)$ -fat Solids . . . . .	153
8.4	Hierarchy Construction using Rectangle Swept Spheres . . . . .	156
8.5	Experimental Results . . . . .	157
8.6	Discussion . . . . .	159
8.7	Future Work . . . . .	162
	8.7.1 Faster Bottom-Up Sphere Hierarchies . . . . .	162
	8.7.2 Interruptible Collision Detection . . . . .	163
	8.7.3 Sphere-Based Shape Primitives . . . . .	164
	8.7.4 Probabilistic Hierarchy Construction . . . . .	166
9	Thesis Summary and Conclusion . . . . .	168
	REFERENCES . . . . .	173



## LIST OF TABLES

<u>Table</u>		<u>page</u>
2–1	Statistics for the computation of medial points in Figure 2–11. . . . .	58
2–2	Parameters and their definitions for the medial surface approximation and partitioning algorithms. . . . .	61
5–1	Mean ( $\mu$ ) and standard deviation ( $\sigma$ ) of the 3 different types of curvature on a section of a cylindrical cup wall (shown in Figure 5–5) consisting of 551 vertices obtained using the method of [100] (Top) and our method (Bottom). We estimate the true positive curvature to be approximately 0.069898 and the true negative curvature to be approximately $-0.0875045$ .	118
7–1	Sphere sets $S_D$ are generated with our method; sphere sets $S_V$ are generated with the method of [23]. Timings are in seconds. . . . .	138
8–1	Timing and error results for separation distance tests for our grown spheres $S_D^+$ and spheres $S_V$ of [23] for a variety of polyhedral models. Timings are shown in milliseconds. Spheres are used as bounding volumes in the <i>SH</i> column and <i>RSS</i> s are used in the <i>RSSH</i> column. The maximum dimension of bounding boxes for all models is 10. . . . .	160

## LIST OF FIGURES

<u>Figure</u>	<u>page</u>
1–1 A ‘bunny’ shape represented using (left) a polygon mesh having 5110 triangle faces, and (centre) a union of 1559 spheres. Also shown is the detail of the bunny’s knee in the union of spheres representation. . . . .	3
1–2 Several implicit surfaces obtained by varying the distance between pairs of generating primitives. Image adapted from [15]. . . . .	7
1–3 An example of a non-solid. . . . .	9
1–4 Top: The medial axis (black) of a 2D solid (orange), overlaid. Bottom: A 3D solid (left) and its medial surface (right). . . . .	11
1–5 Left: Both the interior and the exterior medial axis, cropped to a rectangle. Right: An illustration of Property 1.1. . . . .	12
1–6 Different classes of points that compose the medial surface of a smooth 3D object [58]. $A_1^2$ points are smooth medial points, $A_1^3$ are junction medial points and $A_3$ are edge or rim medial points. Adapted from [106].	13
1–7 The edges of the Voronoi diagram of set of points (black) in the plane are shown in grey. They dual of the Voronoi diagram is the Delaunay Triangulation, whose edges are shown in blue. . . . .	16
1–8 (a) A union of a set of disks; (b) its power diagram overlaid; (c) the union of the set of disks decomposed using the power diagram; (d) the dual of this decomposition. . . . .	17
2–1 A polyhedron (left) and its medial surface (middle). The medial sheets that do not touch the boundary are shown on the right. Figure adapted from [36]. Some corresponding points on the polyhedron and its medial surface are shown. . . . .	25

2-2	Medial axes, in black, for two solids with red boundaries. The boundary on the right is obtained by smoothing the boundary on the left. Produced with [84]. . . . .	26
2-3	Left: The medial surface of a box, with each sheet shown in a different colour. The object angle $\theta$ at the selected medial point is $\pi/4$ and is half the angle between the two spoke vectors shown with black arrows. The object angle for points on the maroon sheet is $\pi/2$ . Right: In this 2D example, the boundary of a 2D solid is shown in black and its medial axis in red. The object angle of medial point $m_1$ ( $\angle A_1 m_1 B_1$ ) is greater than that of medial point $m_2$ ( $\angle A_2 m_2 B_2$ ). . . . .	27
2-4	The edges Voronoi diagram of a set of blue points is shown in grey. The dark edges are the subset of the Voronoi diagram that approximates the inner medial axis. Produced with [84]. . . . .	29
2-5	Arrows show $\nabla D$ , the directions to nearest locations on the boundary to points on $R$ . In this example, the medial axis intersects the line segment $(b, opp(b))$ because $\nabla D(b) \neq \nabla D(opp(b))$ . . . . .	37
2-6	The objects of interest in the proof of Lemma 2.4. . . . .	45
2-7	The boundary of a 2D solid is shown in black and its medial axis in red. In this case $(a, b = a + \gamma \nabla D(a))$ intersects the boundary and although $\nabla D(a) = \nabla D(b)$ , a medial point lies on $(a, b)$ . . . . .	47
2-8	A dragon polyhedron (left) approximated with 1324 spheres (centre) and 4102 spheres (right). Medial points were found in voxels intersected by the boundary $\mathcal{B}$ on the right, but not on the left. . . . .	52
2-9	(a) Sampled points on the red sphere, except $A$ , are medial points of the envelope of the black circles shown. (b) The medial axis of a regular polygon with low object angle medial points shown in black and a high object angle medial point $M$ shown in red. . . . .	53
2-10	An example of an object with boundary $B$ whose medial surface is a single medial sheet $MS$ where neighbouring smooth medial points have different surface normals (two surface normals are shown). Adopted from [39]. . . . .	55

2–11	Left: The original polyhedra; Centre: the union of spheres computed with our method; Right: approximate medial points grouped into smooth sheets. . . . .	57
2–12	Medial points computed for a solid with varying voxel side length $\sigma$ . Shown are the results for voxel lengths equal to $1/10^{th}$ , $1/20^{th}$ , $1/40^{th}$ and $1/100^{th}$ of the maximum dimension of the object bounding box, from left to right, respectively. . . . .	59
2–13	Left: A bumpy sphere model. Centre: Its medial surface approximation when $\theta = 0.6$ and no radius threshold is used. Right: The medial surface approximation when $\theta = 0.6$ and a radius threshold of 10% of the maximum dimension of the bounding box is used. . . . .	60
2–14	Voronoi diagram of three lines cropped to a cube. . . . .	60
3–1	When $\mathcal{B}$ consists of two points outside the sphere, the medial axis is shown as a dashed line. Points $\Phi$ are big dots on the sphere. . . . .	65
3–2	Objects of interest in the proof of Lemma 3.1. . . . .	66
3–3	If a medial point $m$ is in the convex hull of the 6 sampled points on the boundary of the dark disk with nearest boundary points $A$ , $B$ , and $C$ , then the nearest boundary points to $m$ are inside the green disk and outside the grey disks. The dashed lines are the bisectors of $A$ , $B$ and $C$ . . . . .	69
4–1	$\nabla D(A)$ , $\nabla D(B)$ point below the plane through $A$ , $B$ and $C$ , while $\nabla D(C)$ points above this plane. The nearest boundary points to $A$ , $B$ and $C$ lie on an edge of polyhedron. The medial surface does not necessarily intersect $\triangle ABC$ . . . . .	82
4–2	Cases A (Left) and B (Right) of Lemma 4.2. . . . .	83
4–3	Situations when Algorithm 4 returns ‘Undetermined’ while a section of the medial axis passes through circle $C$ . Medial points $p$ and $q$ are equidistant from points $A$ and $B$ on the boundary. Point $p$ has a small object angle, while the distance $d$ from point $p$ to the boundary is small. . . . .	85
4–4	Illustrations for error analysis in Case 1. . . . .	90
4–5	Angle $\beta$ is an upper bound for $\alpha$ , which is twice the object angle. . . . .	92

4–6	The angles of interest in Case 2. . . . .	94
4–7	When $(p, P)$ is orthogonal to $(x, y)$ and $dip = L, d = 2L$ , the $y$ -coordinate of $Y$ is $\approx -0.21L$ . . . . .	95
5–1	An illustration of $S_{rad}^a(v)$ for a given $U^a$ (based on a figure in [37]). . . . .	105
5–2	The case when $\kappa_{ri} = -\text{proj}_U \frac{\partial U_1}{\partial v_1} = \frac{1}{r}$ . . . . .	109
5–3	The medial axis of this object consists of low-object angle segments (dashed lines) and a high-object angle segment (bold line). When approximating the medial axis with a set of points, we only retain points on high-object angle segments of the medial axis. The boundary of the union of the associated medial circles approximates the original object. Medial point $m$ is equidistant from points $A$ and $B$ on the boundary. Surface curvatures at points $A$ and $B$ may be found using the radial shape operator, while for point $C$ a different strategy is used. . . . .	113
5–4	Left: Approximate medial points coloured according to principal curvatures on the solid’s boundary. Centre: Projection of curvature values at the approximate medial points to the boundary of the solid. Right: Curvature values obtained using the method in [100]. The colourmap used is shown in the bottom right corner. See the associated text for a discussion of these results. . . . .	117
5–5	The principal curvature directions recovered on the surface of a cup. . . . .	118
5–6	Top: An oblate spheroid coloured by principal boundary curvature $\kappa_1$ (left) and $\kappa_2$ (right). Middle: Points near the medial surface of the oblate spheroid coloured by principal boundary curvature. Bottom: Respective principal curvature estimates $\kappa'_1$ and $\kappa'_2$ shown on the medial surface. The mean absolute error of the estimation of $\kappa_1$ and $\kappa_2$ is shown below each column. . . . .	122
6–1	(a) A union of the set of balls corresponding to three medial sheets decomposed using the power diagram; (b) the dual of the decomposition of the set of balls. . . . .	126
6–2	A pear model approximated using a progressively larger number of sheets of medial balls. Most of the shape is covered using 4 sheets only. . . . .	127

6–3	A head model approximated using a progressively larger number of sheets of medial balls. Most of the shape is covered using 6 sheets only. . . . .	128
7–1	A polyhedron (Left) and two distributions of sphere centres for a sphere-based approximation this object. (Centre) The subset of the Voronoi vertices of points sampled on the polyhedron boundary that lies near its medial surface. (Right) Sphere centres computed with our method, developed in Chapter 2. . . . .	135
7–2	Consider the solid that is the union of the two grey balls, whose medial surface is the dashed red line. Increasing the voxel resolution (from blue to green) decreases the number of medial spheres generated for the solid.	136
7–3	Error and timing results for the pear (5154 triangles) and triple torus (16000 triangles) models for our method ( $SD$ ) and the AMAA method ( $SV$ ). Note that true $SV$ error may be higher than plotted. . . . .	141
7–4	(a) Reference polyhedron (top) approximated using internal spheres tangent to the polyhedron (bottom). (b)-(d) This sphere set is updated to approximate deformed versions of the reference polyhedron. The fraction of the time to compute the sphere set compared to computing the reference sphere set is shown below each pose. Also shown is the range of the volumetric error normalized by polyhedron volume. . . . .	148
8–1	The envelope of the dark circles $s_1$ and $s_2$ is non-differentiable at point $v$ , whose nearest point on the boundary is $m$ . The radii of both internal circles are increased to create circles that pass through $m$ . . . . .	151
8–2	Comparison of the $S_D^+$ and $S_V$ approximations for various models. (Centre) Histograms of signed distances from points sampled on the boundary of polyhedra to points on the envelopes of $S_D^+$ and $S_V$ . (Right) Histograms of signed distances from points sampled on meshes approximating the envelopes of $S_D^+$ and $S_V$ to the boundary of polyhedra. Stars denote $S_V$ spheres, while circles denote $S_D^+$ spheres. Means of the absolute distance are shown as a dashed line for $S_V$ and a dotted line for $S_D^+$ . . . . .	154
8–3	A rectangle swept sphere. . . . .	156

- 8–4 A fat triangle, or the convex hull of 3 spheres. The spheres contribute red patches to the hull; pairwise sphere interpolants are cylindrical sections shown in blue; the linear section of the fat triangle is shown in yellow. . 165

# Chapter 1

## Introduction

This chapter provides background for this thesis. Section 1.1 introduces the problem of shape representation, states criteria for a useful shape representation, and provides an overview of the existing work. Section 1.2 formally introduces the mathematical objects that are central to this thesis. Section 1.3 explains the objectives of the thesis, states the main contributions, and lists the publications arising from this work.

### 1.1 Shape Representation

Advances in data acquisition are making an unprecedented number of 3D models available for processing, obtained from various imaging modalities, including laser scans, multiple cameras, and biological imaging. Further, sophisticated modeling software enables artists and engineers to create novel 3D shapes of great complexity. Due to the growing number of available 3D models, the need to study appropriate shape representations is greater than ever.

The term *shape* will be used to mean a geometric entity, unique up to transformations due to scale, rotation, and translation, as well as invariant to the choice of appearance properties of its material such as transparency, reflectance, colour, and texture properties. A *shape representation* is a mathematical description of a shape that can be stored, visualized, manipulated, and transmitted by a computer.



The choice of an appropriate shape representation for a given shape and application can be motivated by the following criteria:

- i. ease of generation
- ii. ability to capture meaningful, with respect to the application considered, shape information
- iii. efficiency and accuracy of geometric operations

This thesis introduces the *set of well-spaced medial spheres* as an alternative shape representation and discusses algorithms for its computation (Part I), investigates the descriptive power of this shape representation (Part II), and examines some applications that benefit from this representation (Part III). Specifically, we<sup>1</sup> will consider the set of medial spheres such that the coordinates of each sphere centre, snapped to a grid, are unique.

In this thesis, we will be concerned with shapes that are 3D solids. In this section, we will discuss common types of shape representations for 3D solids and discuss each type with respect to the criteria stated above. Most existing shape representations can be classified as either parametric or implicit. In a *parametric shape representation*, the object's surface is described by a mapping  $f : U \rightarrow Q$  from the local parameter domain  $U$  to the surface  $Q$ . In an *implicit shape representation*, each point of the space in which the object resides is assigned a weight by the mapping  $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ . The object's surface is then given as a level set of  $F$ , usually the zero level set. These two representations differ

---

<sup>1</sup> I, the author, would like to invite you, the reader, to be an active participant in the subsequent discussion. Throughout the thesis, the pronoun 'we' refers to both the author and the reader.




Figure 1–1: A ‘bunny’ shape represented using (left) a polygon mesh having 5110 triangle faces, and (centre) a union of 1559 spheres. Also shown is the detail of the bunny’s knee in the union of spheres representation.

in the following fundamental way. Parametric shape representations explicitly describe the surface of the object, while its interior must be inferred. On the other hand, implicit shape representations describe the volume occupied by the object, while a description of its boundary must be inferred. While the terminology of ‘parametric’ and ‘implicit’ shape representations is standard in the computer graphics community, in the computer vision community these representations are instead referred to as ‘boundary’ and ‘interior’ representations, respectively.

In the following discussion, we will describe each of these types of representations in turn, while commenting on their qualities with respect to the above mentioned criteria. Shape representations based on a union of balls will be discussed in the section on implicit surface representations (Section 1.1.2) because the surface of the union of balls must be derived from this representation. Figure 1–1 shows an example of a polygon mesh (parametric) representation and a union of balls (implicit) representation of the same shape.

### 1.1.1 Parametric/Boundary Representations

In this type of shape representation, an object's surface is approximated using simple patches, each parametrizable. The boundary is thus piecewise parametric, and the individual patches can be selected and assembled so as to provide varying degrees of smoothness. The most popular kind of parametric representation is a triangle mesh, where the boundary is approximated using triangle patches. Such a boundary is thus  $C^0$ -continuous, and often is designed to be a manifold mesh (locally topologically equivalent to a disk). The reader is referred to [22] for a discussion of mesh manifoldness. Alternative patches based on splines offer higher-order continuity and faster convergence to the object being modeled.

Polygon meshes are the de facto standard in the computer graphics and geometry processing communities for representing 3D objects. Such meshes may be generated by connecting points sampled on the shape boundary. The area of 'surface reconstruction' is concerned with the generation of qualitatively useful meshes from point samples. They may also be generated by artists using modeling tools. Polygon meshes offer sufficient flexibility to represent both the coarse and the fine levels of detail of the shape being modeled. There exist algorithms that enable efficient operations on meshes, such as simplification, smoothing, and deformation. Efficient rendering of polygon meshes is made possible in practice via hardware acceleration on modern graphics cards. However, enforcing mesh quality, especially as the mesh is altered, can be challenging. Further, because the interior of the object is not represented explicitly and the number of boundary patches can be very large, point-location queries and distance tests can be expensive. For a thorough discussion on advances in polygon mesh processing, please see [22].

### **1.1.2 Implicit/Interior Representations**

In this type of shape representation, the object interior is represented using volumetric primitives. Among the different types of primitives that we will discuss are voxels, spheres, and metaballs. This representation is less flexible than the boundary mesh representation, in that a substantial number of volumetric primitives may be required to capture small features of the boundary. On the other hand, the surface of an object represented implicitly is non self-intersecting by definition. This is in contrast with parametric representations, where care must be taken to ensure such correctness, especially as objects deform. Further, efficient point-location queries can be made by checking if the point lies in any of the volumetric regions. For points outside the object, the distance to the object represented implicitly is the minimum of the distances to each of the volumetric primitives. When the distance between a point and a volumetric primitive can be computed efficiently and when the number of volumetric primitives is small, implicit representations offer substantial efficiency improvements over parametric representations for distance computations. Further, the given partition into volumetric primitives can serve to provide a meaningful part-based description of the object for such applications as registration, morphing, and matching. Obtaining a part-based object description for objects described parametrically can be difficult. However, rendering implicit representations is usually done by either conversion to a polygon mesh representation of potentially very large size or by ray tracing, and tends to be slow.

#### **Cellular Subdivision Representations**

When space is subdivided into regular-size cubes, each such cube is called a voxel. Voxel representations identify those cubes in a regular partition of space into cubes that

represent a shape. Octrees use axis-aligned cells of mixed size: coarse cells containing geometry are subdivided until a desired resolution is reached. Because each voxel has at most 26 voxel neighbours, voxel data sets can be analyzed by exploring these digital neighbourhoods. Storage requirements of high resolution voxel data sets are large. Octrees have smaller storage requirements than voxel representations when a small number of cells in a volume is ‘on’.

### Union of Balls Representations

Given a set of balls  $B = \{(c_i, r_i) | c_i \in \mathbb{R}^3, r_i \in \mathbb{R}\}$  with centres  $c_i$  and radii  $r_i$ , the surface of the union of balls is the zero level set of the function

$$F : \mathbb{R}^3 \rightarrow \mathbb{R},$$

$$x \mapsto \min_i \{d(x, c_i) - r_i\}.$$

This shape representation is advocated in the work of Alain Fournier [97] as a superior shape representation to triangle meshes because of its stability, as shapes represented using the same or similar sets of spheres can be described by completely different triangle meshes in terms of the location of the triangle vertices and their adjacency, and because of the relative invariance of union of spheres representations to noise in the input.

The generation of union of balls representations that offer advantages over alternative representations is not trivial. Chapter 7, Section 7.1, surveys existing approaches for approximating a shape with balls.

Union of balls shape representations can provide advantages against triangle mesh representations in those applications that do not need to represent small-scale, sharp, and

narrow boundary features, as prohibitively many spheres may be needed to faithfully capture such shapes. These representations are particularly appropriate for modeling molecular surfaces [47].

In recent years, a number of applications have made use of ball-based or sphere-based approximations. Particularly, hierarchies of spheres are used for collision detection in [23, 92, 63] because sphere-sphere intersection tests are fast and simple. Sphere approximations are used in the application of soft shadow generation [98], where a low-frequency object representation is sufficient. Additionally, sphere-based representations are used for efficient level-of-detail rendering [101], registration [125], penetration-depth computation [121], shape morphing [104, 32], and shape deformation [126].

Rendering these representations is expensive, as many triangles are needed to generate qualitatively good approximations to unions of balls. Existing techniques for meshing the envelope of balls include [113, 48, 75] and are discussed in Section 5.4.6.

### **Bloppy Representations**




Figure 1–2: Several implicit surfaces obtained by varying the distance between pairs of generating primitives. Image adapted from [15].

This class of shape representations considers density fields about particles. A level set of the sum of the density fields of the particles in the representation gives the surface of

the shape. Blinn [15] proposed using a truncated Gaussian density field (*cf.* Figure 1–2). Simultaneously, Omura proposed ‘metaballs’, where the density is given by a polynomial function. Wyvill *et al.* [123] introduced ‘soft objects’, where the density function is a polynomial expression that is cheaper to evaluate than that of metaballs. To facilitate modeling shapes that are locally flat, Bloomenthal [16] proposed convolving a piecewise-planar 2D skeleton with a density function.

These ‘blobby’ representations are generated by artists using modeling tools. They are particularly appropriate for modeling amorphous objects, such as raindrops, mud and dough. Objects of high artistic value for computer graphics can be designed using a small number of primitives. The smoothness properties of these representations provide the objects with either an organic or a plasticine, cartoon-like appearance. These representations are an important addition to an artist’s toolbox for shape generation. As with other implicit representations, efficient rendering of such objects is challenging. For additional information, the reader is referred to a survey in [88].

## 1.2 Preliminaries

Section 1.2.1 will formally introduce the medial surface transform and its properties. Its discrete approximation is the ‘union of medial spheres’ shape representation, which is the topic of this thesis. Section 1.2.2 provides an overview of basic space partitions that will be referred to throughout the thesis.

### 1.2.1 Medial Surface Transform

In this thesis, we will propose a shape representation for a shape that is a solid. As a formal definition of a solid, we follow [2, ch. 5], and say that a subset  $X$  of  $\mathbb{R}^n$  is a *solid* if the closure of the interior of  $X$  is  $X$ . In this definition, for example, the union of a cube

and a line segment cannot be a solid, see Figure 1–3. A 3D solid is a solid subset of  $\mathbb{R}^3$ , while a 2D solid is a solid subset of  $\mathbb{R}^2$ .




Figure 1–3: An example of a non-solid.

Consider a 3D solid  $\Omega \subset \mathbb{R}^3$  with boundary  $\mathcal{B}$ .

**Definition 1.1.** A closed ball  $B \subset \mathbb{R}^3$  is a maximal inscribed ball in  $\Omega$  if  $B \subset \Omega$  and  $B$  is not contained in any other ball  $B'$  such that  $B' \subset \Omega$ . A maximal inscribed ball is also called a medial ball.

From this definition, we can define the primitive for our shape representation, a medial sphere:

**Definition 1.2.** A medial sphere  $S$  of  $\Omega$  is the boundary of a medial ball  $B$  in  $\Omega$ .

The centre of a medial sphere has a special name:

**Definition 1.3.** The centre of a medial sphere is a medial point.

Consider a medial ball centred at medial point  $p$ . Because a medial ball is maximal and inscribed, it is tangent to the boundary  $\mathcal{B}$  of  $\Omega$ . These points of tangency are the closest points on  $\mathcal{B}$  to  $p$ . The directions from  $p$  to the points of tangency have a special significance:



**Definition 1.4.** *The vectors from a medial point  $p$  to its closest points on  $\mathcal{B}$  are the spoke vectors of  $p$ .*<sup>2</sup>

The set of all medial points defines the following object:

**Definition 1.5.** *The medial surface  $\mathcal{MS}$  of  $\Omega$  is the closure of the set of medial points of  $\Omega$ .*

We consider the closure in the above definition because the set of medial points may not be closed; see [31] for a 3D example.

When  $\Omega$  is a 2D solid, the 2D counterpart of the medial surface is called the *medial axis* and is the closure of the set centres of maximal inscribed disks in  $\Omega$ . Many authors use the term ‘medial axis’ to refer to the medial surface of a 3D solid, as well as higher-dimensional variants. Figure 1–4 presents examples of a medial axis and the medial surface of non-trivial 2D and 3D solids, respectively.

The medial surface we have thus defined is also known as the *interior medial surface* of  $\Omega$ . The *exterior medial surface* of  $\Omega$  is the closure of the set of centres of all maximal balls exterior to  $\Omega$ , *i.e.*, it is the medial surface of the closure of the complement of  $\Omega$ . This object can be infinite (*e.g.*, for non-convex  $\Omega$  with no cavities) and empty (for convex  $\Omega$  with no cavities). The exterior and interior medial axes are defined analogously for 2D solids. Figure 1–5(left) shows an example of both the interior and the exterior medial axis of a 2D solid.

---

<sup>2</sup> We follow the terminology of [106]. These vectors are also called pannormals in the work of Harry Blum [18].




Figure 1–4: Top: The medial axis (black) of a 2D solid (orange), overlaid. Bottom: A 3D solid (left) and its medial surface (right).

Consider the set of all medial spheres of  $\Omega$ . These spheres have the following very important property [17, 18]:

**Property 1.1.** *The envelope of the medial spheres of  $\Omega$  is the boundary  $\mathcal{B}$  of  $\Omega$ .*

Figure 1–5(Right) provides an example of a finite set of medial spheres such that its envelope approximates the boundary of the 2D solid. Consider the following definition:

**Definition 1.6.** *The medial surface transform of  $\Omega$  is the set of all medial spheres of  $\Omega$ .*

The medial surface transform provides an alternative description of the shape of a solid that captures its local width with respect to its ‘skeleton’: a lower-dimensional representation of the volume of the solid that makes explicit its symmetries. The medial




Figure 1-5: Left: Both the interior and the exterior medial axis, cropped to a rectangle. Right: An illustration of Property 1.1.

surface transform was proposed in the 1960s by Harry Blum [17, 18] as a shape descriptor that makes explicit perceptually significant shape properties. It has since been studied extensively in the computer vision community as it allows abstraction sufficient to guide shape matching, in robotics as it offers a collision avoiding path, in computer graphics to generate deformations of models, in medical imaging to capture local shape variation, in geography to simplify the shape of rivers, among some examples. See [106, Ch. 11] for a detailed overview of applications of the medial surface transform.

For all solids (except those that are disjoint balls) the medial surface transform consists of an infinite set of medial spheres. **The shape representation proposed in this thesis is a union of a finite set of medial spheres and is thus a discretization of the medial surface transform.**




Figure 1–6: Different classes of points that compose the medial surface of a smooth 3D object [58].  $A_1^2$  points are smooth medial points,  $A_1^3$  are junction medial points and  $A_3$  are edge or rim medial points. Adapted from [106].

Here we survey some important properties of the continuous representation that will be important in our discussion of our algorithm to approximate the medial surface transform of a solid in Part I and the analysis of the quality of our approximation in Part II.

For a smooth solid, medial points (generically) fall into a small number of classes [58]. The majority are equidistant from exactly two distinct boundary locations, and are called *smooth medial points*. Contiguous regions of smooth medial points form *medial sheets* (these are manifolds with boundary) [58]. These sheets intersect at *junction curves* made of medial points equidistant from exactly three distinct boundary locations and are bounded by *edge or rim curves* (see Figure 2–3 (right)).

Objects that have tubular structure, *i.e.*, ones that are circular in cross section, are *non-generic* because their medial surface consists of curves, not sheets. Using this definition, the following property of the medial surface [58] summarizes one of its attractive features:

**Property 1.2.** *The medial surface generically has a natural part decomposition into medial sheets.*

Not only does the medial surface transform of the solid  $\Omega$  reconstruct  $\Omega$  completely, but also the medial surface and  $\Omega$  have the same topology: that is,  $\Omega$  and  $\mathcal{MS}$  have the same number of connected components, tunnels, and cavities, and moreover, these have the same relationship in the medial surface  $\mathcal{MS}$  as they do in the solid  $\Omega$ . For example, this is true of the medial axis and medial surface examples in Figure 1–4.

Let us introduce the topological concept of *homotopy equivalence* that formalizes this relationship. Following [35],

**Definition 1.7.** *Let  $X$  and  $Y$  be topological spaces and  $f, g$  continuous functions from  $X$  to  $Y$ . Then  $f$  is homotopic to  $g$  means that there is a continuous function  $H : X \times [0, 1] \rightarrow Y$ , such that for all  $x \in X$ ,  $H(x, 0) = f(x)$  and  $H(x, 1) = g(x)$ .*

Using this definition, *homotopy equivalence* can be defined as follows:

**Definition 1.8.** *Let  $X$  and  $Y$  be topological spaces. Then  $X$  and  $Y$  are homotopy equivalent provided that there exist continuous maps  $f : X \rightarrow Y$  and  $g : Y \rightarrow X$  for which the composite maps  $gf$  and  $fg$  are homotopic to the identity maps on  $X$  and  $Y$ , respectively.*

Lieutier [79] showed the following relationship between the medial surface and  $\Omega$ :

**Property 1.3.** *The complement of  $\Omega$  is homotopy equivalent to the interior of the medial surface of  $\Omega$ .*

Additional properties of the medial surface transform are discussed in [11].

### 1.2.2 Space Partitions

We now define two important space partitions: the Voronoi diagram and the power diagram. We will show how the power diagram can be used to compute the volume of a union of spheres.

## Voronoi Diagram

The Voronoi diagram, and its dual, the Delaunay Triangulation, are important geometric data structures with applications to most areas of science and engineering. In particular, they help approximate the medial surface transform, as we will discuss in Chapter 2 and Chapter 7.

**Definition 1.9.** *Given a set of point sites  $P = \{p_1, p_2, \dots, p_n\}$ ,  $p_i \in \mathbb{R}^3$ , the Voronoi diagram  $VD$  of  $P$  is a partition of  $\mathbb{R}^3$  into Voronoi cells  $V(p_i)$  with respect to the Euclidean distance  $d_E$ , such that*

$$V(p_i) = \{x \in \mathbb{R}^3 \mid d_E(p_i, x) \leq d_E(p_j, x), \forall j \neq i\}. \quad (1.1)$$

A Voronoi cell of a site  $p_i$  is, hence, the set of points that are closer in Euclidean distance to the site  $p_i$  than to any other site in  $P$ .

Let us assume that  $P$  is in general position, such that at most 4 points are co-spherical. A vertex of  $V(p_i)$  is called a *Voronoi vertex* and is the centre of a *Voronoi sphere* that touches 4 sites of  $P$ , but does not contain any point in its interior. A vertex of  $V(p_i)$  is, hence, equidistant from 4 sites in  $P$ . A point on an edge of  $V(p_i)$  is equidistant from 3 sites in  $P$ , and a point on a face of  $V(p_i)$  lies on a bisector of a pair of sites of  $P$ .

The dual of the Voronoi diagram is called the *Delaunay Triangulation* of  $P$ ,  $DT(P)$ . The vertices of  $DT(P)$  are the points in  $P$ , the edges connect those vertices of  $DT(P)$  whose Voronoi cells share a face, triangular faces connect those triples of vertices in  $DT(P)$  that define an edge of  $VD(P)$ , while tetrahedra in  $DT(P)$  connect those quadruples of vertices that define a Voronoi vertex. The circumscribing sphere of a tetrahedron in  $DT(P)$  is the Voronoi sphere.




Figure 1–7: The edges of the Voronoi diagram of set of points (black) in the plane are shown in grey. They dual of the Voronoi diagram is the Delaunay Triangulation, whose edges are shown in blue.

Figure 1.2.2 shows an example of the Voronoi diagram of a set of point sites in the plane and its dual.

### Power Diagram

Consider a ball  $b = (c, r)$  with centre  $c \in \mathbb{R}^3$  and radius  $r \in \mathbb{R}$ . The *power distance* between a point  $x$  and a ball  $b = (c, r)$ ,  $d_P(x, b)$ , is given as

$$d_P(x, b) = d_E^2(x, c) - r^2. \quad (1.2)$$

The power diagram is a space partition with respect to the power distance, as follows:

**Definition 1.10.** Given a set of balls,  $B = \{b_1 = (c_1, r_1), b_2 = (c_2, r_2), \dots, b_n = (c_n, r_n)\}$ , with centres  $c_i \in \mathbb{R}^3$  and radii  $r_i \in \mathbb{R}$ , the power diagram of  $B$ , denoted  $PD(B)$ , is a partition of  $\mathbb{R}^3$  into convex power cells  $P(b_i)$ , such that

$$P(b_i) = \{x \in \mathbb{R}^3 \mid d_P(x, b_i) \leq d_P(x, b_j), \forall j \neq i\}. \quad (1.3)$$

Power cells (Eq. 1.3) are identical to Voronoi cells (Eq. 1.1) when all the ball radii are the same. Figure 6–1(b) shows a 2D example of the power diagram of a set of disks.




Figure 1–8: (a) A union of a set of disks; (b) its power diagram overlaid; (c) the union of the set of disks decomposed using the power diagram; (d) the dual of this decomposition.

A union of balls can be decomposed into convex cells by intersecting each ball with its power cell. The dual of the decomposition of a union of the balls using its power diagram is known as the *dual complex*  $DC(B)$ . The dual complex is a simplicial complex that describes the combinatorial structure of the union of the balls in  $B$  [47]. This complex is also known as the zero-shape or the zero-alpha complex of  $B$ .  $DC(B)$  contains a vertex  $i$  for each ball  $b_i$ , an edge  $(i, j)$  whenever balls  $b_i$  and  $b_j$  share a face of  $PD(B)$ , a triangle  $(i, j, k)$  whenever balls  $b_i, b_j, b_k$  share an edge of  $PD(B)$ , and a tetrahedron  $(i, j, k, l)$  whenever spheres  $b_i, b_j, b_k, b_l$  share a vertex of  $PD(B)$ . As shown in [47], to find the total volume of the union of balls  $b_i$ ,  $\text{vol}(\cup_i b_i)$ , one need only consider the balls corresponding to vertices, edges, triangles and tetrahedra of  $DC(B)$ :

$$\begin{aligned}
 \text{vol}(\cup_i b_i) &= \sum_{i \in DC(B)} \text{vol}(b_i) - \sum_{(i,j) \in DC(B)} \text{vol}(b_i \cap b_j) \\
 &+ \sum_{(i,j,k) \in DC(B)} \text{vol}(b_i \cap b_j \cap b_k) \\
 &- \sum_{(i,j,k,l) \in DC(B)} \text{vol}(b_i \cap b_j \cap b_k \cap b_l).
 \end{aligned} \tag{1.4}$$



The number of simplices of  $DC(B)$  is  $O(|B|^2)$  and is much smaller for typical ball distributions, such as ones that we will consider in this thesis [47]. Use of this simple formula makes the computation of the volume of a union of balls efficient.

## 1.3 Thesis Overview

In this section, the main objectives of the thesis and the order in which they will be addressed are described. We will then list the novel contributions of this thesis. We will also specify which contributions appear in published or submitted work.

### 1.3.1 Objective and Outline

This thesis introduces a new shape representation: the *set of well-spaced medial spheres*, that is, a union of medial spheres the coordinates of whose centres, when rounded to the nearest integer, are unique integer triples. This distribution of medial spheres is well-spaced because there is at most one sphere centre per cubic region of space. The goal of this thesis is to investigate the strengths of this shape representation in terms of the three desiderata of shape representations mentioned in Section 1.1: 1) ease of generation, 2) ability to capture meaningful shape information, and 3) efficiency of geometric operations. The thesis consists of three parts, each of which addresses the desiderata above, in that order.

Part I introduces algorithms for converting from an explicit boundary representation of a solid to a set of well-spaced medial spheres by approximating the medial surface of a solid. When the boundary of the solid is a polygonal mesh, practical algorithms for the efficient computation of the union of medial spheres shape representation are developed. Correctness and completeness issues of the proposed algorithms are discussed in each of the three chapters of Part I.

Part II investigates the descriptive power of the set of well-spaced medial spheres shape representation. When spoke vector estimates are given for each medial point, Chapter 5 shows how boundary curvature can be recovered from this representation. We also show that the proposed shape representation allows us to easily evaluate the significance and simplify the parts of a solid in Chapter 6.

Part III shows that the set of well-spaced medial spheres provides a fast and tight volumetric approximation to a solid in Chapter 7 and also studies advantages offered by this representation when the solid is deformable. Chapter 8 demonstrates how this shape representation can offer improvement over existing union of spheres shape representations for performing proximity queries.

Each of the three parts includes an overview of related previous work, as well as a description of future work arising from those parts.

### **1.3.2 List of Novel Contributions**

The novel contributions of this thesis are as follows:

- Section 2.4 describes an algorithm to obtain a dense set of approximate medial points within a user-specified distance of a solid’s medial surface, along with the spoke vector estimates for the medial points. In Section 2.5, we show how points that lie near distinct smooth medial sheets can be grouped by studying the spoke vector estimates at each approximate medial point.
- Chapter 3 initiates the theoretical analysis of the completeness of this algorithm: we make progress towards a characterization of those medial points that are undetected by our algorithm. Section 3.2 discusses the locations of query points to be considered by the algorithm. Section 3.3 establishes the locations of nearest boundary

points to medial points that are not detected by our algorithm for a finite sampling rate. Section 3.4 shows how the quality of the undetected medial points can be measured under certain conditions.

- Chapter 4 looks at the significantly easier 2D case and considers theoretical properties of an algorithm for computing the medial axis of a 2D solid, based on tools used in the 3D case. We develop additional tools to the ones used in the 3D case for detecting medial points in Section 4.1. Section 4.3 gives a complete geometric description of those medial points that are undetected by these tools.
- Section 5.4 presents a numerical method that allows us to estimate the principal curvatures and principal curvature directions at the two implied boundary patches on either side of each approximate medial point, given a dense sampling of approximate medial points and their spoke vector estimates.
- Sections 6.2 and 6.3 present a method by which the individual sheets of the discrete medial surface approximation may be ordered by significance. This allows us to substantially reduce the size of the resulting part-based representation and to simplify the shape of the solid.
- Chapter 7 looks at the application of our shape representation to the problem of quickly generating a well-fitting sphere-based approximation to a polyhedron. We show how volumetric error of sphere-based shape representations can be evaluated in Section 7.3. Compared to a state-of-the-art method for approximating polyhedra with spheres, we show that our method is significantly faster and provides a tighter fit in terms of volumetric error in Section 7.3. When a model undergoes local feature size preserving deformation, we show in Section 7.4 how the sphere approximation

can be quickly updated and how the volumetric error of the new sphere sets can be evaluated.

- We then use our sphere-based representation to compute approximate separation distance between pairs of polyhedra in Chapter 8. To allow this application, we propose a method to improve the coverage of the polyhedron’s boundary with spheres in Section 8.2 and describe how an efficient bounding volume hierarchy of the sphere sets can be built to accelerate distance tests in Section 8.4.

### **1.3.3 Publications Arising from This Thesis**

We now describe which portions of this thesis have been previously published. The list of such publications, indexed by a capital letter, can be found on the next page.

The majority of the results in Chapter 2 can be found in [E] and its extended version, [C]. The results of Section 3.2 and 3.3 appear in [B]. The results of Chapter 4 appear in [F]. The contributions of Chapter 5 appear in [C, E]. The results of Chapter 6 appear in [C]. The results of Section 7.3, and a preliminary version of the results in Chapter 8, appear in [D, A]. Also included in [A] are the remaining novel contributions of Chapters 7 and the contributions in Chapter 8.

## PUBLICATIONS

- [A] S. Stolpner, P. Kry, K. Siddiqi. Medial Spheres for Shape Approximation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, in press.
- [B] S. Stolpner, S. Whitesides, K. Siddiqi. Approximating the Medial Axis by Shooting Rays: 3D Case. In *Canadian Computational Geometry Conference*, Toronto, Canada, 2011.
- [C] S. Stolpner, S. Whitesides, K. Siddiqi. Sampled Medial Loci for Shape Representation. *Computer Vision and Image Understanding*, 115(5):695–706, 2011.
- [D] S. Stolpner, P. Kry, K. Siddiqi. Medial Spheres for Shape Approximation. In *Symposium on Brain, Body, and Machine*, Montreal, Canada, 2010.
- [E] S. Stolpner, S. Whitesides, K. Siddiqi. Sampled Medial Loci and Boundary Differential Geometry. In *ICCV Workshop on 3D Digital Imaging and Modeling*, pages 1855-1862, Kyoto, Japan, 2009. (selected as oral presentation)
- [F] S. Stolpner, S. Whitesides. Medial Axis Approximation with Bounded Error. In *International Symposium on Voronoi Diagrams*, pages 171-180, Copenhagen, Denmark, 2009.

## **Part I**

### **Medial Spheres – Construction**

## Chapter 2

# Medial Surface Approximation

In this chapter, we present a new algorithm for the approximation of the medial surface of a 3D solid with a set of points, with certain quality guarantees. Section 2.1 explains the nature of medial surfaces of 3D polyhedral solids and motivates the need for their approximation and simplification. Section 2.2 presents an overview of previous work on the computation and simplification of medial surfaces. Section 2.3 describes an approach for the fast detection of cubic regions of space (voxels) that are intersected by the medial surface. Section 2.4 shows how this approximation to the medial surface with voxels can be refined to generate a point cloud that lies near a desirable subset of the medial surface, with one point per voxel. The set of spheres thus generated gives a well-spaced union of medial spheres representation of the solid. We also explain how the medial point cloud can be partitioned into points lying near distinct smooth medial sheets in Section 2.5. Section 2.6 presents experimental results and a discussion of performance.

### 2.1 Background

When  $\Omega$  is a polyhedron, its medial surface  $\mathcal{MS}$  is composed of bisectors of the faces, edges and vertices of the polyhedron boundary  $\mathcal{B}$ . The bisector of two such elements is a quadric surface and these surfaces intersect along curves of higher algebraic degree. Culver *et al.* [36] presents a detailed discussion of the nature of medial sheets




Figure 2–1: A polyhedron (left) and its medial surface (middle). The medial sheets that do not touch the boundary are shown on the right. Figure adapted from [36]. Some corresponding points on the polyhedron and its medial surface are shown.

of a polyhedron. Figure 2–1 shows a relatively simple polyhedron and its medial surface, which consists of a number of curved sheets. For polyhedra with a large number of faces, the algebraic complexity of the medial sheets and their potentially quadratic number prohibits exact computation of the medial surface. For polyhedral models having a large number of faces, typically one only attempts to approximate the medial surface rather than compute it exactly. For higher-order boundary representations, the algebraic complexity of the medial sheets and their junctions increases.

For each pair of adjacent faces of a polyhedral boundary that meet convexly, the complete medial surface of the polyhedron contains portions of the bisector of these faces. This property results in a large number of medial sheets, but not all of these are deemed “significant”. Small modifications to the object boundary can have a significant effect on the medial surface of the object. Figure 2–2 presents a 2D example. When approximating the medial surface, one often seeks to remove less significant portions in order to compute a simpler medial surface that continues to provide a nearly complete representation of the






Figure 2–2: Medial axes, in black, for two solids with red boundaries. The boundary on the right is obtained by smoothing the boundary on the left. Produced with [84].

original object. The nature of significance, or “salience”, is application dependent: for motion planning, an exact medial surface is sometimes necessary, while for matching and animation, a medial surface of minimal geometric complexity that describes the object shape is sought.

Two common measures of significance of a medial point include

- the *radius* of the associated medial sphere, and
- the *object angle* of the medial point, defined as follows:

**Definition 2.1.** *For a smooth medial point  $p$ , the object angle  $\theta$  is the angle between the vector from  $p$  to either of its two closest points on  $\mathcal{B}$  and the tangent plane to  $\mathcal{MS}$  at  $p$ .*

Figure 2–3(Left) gives an example of how the object angle  $\theta$  is evaluated. As shown in [53], removal of medial points that have a small object angle has a small impact on the volume of the union of medial balls (refer to Figure 2–3(Right)). When it is desirable to preserve the volume of the solid being approximated, eliminating medial points with small object angle is an appropriate simplification measure.

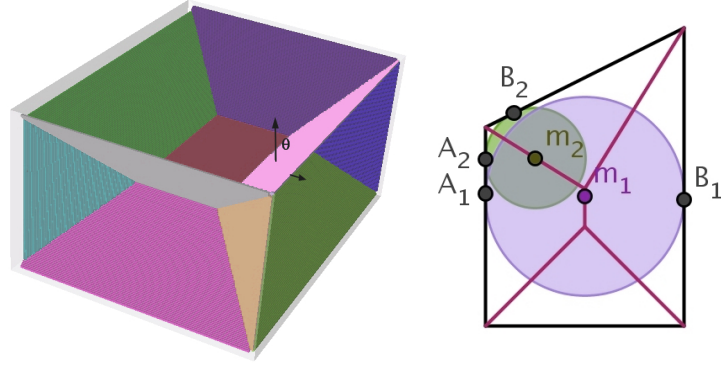


Figure 2–3: Left: The medial surface of a box, with each sheet shown in a different colour. The object angle  $\theta$  at the selected medial point is  $\pi/4$  and is half the angle between the two spoke vectors shown with black arrows. The object angle for points on the maroon sheet is  $\pi/2$ . Right: In this 2D example, the boundary of a 2D solid is shown in black and its medial axis in red. The object angle of medial point  $m_1$  ( $\angle A_1 m_1 B_1$ ) is greater than that of medial point  $m_2$  ( $\angle A_2 m_2 B_2$ ).

## 2.2 Previous Work

In this section, we provide an overview of several notable methods for the approximate and exact computation of the medial surface. We will be concerned only with 3D algorithms whose input is either the boundary of a solid or a cloud of points. There exists a large body of work that processes ‘digital’ inputs, that is, inputs defined on a discrete grid. We will not address such methods in this overview, as the issues involving (truncated) real-valued inputs, considered throughout this thesis, are significantly different than those involving inputs defined on a small integer grid. Often, these methods proceed by an iterative removal of voxels. For a survey of advances in the generation of medial surfaces of digital inputs, the reader is directed to [106, Chs. 4,5]. The most successful methods for the computation of the medial surface for non-digital inputs, at present, belong to these

three categories: Voronoi methods, spatial subdivision methods, and tracing methods. We now describe each type of method in turn.

### 2.2.1 Voronoi Methods

Given a set of points sampled on a surface, the Voronoi diagram of the sampled points is a powerful tool for computing the medial surface of the solid bounded by this surface. Voronoi methods use a subset of the Voronoi faces, edges, and vertices to approximate the medial surface and a subset of the Voronoi edges and vertices to approximate the medial axis. As an example, Figure 2–4 shows the Voronoi diagram of a 2D point set and Voronoi edges that approximate the medial axis are highlighted. For 3D point inputs, the modern CGAL library [1] includes code that computes the Voronoi diagram of a set of points in a manner that is both highly robust and efficient. Issues arising in the implementation of such algorithms are to determine the appropriate sampling rate for the surface points and to determine the right subset of the Voronoi diagram that provides a desirable approximation to the medial surface. Voronoi-based methods can provide theoretical bounds on the quality of a medial surface approximation in terms of homotopy equivalence and geometric proximity to the true medial surface.

The difficulty with using Voronoi diagrams of 3D points to approximate the medial surface is that not all Voronoi vertices lie near the medial surface. Specifically, among the sample points considered, there may exist four nearly co-planar nearby points that determine a small sphere empty of other points, *i.e.*, a Voronoi vertex. This vertex can be arbitrarily far from the medial surface. Amenta *et al.* [7] defines a subset of the Voronoi vertices, called the *poles*, and shows that the set of poles converges to the medial surface




Figure 2–4: The edges Voronoi diagram of a set of blue points is shown in grey. The dark edges are the subset of the Voronoi diagram that approximates the inner medial axis. Produced with [84].

when the sampling density of the surface satisfies certain conditions. Attali and Boissonnat [10] describe conditions when a particular subset of the Voronoi faces lies near the medial surface. When the Hausdorff distance between the sampled points and the surface is less than  $\delta$  (permitting noisy samples on non-smooth surfaces, but requiring many samples if the surface contains small-scale features), Chazal and Lieutier [30] show convergence of those Voronoi vertices, edges and faces, for which the distance between nearest surface samples is at least  $\delta$ , to the subset of the medial surface for which the circumradius of the nearest boundary points is a function of  $\delta$ . This subset of the medial surface is valuable because in [29] Chazal and Lieutier show homotopy equivalence between the medial surface simplified by a parameter  $\lambda$ , the circumradius of the nearest boundary points to medial points, and the complete medial surface, where the choice of the  $\lambda$  parameter is related to the distance from the sampled surface to the medial surface. Giesen *et al.* [61] provide an

algorithm to compute a cell complex that offers homotopy equivalence to the true medial surface, given a dense, but not necessarily uniform, set of sampled surface points.

Consider the Voronoi diagram of a set of sample points on the boundary of a solid. The faces of the Voronoi diagram that are interior to a solid can serve as an approximation to the medial surface of the solid [13, 43]. In order to simplify the set of faces, Attali and Montanvert [13] propose a thinning strategy for removing Delaunay tetrahedra, thus simplifying the dual Voronoi diagram, guided by the dihedral angle of faces of tetrahedra eligible for removal. Attali and Lachaud [12] propose to iteratively remove exterior Voronoi vertices in the medial surface approximation when either an object angle or a radius threshold is not satisfied. Dey and Zhao [43] propose two scale-invariant pruning conditions for Voronoi faces, defined as follows. The furthest interior Voronoi face of a sample point  $p$  is evaluated both by using the angle between the surface normal estimated at  $p$  and the dual Delaunay edge of the Voronoi face, and also by considering the ratio of the length of this dual edge and the circumradii of surface triangles at  $p$  [43].

The “power crust” method for medial surface approximation connects the poles of sampled points on a solid’s surface using the dual of the power diagram of the *polar balls* (Voronoi balls centred at the poles) [8]. The medial surface thus computed is homotopy equivalent to the space enclosed by the sampled surface. It may, however, contain tetrahedra, and thus is not necessarily strictly a 2D cell complex. Numerous techniques have been proposed for pruning this approximation to the medial surface. Tam and Heidrich [112] propose a topology-preserving peeling strategy that first aggregates poles into sheets and then removes those sheets that have a small number of triangles or that correspond to features of small volume (estimated as the volume of the Delaunay tetrahedra associated with

the poles). Aichholzer *et al.* [5, 3] simplify the set of polar balls by growing all balls by a given parameter and then applying a heuristic to solve the set covering problem of finding the minimum size subset of balls that continues to cover the sampled surface points. Miklos *et al.* [60, 85] simplify the set of polar balls by growing each polar ball by a multiplicative factor  $s$ , computing the medial surface of the grown spheres, and scaling the balls back by  $1/s$ .

When a set of points sampled on the surface and normal vectors to each sampled point are given, Ma *et al.* [81] find approximate medial points that lie on a face of the Voronoi diagram of the sample points by iteratively shrinking potential maximal spheres, with parallel computations carried out on the GPU. When the surface is given explicitly, Turkiyyah *et al.* [117] explain how the locations of the Voronoi vertices of points sampled on the surface can be refined to provide a better approximation to the medial surface, using the explicit knowledge of the boundary.

The strengths of Voronoi methods are the substantial body of work on their theoretical properties and the existence of robust software that facilitates computation. However, the density of Voronoi vertices varies depending on the density of the sampled boundary points and the curvature of the boundary, producing dense clusters of Voronoi vertices in certain regions and sparse collections of Voronoi vertices in others. The method we will develop for medial surface approximation in this chapter will generate a set of medial points that are distributed relatively uniformly on the medial surface. The advantage of this distribution for applications in computer graphics will be discussed in Chapters 7 and 8.

### 2.2.2 Tracing Methods

Computing the medial surface by tracing involves identifying points on junction curves of medial sheets and following these to discover the complete set of junctions of medial sheets. When the medial surface of a polyhedron must be computed exactly, Culver *et al.* [36] use exact arithmetic to accurately compute the junction curves of the medial surface by tracing. Because of the high computational cost of using exact arithmetic, this method is only feasible for polyhedra having a small number of faces. Leymarie and Kimia [78] consider as input a cloud of points and trace medial sheet junctions from sources detected using exact bisector computations between clusters of surface points. Interacting clusters are found quickly by taking advantage of spatial relationships between clusters. Seams are traced in the direction of increasing medial sphere radius. Simplification is achieved by removing junctions close to the boundary.

Tracing methods provide a rich description of the medial surface by explicitly locating sheet junctions in an efficient manner. However, these descriptions are sensitive to noise or require sophisticated regularization steps [28] to be useable for the application of shape matching.

### 2.2.3 Spatial Subdivision Methods

These types of algorithms partition space into disjoint cells and study the nearest boundary elements to these cells to locate the medial surface. Many compute the generalized Voronoi diagram of a solid, rather than the medial surface of a solid. For a polyhedron, the generalized Voronoi diagram is a superset of the medial surface that additionally contains bisectors of faces and their incident reflex edges and vertices. When the solid is a

set of disjoint convex sites, the exterior medial surface and the generalized Voronoi diagram are equivalent. Unlike Voronoi methods, spatial subdivision methods rely on a fine partition of space to localize the medial surface, and not on a dense boundary sampling.

Etzion and Rappaport [51] perform octree spatial subdivision until each octree cell has at most 4 nearest vertices, edges and faces of a polyhedron, *i.e.*, it contains at most one vertex of the generalized Voronoi diagram of the polyhedron. This proximity structure is then used to establish the adjacency of the cells containing the Voronoi vertices in the Voronoi diagram. Boada *et al.* [19] approximate the generalized Voronoi diagram using recursive spatial subdivision, where cells whose corners have the same nearest boundary element are not subdivided. The medial surface is approximated using those cell facets whose edges have different nearest boundary elements to their endpoints. Vleugels and Overmars [119] use octree spatial subdivision to approximate the medial surface of a set of disjoint convex polytopes in any dimension with octree cells by considering the nearest boundary elements to the octree cell vertices. Connectivity and convergence of the approximation is guaranteed for sufficiently small cells. Teichmann and Teller [116] compute the generalized Voronoi diagram of a set of triangles in 3D by dividing space into tetrahedral cells and determining the exact set of Voronoi regions intersecting a given cell. Cells of interest are subdivided and a polygonal approximation to the Voronoi diagram is output, which converges to the true Voronoi diagram for sufficiently small cells.

Foskey *et al.* [53] approximate the medial surface of a polyhedron by considering the directions to nearest boundary points for points on either a uniform or an adaptive grid. If a pair of points have divergent directions to nearest boundary points, the medial surface of a polyhedron intersects the line segment connecting these points. The object



angle is estimated using the angle between these vectors. The medial surface simplified by object angle is approximated by axis-aligned facets. When the sampled points are grid points, the computation of directions to nearest boundary elements can be accelerated by graphics hardware. Because simplification by object angle can disconnect the medial surface, Sud *et al.* [110] describe how homotopy equivalence between the polyhedron and its simplified medial surface approximation can be enforced by iteratively removing eligible medial sheets. The order of sheet removal for eligible sheets is given by an upper bound on the object angle of a sheet.

Of the spatial subdivision algorithms cited above, those algorithms that offer guarantees of convergence to the true medial surface require potentially a very fine spatial subdivision in order to accurately establish the location of medial points. To make guarantees at a coarse resolution, they must compute the complete set of boundary primitives that are closest to a cell. This computation may not be practical for large inputs. In the present work, we also design a spatial subdivision method for medial surface approximation, but we will consider only a fixed number of nearest neighbour queries to a given cell of an arbitrary size in order to determine if it contains a medial point. While this approach will not compute the complete medial surface, we argue that the quality of the approximation as a function of the object angle and other parameters can be bounded in the 2D case in Chapter 4, and provide tools for measuring the quality of the approximation in 3D in Chapter 3.

## 2.3 Voxel-Based Approximation

In this section, we will consider a regular partition of space into cubes, or voxels. We now overview a method for the detection of regions of space that are intersected by

the medial surface of a polyhedron described in [108] (the author’s Master’s thesis) to approximate the medial surface of a polyhedron with voxels. Based on values of the gradient of the Euclidean distance transform of a solid in a voxel, we will describe how to decide if this voxel is intersected by the medial surface of the solid. The present work, as well as the work in [108], builds on the work of Siddiqi *et al.* described in [105, 44] and [106](Ch. 4).

### 2.3.1 Foundations

Consider an 3D solid  $\Omega$  with boundary  $\mathcal{B}$ .

**Definition 2.2.** *The Euclidean distance transform of a solid  $\Omega$  with boundary  $\mathcal{B}$  is given by  $D(p) = -\inf_{q \in \mathcal{B}} d(p, q)$ , where  $p \in \Omega$  and  $d(p, q)$  denotes Euclidean distance.*

The gradient of  $D$ ,  $\nabla D : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  has the following properties. For a point  $p$  that is not on the medial surface or the boundary  $\mathcal{B}$ ,  $\nabla D$  is a unit vector field that assigns each point  $p$  the direction to its nearest point  $P$  on the boundary  $\mathcal{B}$  [52, 4.8(3)]. When  $p$  is inside  $\Omega$ ,  $\nabla D(p)$  is the unit outward normal to  $P$ , while when  $p$  is outside  $\Omega$ ,  $\nabla D(p)$  is the unit inward normal to  $P$ . The vector field  $\nabla D$  is uniquely defined everywhere except on the medial surface of  $\Omega$  and the boundary  $\mathcal{B}$  of  $\Omega$ . For points on the medial surface,  $\nabla D$  is multi-valued: there are two or more nearest boundary locations. This property is the basis for our method that locates medial points: we will look for regions where  $\nabla D$  is multi-valued.

Consider a sphere  $S$  interior to  $\mathcal{B}$  with outward normal vector  $N_S$ .<sup>1</sup> Then, as discussed in [38, 105, 45], in the continuum, the value of the average outward flux of  $\nabla D$  through a shrinking region  $S$ ,  $\mathcal{AOF}_S(\nabla D)$ , can be used to decide the presence and nature of the medial surface in  $S$  as the surface area of  $S$  tends to zero. Precisely,

$$\mathcal{AOF}_S(\nabla D) = \frac{\iint_S \nabla D \cdot N_S dS}{\iint_S dS}. \quad (2.1)$$

If the medial surface does not intersect  $S$ ,

$$\lim_{\text{area}(S) \rightarrow 0} \mathcal{AOF}_S(\nabla D) = 0. \quad (2.2)$$

However, if  $S$  contains a medial point with object angle  $\theta$ ,

$$\lim_{\text{area}(S) \rightarrow 0} \mathcal{AOF}_S(\nabla D) = \frac{1}{2} \sin(\theta). \quad (2.3)$$

It is important to note that this relationship holds only as the area of  $S$  shrinks to zero.

We now describe how to use a similar measure to decide the presence of medial points in a sphere  $S$  of positive surface area when  $\Omega$  is a polyhedron, as shown in [108]. In this special case, the  $\nabla D$  vector field over the surface of a sphere  $S$  can be partitioned into a small number of classes. This is possible because the nearest locations on the boundary  $\mathcal{B}$  of  $\Omega$  to points  $p$  inside  $\Omega$  lie either (1) in the interior of a face, (2) in the interior of an edge, or (3) on a vertex of  $\Omega$ .

---

<sup>1</sup> The choice of spherical region, and not a region of a different shape, simplifies the analysis.




Figure 2–5: Arrows show  $\nabla D$ , the directions to nearest locations on the boundary to points on  $R$ . In this example, the medial axis intersects the line segment  $(b, opp(b))$  because  $\nabla D(b) \neq \nabla D(opp(b))$ .

The following notion will be important for our subsequent analysis of  $\nabla D$  through  $S$ .

**Definition 2.3.** Consider a point  $p$  on the surface of a sphere  $S$ . Let  $l$  be a line through a point  $p$  with direction  $\nabla D(p)$ . Then the opposite of  $p$ ,  $opp(p)$  is the other point of intersection of  $l$  with  $S$ . In case  $\nabla D(p)$  is tangent to  $S$  at  $p$ ,  $opp(p) = p$ .

In fact, for any convex region, the notion of an opposite is uniquely defined. Figure 2–5 shows examples of points and their opposites. We now define a special class of  $\nabla D$  vectors for a polyhedral boundary  $\mathcal{B}$ , which will prove useful shortly.

**Definition 2.4.** A two-sided fan for a polyhedron with boundary  $\mathcal{B}$  is a set of locations  $p$  such that  $p$  lies on an edge or a vertex of  $\mathcal{B}$  and  $\nabla D(opp(p)) = \nabla D(p)$ .

The following theorem, proved in [108], uses the notion of two-sided fans to show how an analysis of the average outward flux through a region of positive surface area can be useful for detecting the medial surface of a polyhedron.

**Theorem 2.1.** *Consider a sphere  $S$  inside a 3D polyhedral solid  $\Omega$  with boundary  $\mathcal{B}$ . Let  $S^* \subset S$  be those locations on  $S$  that are not two-sided fans for  $\mathcal{B}$ . The medial surface of  $\Omega$  passes inside  $S$  if and only if  $\mathcal{AOF}_{S^*}(\nabla D)$  is positive.*

The magnitude of  $\mathcal{AOF}_{S^*}(\nabla D)$  remains proportional to the object angle of points on the medial surface that passes through  $S$ . In addition, it depends on the area of the intersection of the medial surface of the polyhedron with  $S$  since the area of intersection of a medial sheet with a constant object angle is reflected in the magnitude of  $\mathcal{AOF}_S(\nabla D)$  [38].

Because the integral in the numerator of Equation 2.1 is hard to compute exactly, it is desirable to approximate it. Using a simple quadrature rule, a discrete version of  $\mathcal{AOF}_S(\nabla D)$  then is

$$AOF_S(\nabla D) = \frac{\sum_{i=1}^N \nabla D(p_i) \cdot N_S(p_i)}{N}, \quad p_i \in S. \quad (2.4)$$

Theorem 2.1 is only true for an infinite sampling rate  $N$ . This presents problems for the correctness and completeness of algorithms using the  $AOF_{S^*}(\nabla D)$  measure to decide the presence of medial points using a finite sampling rate  $N$ . For example, Figure 4–3 shows situations when  $AOF_{S^*}(\nabla D)$  is zero, yet the medial surface intersects  $S$ . Chapters 3 and 4 study correctness issues of algorithms for the detection of medial points based on the analysis of nearest boundary locations to a finite number of samples in 3D and 2D, respectively.

### 2.3.2 Approximating the Medial Surface with Voxels

Here we use the foundations presented above to describe the algorithm in [108] that we use to efficiently approximate the medial surface of  $\Omega$  with voxels of a user-chosen

size. The algorithm we will describe repeatedly computes the  $AOF_{S^*}(\nabla D)$  measure for overlapping spheres interior to  $\Omega$ . The algorithm takes the following 3 steps:

1. We overlay  $\Omega$  with a coarse voxel grid where each voxel has side length  $l$  and record the position of those voxels such that a sphere  $S$  of radius  $l$  centred at the voxel center lies inside  $\Omega$ . We then repeat the following steps for each such voxel:
2. We distribute  $N$  points approximately uniformly on the sphere  $S$  centred at the voxel centre and with radius equal to the voxel side length. We then compute  $AOF_{S^*}(\nabla D)$ , where  $S^*$  is the subset of  $S$  that does not contain any two-sided fans for  $\nabla D$ .
3. When  $AOF_{S^*}(\nabla D)$  is positive and we have reached the desired voxel resolution, we output this voxel. Otherwise, we subdivide the voxel considered into 8 equal size voxels and repeat steps 2-3 for the smaller voxels.

To distribute  $N$  points approximately uniformly on the boundary of a sphere, we will use the popular ‘spiral distribution’ described in [102]. A larger sampling rate  $N$  is used for top-level voxels than lower level voxels.

The result is a set of fine resolution voxels that are intersected by the medial surface of  $\Omega$  as  $N \rightarrow \infty$ . Let us call the process of subdividing a voxel outlined in steps 1-3 ‘zooming’. This coarse-to-fine procedure is efficient because it allows early dismissal of large regions that do not contain medial points. In order to consider as much of the interior of  $\Omega$  as possible, we repeat this procedure for any smaller resolution voxels near the boundary that were not interior to any coarse-level voxels. This new step finds additional voxels near the object boundary that are intersected by the medial surface.

To ensure the efficiency of this algorithm, the computation of exact  $\nabla D$  values, *i.e.*, finding the nearest point on a mesh boundary to a query point, needs to be done quickly. To this aim, we will represent the polyhedral boundary as a rectangle swept sphere hierarchy using the proximity query package PQP of [76]. This hierarchy enables one to quickly find the nearest mesh triangle to a query point.

For the purposes of the current work, the following improvements have been made to the algorithm in [108], to further improve efficiency. The PQP package was designed to compute the exact shortest distance between pairs of triangle meshes. This package has been modified to work more efficiently in the case when one of the inputs is a point by simplifying various geometric and logical operations. As a heuristic, when ‘zooming’, we will save the nearest boundary triangles to points on the sphere bounding the top-level coarse voxel. When computing  $\nabla D$  values for points inside this voxel, we will then reduce the search for the nearest mesh triangle to these saved triangles. When these saved triangles form a superset of nearest mesh triangles to points inside this voxel, the  $\nabla D$  values are computed exactly.

When describing our approximation, we use  $\sigma$  to denote the side length of the finest resolution voxel considered in this procedure. This value measures the density of the approximation to the medial surface.

## 2.4 Point-Based Approximation

In the previous section, we showed how voxels that are potentially intersected by the medial surface can be identified. In this section, we will explain how to refine the estimates for medial point locations to obtain a dense collection of points within a user-specified constant of the salient medial sheets. We will discuss properties of this approximation.

### 2.4.1 Precise Medial Points

Given a set of voxels that are believed to be intersected by the medial surface, we would like to obtain more precise estimates of the location of the medial surface in these voxels. To help us design an algorithm that generates these precise estimates of locations of medial points, we introduce and prove several properties of the Euclidean distance transform  $D$ .

Let  $(a, b)$  denote the closed line segment with endpoints  $a$  and  $b$ . Let  $d(a, b)$  denote the length of  $(a, b)$ .

**Lemma 2.1.** *Let  $a$  be a point with a unique nearest point  $A$  on the boundary  $\mathcal{B}$ . Then  $A$  is also the nearest boundary point to any point  $b$ ,  $b \in (a, A)$ .*

*Proof.* Since  $A$  is the nearest boundary point to  $a$ ,  $A = a + \gamma \nabla D(a)$ , for some positive constant  $\gamma$ . Suppose, for a contradiction, that there exists a boundary point  $B$  on  $\mathcal{B}$  such that  $B \neq A$  and

$$d(b, B) < d(b, A), \quad (2.5)$$

for some point  $b \in (a, A)$ . As  $A$  is the closest point to  $a$ ,

$$d(a, B) > d(a, A). \quad (2.6)$$

By the triangle inequality,

$$\begin{aligned} d(a, B) &\leq d(a, b) + d(b, B) \\ &\leq d(a, b) + d(b, A) \\ &= d(a, A), \end{aligned}$$



where the second inequality follows from (2.5) and the last because  $b \in (a, A)$ . Therefore,  $d(a, B) \leq d(a, A)$ , contradicting (2.6).  $\square$

**Corollary 2.1.** *Let  $a$  be a point with a unique nearest boundary point  $A$ . Let  $b = a + \gamma \nabla D(a)$  for a scalar  $\gamma$ . Let  $B$  be the unique nearest boundary point to  $b$ . If  $A = B$ , then no point of the medial surface lies on  $(a, b)$ .*

*Proof.* Suppose that  $(a, b)$  does not intersect  $\mathcal{B}$ . Consider that  $a$  is farther from the boundary  $\mathcal{B}$  than  $b$ . By Lemma 2.1, the unique boundary point  $A$  to  $a$  is also the unique boundary point for any point on  $(a, b)$ . As the nearest boundary point to any point on  $(a, b)$  is unique, no medial points lie on  $(a, b)$ . If  $b$  is farther from  $\mathcal{B}$  than  $a$ , a similar argument holds.

Now suppose that  $(a, b)$  intersects  $\mathcal{B}$ . Since  $A = B$ ,  $(a, b)$  intersects  $\mathcal{B}$  at a single point  $C$ . Then  $C = A = B$ . But by Lemma 2.1, the nearest boundary point to points on  $(a, C)$  and  $(C, b)$  is  $C$ . Since the nearest boundary point to all points on  $(a, b)$  is unique, no medial point lies on  $(a, b)$ .  $\square$

In particular, Corollary 2.1 holds for  $b = \text{opp}(a)$  (see Figure 2–5). Given the boundary  $\mathcal{B}$  of a solid  $\Omega$ , let us denote by  $\mathcal{B}(p)$  the nearest point to  $p$  on  $\mathcal{B}$ , whenever this point is uniquely defined (*i.e.*, whenever  $p$  is not a medial point).

**Lemma 2.2.** *Consider two points  $a$  and  $b = a + \gamma \nabla D(a)$  such that  $\gamma$  is a non-zero scalar. If no points of the interior or exterior medial surface of  $\Omega$  lie on  $(a, b)$ , then  $\mathcal{B}(a) = \mathcal{B}(b)$ .*

*Proof.* Consider the case that  $(a, b)$  does not intersect  $\mathcal{B}$  at a point  $c \notin \{a, b\}$ . Suppose that  $b$  is closer to  $\mathcal{B}$  than  $a$ . Then by Lemma 2.1, the unique closest boundary point to  $a$  is also the unique closest boundary point to  $b$ . Then  $\mathcal{B}(a) = \mathcal{B}(b)$ . Now suppose that  $a$  is closer than  $b$  to  $\mathcal{B}$  and no medial points lie on  $(a, b)$ . We now use some tools from

the theory of geometric curve evolution. The medial surface is the set of shocks of the grassfire flow applied to  $\mathbb{R}^3 \setminus \mathcal{B}$ , as explained in [18]. The grassfire flow evolves a moving front  $W(t)$  according to the equation  $\frac{\partial W}{\partial t} = -n$ , where  $n$  is the unit outward normal to  $W$  with the initial condition  $W(0) = \mathcal{B}'$ , where  $\mathcal{B}'$  is the boundary of  $\Omega$  if  $(a, b) \subset \Omega$  and is the boundary of the closure of  $\mathbb{R}^3 \setminus \Omega$  otherwise. Suppose  $W(t_0)$  passes through  $a$ . Since  $(a, b)$  contains no medial points, point  $a$  on  $W(t_0)$  after time  $\gamma$  becomes point  $b = a + \gamma \nabla D(a)$  on  $W(t_0 + \gamma)$ . The outward normal to  $W(t_0 + \gamma)$  at  $b$  is  $\nabla D(a)$ , a known property of the grassfire flow [72]. Thus,  $\nabla D(b) = \nabla D(a)$ . Since  $a \in (b, \mathcal{B}(a))$  and  $(b, \mathcal{B}(a))$  does not intersect  $\mathcal{B}$  by our assumption,  $\mathcal{B}(b) = \mathcal{B}(a)$ .

Now consider the case that  $(a, b)$  does intersect  $\mathcal{B}$  at point  $c \notin \{a, b\}$ . Let  $C$  be the intersection point of  $\mathcal{B}$  with  $(a, b)$  that is closest to  $a$ . By applying the earlier argument to  $(a, c)$ , we can show that  $C = \mathcal{B}(c) = \mathcal{B}(a)$ . Now consider the segment  $(c, b)$ . By repeating the argument in this proof, we can show that  $\mathcal{B}(c) = \mathcal{B}(b)$ .  $\square$

Let  $\overline{\Omega}$  denote the closure of the complement of  $\Omega$ . Corollary 2.1 together with Lemma 2.2 imply the following:

**Lemma 2.3.** *Let  $p$  be a point in  $\mathbb{R}^3$  that is not a medial point of  $\Omega$  or  $\overline{\Omega}$ . Let  $q = p + \gamma \nabla D(p)$ , such that  $\gamma$  is a scalar and  $q$  is not a medial point. A medial point of  $\Omega$  or  $\overline{\Omega}$  lies on  $(p, q)$  if and only if  $\mathcal{B}(p) \neq \mathcal{B}(q)$ .*

Using this lemma, we design the subroutine RETRACT, shown as Algorithm 1, that performs binary search to estimate the intersection of the medial surface of  $\Omega$  or  $\overline{\Omega}$  with a line segment to a desired accuracy, where the line segment has the form  $q = p + \gamma \nabla D(p)$ ,

$\mathcal{B}(p) \neq \mathcal{B}(q)$ . By Lemma 2.3, Algorithm 1 finds a point on  $(p, q)$  that is within a user-specified tolerance  $\epsilon$  of the medial surface of  $\Omega$  or  $\bar{\Omega}$ . Similar algorithms have been proposed in the literature [44, 90, 53, 122] for localizing medial points, but none have offered a similarly rigorous correctness analysis.

---

**Algorithm 1** RETRACT( $p, q, \mathcal{B}, \epsilon$ )

---

**Input:** Non-medial points  $p, q$  such that  $q = p + \gamma \nabla D(p)$  and  $\mathcal{B}(p) \neq \mathcal{B}(q)$ , tolerance  $\epsilon$ .

**Output:** A point within  $\epsilon$  of the medial surface of  $\Omega$  or  $\bar{\Omega}$  and its two spoke vectors.

```

1: while  $d(p, q) > \epsilon$  do
2:    $m = \frac{1}{2}(p + q)$ 
3:   if  $m$  is equidistant from two points on  $\mathcal{B}$  then
4:     Return  $m$  and the directions to two nearest points on  $\mathcal{B}$  to  $m$ .
5:   end if
6:   if  $\mathcal{B}(m) \neq \mathcal{B}(p)$  then
7:      $q = m$ 
8:   else
9:      $p = m$ 
10:  end if
11: end while
12: Return  $p, \nabla D(p), \nabla D(q)$ .
```

---

The point  $p$  returned by Algorithm 1 is an approximate location of a medial point of  $\Omega$  or  $\bar{\Omega}$ , which we will also refer to as an *approximate medial point*. The radius of the medial sphere at  $p$  is the distance from  $p$  to its the nearest point on  $\mathcal{B}$ . By definition, this sphere is tangent to  $\mathcal{B}$ . The spoke vector  $\nabla D(p)$  is one of the true spoke vectors of a medial point within distance  $\epsilon$  from  $p$  along  $\nabla D(p)$ , while the spoke vector  $\nabla D(q)$  is an estimate for the second spoke vector at  $p$ .

Let  $(p, \nabla D(p), \nabla D(q))$  be the output of RETRACT( $p, q, \mathcal{B}, \epsilon$ ). Suppose we assume that  $p$  is a medial point having object angle  $\arccos\left(\frac{\nabla D(p) \cdot \nabla D(q)}{2}\right)$ , that the radius of the medial sphere at  $p$  is  $d(p, \mathcal{B}(p))$ , and that the two nearest boundary points to this medial

point are  $P = \mathcal{B}(p)$  and  $Q' = p + d(p, \mathcal{B}(p))\nabla D(q)$ . If  $q = p + \gamma\nabla D(p)$  for some constant  $\gamma \in [0, \epsilon]$  and if  $\mathcal{B}(q) = Q$ , how far away is  $Q'$  from  $Q$ ? The following lemma answers this question.

**Lemma 2.4.** *If  $q = p + \gamma \nabla D(p)$  for some constant  $\gamma \in [0, \epsilon]$ , then the estimated second point of tangency of the approximate medial sphere at  $p$  with radius  $d(p, \mathcal{B}(p))$  is at most  $2\epsilon$  from  $\mathcal{B}(q)$ .*




Figure 2–6: The objects of interest in the proof of Lemma 2.4.

*Proof.* Consider the plane spanned by  $\nabla D(p)$  and  $\nabla D(q)$ . Let  $Q = \mathcal{B}(q)$  and  $P = \mathcal{B}(p)$ . Let  $d(p, P) = r_p$  and  $d(q, Q) = r_q$ . Let  $S_p$  be the sphere at  $p$  with radius  $r_p$ . Let point  $Q'$  be found by intersecting  $S_p$  with a ray through  $p$  with direction  $\nabla D(q)$ .  $Q'$  is the estimate for the second point of tangency of  $S_p$  with  $\mathcal{B}$ . Figure 2–6 presents several examples of configurations of the points mentioned. Let  $r$  be  $q + r_p \nabla D(q)$ . Consider the

quadrilateral  $Q'pqr$ . By construction, it is a parallelogram. Therefore,  $d(Q', r) \leq \epsilon$ . Also by construction,  $d(r, Q) = |d(q, Q) - d(q, r)|$ .

Since  $d(q, r) = d(p, P)$ ,  $d(r, Q) = |d(q, Q) - d(p, P)|$ . However, since  $Q$  is no farther from  $q$  than  $P$ ,  $d(q, Q) \leq d(q, P)$ . On the other hand,  $d(q, P) \leq d(p, P) + \epsilon$ . Therefore,

$$d(q, Q) - d(p, P) \leq \epsilon. \quad (2.7)$$

Also, since  $P$  is no farther from  $p$  than  $Q$ ,  $d(p, P) \leq d(p, Q) \leq d(p, q) + d(q, Q) \leq \epsilon + d(q, Q)$ . Therefore,

$$d(p, P) - d(q, Q) \leq \epsilon. \quad (2.8)$$

By Equations 2.7 and 2.8,  $d(r, Q) = |d(q, Q) - d(p, P)| \leq \epsilon$ . Then, by the triangle inequality,  $d(Q, Q') \leq d(Q', r) + d(r, Q) \leq 2\epsilon$ .  $\square$

Lemma 2.4 implies that not only is the approximate medial point output by Algorithm 1 within  $\epsilon$  of the true medial surface, but also there is boundary point within  $2\epsilon$  of the estimated boundary point of tangency. When  $\epsilon$  is small, Algorithm 1 provides a reasonable estimate of both the location and the spoke vectors of a medial point on  $(p, q)$ . If necessary, the exact location of a medial point on  $(p, q)$  and the exact spoke vectors can be found by finding the other element of  $\mathcal{B}$  (polyhedron face, edge or vertex) closest to  $p$  and computing the intersection of the bisector of that element and point  $P$  and the line segment  $(p, q)$ .

When a line segment  $(a, b)$  does not cross  $\mathcal{B}$ ,  $\nabla D(a) = \nabla D(b) \Leftrightarrow \mathcal{B}(a) = \mathcal{B}(b)$ . However, when  $(a, b)$  crosses  $\mathcal{B}$ ,  $\nabla D(a) = \nabla D(b) \nRightarrow \mathcal{B}(a) = \mathcal{B}(b)$  and a medial point can lie on  $(a, b)$ . Figure 2–7 presents an example. If it is not important to consider those

$(a, b)$  that cross  $\mathcal{B}$ , and further, if  $\nabla D(a)$  is easier to compute than  $\mathcal{B}(a)$ ,  $\nabla D(a)$  may be used instead of  $\mathcal{B}(a)$  in Algorithms 1 and 2.




Figure 2–7: The boundary of a 2D solid is shown in black and its medial axis in red. In this case  $(a, b = a + \gamma \nabla D(a))$  intersects the boundary and although  $\nabla D(a) = \nabla D(b)$ , a medial point lies on  $(a, b)$ .

#### 2.4.2 Well-Spaced Salient Medial Points

We wish to find a set of points that lies near some salient subset of the medial surface. We want the salient subset to have the property that the solid may be accurately reconstructed from the computed medial spheres. As discussed in Sections 2.1 and 2.2, a popular measure for the simplification of the medial surface are the object angle and radius of the medial points. However, simplification by object angle alone can disconnect the medial surface when low object angle medial points lie deep inside  $\Omega$ . In an effort to maintain the connectivity of the medial surface and to achieve adequate simplification, we will retain approximate medial points that are estimated as having a large object angle *or* a large radius.

For each voxel of interest, Algorithm 2 looks for salient medial points inside this voxel and outputs one approximate medial point inside this voxel if it is deemed salient in terms of the object angle and radius parameters.

We described in Section 2.3 how voxels that are possibly intersected by the medial surface of a polyhedron can be identified based on the measure of  $AOF_{S^*}(\nabla D)$  for a sphere  $S$  circumscribing a voxel. When  $\Omega$  is a polyhedron, we will use the  $AOF_{S^*}(\nabla D)$  measure to pre-select those voxels likely to be intersected by medial sheets containing points of high object angle and radius.

Recall that the value of  $AOF_{S^*}(\nabla D)$  is proportional to the area of intersection of the medial surface  $\mathcal{MS}$  with  $S$  as well as the object angle of the portion of  $\mathcal{MS}$  in  $S$ . Consider a voxel  $v$  with side length  $l$  circumscribed by a sphere  $S$  with radius  $l$ . We will use the  $AOF_{S^*}(\nabla D)$  value of this voxel to provide an estimate of the object angle of medial points in this voxel, assuming the medial sheet has a certain minimal area of intersection with  $S$ . We will ignore voxel  $v$  when  $AOF_{S^*}(\nabla D)$  is below that of a planar medial sheet with object angle  $\theta$  intersecting  $S$  in a circle with radius  $\sqrt{3}l/2$ , as happens when a medial sheet passes on the border of 2 voxels and falls directly on a face of the voxel  $v$ . As shown in [108], the value of  $AOF_{S^*}(\nabla D)$  in this case is  $-2 \sin(\theta) \frac{\pi(\sqrt{3}l/2)^2}{4\pi l^2}$ .

To avoid missing medial points with radii above the radius threshold  $r$ , we also consider voxel  $v$  with side length  $l$  for which the maximum distance to the boundary from points sampled on a sphere of radius  $l$  circumscribing  $l$  exceeds  $r - 2l$ .

For a voxel  $v$  of interest, given a set  $\Phi$  of points sampled on a sphere circumscribing  $v$ , Algorithm 2 uses subroutine RETRACT to find locations inside  $v$  near the medial surface of a solid  $\Omega$ . Of these, Algorithm 2 returns a single approximate medial point if the estimates of the radius and object angle of this point satisfy the user's requirements of saliency.

---

**Algorithm 2** SALIENTMEDIALPOINT( $\Phi, v, \mathcal{B}, \theta, r, \epsilon$ )

---

**Input:** A set of non-medial points  $\Phi$  sampled on a sphere  $S$ , voxel  $v$  inside  $S$ , boundary  $\mathcal{B}$  of solid  $\Omega$ , thresholds  $\theta, r, \epsilon$ .

**Output:** An approximate location of a medial point  $m$  of  $\Omega$  inside  $v$  within  $\epsilon$  of the medial surface of  $\Omega$ , deemed as having object angle greater than  $\theta$  or radius greater than  $r$ , along with its spoke vector estimates.

```
1: for all  $\phi_i \in \Phi$  do
2:   if  $\mathcal{B}(\phi_i) \neq \mathcal{B}(\text{opp}(\phi_i))$  then
3:     Let  $(m, \nabla D(m), \nabla D(q)) = \text{RETRACT}(\phi_i, \text{opp}(\phi_i), \mathcal{B}, \epsilon)$ .
4:     if  $m$  is inside  $v$  and inside  $\Omega$  then
5:       if  $\arccos(A \cdot B)/2 > \theta$  then
6:         Return  $(m, \nabla D(m), \nabla D(q))$ .
7:       end if
8:       if  $d(m, \mathcal{B}(m)) > r$  then
9:         Return  $(m, \nabla D(m), \nabla D(q))$ .
10:      end if
11:    end if
12:  end if
13: end for
```

---



Using Algorithm 2, at most one approximate location of a medial point is found per voxel. Thus, the density of the point-based approximation to the medial surface is controlled by the resolution of the voxels considered. Given a boundary representation of a solid, the spheres generated by Algorithm 2 make up the *well-spaced medial spheres* shape representation of that solid.

Consider the solid  $\Omega'$  that is the union of balls associated with the spheres we compute. Are the spheres maximal with respect to  $\Omega'$ ? A ball  $b$  with centre  $m$  of  $\Omega'$  contains on its boundary a point  $M$ , which is the closest point on the boundary of  $\Omega$  to  $m$ . Point  $M$  may also lie on the boundary of those balls of  $\Omega'$  whose centres are closer than  $\epsilon$  to  $m$  along  $\nabla D(m)$ . To simplify the set of spheres produced with our method, approximate medial points within  $\epsilon$  of a given approximate medial point can be removed. When  $\epsilon \ll \sigma$ , this simplification step is local and removes at most 7 points from neighbouring voxels for each point retained. The resulting set of balls is maximal in  $\Omega'$ , and inscribed in  $\Omega'$  by definition. Hence the set of spheres we produce are medial with respect to the solid  $\Omega'$ , the union of the associated balls, after simplification.

### 2.4.3 Handling Narrow Features

When using the  $AOF_{S^*}(\nabla D)$  measure to pre-select voxels intersected by the medial surface, the sphere of radius  $\sigma$  circumscribing a voxel  $v$  with side length  $\sigma$  has to be inside  $\Omega$  for the flux analysis to be valid. We will explain how the minimum voxel side length  $\sigma$  considered determines the smallest radius of a medial sphere that can be found, when interior spheres  $S$  are used.

Consider a medial sphere with radius  $r$  whose centre is inside voxel  $v$ . If the nearest boundary point to the centre of this sphere is inside the circumscribing sphere for the voxel  $v$ , the medial sphere fails to be detected. We make the following observations:

- Medial spheres with radius  $r < \sigma - \frac{\sqrt{3}}{2}\sigma$  are never found by our algorithm.
- Medial spheres with radius  $r < \sigma + \frac{\sqrt{3}}{2}\sigma$  are not guaranteed to be found by our algorithm.

Therefore, small spheres with respect to the voxel resolution may be undetected by our algorithm. Although inability to detect a certain kind of spheres can be viewed as a shortcoming, this feature actually accounts for stability of our algorithm: removal of small radius medial spheres near the boundary is a common simplification step that preserves medial surface connectivity and removes insignificant medial points for shape reconstruction.

On the other hand, solids with narrow portions or small-scale features require a very large voxel resolution  $\sigma$  in order for our algorithm to detect the small radius medial spheres in such regions. Using a very large value for  $\sigma$  produces a large number of approximate medial points, which may not be desirable, and is computationally expensive. However, note that because Algorithm 1, RETRACT, detects medial points on a line segment  $(p, q)$  even when this line segment is intersected by the boundary  $\mathcal{B}$ , Algorithm 2 outputs approximate medial points even for voxels whose circumscribing spheres are intersected by  $\mathcal{B}$ . Figure 2–8 presents an example of how the approximation of a solid with many narrow parts and small-scale features is improved by additionally considering voxels intersected by the solid’s boundary.




Figure 2–8: A dragon polyhedron (left) approximated with 1324 spheres (centre) and 4102 spheres (right). Medial points were found in voxels intersected by the boundary  $\mathcal{B}$  on the right, but not on the left.

When it is desirable to approximate narrow regions and small-scale features with respect to the minimum voxel resolution  $\sigma$  considered, we run Algorithm 2 on all voxels whose circumscribing spheres are intersected by  $\mathcal{B}$ . Note that not all such voxels will contain salient medial points. Further, in order to output points near the medial surface of  $\Omega$  and not  $\overline{\Omega}$ , we must be able to determine if a point lies inside or outside  $\Omega$ . These operations add an extra expense to the process of detecting medial points. When approximating a solid with a small number of medial spheres, the extra expense of using Algorithm 2 on voxels whose circumscribing spheres are intersected by  $\mathcal{B}$  may be justifiable as it allows us to find medial points near the boundary.

#### 2.4.4 Degeneracy Issues

In Algorithm 2, we need to be able to find non-medial points on a sphere  $S$ . As the medial surface of  $\Omega$  is of lower dimensionality than  $\Omega$ , finding such points is typically not a problem. However, for very special inputs the medial surface and  $S$  coincide. Figure 2–9

(left) shows 20 spheres whose envelope approximates the boundary of a cardioid. The medial surface of this cardioid consists of all the points on the red sphere, except the point closest to  $A$ . It would not be possible to find more than one non-medial point to sample on  $S$  in this case.




Figure 2-9: (a) Sampled points on the red sphere, except  $A$ , are medial points of the envelope of the black circles shown. (b) The medial axis of a regular polygon with low object angle medial points shown in black and a high object angle medial point  $M$  shown in red.

For solids such as a cylinder, the medial surface is degenerate: it contains curves, not only surfaces. When a cylinder is approximated by a polyhedron, it may happen that, in cross section, the polyhedron is a regular polygon (see Figure 2-9 (right)). The medial surface of such a polyhedron will consist of low object angle branches corresponding to bisectors of adjacent boundary faces meeting at high object angle medial points. In order to detect the high object angle medial points using Algorithm 2, sample points  $\Phi$  on a sphere  $S$  must be chosen from a finite set of locations, such that the high object angle portion of the medial surface can be found by following the  $\nabla D$  directions from these locations. We say that a sampled point  $p$  *retracts* to a medial point if the medial point lies on a line segment  $(p, p + \gamma \nabla D(p))$ , for a scalar  $\gamma$ . In Figure 2-9 (right), there are at

most 12 locations on a given sphere that retract to  $M$ . It is unlikely that sample points  $\Phi$  will be chosen among this set of locations. In order to detect salient medial points using Algorithm 2 for degenerate solids, one should perturb the boundary of the solid slightly (for example, by translating each vertex by a small random value) so that its medial surface ceases to be degenerate and so that the set of points that retracts to salient medial points is enlarged.

## 2.5 Grouping Medial Points into Medial Sheets

As discussed in [58], medial points with two distinct points of tangency lie on smooth medial sheets. In this section, we will explain how, given a sufficiently dense set of approximate medial points and their spoke vector estimates, the set of medial points can be partitioned into those belonging to distinct medial sheets. This partition gives additional power to our union of medial spheres shape representation as it enables a part-based analysis. Given a cloud of points sampled near the medial surface along with their spoke vector estimates, we will show how to partition the set of approximate medial points into sets of points that belong to distinct smooth medial sheets.

For a smooth medial sheet, a small spatial region can contain medial points that all belong to the same sheet, yet the normals to the medial sheet at these points can be very different, as illustrated with the example in Figure 2–10, adopted from [39] (in this case, a medial sheet locally self-intersects). When designing a partitioning algorithm for medial points, we need to make sure that medial points that belong to such sheets are grouped into a single sheet.

Smooth medial sheets can be characterized as follows: for a pair of medial points  $a$  and  $b$  on a smooth medial sheet, there is a path along the sheet connecting  $a$  and  $b$ , along




Figure 2–10: An example of an object with boundary  $B$  whose medial surface is a single medial sheet  $MS$  where neighbouring smooth medial points have different surface normals (two surface normals are shown). Adopted from [39].

which the normals to the medial sheet vary smoothly. We use this criterion to design a method for partitioning the set of medial points into distinct sheets.

To start, we need a notion of a surface normal to a medial point. It turns out that this information is encoded by the spoke vectors of a medial point.

Consider a medial point  $x_0$  and its two spoke vectors  $U^a(x_0)$  and  $U^b(x_0)$ . The spoke vectors make an equal angle with the tangent plane to the medial surface  $MS$  at  $x_0$ ,  $T_{x_0}MS$ . Let the normal to  $T_{x_0}MS$  be  $N(x_0)$ . Then for  $x \in T_{x_0}MS$ ,

$$U^a(x_0) \cdot (x_0 - x) = U^b(x_0) \cdot (x_0 - x) \quad (2.9)$$

$$\Rightarrow (U^a(x_0) - U^b(x_0)) \cdot (x_0 - x) = 0. \quad (2.10)$$

Therefore,

$$N(x_0) = U^a(x_0) - U^b(x_0) \quad (2.11)$$

is a normal direction to  $T_{x_0}MS$ . Based on estimates of spoke vectors at our approximate medial points, we can compute normal estimates to the medial surface at these points using Equation 2.11.

We will group neighbouring approximate medial points that have locally consistent normals into sheets. As each approximate medial point  $q$  lies inside a distinct voxel  $v_q$ , those medial points found in neighbouring voxels of  $v_q$  are considered the neighbours of  $q$ . We will add to the sheet of point  $q$  all medial points  $p$  inside voxels in the 26-neighbourhood of  $v_q$  whose estimates of the normal to the medial surface are within an allowed tolerance  $\tau$  of the normal estimate at  $q$ . The result is a grouping of the medial points into medial points belonging to distinct smooth medial sheets.

## 2.6 Experimental Results and Discussion

In this section experimental results of applying the medial surface approximation and segmentation algorithms we have presented are shown.

Figure 2–11 presents the approximate medial points computed for four polyhedra of substantial geometric complexity. The object angle threshold used is 0.3 radians and the radius threshold is 0.25 times the maximum dimension of the bounding box of each polyhedron. The approximate medial points of the four polyhedra are shown segmented into medial sheets in Figure 2–11 (Right), where the allowed tolerance  $\tau$  between adjacent normal estimates for grouping is 5.7 degrees. Figure 2–11 (Centre) illustrates the union of the spheres corresponding to the approximate medial points.

In Section 1.2.1, we mentioned that the medial surface transform of  $\Omega$  has several important properties. Of these, let us consider completeness (Property 1.1) and part-structure (Property 1.2). As the envelopes of spheres computed with our method bear a close resemblance to the boundaries of polyhedra, our experimental results support the claim that the approximate medial points computed with our method are a good approximation to the medial surface in terms of providing completeness of shape description. Our ability




Figure 2–11: Left: The original polyhedra; Centre: the union of spheres computed with our method; Right: approximate medial points grouped into smooth sheets.



to organize medial points into distinct smooth medial sheets speaks of the success of our approximation to the medial surface at capturing part-structure.

Computation of smooth medial sheets that provably lie near salient portions of the medial surface is a challenge of all current methods for medial surface approximation. In this respect, our method offers an approximation to the medial surface of comparable quality as that produced with a leading Voronoi-based method of [85]. In contrast to alternative spatial subdivision methods which consider all the nearest boundary triangles to a sphere [51, 19, 119, 116], our method only considers a sampling of points on these triangles. As we will see in the following chapter, this approach may not detect all medial points. However, we will see that it is also less sensitive to noise as it may not detect spheres that are intersected by low object angle medial points, which are typically subject to pruning.

Model	# Triangles	Final Voxel Resolution	# Medial Points	Time (min)
Torus	16,000	3,033,136	64,195	5.93
Pear	86,016	5,177,526	18,147	22.24
Head	6,816	5,244,401	74,853	6.22
Venus	22,688	7,739,924	116,793	12.82

Table 2–1: Statistics for the computation of medial points in Figure 2–11.

Statistics for the computation of medial points in Figure 2–11 are given in Table 2–1. The final voxel resolution is the number of finest-resolution voxels in the interior of the model. The number of medial points is the number of salient medial points found, at most one per finest-resolution voxel. We used 4 levels of recursive voxel subdivision, and the smallest voxel considered had length  $\sigma$  equal to  $1/640^{th}$  of the maximum dimension of the

object bounding box. The sampling rate  $N$  for points on spheres circumscribing voxels was 400, 200, 100 and 50 for each of the 4 levels of subdivision, with smaller  $N$  for smaller voxels. Timings are shown for a single 3.4 GHz Pentium IV processor with 3 GB of RAM. Partitioning into medial sheets took a few seconds to complete. To accelerate computation, one should consider using a larger final voxel side length and coarser resolution polyhedra.

Figure 2–12 illustrates the effect of decreasing the smallest voxel side length  $\sigma$  on the density and number of medial points produced. We will see in Part III of this thesis that the ability to output a set of points that are arbitrarily close to the medial surface but whose density is user-chosen, is a special feature of our algorithm.




Figure 2–12: Medial points computed for a solid with varying voxel side length  $\sigma$ . Shown are the results for voxel lengths equal to  $1/10^{th}$ ,  $1/20^{th}$ ,  $1/40^{th}$  and  $1/100^{th}$  of the maximum dimension of the object bounding box, from left to right, respectively.

Figure 2–13 shows that using the object angle threshold alone can produce a set of medial points whose voxels are disconnected. Additionally computing medial points having large radiusf improves connectivity, providing support for our salience measure which incorporates both object angle and radius measures.

Given a method to evaluate  $\mathcal{B}(p)$  for any point  $p \in \mathbb{R}^3$ , Algorithm 2 can be used to approximate the medial surface for other inputs  $\mathcal{B}$  than boundaries of polyhedra. Moreover, the  $\Omega$  does not have to be connected. When  $\Omega$  is a collection of convex sites, their external




Figure 2–13: Left: A bumpy sphere model. Centre: Its medial surface approximation when  $\theta = 0.6$  and no radius threshold is used. Right: The medial surface approximation when  $\theta = 0.6$  and a radius threshold of 10% of the maximum dimension of the bounding box is used.

medial surface corresponds to the generalized Voronoi diagram of the sites. For non-convex sites, pruning of portions of the medial surface produces the generalized Voronoi diagram of the sites. Figure 2–14 shows our approximation of the Voronoi diagram of three lines.




Figure 2–14: Voronoi diagram of three lines cropped to a cube.

Our method for approximating the medial surface with a sampling of points and for grouping these points into distinct medial sheets uses a number of parameters. Table 2–2

lists these parameters, along with their functions. The values of these parameters have been selected empirically. It is important to establish theoretical results giving appropriate settings of these parameters, such that desirable results are guaranteed. Among desirable results are the guaranteed generation of a set of spheres whose union is homotopy equivalent to the original solid, and the guaranteed detection of salient medial points when some lie in a voxel. The next two chapters are concerned with establishing the relationship between the number  $N$  of sample points and the ability to detect salient medial points. Future work is discussed in both of these chapters.

Parameter	Definition
$N$	Number of points to generate on a sphere circumscribing a voxel
$\sigma$	Smallest voxel side length considered
$\epsilon$	Maximum distance from an approximate medial point to the medial surface
$\theta$	Object angle threshold for medial points
$r$	Radius threshold for medial points
$\tau$	Maximum difference in normals to the medial surface for points on the same medial sheet

Table 2–2: Parameters and their definitions for the medial surface approximation and partitioning algorithms.

## 2.7 Chapter Summary

In this chapter, we have presented algorithms for the approximation of the medial surface of a 3D solid whose boundary is given explicitly with a collection of points within a user-chosen threshold of medial points deemed salient, such that at most one such point is output per regular grid cell. We have discussed correctness issues of the algorithms presented, efficiency considerations, as well as issues with degenerate inputs. The

availability of grid neighbourhoods, as well as an easily computable surface normal estimate at each approximate medial point, allows grouping those approximate medial points sampled near the same smooth medial sheet. Experimental results support the validity of our approach for approximating the medial surface and segmenting the approximation into points belonging to distinct sheets. The next chapter studies a simplified version of Algorithm 2 and presents an analysis of properties of medial points that are not detected by this algorithm.

The union of the set of spheres, where each sphere is centred on an approximate medial point we compute and whose radius is the shortest distance from the centre to the boundary of the solid, is the “set of well-spaced medial spheres” shape representation of the solid. The quality of shape description offered by this representation will be discussed in Part II, while the utility of this shape representation for applications in computer graphics will be discussed in Part III.

## Chapter 3

# Theoretical Analysis: 3D

In the previous chapter, in considering the medial surface of a polyhedron, we described how a measure based on the average outward flux of the gradient of the signed Euclidean distance transform can be used to select voxels that are potentially intersected by a portion of the medial surface of high object angle. This step is not supported by guarantees that all voxels that are intersected by salient portions of the medial surface are detected, nor that the voxels output are indeed intersected by the medial surface.

In this chapter, we will propose a basic algorithm, derived from Algorithm 2, to detect convex regions intersected by the medial surface by analyzing the nearest boundary points to a set of points sampled on the surface of a convex region in Section 3.1. However, this algorithm may not detect all such regions. In Section 3.2 we will discuss which points should be sampled. Sections 3.3 and 3.4 present an error analysis for the algorithm in Section 3.1.

### 3.1 Basic Algorithm

Consider Algorithm 3, whose correctness follows by Lemma 2.2.

When the samples in  $\Phi$  are dense enough and there is a medial point inside  $S$ , one of the points  $\phi_i \in \Phi$  will lie on one of the two spoke vectors of the medial point, and Algorithm 3 will return ‘True’. However, for a finite sampling  $\Phi$ , it is not guaranteed that such a point  $\phi_i$  will be sampled. Ideally, if Algorithm 3 returns ‘Undetermined’ and a

---

**Algorithm 3** DECIDEMS( $\mathcal{B}, S, \Phi$ )

---

**Input:** Boundary  $\mathcal{B}$  of a solid  $\Omega$ , sphere  $S$  inside  $\Omega$ , set of non-medial points  $\Phi$  distributed on  $S$ .

**Output:** ‘True’ if  $S$  contains a medial point, ‘Undetermined’ if no such conclusion can be drawn.

```
1: for all  $\phi_i \in \Phi$  do
2:   if  $\mathcal{B}(\phi_i) \neq \mathcal{B}(\text{opp}(\phi_i))$  then
3:     Return ‘True’
4:   end if
5: end for
6: Return ‘Undetermined’
```

---

medial point  $m$  does lie inside  $S$ , then  $m$  should not be a “significant” medial point. The next section describes a situation when an arguably significant medial point lies inside  $S$ , while Algorithm 3 returns ‘Undetermined’.

### 3.2 Deep Samples

Suppose that DECIDEMS( $\mathcal{B}, S, \Phi$ ) returns ‘Undetermined’. Let  $\Theta = \Phi \cup \{\text{opp}(\phi_i) | \phi_i \in \Phi\}$  be the set of all sampled points considered on  $S$ . It may happen that none of the line segments  $(\phi, \text{opp}(\phi))$  are long enough to penetrate deeply into  $S$  and none intersect the medial surface. As a result, it is possible to fail to detect medial points in  $S$ , as shown in Figure 3–1. The medial points missed in this example are of the highest object angle possible ( $\pi/2$  for the medial point at the sphere centre). Further, as the radius of  $S$  can be chosen to be arbitrarily large, the medial points missed have arbitrarily large radius.

In order to improve the ability of Algorithm 3 to detect significant medial points (in terms of object angle and radius), let us consider two additional query points  $c_{in}$  and  $c_{out}$ , defined as follows. Let the centre of  $S$  be  $c$ . Let the nearest point on the boundary  $\mathcal{B}$  to  $c$  be  $C$ , which is outside  $S$  by the assumption that  $S$  does not intersect  $\mathcal{B}$ . Define




Figure 3–1: When  $\mathcal{B}$  consists of two points outside the sphere, the medial axis is shown as a dashed line. Points  $\Phi$  are big dots on the sphere.

$\{c_{in}, c_{out}\} \in S$ , where  $c_{out}$  is the intersection of  $S$  and the ray at  $c$  with direction  $\overrightarrow{(c, C)}$  and  $c_{in}$  is the intersection of  $S$  and the ray at  $c$  with direction  $\overrightarrow{(C, c)}$  (see Figure 3–2). Line segment  $(c_{in}, c_{out})$  is the longest line segment possible connecting a pair of points on  $S$ . In the example in Figure 3–1,  $\nabla D(c_{in}) \neq \nabla D(c_{out})$ , where  $c_{in}$  and  $c_{out}$  are two points on the sphere that are closest to the two boundary points. Therefore, in this example, by including  $c_{in}$  and  $c_{out}$  among the sampled points on  $S$ , we are guaranteed to detect a medial surface point in  $S$ . If we still do not detect a medial point in  $S$ , Lemma 3.1 characterizes where the set of nearest boundary points to points sampled on  $S$  must lie.

In the proof of the following lemma and in our subsequent discussion, we will use  $B(a, A)$  to denote a closed ball centred at point  $a$  and having point  $A$  on its boundary.

**Lemma 3.1.** *If  $\nabla D(c_{in}) = \nabla D(c_{out})$  and DECIDEMS  $(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’, then all the nearest points on  $\mathcal{B}$  to points in  $\Theta$  lie above the plane  $\pi$  through  $c_{in}$  with normal  $\overrightarrow{(c, C)}$ .*





of the line  $l$  with  $\pi_1$ . Since  $\angle Cpc_{in} > \pi/2$  and  $l$  is orthogonal to  $\pi_1$ ,  $p''$  is left of  $p$  on  $l$ . Hence,  $p''$ , just like  $p$ , is above  $\pi$ . Since  $\pi_2$  is tangent to  $S$  at  $p$  and since  $p''$  is left of  $p$  on  $l$ ,  $p'$  is above  $l$  on  $\pi_2 \cap \rho$  and hence, above  $\pi$ .  $\square$

In Figure 3–1, since the two detected boundary points are to different sides of the plane through  $c_{in}$  with normal  $\overrightarrow{(c, C)}$ , it must be that  $\nabla D(c_{in}) \neq \nabla D(c_{out})$ .

Lemma 3.1 explains how using the sample points  $c_{in}$  and  $c_{out}$  restricts the situations where Algorithm 3 returns ‘Undetermined’. The next section explains how the set of all possible locations of the two nearest boundary points to a medial point missed by Algorithm 3 can be computed.

### 3.3 Building the Foam

Suppose that Algorithm 3,  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$ , returns ‘Undetermined’. Consider the convex hull of the points  $\Theta = \Phi \cup \{\text{opp}(\phi_i) | \phi_i \in \Phi\}$ ,  $CH(\Theta)$ . Suppose that there is a medial point  $m$  inside  $CH(\Theta)$ . We will now seek to know the locations of  $m$ ’s nearest points on  $\mathcal{B}$ .

Recall that  $B_a = B(a, A)$  is a closed ball with centre  $a$  having point  $A$  on its boundary and let  $d(a, b)$  be the Euclidean distance between points  $a$  and  $b$ . The following tool will prove helpful in locating the nearest boundary points to  $m$ :

**Lemma 3.2.** *Consider two closed balls  $B_a = B(a, Y)$  and  $B_b = B(b, Y)$ . Then for any ball  $B_c = B(c, Y)$ ,  $c \in (a, b)$ ,  $B_c \subseteq B_a \cup B_b$ .*

*Proof.* <sup>1</sup> Let  $x$  be the intersection of line segment  $(a, b)$  with  $B_a \cap B_b$  (a disk). Let  $I$  be the boundary of  $B_a \cap B_b$  (a circle). We want to show that the distance from  $c$  to  $I$  is less than or equal to  $d(c, Y)$ . Let  $X \in I$ . Then  $d(c, X)^2 = d(c, x)^2 + d(x, X)^2$ . Also  $d(c, Y)^2 = d(c, x)^2 + d(x, Y)^2$ . However, note that  $d(x, X)^2 = d(x, Y)^2$ , since  $X$  and  $Y$  both lie on the circle  $I$  with centre  $x$ . It follows that  $d(c, Y)^2 = d(c, X)^2$  and  $d(c, X) = d(c, Y)$ . Thus,  $B_c$  is contained in the union of  $B_a$  and  $B_b$ .  $\square$

Let  $\mathbb{B}$  be the set of closest points on  $\mathcal{B}$  to  $\Theta$ . Let  $N = |\Phi|$ . We will assume that there are exactly  $N$  distinct points in  $\mathbb{B}$  (this holds when the boundary  $\mathcal{B}$  is  $C^1$ ). We now explain how to construct a region of  $\mathbb{R}^3$  that contains all the possible nearest boundary points to a medial point  $m$  inside  $CH(\Theta)$ . This region will be found by subtracting the “empty foam” from the “full foam”, which we define and explain how to compute in the following discussion.

**Empty Foam** For each point  $p \in \Theta$ , if  $P \in \mathbb{B}$  is the nearest boundary point to  $p$ , then ball  $B_p = (p, P)$  has an empty interior and the only point on its boundary is  $P$ . Let  $\mathcal{F}_e = \bigcup B_p \setminus \mathbb{B}$  be the union of balls hereafter called the *empty foam*.

**Full Foam** Consider the Voronoi diagram of  $\mathbb{B}$ ,  $VD(\mathbb{B})$ . Since  $m$  is a medial point, it is not one of the points in  $\Theta$ . Suppose that  $m$  is in  $A$ 's Voronoi region,  $V(A)$ ,  $A \in \mathbb{B}$ . Then  $m$ 's nearest point on  $\mathcal{B}$  is no further than  $d(m, A)$ , *i.e.* its nearest boundary point is on or inside the ball  $B_m = B(m, A)$ . Using the information about  $A$ 's Voronoi neighbours, we will find the region of space that contains  $B_m$ . This region of space will be a union of

---

<sup>1</sup> I would like to thank Nina Amenta for the idea behind this proof.




Figure 3–3: If a medial point  $m$  is in the convex hull of the 6 sampled points on the boundary of the dark disk with nearest boundary points  $A$ ,  $B$ , and  $C$ , then the nearest boundary points to  $m$  are inside the green disk and outside the grey disks. The dashed lines are the bisectors of  $A$ ,  $B$  and  $C$ .

balls, which we call the *full foam of A*,  $\mathcal{F}_f^A$ . The set  $\mathcal{F}_f = \{\bigcup \mathcal{F}_f^P | P \in \mathbb{B}\}$  is called the *full foam*.

We now explain how the full foam of A can be computed.

Let  $\{a, opp(a)\} \subset \Theta$  be the points in  $\Theta$  that have  $A$  as their nearest boundary point. Let  $a'$  be the nearest point on the line segment  $(a, opp(a))$  to  $m$ . It can be easily shown that the nearest boundary point to  $a'$  is  $A$ . Consider the ray at  $a'$  with direction  $\overrightarrow{(a', m)}$ . Let  $m'$  be the intersection of this ray with either the boundary of  $V(A)$ , or  $CH(\Theta)$ , whichever

occurs first. Let  $B_{m'} = B(m', A)$ . Let  $B_{a'} = B(a', A)$ . By Lemma 3.2,  $B_m \subset B_{a'} \cup B_{m'}$ . Let  $B_a = B(a, A)$  and  $B_{opp(a)} = B(opp(a), A)$ . By Lemma 3.2,  $B_{a'} \subset B_a \cup B_{opp(a)}$ . We add  $B_a$  and  $B_{opp(a)}$  to  $\mathcal{F}_f^A$  and now proceed to find spheres that contain  $B_{m'}$ .

There are several cases: (1)  $m'$  is on a Voronoi face, (2)  $m'$  is on a Voronoi edge, (3)  $m'$  is on a Voronoi vertex, or (4)  $m'$  is on  $CH(\Theta)$ . We consider each case in turn.

**Case 1 [ $m'$  is on a Voronoi face]** Suppose that the Voronoi face is a bisector of points  $A$  and  $B$  in  $\mathbb{B}$ . Let  $bis(A, B)$  denote the bisector of points  $A$  and  $B$ . It is a plane. Consider a plane at  $m'$  with normal direction  $(a, A)$ . This plane necessarily intersects  $bis(A, B)$  as  $\overrightarrow{(a', m')}$  is necessarily not parallel to  $\overrightarrow{(a, A)}$  and the intersection is a line on  $bis(A, B)$  passing through  $m'$ . By following this line, we will find two points  $m_1^*$  and  $m_2^*$ , where each point either lies on an edge of  $V(A)$  (Case 2), a vertex of  $V(A)$  (Case 3), or on the  $CH(\Theta)$  (Case 4). Define  $B_{m_1^*} = B(m_1^*, A)$  and  $B_{m_2^*} = B(m_2^*, A)$ . Then  $B_{m'} \subset B_{m_1^*} \cup B_{m_2^*}$  by Lemma 3.2. We now proceed to the respective cases to find balls containing  $B_{m_1^*}$  and  $B_{m_2^*}$ .

**Case 2 [ $m'$  is on a Voronoi edge]** Suppose that the Voronoi edge of  $V(A)$  is a trisector of points  $A, B$  and  $C$  in  $\mathbb{B}$ . Starting from a point  $m'$  on the edge, we will move up and down this edge until either we hit a Voronoi vertex of  $V(A)$  (Case 3), or we hit the convex hull of  $\Theta$  (Case 4) at points  $v_1$  and  $v_2$ . Then  $m' \in (v_1, v_2)$ . Let  $B_{v_1} = B(v_1, A)$  and  $B_{v_2} = B(v_2, A)$ . Then  $B_{m'} = B(m', A)$  is contained in  $B_{v_1} \cup B_{v_2}$  by Lemma 3.2. We add  $B_{v_1}$  and  $B_{v_2}$  to the full foam of  $A$   $\mathcal{F}_f^A$ .

**Case 3 [ $m'$  is on a Voronoi vertex]** Any Voronoi vertex  $v$  of  $V(A)$  inside  $CH(\Theta)$  defines a ball  $B_v = B(v, A)$  which we add to the full foam of  $A$   $\mathcal{F}_f^A$ .

**Case 4 [ $m'$  is on  $CH(\Theta)$ ]** Suppose that  $m'$  is a vertex of  $CH(\Theta)$ . This vertex cannot be  $a$  or  $opp(a)$  because we reached it by following the direction  $\overrightarrow{(a', m')}$  from  $a'$ . Any

other point in  $\Theta$  is outside of  $V(A)$ , and so is this vertex. But then we would have hit the boundary of  $V(A)$  before hitting this vertex when following the ray  $\overrightarrow{(a', m)}$  from  $a'$ . Therefore,  $m'$  is not a vertex and lies on an edge of  $CH(\Theta)$  (Case 4.1) or on the interior of some triangle of  $CH(\Theta)$  (Case 4.2).

**Case 4.1** [ $m' \in V(A)$  lies on an edge  $e$  of  $CH(\Theta)$ ] In case  $e$  is  $(a, opp(a))$ , then  $B_{m'} = B(m', A)$  is contained in  $B_a = B(a, A)$  and  $B_{opp(a)} = B(opp(a), A)$  and these balls have already been added to  $\mathcal{F}_f^A$ . Suppose edge  $e$  is  $(a, b)$  or  $(opp(a), b)$  for some point  $b \in \Theta$  outside  $V(A)$ . Then  $V(A)$  intersects  $(a, b)$  at some point  $x$ . By Lemma 3.2, either  $B_{m'} \subset B_a \cup B_x$  or  $B_{m'} \subset B_{opp(a)} \cup B_x$ , where  $B_x = B(x, A)$ . In this case, we add  $B_x$  and either  $B_a$  or  $B_{opp(a)}$  to  $\mathcal{F}_f^A$ . Now suppose edge  $e$  is  $(b, c)$ , which is intersected by  $V(A)$ , for some pair of points  $b, c \in \Theta$  outside of  $V(A)$ . In this case there are two points  $v_1$  and  $v_2$  on  $(b, c)$  that are the intersections of  $V(A)$  with  $(b, c)$ , such that  $m' \in (v_1, v_2)$ . If  $B_{v_1} = B(v_1, A)$  and  $B_{v_2} = B(v_2, A)$ , then  $B_{m'} \subset B_{v_1} \cup B_{v_2}$ . We add  $B_{v_1}$  and  $B_{v_2}$  to  $\mathcal{F}_f^A$ .

**Case 4.2** [ $m' \in V(A)$  lies on the interior of a triangle  $t$  of  $CH(\Theta)$ ] At least one vertex of triangle  $t$  is  $a$  or  $opp(a)$ . Suppose it is  $a$ . Then by following direction  $\overrightarrow{(a, m')}$ , we will hit either (4.2.1) an edge of  $t$  at point  $m''$ , or (4.2.2) the boundary of  $V(A)$  at point  $m''$ . Ball  $B_{m'} = B(m', A)$  is contained in  $B_a = B(a, A)$  and  $B_{m''} = B(m'', A)$ . In case 4.2.1, we proceed to Case 4.1 for point  $m''$  (recalling that  $B_a$  is already in  $\mathcal{F}_f^A$ ). In case 4.2.2, if  $m''$  is on an edge or vertex of  $V(A)$ , then we add  $B_{m''}$  to  $\mathcal{F}_f^A$  (recalling that  $B_a$  is already in  $\mathcal{F}_f^A$ ). Otherwise, if  $m''$  is on a face of  $V(A)$ , then the intersection of this face and  $t$  is a line segment  $(v_1, v_2)$ , where  $v_1$  and  $v_2$  are either on a Voronoi edge or vertex, or on an edge of  $t$ . If we define  $B_{v_1} = B(v_1, A)$  and  $B_{v_2} = B(v_2, A)$ , then  $B_{m'} \subset B_{v_1} \cup B_{v_2}$ . In this case, we add  $B_{v_1}$  and  $B_{v_2}$  to  $\mathcal{F}_f^A$  (recalling that  $B_a$  is already in  $\mathcal{F}_f^A$ ).

**Summary of balls in  $\mathcal{F}_f^A$ :** In this argument, for a medial point  $m$  in  $CH(\Theta)|V(A)$ , we have added balls to  $\mathcal{F}_f^A$  passing through  $A$  centred at the following types of points  $q$ :

- **Type 1:**  $q \in (a, opp(a))$
- **Type 2:**  $q$  is a vertex of  $V(A)$  inside or on  $CH(\Theta)$
- **Type 3:**  $q$  is an intersection of an edge of  $V(A)$  with  $CH(\Theta)$
- **Type 4:**  $q$  is an intersection of a face of  $V(A)$  with edges of  $CH(\Theta)$ .

By the argument above, which uses multiple invocations of Lemma 3.2 to create a set of spheres that contain  $B_m = B(m, A)$  for an arbitrarily positioned  $m \in CH(\Theta)|V(A)$ , it follows that  $B_m \subset \mathcal{F}_f^A$ . Starting with an arbitrary point  $m \in V(A)$ , we can construct the full foam of  $A$  by taking the union of the four types of balls described above.

The union of the full foams of each boundary point  $P \in \mathbb{B}$  gives the full foam:

$$\mathcal{F}_f = \{\bigcup \mathcal{F}_f^P | P \in \mathbb{B}\}.$$

Recall that the empty foam is

$$\mathcal{F}_e = \{\bigcup B_p | p \in \Theta\} \setminus \mathbb{B},$$

where  $B_p = B(p, P)$ , and  $P \in \mathbb{B}$  is the nearest boundary point to  $p$ . The region  $\mathcal{F}_e$  does not contain any points in  $\mathbb{B}$ . We have shown the following lemma:

**Lemma 3.3.** *Suppose that  $\text{DECIDEMS}(\mathcal{B}, S, \Theta)$  returned ‘Undetermined’. Let  $\mathbb{B}$  be the set of nearest boundary points to  $\Phi$  and let  $\Theta = \Phi \cup \{opp(\phi_i) | \phi_i \in \Phi\}$ . If there are exactly  $|\Phi|$  distinct points in  $\mathbb{B}$ , then for any medial point  $m \in CH(\Theta)$ , its nearest boundary points lie in  $\mathcal{F}_f \setminus \mathcal{F}_e$ .*

When computing the full foam, consider the following. Since balls of type 1 in  $\mathcal{F}_f$  are contained in  $\mathcal{F}_e \cup \mathbb{B}$ , we do not need to consider them when building  $\mathcal{F}_f \setminus \mathcal{F}_e$ . Balls of types 2-4 are on the border of Voronoi regions and hence are equidistant from their respective nearest points in  $\mathbb{B}$ . Therefore, when constructing  $\mathcal{F}_f \setminus \mathcal{F}_e$ , we do not need to compute these balls multiple times for each  $A \in \mathbb{B}$ . Thus, if we want to construct  $\mathcal{F}_f \setminus \mathcal{F}_e$ , we need only consider the vertices of  $VD(\mathbb{B})$  inside or on  $CH(\Theta)$ , the intersections of edges of  $VD(\mathbb{B})$  with  $CH(\Theta)$ , and the intersection of faces of  $VD(\mathbb{B})$  with edges of  $CH(\Theta)$ .

We can easily compute all the potential nearest boundary points to a medial point  $m$  in  $CH(\Theta)$  by finding the intersection of the boundary  $\mathcal{B}$  with all the balls of the full foam of type 2-4.

### 3.4 Measuring Missed Medial Point Quality

Suppose that a medial point  $m$  is inside  $S$ , yet  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’. In the previous section, we described how the set of potential boundary points that are closest to the medial point  $m$  can be computed. In this section, we discuss how bounds on the quality of the missed medial point  $m$  can be computed.

If  $m$  is equidistant from boundary points  $X$  and  $Y$ , then  $w = d(X, Y)$  provides a measure of the quality of this medial point, called the *width*. Both the radius of a medial sphere at  $m$ ,  $d = d(m, X) = d(m, Y)$ , and the object angle,  $\theta = \angle X m Y / 2$ , as well as width,  $w = d(X, Y)$  have been suggested as reasonable measures of quality of a medial point  $m$  in the literature [7, 42, 45, 53, 12, 30]. Since

$$w = 2d \sin(\theta), \tag{3.1}$$



a small value of  $w$  implies a small value of either  $d$  or  $\theta$ , and suggests that  $m$  is not a significant medial point.

As argued in the previous section, if  $m \in V(A)$ , then  $\mathcal{F}_f^A \setminus \mathcal{F}_e$  contains all the possible nearest boundary points to  $m$ . By considering all the possible pairs of nearest boundary points in  $\mathcal{F}_f^A \setminus \mathcal{F}_e$ , we can compute an upper bound on width for the missed medial point  $m$ . However, the width of a missed medial point can be very large. Consider, for example, the region  $\mathcal{F}_f \setminus \mathcal{F}_e$  in Figure 3–3. Here it is possible to pick two points in  $\mathcal{F}_f \setminus \mathcal{F}_e$  that lie in the same Voronoi region, such that their distance is very large (consider, for example, the two red points).

By introducing an additional measure of medial point quality, besides width, we can argue that some missed medial points having large width are not significant. Consider the distance from a medial point  $m$  to the boundary of a sphere  $S$ . When this distance is small, medial point  $m$  can be viewed as not important when using spheres  $S$  which overlap a large amount.

When  $S$  intersects the boundary  $\mathcal{B}$  of the solid  $\Omega$ , the width of a missed medial point of  $\Omega$  is at most the diameter of  $S$ . By using small spheres  $S$  near the boundary  $\mathcal{B}$ , one can ensure that the width of missed medial points in such regions is small.

Now, let us consider the case when  $S$  does not intersect the boundary  $\mathcal{B}$ . Recall that the  $|$  operator denotes restriction. Let  $\Theta = \Phi \cup \{\text{opp}(\phi_i) | \phi_i \in \Phi\}$ . Suppose that  $m \in V(A)|CH(\Theta)$ ,  $A \in \mathbb{B}$ , is a missed medial point. As shown in the previous section,  $\mathcal{F}_f^A \setminus \mathcal{F}_e$  contains all the possible nearest boundary points to  $m$ . We now make the following assumption.

**Assumption**  $V(A)$  intersects the boundary of  $CH(\Theta)$  in two closed polygonal chains.

We will show that either the distance from  $m$  to  $S$  can be bounded, or that a new set of possible locations of nearest boundary points to  $m$  can be computed by shrinking some of the balls of the full foam.

Let the intersection of  $V(A)$  and the boundary of  $CH(\Theta)$  be the two closed polygonal chains  $c_1$  and  $c_2$ . These chains divide  $CH(\Theta)$  into three regions. Let  $a_{in}$  be the further point from  $A$  among  $\{a, opp(a)\}$  (defined since  $\mathcal{B}$  does not intersect  $S$ ). As  $a_{in} \in S$ , it is on the boundary of  $CH(\Theta)$ . Suppose that chain  $c_1$  bounds the region containing  $a_{in}$ . Consider those vertices  $R$  of  $c_1$  that are on edges of  $V(A)$  (ball centres of type 3). Consider a triangulation  $T$  of  $R$  and  $a_{in}$ , such that for each pair of vertices  $\{v_1, v_2\} \in R$  on the same Voronoi face, there is a triangle with edges  $(v_1, v_2)$ ,  $(v_1, a_{in})$ , and  $(v_2, a_{in})$ . The triangles in  $T$  partition  $V(A)|CH(\Theta)$  into two regions. Let the region containing  $opp(a_{in})$  be called the region *above*  $T$ .

Consider each ball  $B_x = B(x, A)$  in  $\mathcal{F}_f^A \setminus \mathcal{F}_e$ . As we argued earlier, balls of type 1 are redundant in  $\mathcal{F}_f^A \setminus \mathcal{F}_e$ . Balls of type 2 are not in  $\mathcal{F}_f^A \setminus \mathcal{F}_e$  by the assumption about the nature of the intersection of  $CH(\Theta)$  and  $V(A)$ . Therefore, balls  $B_x$  are of types 3 and 4. Let  $\mathcal{F}_f^{A\dagger}$  be those balls in  $\mathcal{F}_f^A$ , save for those of type 4 whose centres are on the polygonal chain  $c_1$ . Consider a medial point  $m$  that lies in  $V(A)|CH(\Theta)$  and above  $T$ . By a similar argument as that used in Section 3.3, balls  $B_m = B(m, A)$  are contained in  $\mathcal{F}_f^{A\dagger}$ .

Consider the balls  $B^* \subset \mathcal{F}_f^{A\dagger}$ ,  $B^* = \{B(x, A) | x \in R\}$ . The balls in  $B^*$  are either all connected in  $\mathcal{F}_f^{A\dagger} \setminus \mathcal{F}_e$ , or at least one ball is disconnected. If all the balls are connected, then either the width of a missed medial point  $m$  is at most the maximum distance between pairs of points in  $\mathcal{F}_f^{A\dagger} \setminus \mathcal{F}_e$  or the distance from  $m$  to  $S$  is at most the maximum distance from  $T$  to  $S$ .

Now suppose that at least one ball  $B_x \in B^*$  is disconnected in  $\mathcal{F}_f^A \dagger \setminus \mathcal{F}_e$ . The ball centre,  $x$ , is on an edge of  $V(A)$  and is equidistant from  $A$  and two other points in  $\mathbb{B}$ :  $B$  and  $C$ . We now shrink  $B_x$  to make a new ball  $B'_x = (x', A)$ , such that  $x'$  is on the same edge of  $V(A)$  as  $x$  and  $B'_x \setminus \mathcal{F}_e$  is connected. When all the spheres in  $\mathcal{F}_f^A \dagger$  with centres in  $R$  are made connected in this way, let  $R'$  be the new set of sphere centres. Now consider triangulation  $T'$  of  $R'$  and  $a_{in}$ , made in the same way as  $T$ . Let  $\mathcal{F}_f^A \ddagger$  be the spheres in  $\mathcal{F}_f^A \dagger$ , where spheres with centres in  $R$  are replaced with their shrunken versions. By the same argument as before, any ball  $B_m = B(m, A)$ , such that its centre  $m$  lies in the region  $V(A) \setminus CH(\Theta)$  and further above  $T'$ , is contained in  $\mathcal{F}_f^A \ddagger$ . Then either the width of a medial point  $m$  is at most the maximum distance between pairs of points in  $\mathcal{F}_f^A \ddagger \setminus \mathcal{F}_e$ , or the distance from  $m$  to  $S$  is at most the maximum distance from  $T'$  to  $S$ .

It should be possible to construct an argument showing that either the width (distance between pairs of boundary points in  $\mathcal{B}_f^A$ ) or the distance to  $S$  of a missed medial point  $m \in V(A)$  decreases as more samples  $\Theta$  are considered by Algorithm 3, but this remains to be demonstrated. The argument presented in this section, describing how the distance from a missed medial point to the sampling region  $S$  can be measured, should provide a building block for such an argument. As more samples are considered, it should be possible to show that those Voronoi regions that do not intersect the boundary of  $CH(\Theta)$  in two closed polygonal chains will contain points whose distance to  $S$  is small.

### 3.5 Chapter Summary and Future Work

This chapter has presented a preliminary study of an algorithm for the detection of spheres containing medial points, based on an analysis of nearest boundary points to a set of points sampled on spheres, where some important properties of the nature and quality

of the medial points that are not detected by the algorithm have been established. Much remains to be done in order to understand what theoretical guarantees can be offered by such an algorithm in 3D and higher dimensions.

Below we list several open problems. Medial point quality may be measured using either object angle, radius or distance from the medial point to the query region.

1. Show that if  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’, the quality of medial points present in sphere  $S$  decreases as the density of  $\Phi$  increases.
2. Suppose that  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’. For each sphere in  $\mathcal{F}_f$ , we add its centre to the set of query points  $\Phi$  to create a set of query points  $\Phi'$ . Next, execute  $\text{DECIDEMS}(\mathcal{B}, S, \Phi')$ . By carrying out this operation repeatedly, what can be said about the quality of the missed medial points as more iterations are considered?
3. Design other rules based on the analysis of a finite number of nearest boundary points to query points on a sphere that enable detection of medial points inside the sphere, such that the quality of the missed medial points can be shown to decrease by increasing the density of query points used, for a fixed-size query region  $S$ .

Our method for identifying spheres whose interior is intersected by the medial surface of a solid generates samples on the boundary of the solid, which are the nearest boundary points to query points on spheres, and checks whether these samples are nearest boundary points to medial points inside spheres. Voronoi methods for medial surface approximation also identify medial points generated by boundary samples. The theoretical guarantees offered by Voronoi methods usually stem from the assumption that the Hausdorff distance between the boundary samples considered and the true surface of the solid is bounded,

with additional assumptions about the smoothness of the solid boundary. In our work, we have not made any assumptions about the Hausdorff distance of the boundary samples we generate and the solid boundary.

Regions intersected by the medial axis of a 2D solid can also be detected based on an analysis of nearest boundary points to a set of query points on the boundary of a convex 2D region. The next chapter considers the theoretical guarantees that can be offered in the 2D case. This 2D case is significantly easier than the 3D case. For this reason, it has been possible to make more progress in the 2D case than in the 3D case considered in the current chapter. The 2D analysis offers a more complete perspective on the algorithms we are discussing. Further, the difference between approaches possible in 2D and 3D highlights the challenges of the 3D case.

## Chapter 4

# Theoretical Analysis: 2D

In this thesis, we are concerned with the design and analysis of algorithms for the approximation of medial surfaces of 3D solids based on the analysis of nearest boundary points to query points. However, in the 3D case, we have not yet been able to establish an exact relationship between the distribution of the query points and the nature of medial points that remain undetected. For the easier 2D case, this chapter establishes an explicit relationship between the density of query points on a circle and the quality of medial points inside the circle that remain undetected.

We develop additional tools for identifying medial points based on the analysis of the  $\nabla D$  vector field in 2D in Section 4.1. Based on these tools, we design algorithm DECIDEMA for approximating the medial axis of a 2D solid in Section 4.2. In Section 4.3, we give a geometric description of the medial points undetected by this 2D algorithm. We argue that the quality measures used to describe missed medial points are consistent with current practice. In Section 4.4, we use DECIDEMA to design an algorithm for medial axis approximation.

We have provided a detailed overview of algorithms that compute the medial surface (of 3D solids) in Section 2.2. For an overview of algorithms that compute the medial axis (of 2D solids), we invite the reader to consult the thorough review in [44]. We will just mention that the medial axis of a simple polygon can be computed in linear time [33]. In

terms of practical implementations, the CGAL library [1] has a very efficient and robust implementation of the incremental algorithm by Karavelas [69] to compute the Voronoi diagram of line segments, which can be easily simplified to obtain the medial axis of a polygon.

In this chapter, let  $\Omega$  be a 2D solid with boundary  $\mathcal{B}$ . Let  $C$  be a circle inside  $\Omega$ .

## 4.1 Tools for Identification of Medial Points

In this section, we establish some tools for the detection of medial points based on the analysis of  $\nabla D$  vectors in 2D. These tools will provide the basis for the design of our algorithm in Section 4.2.

Algorithm 3 for identifying spheres that contain medial points is based on Lemma 2.3. This lemma holds in 2D as well. We now introduce additional properties of the  $\nabla D$  vector field that enable detection of medial points.

Recall that  $(a, b)$  is a closed line segment with endpoints  $a$  and  $b$ . We start with the following consequence of Lemma 2.1:

**Corollary 4.1.** *Let  $b = a + \gamma \nabla D(a)$  and  $d = c + \delta \nabla D(c)$ , where  $\gamma, \delta \in \mathbb{R}$ , for distinct  $a, b, c, d$ . If  $\nabla D(a) = \nabla D(b)$  and  $\nabla D(c) = \nabla D(d)$ , then  $(a, b)$  and  $(c, d)$  do not cross.*

*Proof.* Suppose  $(a, b)$  and  $(c, d)$  cross at a point  $e$ . Then  $\nabla D$  is multi-valued at  $e$ . However, by Lemma 2.1, the  $\nabla D$  direction is unique for points on  $(a, b)$  and  $(c, d)$ .  $\square$

Although this lemma holds in 3D as well, the likelihood of two arbitrary 3D line segments crossing is very small. However, in 2D, we will see that this property is very useful for “trapping” the medial axis.

Let  $l(a, b)$  denote the line through points  $a$  and  $b$ . We will make use of the following definition.

**Definition 4.1.** Consider two vectors  $\nabla D(p)$  and  $\nabla D(q)$  and let  $n(p, q)$  be a vector normal to the line through  $p$  and  $q$ ,  $l(p, q)$ . Then  $\nabla D(p)$  and  $\nabla D(q)$  are to the same side of  $l(p, q)$  if  $\text{sgn}(\nabla D(p) \cdot n(p, q)) = \text{sgn}(\nabla D(q) \cdot n(p, q))$ , and to different sides of  $l(p, q)$  otherwise.

The notion of an *opposite* in 2D is the same as that in 3D: for  $a \in C$ ,  $\text{opp}(a)$  is the other point of intersection of the line with direction  $\nabla D(a)$  with  $C$ .

**Lemma 4.1.** Suppose  $a$  and  $b$  are two points on  $C$ ,  $b \neq \text{opp}(a)$ , and there are no medial points along  $(a, b)$ . Then  $\nabla D(a)$  and  $\nabla D(b)$  are to the same side of  $l(a, b)$ .

*Proof.* Suppose that there are no medial points along  $(a, b)$ , yet  $\nabla D(a)$  and  $\nabla D(b)$  are to different sides of  $l(a, b)$ . Since there are no medial points along  $(a, b)$ , then  $\nabla D$  varies smoothly for points along  $(a, b)$  [74]. Thus, at some point  $c$ ,  $\nabla D(c)$  is aligned with  $(a, b)$ . Since there is no medial point along  $(a, b)$ , by Lemma 2.3, either  $\nabla D(a) = \nabla D(c)$  or  $\nabla D(b) = \nabla D(c)$ . Then  $b = \text{opp}(a)$ , a contradiction.  $\square$

**Corollary 4.2.** If  $\nabla D(a)$  and  $\nabla D(b)$  are to different sides of  $l(a, b)$ , then there is a medial point on  $(a, b)$ .

Note that Corollary 4.2 does not generalize to 3D: three  $\nabla D$  vectors may point to different sides (above and below) of a plane through their tails, yet no medial point may lie on the triangle through their tails. Figure 4–1 shows an example of this situation in 3D.

For a circle  $C$  inside  $\Omega$ , let  $\Phi$  be a set of  $N$  points sampled on  $C$ , such that the maximum distance between neighbouring points in  $\Phi$  on  $C$  is small. Let  $\Psi = \{\text{opp}(\phi_i), \phi_i \in$






Figure 4–1:  $\nabla D(A)$ ,  $\nabla D(B)$  point below the plane through  $A$ ,  $B$  and  $C$ , while  $\nabla D(C)$  points above this plane. The nearest boundary points to  $A$ ,  $B$  and  $C$  lie on an edge of polyhedron. The medial surface does not necessarily intersect  $\triangle ABC$ .

$\Phi\}$ . Let  $\Theta = \Phi \cup \Psi$ . For two points  $a$  and  $b$  in  $\Theta$ , let  $\widehat{(a, b)}$  denote the closed arc of  $C$  connecting  $a$  and  $b$ , where the choice of arc is the one empty of other points in  $\Theta$ , unless otherwise specified. Points  $a$  and  $b$  are *nearest neighbours* in  $\Theta$  if no other point in  $\Theta$  is in the shorter arc of  $C$  with endpoints  $a$  and  $b$ .

**Lemma 4.2.** *Let  $a$  and  $b$  be points in  $\Theta$  that are nearest neighbours along  $C$ . Suppose that  $\nabla D(a)$  and  $\nabla D(b)$  are to the same side of  $l(a, b)$  and also that  $\nabla D(a) = \nabla D(\text{opp}(a))$  and  $\nabla D(b) = \nabla D(\text{opp}(b))$ . Suppose that point  $c \in \Theta, c \notin \{\text{opp}(a), \text{opp}(b)\}$ , lies in  $(\text{opp}(a), \text{opp}(b))$  (where the choice of arc is the one not containing  $a$  or  $b$ ). Then there is a medial point inside  $C$  on  $(a, c)$  or  $(b, c)$ .*

*Proof.* Suppose that there is no medial point inside  $C$ . Consider  $(a, c)$  and  $(c, b)$ . Since there is no medial point inside  $C$ , by Lemma 4.1,  $\nabla D(c)$  is to the same side of  $l(a, c)$  as




Figure 4-2: Cases A (Left) and B (Right) of Lemma 4.2.

$\nabla D(a)$  and also to the same side of  $l(b, c)$  as  $\nabla D(b)$ . By Corollary 4.1,  $(a, opp(a))$  and  $(b, opp(b))$  do not cross.

For the remainder of the proof, we will orient  $(a, b)$  horizontally, such that  $a$  is to the left of  $b$ , as in Figure 4-2. The next two paragraphs handle the two cases: (1)  $\nabla D(a)$  and  $\nabla D(b)$  both point above  $l(a, b)$  and (2)  $\nabla D(a)$  and  $\nabla D(b)$  both point below  $l(a, b)$ .

If  $\nabla D(a)$  and  $\nabla D(b)$  both point above  $l(a, b)$ , then  $\nabla D(a)$  is to the right of  $(a, c)$  and  $\nabla D(b)$  is to the left of  $(b, c)$  because  $c \in (\widehat{opp(a), opp(b)})$ . Then  $\nabla D(c)$  is to the right of  $l(a, c)$  and to the left of  $l(b, c)$ . Thus  $opp(c) \in (\widehat{a, b})$ , an impossibility since  $(\widehat{a, b})$  is empty of other points in  $\Theta$  and  $opp(a) \neq c$  and  $opp(b) \neq c$ . Refer to Figure 4-2 (Right).

If  $\nabla D(a)$  and  $\nabla D(b)$  are both below the line through  $(a, b)$ , then  $\nabla D(a)$  is to the left of  $l(a, c)$  and  $\nabla D(b)$  is to the right of  $l(b, c)$ . Then  $\nabla D(c)$  is to the left of the line through  $(a, c)$  and to the right of the line through  $(b, c)$ . Thus  $opp(c)$  is again in  $(\widehat{a, b})$ , an impossibility as we argued above. Refer to Figure 4-2 (Left).

Therefore, if there exists a point  $c \in \Theta$  in the arc  $(\widehat{opp(a), opp(b)})$ , then there is a medial point in  $C$ . In particular, there is a medial point on  $(a, c)$  or  $(b, c)$ .  $\square$

## 4.2 Algorithm

Using the tools developed in the previous section, we now propose Algorithm 4, DECIDEMA, for the detection of regions of 2D space that contain medial points.

---

**Algorithm 4** DECIDEMA( $\mathcal{B}, C, \Phi$ )

---

**Input:** Boundary  $\mathcal{B}$  of a 2D solid  $\Omega$ , circle  $C$  inside  $\Omega$ , set of points  $\Phi$  distributed on  $C$ .

**Output:** ‘True’ if  $C$  contains a medial point, ‘Undetermined’ if no such conclusion can be drawn.

```

1: for all  $\phi_i \in \Phi$  do
2:   if  $\nabla D(\phi_i) \neq \nabla D(\text{opp}(\phi_i))$  then
3:     Return ‘True’
4:   end if
5: end for
6: Let  $\Theta = \Phi \cup \{\text{opp}(\phi_i) | \phi_i \in \Phi\}$ 
7: for all pairs of nearest neighbours  $\theta_i$  and  $\theta_j$  in  $\Theta$  do
8:   if  $\nabla D(\theta_i)$  and  $\nabla D(\theta_j)$  are to different sides of  $l(\theta_i, \theta_j)$  then
9:     Return ‘True’
10:  end if
11: end for
12: for all pairs of nearest neighbours  $\theta_i$  and  $\theta_j$  in  $\Theta$  do
13:   if  $(\text{opp}(\theta_i), \text{opp}(\theta_j))$  contains points in  $\Theta$  then
14:     Return ‘True’
15:   end if
16: end for
17: Return ‘Undetermined’

```

---

The following lemma attests to the correctness of Algorithm 4.

**Lemma 4.3.** *If DECIDEMA( $\mathcal{B}, C, \Phi$ ) returns ‘True’ then  $C$  contains a medial point.*

*Proof.* If DECIDEMA returns ‘True’ on line 3, then by Lemma 2.3, there is a medial point in  $C$ . If DECIDEMA returns ‘True’ on line 8, then by Corollary 4.2, there is a medial point in  $C$ . If DECIDEMA returns ‘True’ on line 13, then by Lemma 4.2, there is a medial point in  $C$ . □

Algorithm DECIDEMA performs  $O(N)$  distance queries and takes  $O(N \log N)$  time to establish the ordering of points in  $\Theta$  along  $C$ , when it is not known a priori.

### 4.3 Missing a medial point

In this section, we consider the possibility that  $\text{DECIDEMA}(\mathcal{B}, C, \Phi)$  returns ‘Undetermined’, yet the convex hull of  $\Theta$  contains a medial point  $p$ . Some situations when this happens are illustrated in Figure 4–3.

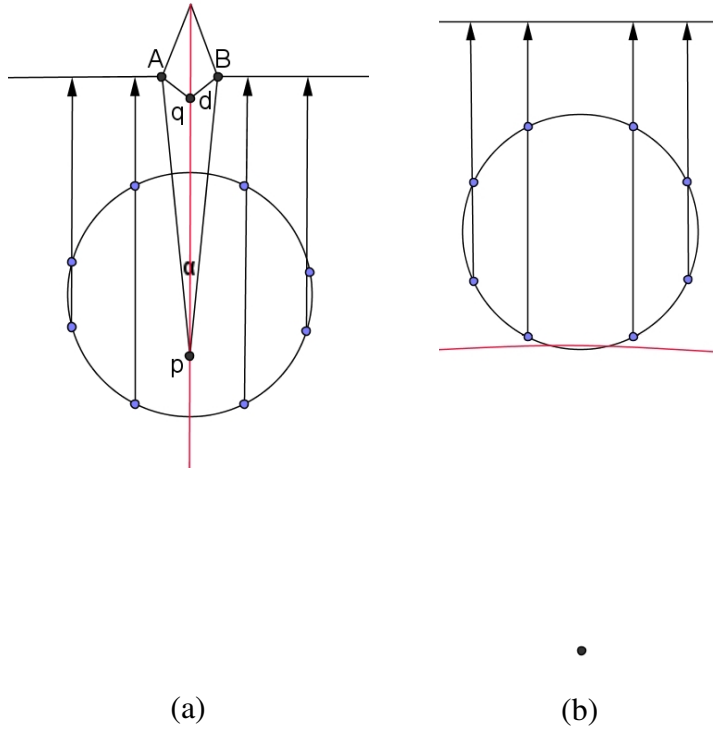


Figure 4–3: Situations when Algorithm 4 returns ‘Undetermined’ while a section of the medial axis passes through circle  $C$ . Medial points  $p$  and  $q$  are equidistant from points  $A$  and  $B$  on the boundary. Point  $p$  has a small object angle, while the distance  $d$  from point  $p$  to the boundary is small.

Although our algorithm fails to identify medial points in these examples, in this section we will show that whenever circle  $C$  contains a medial point  $p$  and DECIDEMA

returns ‘Undetermined’, when sample points  $\Phi$  on  $C$  are sufficiently dense, either (1)  $p$  has a small object angle, (2)  $p$  is close to  $\mathcal{B}$ , (3)  $C$  is close to  $\mathcal{B}$ , or (4) the medial point  $p$  is shallow inside  $C$ . The medial points  $p$  and  $q$  in Figure 4–3(a) are equidistant from points  $A$  and  $B$  on  $\mathcal{B}$ . Point  $p$  has a small object angle  $\angle ApB$ , while point  $q$  is close to  $\mathcal{B}$ . Figure 4–3(b) shows an example of missing medial points that are close to  $C$ .

Suppose that a medial point  $p$  is on or inside the convex hull of  $\Theta$ . Point  $p$  is equidistant from two locations on  $\mathcal{B}$ ; let us call these locations  $A$  and  $B$ . Let  $a$  be the point of intersection of  $(p, A)$  with  $C$ ; define  $b$  similarly for  $B$ . Points  $a$  and  $b$  cannot be in  $\Theta$ , because DECIDEMA ensured that  $\nabla D(a) = \nabla D(\text{opp}(a))$  and  $\nabla D(b) = \nabla D(\text{opp}(b))$ , and by Corollary 2.1,  $(a, \text{opp}(a))$  and  $(b, \text{opp}(b))$  do not contain medial points. Call the two nearest neighbours of  $a$  in  $\Theta$   $x$  and  $y$ ; similarly, the nearest neighbours of  $b$  in  $\Theta$  are  $x'$  and  $y'$ . Also, since DECIDEMA returned ‘True’,  $\nabla D(x)$  and  $\nabla D(y)$  are to the same side of  $l(x, y)$  and  $\nabla D(x')$  and  $\nabla D(y')$  are to the same side of  $l(x', y')$ .

For a medial axis point  $p$  inside the convex hull of  $\Theta$ , let us consider two cases:

- Case 1:  $\nabla D(a)$  is to the same side of  $l(x, y)$  as are  $\nabla D(x)$  and  $\nabla D(y)$ , and also  $\nabla D(b)$  is to the same side of  $l(x', y')$  as are  $\nabla D(x')$  and  $\nabla D(y')$ .
- Case 2: Either  $\nabla D(a)$  is to the other side of  $l(x, y)$  as  $\nabla D(x)$  and  $\nabla D(y)$ , or  $\nabla D(b)$  is to the other side of  $l(x', y')$  as  $\nabla D(x')$  and  $\nabla D(y')$ .

Missed medial points in Figure 4–3 (Left) belong to Case 1, while those missed in Figure 4–3 (Right) belong to Case 2.

Let  $L$  be the longest length of the shortest arc on  $C$  between two neighbouring points sampled on  $C$ . We will require that  $L$  decreases as  $N$  increases. Let  $R$  be the radius of  $C$ .

Let  $N$  be big enough such that  $L < 2R$ . Let  $d = d(p, A) = d(p, B)$ . We will study the nature of the medial points missed as a function of  $L$ ,  $R$  and  $d$ .

#### 4.3.1 Case 1

Suppose that  $C$  contains a medial point  $p$  as described by Case 1, yet DECIDEMA  $(\mathcal{B}, C, \Phi)$  returned ‘Undetermined’.

We first show that, in this case,  $b \in \widehat{(x, y)}$ . Suppose  $b \notin \widehat{(x, y)}$ . Note that  $(x, opp(x))$  and  $(y, opp(y))$  do not cross, by Corollary 4.1. Let  $(x, \widehat{opp(x)})$  be the arc of  $C$  not containing  $y$  and let  $(y, \widehat{opp(y)})$  be the arc of  $C$  not containing  $x$ . Point  $b$  cannot be in  $(x, \widehat{opp(x)})$  or  $(y, \widehat{opp(y)})$ , because  $(x, opp(x))$  and  $(y, opp(y))$  cannot cross  $(p, b)$ , by Corollary 4.1. Thus  $b \in (\widehat{opp(x), opp(y)})$ . Consider  $x'$  and  $y'$ . Since algorithm DECIDEMA returned ‘Undetermined’,  $(\widehat{opp(x), opp(y)})$  does not contain any points of  $\Theta$  in its interior. Therefore,  $x'$  and  $y'$  are  $opp(x)$  and  $opp(y)$ . Vectors  $\nabla D(x)$  and  $\nabla D(y)$  point to the same side of  $l(x, y)$  as  $\nabla D(a)$ , that is, outside of  $C$ . Likewise,  $\nabla D(x')$  and  $\nabla D(y')$  point to the same side of  $l(x', y')$  as  $\nabla D(b)$ , that is, outside of  $C$ . Thus,  $\nabla D(x) = \nabla D(opp(x)) = x - opp(x)$ . Since one of  $\nabla D(x')$  or  $\nabla D(y')$  is  $D(opp(x))$ , one of  $\nabla D(x')$  or  $\nabla D(y')$  points to different sides of  $(x', y')$  as does  $\nabla D(b)$ . Since we are in Case 1, point  $b \notin (\widehat{opp(x), opp(y)})$ . It must lie in  $\widehat{(x, y)}$ .

Suppose that  $(x, y)$  is aligned with the horizontal and that  $x$  is left of  $y$ . Let us call  $d(p, A)$  the *distance*  $d$  from the medial axis point  $p$  to the boundary  $\mathcal{B}$ . Let  $\alpha$  be the angle between  $(p, A)$  and  $(p, B)$ . Angle  $\alpha$  is twice the object angle of medial point  $p$ . Let  $R$  be the radius of  $C$ . Let  $X$  be the nearest point on  $\mathcal{B}$  to  $x$  and let  $Y$  be the nearest point of  $\mathcal{B}$  to  $y$ . We will show that if the missed medial point  $p$  is of the type in Case 1, then either the *distance* or the *object angle* is bounded in terms of  $R$  and  $L$ .

Let  $d(p, A) = d(p, B)$  be  $d$ . When  $d < 2L$ , this is a reasonable bound on the quality of the medial point missed, as we will argue in Section 4.4.

When  $d \geq 2L$ , we look for an upper bound on angle  $\alpha$  given  $L$ ,  $R$  and  $d$ . Consider the isosceles triangle with vertices  $p$ ,  $A$  and  $B$ . We know that  $a$  and  $b$  are in  $\widehat{(x, y)}$ . We now argue that  $\triangle pAB$  must be fully contained in the region right of  $l(x, opp(x))$  and left of  $l(y, opp(y))$ . Suppose that  $A$  lies outside this region, left of  $l(x, opp(x))$ . Then, since  $(x, X)$  and  $(p, A)$  do not cross by Corollary 4.1,  $X$  is closer to  $p$  than is  $A$ , contradicting the fact that  $A$  is a closest point on  $\mathcal{B}$  to  $p$ .

For an upper bound on  $\alpha$ , let us consider an isosceles triangle  $\triangle uvw$  such that  $d(u, w) = d(w, v) = d$  and such that  $\triangle uvw$  is fully contained inside the region right of  $l(x, opp(x))$  and left of  $l(y, opp(y))$ . We look for a placement of vertex  $w$  on or inside the convex hull of  $\Theta$  such that  $\angle uvw$  is maximum. The maximum  $\angle uvw$  is achieved when  $opp(x)$  and  $opp(y)$  coincide at a point  $z$  on  $C$ . The region right of  $l(x, z)$  and left of  $l(y, z)$  is a wedge  $xzy$ . For an upper bound on  $\alpha$ , let the length of arc  $\widehat{(x, y)}$  be  $L$ . The inscribed acute angle  $\phi$  made by the chord  $(x, y)$  with  $z$  is  $\arcsin\left(\frac{L}{2R}\right)$  (see Figure 4–4).

Consider all possible isosceles triangles  $uvw$  that lie inside the wedge  $xzy$ , such that  $d > L$ . We want to find the position of this triangle that gives the maximum  $\angle uvw$ . Suppose that  $\angle uvw$  is maximum when  $w$  lies strictly inside the convex hull of  $\Theta$  (see Figure 4–4(a)). Since  $\angle uvw$  is maximum and  $d > L$ , each of  $u$  and  $v$  lies on one of  $l(z, x)$  and  $l(z, y)$ . Suppose that  $u$  lies on  $l(z, x)$ . Now consider translating the triangle  $uvw$  in the direction parallel to  $(z, x)$ , so that  $u$  moves along  $l(z, x)$  until  $w$  lies on the chord  $(x, y)$  instead of inside the convex hull. Since  $\phi > 0$ ,  $v$  no longer lies on  $l(z, y)$  and, since  $d \geq 2L > d(x, y)$ ,  $\angle uvw$  can be increased so that  $v$  lies on  $l(z, y)$  (see Figure 4–4(b)).

Thus,  $\angle u w v$  is not maximum when  $w$  lies strictly inside the convex hull of  $\Theta$ . Therefore, to maximize  $\angle u w v$ , let us place  $w$  on the chord  $(x, y)$  of  $C$ .

We know that the maximum value of  $\angle u w v$  occurs when  $w$  lies on the chord  $(x, y)$ . Now we want to know an upper bound on  $\angle u w v$  when  $w$  lies on the chord  $(x, y)$ . We measure the *height* of  $w$  with respect to  $z$  as follows. Let  $x^*$  and  $y^*$  be the intersections of a line through  $w$  with  $l(x, z)$ ,  $l(y, z)$ , such that  $\angle z x^* y^* = \angle z y^* x^*$ . Then  $d(x^*, z) = d(y^*, z) = h$  is the height of  $w$  in the wedge  $x y z$  (see Figure 4–4(c)).

For all possible points  $z$  on  $C$ , either  $h \leq d(z, y)$  or  $h \leq d(z, x)$ . Suppose it is the latter. Since  $(x, z)$  and  $(y, z)$  are chords of  $C$ ,  $d(x, z) \leq 2R$  and  $d(y, z) \leq 2R$ . Since  $d(z, y) \leq 2R$ ,  $h \leq 2R$ . Consider an isosceles triangle  $\triangle x' z y'$  such that  $d(z, x') = d(z, y') = 2R$  and  $\angle x' z y' = \phi$ . Call  $d(x', y')$   $L'$ . The triangle  $u w v$  lies inside the wedge  $x' z y'$ . Then  $h = d(x^*, z) = d(y^*, z) \leq 2R$ . That is,  $w$  lies on or below  $(x', y')$  in the wedge  $x' z y'$ . By earlier arguments,  $w$  can be placed on  $(x', y')$  (see Figure 4–4(c)). Again, by earlier arguments,  $\angle u w v$  can be increased by placing  $u$  on  $l(x, z)$  and  $v$  on  $l(y, z)$  (see Figure 4–4(d)).


We want to find the position of  $w$  on  $(x', y')$  such that  $\triangle w u v$  is inside the wedge  $x' z y'$  that maximizes  $\angle u w v$ , where

$$\angle u w v = \phi + \angle x' u w + \angle w v y'.$$


Also,  $d(x', w) = L' \cdot t$  and  $d(w, y') = L' \cdot (1 - t)$ ,  $t \in [0, 1]$ . Letting  $\angle u w v = \beta(t)$ , we have  $\beta(t) =$

$$\phi + \arcsin \left( \sin(\angle u x' y') \frac{L' t}{d} \right) + \arcsin \left( \sin(\angle v y' x') \frac{L' (1 - t)}{d} \right).$$







(a)



(b)



(c)



(d)

Figure 4–4: Illustrations for error analysis in Case 1.

Taking the derivative with respect to  $t$ ,  $\frac{\partial \beta}{\partial t} =$

$$\frac{\frac{L'}{d} \sin(\angle ux'y')}{\sqrt{1 - (\sin(\angle ux'y') \frac{L't}{d})^2}} - \frac{\frac{L'}{d} \sin(\angle vy'x')}{\sqrt{1 - (\sin(\angle vy'x') \frac{L'(1-t)}{d})^2}}.$$

Setting  $\frac{\partial \beta}{\partial t} = 0$ ,

$$\frac{\frac{L'}{d} \sin(\angle ux'y')}{\sqrt{1 - (\sin(\angle ux'y') \frac{L't}{d})^2}} = \frac{\frac{L'}{d} \sin(\angle vy'x')}{\sqrt{1 - (\sin(\angle vy'x') \frac{L'(1-t)}{d})^2}}.$$

Triangle  $\triangle x'zy'$  is an isosceles triangle. Therefore,  $\angle zx'y' = \angle zy'x'$ ,  $\angle ux'y' = \angle vy'x'$  and hence

$$t^2 = (1 - t)^2,$$

which implies that  $t = \frac{1}{2}$ . Thus when  $t = \frac{1}{2}$ ,  $\beta(t)$  has either a global maximum or a global minimum. We compare the value of  $\beta(\frac{1}{2})$  with the values at the extreme points  $t = 0, 1$ .

$$\begin{aligned} \beta(0) &= \phi + \arcsin \left( \sin(\angle xy'v') \frac{L'}{d} \right), \\ \beta(1) &= \phi + \arcsin \left( \sin(\angle xy'v') \frac{L'}{d} \right), \text{ and} \\ \beta\left(\frac{1}{2}\right) &= \phi + 2 \arcsin \left( \sin(\angle xy'v') \frac{L'}{2d} \right). \end{aligned}$$

As can be shown,  $\arcsin(x/2) < \arcsin(x)/2$  for  $x > 0$ . Since  $\angle xy'v' < \pi$ ,  $\sin(\angle xy'v') \frac{L'}{d} > 0$  and hence,  $\beta(\frac{1}{2}) < \beta(0) = \beta(1)$ . Therefore,  $\beta(\frac{1}{2})$  is a global minimum for the function  $\beta(t)$  and the maximum occurs when  $t = 0$  or  $t = 1$ .

Thus, the maximum value of  $\angle u w v$ ,  $\beta$ , is obtained when  $w$  is placed at either one of  $x'$  or  $y'$ . Since  $d \geq 2L$  and  $\beta(0) = \beta(1)$ , the geometry of this configuration is as shown in Figure 4–5. In particular,  $\beta = \phi + \gamma$ , where  $\phi = \arcsin\left(\frac{L}{2R}\right)$  and  $\gamma = \arcsin\left(\frac{2R}{d} \sin \phi\right)$ . Thus,  $\gamma = \arcsin\left(\frac{2R}{d} \frac{L}{2R}\right) = \arcsin\left(\frac{L}{d}\right)$  and  $\beta = \arcsin\left(\frac{L}{2R}\right) + \arcsin\left(\frac{L}{d}\right)$ .

The angle  $\beta = \angle uvw$  gives an upper bound for the angle  $\alpha$ , which is twice the object angle of the missed medial point inside the convex hull of  $\Theta$ .

Thus, if  $d \geq 2L$ ,  $\alpha \leq \beta = \arcsin\left(\frac{L}{2R}\right) + \arcsin\left(\frac{L}{d}\right)$ . For a fixed  $L$  and  $R$ ,  $\alpha$  decreases as  $d$  increases. So, if one wants to find medial points with small  $d \geq 2L$ , one should choose  $L \ll d$ .




Figure 4–5: Angle  $\beta$  is an upper bound for  $\alpha$ , which is twice the object angle.

#### 4.3.2 Case 2

In this case, we consider the scenario when either  $\nabla D(a)$  is to the other side of  $l(x, y)$  as  $\nabla D(x)$  and  $\nabla D(y)$ , or  $\nabla D(b)$  is to the other side of  $l(x', y')$  as  $\nabla D(x')$  and  $\nabla D(y')$ . Without loss of generality, assume that  $\nabla D(x)$  and  $\nabla D(y)$  are to the other side of  $l(x, y)$  as is  $\nabla D(a)$ . Assume that  $l(x, y)$  is oriented horizontally and that  $\nabla D(x)$  and  $\nabla D(y)$  are above  $l(x, y)$ , while  $\nabla D(a)$  is below.

Let  $X, Y$  and  $P$  be the nearest points on  $\mathcal{B}$  to  $x, y$  and  $p$ , respectively. We know that

$$d(x, X) < d(x, P)$$

$$d(y, Y) < d(y, P)$$

$$d(p, P) \leq d(p, X)$$

$$d(p, P) \leq d(p, Y).$$

Let  $b(a, b)$  be the bisector of points  $a$  and  $b$ . Then  $x$  is above  $b(X, P)$  and  $y$  is above of  $b(Y, P)$ , while  $p$  can be on or below  $b(X, P)$ .

Recall that  $d(x, y) \leq L$ . Let  $p'$  be the point of intersection of  $(p, P)$  with  $(x, y)$ . We want to know the maximum possible value of  $d(p, p')$ . We call  $d(p, p')$  the *dip* and investigate the relationship between the dip and  $L$ .

Let  $d = d(p, P)$ . When  $d < 2L$ , then we have a bound on  $d$  as a function of  $L$ , as desired.

Consider the case  $d \geq 2L$ . Suppose that  $dip \geq L$  in this case. Since we are in Case 2,  $\nabla D(x)$  and  $\nabla D(y)$  are to a different side of  $l(x, y)$  as is  $\nabla D(a)$ . We will show that certain conditions must be met when Case 2 holds. To this aim, we will construct an extremal scenario where  $X$  and  $Y$  are maximally above  $l(x, y)$ .

Let  $x^*$  be the point of intersection of  $b(X, P)$  with  $(x, y)$  and let  $y^*$  be the point of intersection of  $b(Y, P)$  with  $(x, y)$ . To construct a scenario where  $X$  and  $Y$  are maximally above  $l(x, y)$ , we want to maximize  $\min(\angle xy^*Y, \angle yx^*X)$ . This value is largest when  $\angle xy^*Y = \angle yx^*X$ . Given  $b(Y, P)$  and  $P$ , point  $Y$  is found by dropping a perpendicular from  $P$  to  $b(Y, P)$  of length  $l$  and then traveling distance  $l$  in the direction of this perpendicular. Therefore,  $\angle xy^*Y$  is largest when the slope of the bisector  $b(Y, P)$




Figure 4-6: The angles of interest in Case 2.

is smallest, *i.e.* when  $\angle py^*x$  is smallest. A lower bound on the slope of  $b(Y, P)$  is achieved by aligning it with  $p$  and  $y$ . Similarly, we achieve a lower bound on the slope of  $b(X, P)$  by aligning  $b(X, P)$  with  $x$  and  $p$ . Since  $\angle xy^*Y = \angle yx^*X$ ,  $b(X, P)$  and  $b(Y, P)$  have the same slope, *i.e.*,  $\angle pxy = \angle pyx$ . This implies that  $d(x, p) = d(p, y)$ . Refer to Figure 4-6. Let  $b_{XP}$  be the intersection of  $b(X, P)$  with  $l(X, P)$ . Define  $b_{YP}$  similarly. By construction,  $\angle Xxb_{XP} = \angle b_{XP}xP$  and also  $\angle Pyb_{YP} = \angle b_{YP}yY$ . Since  $\angle xyY = \angle yxX$ ,  $2\angle b_{XP}xP + \angle Pxy = 2\angle b_{YP}yP + \angle Pyx$ . Combining this with the fact that  $\angle b_{XP}xP + \angle Pxy = \angle b_{YP}yP + \angle Pyx$ , we get that  $\angle b_{XP}xP = \angle b_{YP}yP$ . Hence  $\angle Pxy = \angle Pyx$  and  $d(P, y) = d(P, x)$ . Since  $d(P, y) = d(P, x)$  and  $d(x, p) = d(y, p)$ , the quadrilateral  $xpyP$  has the property that its diagonals  $(x, y)$  and  $(p, P)$  intersect at right angles. Their point of intersection is  $p'$  and  $d(x, p') = d(p', y)$ .

An upper bound for  $d(x, p') = d(p', y)$  is  $L/2$ . Setting  $d(x, p') = d(p', y)$  to the largest possible value  $L/2$  makes the slope of  $b(Y, P)$  smallest. By the arguments in the previous paragraph, when the slope of  $b(Y, P)$  is smallest, segment  $(p, p')$  is orthogonal to  $(x, y)$  and  $b(Y, P)$  passes through  $p$  and  $y$ . Now suppose that  $d = d(p, P) = 2L$ ,  $dip = d(p, p') = L$  and  $(x, y)$  is aligned with the horizontal, as shown in Figure 4-7.

[illegible]

Therefore, when  $d \geq 2L$  and  $d(p', P) \geq L$ , it is not possible that  $dip \geq L$  when we are in case 2. Thus, when  $d \geq 2L$  and  $d(p', P) \geq L$ ,  $dip < L$ . So, in summary for case 2, either  $dip < L$  or  $d < 2L$  or  $d(p', P) < L$ . In the future, when we say that medial point  $p$  has nearest points on  $\mathcal{B}$  near  $C$ , we mean that  $d(p', P)$  is small.

Our discussion of the error in cases 1 and 2 serves to prove the following theorem:

95

$R$  be the radius of the circle  $C$ . Let two of the nearest points on  $\mathcal{B}$  to  $p$  be  $A$  and  $B$ . Let  $d = d(p, A)$ ,  $\alpha = \angle ApB$ ,  $a'$  be the point of intersection of  $(p, A)$  with the convex hull of  $\Theta$ ,  $b'$  be the point of intersection of  $(p, B)$  with the convex hull of  $\Theta$ . Then either

1.  $d < 2L$ ,
2.  $\alpha \leq \arcsin\left(\frac{L}{2R}\right) + \arcsin\left(\frac{L}{d}\right)$ ,
3.  $d(p, a') < L$ ,
4.  $d(p, b') < L$ ,
5.  $d(a', A) < L$ , or
6.  $d(b', B) < L$ .

Therefore, when  $L$  is small, an undetected medial point either has small distance to the boundary, small object angle, small dip, or has nearest points on  $\mathcal{B}$  near  $C$ . In the next section, we shall evaluate the significance of these error measures and also show how algorithm DECIDEMA can be used to build an algorithm that approximates the medial axis with a dense set of medial points.

#### 4.4 Medial Axis Approximation

As discussed in Section 2.1, small perturbations of  $\mathcal{B}$  can result in significant perturbations of the medial axis. For this reason, it is common to approximate the medial axis using a stable subset of medial points. Recall that popular simplification schemes favour medial points with large object angle [7, 42, 45, 53, 12] and large distance to  $\mathcal{B}$  [7, 30, 12]. Thus, object angle and distance to  $\mathcal{B}$  of any undetected medial points are reasonable quantities to bound when evaluating quality of Algorithm DECIDEMA.

In order to use DECIDEMA to find a dense collection of medial points on the medial axis of a solid  $\Omega$ , we cover  $\Omega$  with overlapping disks of varying radii. By using overlapping

disks, we ensure that if a medial point has a small dip in one disk, it will have a bigger dip in a neighbouring disk. Recall that when a circle  $C$  is near  $\mathcal{B}$ , it is possible that DECIDEMA fails to detect medial points in  $C$ . We overcome this limitation by using disks of varying radii. If we place small disks near  $\mathcal{B}$  and if medial points fail to be detected in these, then such medial points will have small  $d$  and will be candidates for pruning. Proceeding in this way, the medial points undetected in overlapping disks of varying radii are ones that have either a small object angle or small distance to the boundary. Thus, it is possible to design an approximation method that, in 2D, provably detects those regions of space that contain medial points of large object angle and large distance to the boundary.

Finally, to locate medial points inside a circle  $C$  that is chosen by Algorithm DECIDEMA as containing medial points, we can perform binary search along appropriate line segments  $(\theta_i, \theta_j)$ , where  $\theta_i, \theta_j \in \Theta$ , to find the approximate location where a  $\nabla D$  vector changes or crosses a line between two sampled points, similar to Algorithm 1.

## 4.5 Chapter Summary and Discussion

We have presented a theoretical analysis of the soundness and completeness of an algorithm for the detection of 2D regions intersected by the medial axis of a 2D solid, based on the analysis of the gradient of its Euclidean distance transform. Such an analysis is important because methods based on the analysis of the gradient of the Euclidean distance transform are proposed in the literature [45] and used in practice to approximate the medial axis. Further, we are not aware of a comparable quantitative analysis of the quality of the medial axis approximation as a function of the density of the query points considered. Thus, we believe our analysis of the quality of the approximation in this 2D case is both useful and novel.



Nonetheless, the error bounds we have presented here are not tight: we do not have examples of missed medial points displaying the qualities described in Theorem 4.1. Therefore, there is room for improvement of the bounds presented. Further, new rules for medial point detection could produce algorithms with improved approximation quality.

In Part I of this thesis, we have presented algorithms for medial surface and medial axis approximation and have studied their correctness in terms of soundness and completeness. Our approximation of the medial surface of a 3D solid generates an alternative shape representation for the solid: that of a set of well-spaced medial spheres. The next part of this thesis will examine the quality of the shape description offered by this representation.

## **Part II**

# **Quality of Shape Description**

## Chapter 5

# Boundary Differential Geometry from Medial Geometry

In Section 1.2.1, we mentioned that the envelope of the (generically) infinite set of medial spheres of a solid reconstructs the boundary of the solid (Property 1.1). In Chapter 2, we have presented an algorithm for the approximation of the medial surface of a solid with a set of points, as well as an estimate of the spoke vectors for each point. It is natural to ask how well a shape representation based on a finite set of medial spheres and their spoke vectors describes the original solid. One natural way to describe the shape of a solid is to measure the amount of bending of its boundary. The area of differential geometry investigates such measures. In this chapter, we first recall the relevant basics of differential geometry in Section 5.1. Next, we explain how boundary curvature has been estimated when the boundary of a solid is not differentiable, but approximated with a triangle mesh or a cloud of points in Section 5.2. In Section 5.3, we explain how the medial surface transform of a solid can be used to measure the boundary curvature of the solid. In Section 5.4, we explain how this result can be used with our discrete approximation to the medial surface transform. The ability to estimate boundary curvature directly from a discrete set of medial points and their spoke vectors, without an explicit boundary reconstruction step, speaks of the quality of this shape representation. The shape representation of a union of medial spheres together with their spoke vectors is closely related to that of *m-reps*, used in medical imaging to perform statistical shape analysis [95, 106].

## 5.1 Differential Geometric Shape Operator

In this section, we will provide a brief overview of the relevant concepts from differential geometry, to be used in our subsequent development.

Consider a differentiable surface  $\mathcal{B} \subset \mathbb{R}^3$  with outward unit normal  $n$ . For a point  $x_0$  on  $\mathcal{B}$  and a tangent vector  $v$  to  $\mathcal{B}$  at  $x_0$ , define

$$S(v) = \frac{\partial n}{\partial v}. \quad (5.1)$$

Choosing an orthonormal basis  $\{v_1, v_2\}$  for the tangent plane  $T_{x_0}\mathcal{B}$  to  $\mathcal{B}$  at  $x_0$ ,

$$\frac{\partial n}{\partial v_i} = -a_{1i} \cdot v_1 - a_{2i} \cdot v_2, i = 1, 2. \quad (5.2)$$

The matrix

$$S = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \quad (5.3)$$

is the differential geometric *shape operator*.

Given a direction  $u$  tangent to  $\mathcal{B}$  at  $x_0$ , the quantity  $u^T S u$  is called the *normal curvature* at  $x_0$  in the direction  $u$ . The plane  $\pi$  spanned by  $n$  and  $u$  cuts the surface  $\mathcal{B}$  in a plane curve, called a *normal section*. A normal curvature of zero means that the normal section is flat. The sign of normal curvature is positive when the normal section bends away from  $n$  at  $x_0$  and negative if it bends towards  $n$ . The radius of the osculating circle fitted to the normal section is the reciprocal of the magnitude of the normal curvature and is called the *radius of curvature*.

The maximum and minimum values of the normal curvature are the *principal curvatures* at  $x_0$  and are the eigenvalues of  $S$ . The tangent vectors that give the two principal

curvatures are called the *principal curvature directions* and are the eigenvectors of  $S$ . The two principal curvatures together with the normal vector at  $x_0$  form an orthonormal basis at  $x_0$ .

For a more in-depth presentation of differential geometry basics, the reader is referred to the introductory text [87].

## 5.2 Discrete Differential Geometry

Due to the abundance of polygonal meshes and 3D point cloud data sets, there is great interest in designing discrete counterparts of differential geometric operators for use on these discrete data sets. Surface curvature measures on meshes help identify salient mesh features [124], are used to generate suggestive contours [41], and facilitate segmentation [77]. The new area of *discrete differential geometry* is concerned with the development of discrete geometric concepts that lead to efficient algorithms for computing differential geometry measures on discretized surfaces. We survey some existing methods for estimating the curvature tensor for triangle meshes and 3D range data. The curvature tensor is a  $3 \times 3$  matrix with eigenvalues being the two principal curvatures and zero, and eigenvectors the two principal curvature directions and the surface normal. It generalizes the shape operator for computing normal curvature in directions expressed in an arbitrary basis.

Taubin [115] finds curvature tensors on vertices of triangle meshes using estimates of normal curvature for each edge incident at a mesh vertex. Meyer *et al.* [83] offers an improvement on the work in [115] for irregular meshes by taking into account the angles in a triangulation and the areas of the Voronoi cells of vertices. Cohen-Steiner and Morvan [34] propose a formula for estimating the curvature tensor that converges when points

are sampled on a particular kind of mesh and satisfy a sampling condition. Cazals and Pouget [25] show that differential quantities evaluated for fitted polynomial surfaces of required degree converge to the true values given that a sampling condition on the boundary is met. The method of Rusinkiewicz [100] estimates the curvature tensor on a per-triangle basis by introducing linear constraints on this tensor stemming from finite differences of vertex normals. The work of Taubin [115] is extended to non-homogeneously scattered and noisy 3D range data in [66].

For a detailed review of the advances in this recent mathematical field of discrete differential geometry, the interested reader is directed towards the recent text in [20]. The detailed survey in [93] also provides an excellent introduction.

### 5.3 Boundary Geometry from Medial Geometry

Since the medial surface transform of a solid reconstructs the solid completely (Property 1.1), it is natural to ask if it is possible to estimate boundary curvature of a solid using the medial surface transform, without reconstructing the solid's boundary. To our knowledge, there has been little work that relates medial differential geometry to boundary geometry, with three exceptions. In [86], formulas are derived for the Gaussian and Mean curvatures for 3D boundaries based on derivatives along medial sheets, but this theory has not yet lead to implementations. For objects with non-branching medial topology, Yushkevich *et al.* [127] fit a single-sheet continuous medial representation (an *m-rep*) to medical image data and derive conditions to compute the implied boundary.

Recent work by differential geometer James Damon has rigorously established the relationship between medial and surface geometry. In this section, we overview some of the important results in [37] that will serve as the basis of our numerical method for the

estimation of boundary differential geometry from a sampling of medial points in Section 5.4.

If  $\mathcal{B}$  is the boundary of a 3D solid  $\Omega$  with unit outward normal field  $n$ , then, as described in the previous section, the rate of change of  $n$  along  $\mathcal{B}$  describes the curvature of  $\mathcal{B}$ . Rather than studying the rate of change of  $n$  as one moves along  $\mathcal{B}$ , consider the rate of change of  $n$  as one moves along the medial surface  $\mathcal{MS}$  of  $\Omega$ . Recall that smooth medial points are equidistant from exactly 2 points on  $\mathcal{B}$ . The vectors from smooth medial points to nearest points on  $\mathcal{B}$ , the spoke vectors, are normal to  $\mathcal{B}$ . Thus, by studying the rate of change of the spoke vectors we obtain information about the rate of change of the normals to the boundary. This observation is the basis for the mathematics we develop in this section. We now proceed to define a special shape operator on the medial surface, Damon's *radial shape operator* [37].

Consider a smooth medial point  $x_0 \in \mathcal{MS}$  and let  $U^a, U^b$  be the two spoke vectors at  $x_0$ . Let  $U_1^a = U^a / \|U^a\|, U_1^b = U^b / \|U^b\|$ . Damon defines [37] the *radial shape operator* as follows:

$$S_{rad}^a(v) = -\text{proj}_{U^a} \left( \frac{\partial U_1^a}{\partial v} \right) \text{ and} \quad (5.4)$$

$$S_{rad}^b(v) = -\text{proj}_{U^b} \left( \frac{\partial U_1^b}{\partial v} \right), \quad (5.5)$$

where  $\text{proj}_U$  denotes projection onto  $T_{x_0}\mathcal{MS}$ , the tangent plane to  $\mathcal{MS}$  at  $x_0$  along  $U$  [37]. Refer to Figure 5–1.

The projection onto the tangent plane to the medial surface is necessary because the derivative of a spoke vector may not necessarily lie in this tangent plane. The derivative of




Figure 5–1: An illustration of  $S_{rad}^a(v)$  for a given  $U^a$  (based on a figure in [37]).

the surface normal, however, does lie in the tangent plane to the surface. For this reason, no projection is required when computing the regular shape operator (see Equation 5.1).

Choosing an orthonormal basis  $\{v_1, v_2\}$  for  $T_{x_0}\mathcal{MS}$ , let  $S_{rad}^a, S_{rad}^b$  denote the matrix representation of the two radial shape operators. Then, as in the case of the standard shape operator, the eigenvectors are called the principal radial curvature directions and the eigenvalues are called the principal radial curvatures. Let  $\kappa_{ri}$  denote the principal radial curvatures. The following condition is required to ensure the smoothness of the boundary  $\mathcal{B}$ .

**Definition 5.1.** [37] *The Radial Curvature Condition is defined as follows. For all smooth medial points with distance  $r$  to the boundary  $\mathcal{B}$ ,*

$$r < \min \left\{ \frac{1}{|\kappa_{ri}|} \right\} \quad (5.6)$$

*for all positive principal radial curvatures  $\kappa_{ri}$ .*

Consider the correspondence between medial points and their closest boundary points given by

$$\psi^a : \mathcal{MS} \rightarrow \mathcal{B}, \psi^b : \mathcal{MS} \rightarrow \mathcal{B}.$$



Denote  $\psi^a(x_0)$  by  $x_0^a$  and  $\psi^b(x_0)$  by  $x_0^b$ . It is shown in [37] that:

**Theorem 5.1.** *Let  $x_0^a$  be a smooth medial point of an object with a smooth boundary  $\mathcal{B}$  that satisfies the Radial Curvature Condition and let  $r$  be the distance from  $x_0^a$  to  $\mathcal{B}$ . The principal curvatures  $\kappa_i^a$  of  $\mathcal{B}$  at  $x_0^a$  and the principal radial curvatures  $\kappa_{ri}$  of  $\mathcal{MS}$  at a smooth medial point  $x_0$  have the following relationship:*

$$\kappa_i^a = -\frac{\kappa_{ri}}{1 - r\kappa_{ri}} \quad (5.7)$$

where  $r = \|U^a\|$ . Furthermore, the principal radial curvature directions corresponding to  $\kappa_{ri}$  can be found by applying the map  $\psi^a$  to the principal curvature directions corresponding to  $\kappa_i^a$ . The case of  $S_{rad}^b$  is symmetric.

Thus, computing the radial shape operator at a smooth medial point of a solid with a smooth boundary, one can find principal curvatures and principal curvature directions on the associated locations on the solid's boundary.

Note that the shape operator of the medial surface needs not be explicitly computed. Damon [37] shows how the radial shape operator may be computed in terms of the shape operator for the medial surface, allowing one to express the boundary shape operator in terms of the shape operator of the medial surface. However, this formulation is needlessly more complex than the formulation of the boundary shape operator in terms of the radial shape operator.

## 5.4 Boundary Geometry from Sampled Medial Geometry

In the same spirit as the body of work on extending continuous mathematics for differentiable surfaces to be used with non-differentiable surfaces (see Section 5.2), we adapt the continuous formulations presented in the previous section to the situation where the

medial surface is approximated by a medial point cloud, and a pair of spoke vector estimates is given at each medial point.

This medial point cloud is sampled near the medial surface, which is a “branched manifold” [37]: a set of manifolds with boundary. Fitting a suitable interpolating surface to such a point set is, to the best of our knowledge, a challenging open problem in the domain of surface reconstruction. Rather than performing computations on a surface that interpolates the medial points, we will work with the set of medial points directly.

Recall that we can estimate the normal to the medial surface at an approximate medial point by considering the difference of its two estimated spoke vectors, as described in Section 2.5. We also compute surface normals to the approximate medial points using the Voronoi-based method in [42]. These normal estimates have consistent signs and we use these estimates to disambiguate between the  $U^a$  and  $U^b$  spoke vectors. Further, the surface normals computed using the method in [42] are set to zero near medial sheet junctions and boundaries, where normal estimation is ill-defined. Performing numerical differentiation for points located near medial sheet junctions or boundaries is not desirable as these are special cases not described by the mathematics we have presented in Section 5.3. We do not apply our numerical method to points whose normals cannot be unambiguously computed using the method in [42].

Our dense collection of approximate medial points with a surface normal estimate attributed to each point is actually a type of shape representation in its own right. The area of *point-based graphics* studies methods to process such data sets [64]. The primitive element considered is a *surfel*, for “surface element”, defined as a point with normal and disk-radius attributes, and an optional colour value. A surfel is basically an oriented disk.

Surfaces can be rendered using surfels without requiring an explicit surface reconstruction step. Rendering tools from the domain of point-based graphics are used to visualize our medial points.

In this section, we describe a numerical method by which boundary curvature may be inferred from a collection of approximate medial points and their spoke vector estimates. Starting with a point-based approximation of the medial surface computed using the method in Chapter 2, we now explain the steps we take to estimate principal curvatures and principal curvature directions on the boundary of a polyhedral solid.

#### 5.4.1 Imposing Smoothness on the Boundary

We are considering a surface  $\mathcal{B}$ , which is the boundary of a polyhedron and is piecewise planar. In order to satisfy the Radial Curvature Condition (Definition 5.1) at smooth medial points, we have to ensure that  $r < \min \left\{ \frac{1}{\kappa_{ri}} \right\}$  for positive values of  $\kappa_{ri}$ . Non-smooth objects may not satisfy the Radial Curvature Condition. For such objects,  $r \geq \min \left\{ \frac{1}{|\kappa_{ri}|} \right\}$  for some medial points. When  $r > \min \left\{ \frac{1}{\kappa_{ri}} \right\}$  for a positive  $\kappa_{ri}$ ,  $\mathcal{B}$  locally self-intersects. When  $r = \min \left\{ \frac{1}{\kappa_{ri}} \right\}$  for a positive  $\kappa_{ri}$ ,  $\mathcal{B}$  has a non-differentiable concavity, *e.g.*, a concave edge or vertex of a polyhedron. When  $\mathcal{B}$  is the boundary of a non-convex polyhedron, the concave edges and vertices may be the nearest boundary points to points sampled on the medial surface of the polyhedron. Moreover, several sampled medial points may have the same nearest boundary point. Figure 5–2 illustrates a scenario where two nearby medial points share the same nearest boundary point on a concave vertex of a piecewise linear boundary  $\mathcal{B}$ . It can be easily verified that in this case, indeed,  $r = \frac{1}{\kappa_{ri}}$ . In order to apply the theory of Section 5.3 to the case where the medial surface is computed for a non-convex polyhedron, a preliminary smoothing step is required.

When working with approximate medial points computed for non-convex polyhedra, we impose smoothness on the two spoke vector estimates  $U^a$  and  $U^b$  of an approximate medial point  $x_0$  as follows.

1. We consider the approximate medial points located within a small radius of  $x_0$ . Some of these neighbouring medial points may lie near a different medial sheet than  $x_0$ . However, since we are considering a small neighbourhood of  $x_0$ , those medial points that lie on a different medial sheet than  $x_0$  will have a different normal to the medial surface than  $x_0$ . We restrict our attention to those neighbouring medial points whose normals differ by less than a constant  $\gamma = 0.8$  radians.
2. We then replace the two spoke vector estimates at  $x_0$  with a local Gaussian weighted average of its neighbours' spoke vectors. Specifically, the weights of the contribution of vectors at medial point  $x_1$  to vectors at medial point  $x_0$  are given by the Gaussian function  $\frac{1}{(2\pi)^{3/2}\sigma} e^{-\frac{\|x_0 - x_1\|_2^2}{2\sigma^2}}$ , with  $\sigma = 0.25$ .

This procedure corresponds to smoothing the boundary implied by the approximate medial points and their spoke vector estimates. The use of this heuristic prohibits incorrectly computing very large curvatures at concave edges and vertices of the polyhedron's boundary. We only perform this smoothing step for non-convex polyhedra.




Figure 5-2: The case when  $\kappa_{ri} = -\text{proj}_U \frac{\partial U_1}{\partial v_1} = \frac{1}{r}$ .

### 5.4.2 Derivatives on Medial Sheets

For an approximate medial point  $x_0$ , we next measure the rate of change of the vectors  $U^a(x_0)$  and  $U^b(x_0)$  in the directions of an estimated orthonormal basis to the tangent plane at  $x_0$ ,  $T_{x_0}\mathcal{MS}$ , by studying the estimated spoke vectors at the approximate medial points in the vicinity of  $x_0$  that have similar normals as  $x_0$  (we considered two unit normals as similar if their dot product was greater than 0.8).

Let  $v$  be a direction orthogonal to  $N_{\mathcal{MS}}(x_0)$ , the normal estimate to  $x_0$ . We find the nearest point  $x_1$  to  $x_0 + \Delta \frac{v}{\|v\|}$ , for a given step size  $\Delta$ . We then check that  $x_1$  lies within distance  $\Delta \pm \epsilon_1 \Delta$  from  $x_0$  (we use  $\epsilon_1 = 0.2$ ). In order to ensure that  $x_1$  lies near the same medial sheet as  $x_0$ , we also check that the angle between normal estimates to  $x_1$  and  $x_0$  is less than  $\gamma$ . If  $x_0 + \Delta \frac{v}{\|v\|}$  lies near a medial sheet boundary or junction, then we may not find another approximate medial point  $x_1$  near the same medial sheet and also a desired distance away from  $x_0$ . To address this issue, when the neighbour  $x_1$  cannot be found, we try reversing the direction  $v_1$  and decreasing the step size  $\Delta$ . If  $x_1$  is found, we consider the spoke vector estimates at  $x_0$  and  $x_1$ , normalized to unit length:  $U_1^a(x_0)$ ,  $U_1^b(x_0)$  and  $U_1^a(x_1)$ ,  $U_1^b(x_1)$ . We set the vector

$$v_1 = \frac{x_1 - x_0}{\|x_1 - x_0\|}. \quad (5.8)$$

We then estimate the partial derivatives of the unit spoke vectors in the direction  $v_1$ .

$$\frac{\partial U_1^a}{\partial v_1} = \frac{U_1^a(x_1) - U_1^a(x_0)}{\|x_1 - x_0\|}, \quad (5.9)$$

$$\frac{\partial U_1^b}{\partial v_1} = \frac{U_1^b(x_1) - U_1^b(x_0)}{\|x_1 - x_0\|}. \quad (5.10)$$

Let

$$v_2 = \frac{N_{\mathcal{MS}}(x_0) \times v_1}{\|N_{\mathcal{MS}}(x_0) \times v_1\|}, \quad (5.11)$$

such that  $v_1, v_2$  and  $N_{\mathcal{MS}}(x_0)$  are orthonormal. Consider the point  $x_2 = x_0 + \|x_1 - x_0\|v_2$ . As it is unlikely to find an approximate medial point at  $x_2$ , we estimate the  $U^a$  and  $U^b$  directions at  $x_2$  as a Gaussian weighted average of the normalized spoke vector estimates of the approximate medial points within a distance of  $\epsilon_2$  of  $x_2$  ( $\epsilon_2$  was set to  $\Delta/4$ ). Once again, to restrict our attention to the approximate medial points on the same medial sheet, we only consider those points that have similar normals as  $N_{\mathcal{MS}}(x_0)$ , when such points are available. We then estimate the partial derivatives of the unit spoke vectors in the direction  $v_1$ . Then the partial derivatives of the unit spoke vectors in the direction  $v_2$  are estimated as:

$$\frac{\partial U_2^a}{\partial v_2} = \frac{U_2^a(x_2) - U_1^a(x_0)}{\|x_2 - x_0\|}, \quad (5.12)$$

$$\frac{\partial U_2^b}{\partial v_2} = \frac{U_2^b(x_2) - U_1^b(x_0)}{\|x_2 - x_0\|}. \quad (5.13)$$

### 5.4.3 Estimating Surface Curvature

When it is possible to compute  $\frac{\partial U_1^a}{\partial v}, \frac{\partial U_1^b}{\partial v}$  estimates for  $v \in \{v_1, v_2\}$ , we proceed to estimate the  $S_{rad}$  operator at approximate medial points using Equations 5.4 and 5.5. We then use its eigenvalues to find the principal curvature estimates  $\kappa_i^a, \kappa_i^b, i = 1, 2$  at the two estimated points of tangency on  $\mathcal{B}$  of the medial spheres corresponding to the approximate medial points, using Equation 5.7.

If we would like to estimate the curvature of the polygonal mesh boundary  $\mathcal{B}$ , we are presented with the obstacle that the set of nearest boundary points to the set of medial

points we have computed does not include points on certain regions of  $\mathcal{B}$ . For a medial sphere that contributes a large spherical patch to the union of medial spheres, we only have curvature information about its two nearest boundary points (refer to Figure 5–3 for a 2D illustration). Because we do not consider any low-object angle medial spheres, certain medial spheres that lie near boundaries of medial sheets may contribute a large sphere patch to the union of the medial spheres. In particular, large parts of the polyhedron’s boundary that are nearly spherical may be only approximated with a single medial sphere patch for which we can only compute surface curvature values at two points.

We now describe the method we use to assign boundary curvature values to the boundary mesh vertices  $V$  based on our estimates of boundary curvature at points  $P$  on the boundary mesh. We first compute the Voronoi diagram of the mesh vertices  $V$ . For each mesh vertex  $v \in V$ , we set each of the principal curvatures of  $v$  to the mean of the principal curvature estimates of the points in  $P$  that lie in  $v$ ’s Voronoi cell. Those vertices  $v' \in V$  that have not been assigned any curvature values, but that are adjacent in a triangulation of the polyhedron’s boundary  $\mathcal{B}$  to vertices that have been assigned curvature values, are assigned the mean of their neighbours’ principal curvatures. This process continues until all vertices in  $V$  are assigned principal curvature values, and essentially propagates boundary curvature values from areas where these values are known to those where they are unknown.

Although effective at assigning curvature values to all polyhedron vertices based on known curvature values of a sampling of points on the surface of a polyhedron, this strategy has limitations. When considering surfaces like a slightly flattened capped cylinder, where one of the principal curvatures changes suddenly from zero to a large positive value,

it may happen that all the approximate medial points we consider have one surface principal curvature of zero. By using this method, the cap of the cylinder will be wrongly assigned a principal curvature of zero.




Figure 5–3: The medial axis of this object consists of low-object angle segments (dashed lines) and a high-object angle segment (bold line). When approximating the medial axis with a set of points, we only retain points on high-object angle segments of the medial axis. The boundary of the union of the associated medial circles approximates the original object. Medial point  $m$  is equidistant from points  $A$  and  $B$  on the boundary. Surface curvatures at points  $A$  and  $B$  may be found using the radial shape operator, while for point  $C$  a different strategy is used.

#### 5.4.4 Sources of Error

We have described a numerical scheme that estimates surface principal curvatures from a set of medial points of a polyhedron, and their spoke vectors, that requires the following assumptions to hold in order to be effective in estimating the principal curvatures on the boundary of a smooth solid approximated by the polyhedron:

1. The polyhedron's boundary is of sufficiently high resolution and is sufficiently smooth such that the variation in neighbouring medial points' spoke vectors captures the rate of change of the surface normals of the smooth solid approximated by the polyhedron.



2. The polyhedron's boundary is of sufficiently high resolution and is sufficiently smooth such that high object angle medial points lie near the medial surface of the smooth solid approximated by the polyhedron.
3. The difference in normal estimates to neighbouring approximate medial points serves as an indication of whether they are sampled near the same medial sheet.
4. The medial surface is non-degenerate, *i.e.*, it consists of sheets, not curves, and the medial sheets are sufficiently wide.
5. The approximation density of the medial surface is fine enough such that approximate medial points can be locally well-fitted with a plane.

Note that items 1, 3, 4 and 5 imply that the step size for the derivatives needs to be big enough to capture variability in the polyhedral surface, while also small enough so that the medial surface can be locally approximated with a planar patch. At the same time, the step size for the derivatives should not be smaller than the maximum distance between two points sampled in adjacent voxels of side length  $\sigma$ ,  $2\sqrt{3}\sigma$ .

When we are able to estimate the principal radial curvatures  $\kappa_{ri}$  at an approximate medial point  $m$  with distance  $r$  to the boundary, we use Equation 5.7 to compute principal curvatures at the two nearest points to  $m$  on the object boundary. We will now show how errors in the estimation of  $\kappa_{ri}$  are reflected in errors in the radius of curvature estimates.

**Lemma 5.1.** *Let  $\kappa'_{ri}$  denote the value for a principal radial curvature at  $m$  computed with our numerical method. Let  $\kappa_{ri}$  be the true principal radial curvature at  $m$ . Suppose that*

*$\frac{\kappa_{ri}}{\kappa'_{ri}} = c$  for some constant  $c$ . Then*

$$\left| \frac{1}{\kappa'_i} - \frac{1}{\kappa_i} \right| = \left| \frac{1-c}{\kappa_{ri}} \right|. \quad (5.14)$$

*Proof.* Equation 5.7 can be rewritten as  $\frac{1}{\kappa_i} = r - \frac{1}{\kappa_{ri}}$ . Then  $\frac{1}{\kappa'_i} = r - \frac{1}{\kappa'_{ri}}$ . Let the error of the approximation of  $\kappa_i$  with  $\kappa'_i$  be the difference of the two radii of curvature,  $\left| \frac{1}{\kappa'_i} - \frac{1}{\kappa_i} \right|$ . Then  $\left| \frac{1}{\kappa'_i} - \frac{1}{\kappa_i} \right| = \left| r - \frac{c}{\kappa_{ri}} - r + \frac{1}{\kappa_{ri}} \right| = \left| \frac{1-c}{\kappa_{ri}} \right|$ .  $\square$

Therefore, for small values of  $|\kappa_{ri}|$ , the ratio  $c = \frac{\kappa_{ri}}{\kappa'_{ri}}$  needs to be close to 1 in order to ensure that the radius of curvature,  $\frac{1}{\kappa_i}$ , is well estimated.

### 5.4.5 Experimental Results and Analysis

We now present numerical results of principal curvature estimation for polyhedra of varying topology and geometric complexity, and with varying medial surface branching topology.<sup>1</sup> To show surface curvatures, we use the colormap in Figure 5–4 (bottom right). Here, red corresponds to a convexity, blue to a concavity, green to a saddle-shaped region, yellow to a cylindrical patch curving toward the object (the non-zero principal curvature is positive), cyan to a cylindrical patch curving away from the object (the non-zero principal curvature is negative) and white to a flat region. Figure 5–4 (left) shows the approximate medial points coloured according to surface curvature estimates and Figure 5–4 (middle) shows these estimates projected onto the boundary  $\mathcal{B}$ . For comparison, Figure 5–4 (right) shows the result of applying Rusinkiewicz’s method for estimating boundary principal curvatures [100] directly on  $\mathcal{B}$ .

Although there are subtle numerical differences, the results obtained by the two methods are qualitatively very consistent. We emphasize that the implementation of our method

---

<sup>1</sup> Models from the Princeton Shape Benchmark, <http://vcg.isti.cnr.it/polycubemaps/models/>, and <http://www.cs.princeton.edu/gfx/proj/sugcon/models/>

works at a fixed spatial resolution on the set of medial points, whereas the method in [100] is aided by the explicit representation of the surface discontinuities of the mesh boundary and the connectivity of the mesh. The majority of surface regions recovered by our method are correctly coloured. Examples include the holes and the sides of the torus; the stem, neck and base of the pear; the eye sockets, neck, chin and nose of the head; and the neck, shoulders, breasts, chest, thighs, and navel of the Venus model. We note that the union of medial spheres is a slightly different object than the original polyhedral object, as object angle simplification removes some small-scale boundary details. This fact should be taken into account when comparing the two different curvature estimates in Figure 5–4.

In order to demonstrate the validity of the principal curvature direction estimates from medial geometry, Figure 5–5 illustrates these on a model of a cup. As explained in Section 5.3, these estimates are obtained from the eigenvectors of  $S_{rad}$ . Note how the directions are orthogonal and correspond to the directions of maximal bending (red) and minimal bending (blue).

Since ground truth curvatures are not available for the polyhedra in Figure 5–4, a quantitative evaluation of the quality of our curvature estimates is not possible. However, for the simple example of a portion of a cylindrical cup wall, such that one side of the wall has positive and zero principal curvatures, while the other side has negative and zero principal curvatures, such a numerical comparison is feasible. Based on estimates of exterior and interior radii of the cylinders defining the cup walls, we obtain an estimate for the positive and negative principal curvatures, respectively. The distribution of the curvatures estimated with our method and that of [100] is shown in Table 5–1. We observe that our relative error is 2.2% for the positive curvature and is 0.4% for the negative curvature.




Figure 5–4: Left: Approximate medial points coloured according to principal curvatures on the solid’s boundary. Centre: Projection of curvature values at the approximate medial points to the boundary of the solid. Right: Curvature values obtained using the method in [100]. The colourmap used is shown in the bottom right corner. See the associated text for a discussion of these results.




Figure 5–5: The principal curvature directions recovered on the surface of a cup.

	Positive	Negative	Zero
$\mu$	0.069834	-0.087302	$1.205 \times 10^{-6}$
$\sigma$	0.000101	0.000148	$8.707 \times 10^{-6}$
$\mu$	0.068380	-0.087144	$-8.973 \times 10^{-7}$
$\sigma$	0.003784	0.002784	$1.428 \times 10^{-5}$

Table 5–1: Mean ( $\mu$ ) and standard deviation ( $\sigma$ ) of the 3 different types of curvature on a section of a cylindrical cup wall (shown in Figure 5–5) consisting of 551 vertices obtained using the method of [100] (Top) and our method (Bottom). We estimate the true positive curvature to be approximately 0.069898 and the true negative curvature to be approximately  $-0.0875045$ .

Figure 5–6 presents results for another quantitative evaluation of our method. For the oblate spheroid considered, we compare the true principal curvatures to our principal curvature estimates. The finite difference approach to estimating the rate of change of the spoke vectors produces exact results when the spoke vector direction changes linearly over the neighbourhood, *i.e.*, the boundary is round, like in the case of the cup model. For the spheroid, the boundary is not round and the error introduced by the approximation is expected to be larger.

It has been observed experimentally that increasing the number of medial points computed (by increasing the voxel resolution), and subdividing and smoothing the boundary of the polyhedra, greatly improves the performance of our algorithm in terms of the quality of principal curvature estimates, as expected.

#### **5.4.6 Discussion**

The quality of our numerical method for boundary curvature estimation from a dense set of approximate medial points of a polyhedron and their spoke vector estimates is not intended to rival that of the discrete differential geometry methods surveyed in Section 5.2. Discrete differential geometry methods face the challenge of estimating boundary curvature of a smooth solid by measuring local variation in surface normals to a non-smooth approximation to the smooth object boundary. Our method also measures the local variation in surface normals to a non-smooth approximation to a smooth solid, but moreover, it must also correctly estimate the tangent plane to a discrete set of points near the medial surface. As we approximate the medial surface of a smooth solid with a sampling of points that lie near the medial surface of an approximation of that solid, and because the medial surface is a branched manifold with boundary, correctly estimating the tangent plane to

this collection of medial points is a non-trivial task. This extra layer of discretization and approximation makes our method less attractive for computing boundary curvature than the existing discrete differential geometry methods that work directly on the object boundary, when this boundary is available.

However, precisely in the scenario when the boundary is not available, *i.e.*, when the shape representation of the solid is a collection of medial points and their spoke vectors, we have shown that we can recover qualitatively meaningful measures of boundary curvature using this representation alone. Therefore, the shape representation of a solid using a set of well-spaced medial spheres, and additionally spoke vectors for each sphere, captures curvature information of the solid’s boundary. In Section 5.5, we will discuss how the spoke vectors can be estimated from the union of medial spheres.

Given a union of spheres representation of a solid, one may ask if it is possible to estimate the boundary curvature of the solid by using existing discrete differential geometry methods on a polygonal mesh approximation to the envelope of spheres. Numerous studies have been devoted to the computation of the envelope of spheres. Tam and Heidrich [113] mesh a dense union of spheres by connecting the vertices of the envelope of spheres. Where smoothness of the boundary is important, Edelsbrunner [48] defines the *skin surface* of a set of weighted points which is tangent-continuous and defined with respect to a single shrinking parameter. Kruithof and Vegter [75] extend the work of Edelsbrunner to permit a better interpolation of spheres by introducing local quality parameters. Given our medial spheres approximation to a solid, we can compute the envelope of spheres using, for example, the method of [48] with a zero shrinking parameter to produce a mesh of the envelope of spheres, whose implementation is available in the CGAL library [1].

However, existing discrete differential geometry methods interpret surface discontinuities as an indication of a change of curvature, and vertices in the union of balls envelope will be interpreted as dimples in the surface. Therefore, it is not clear how useful boundary curvature measures can be estimated from the envelope of a union of spheres.

## 5.5 Chapter Summary and Future Work

We described a method that starts with a union of medial spheres of a solid, together with the spoke vectors for each sphere, and estimates the boundary shape operator at the spoke vector tips located on the boundary of the solid.

As we would like a shape representation to be as minimal as possible, we ask if it is possible to estimate the spoke vectors for each medial sphere from the contribution of each sphere to the envelope of the union of medial spheres. Consider the medial surface of a union of medial balls corresponding to these medial spheres. Those medial balls that are internal to medial sheets in this medial surface contribute two disjoint regions to the surface of the union of the medial balls. Using the power diagram of the medial balls, it is possible to compute these regions for each medial sphere. A single point chosen appropriately on each of these two regions can be used as an estimate of each of the two spoke vector tips of a medial sphere.

Consider the following research direction:

*Given a union of medial spheres, compute an estimate for the spoke vectors for each medial sphere, when possible. Use these estimates to, in turn, provide an estimate of the boundary differential geometry of the solid represented using this union of medial spheres, using tools developed in this chapter.*



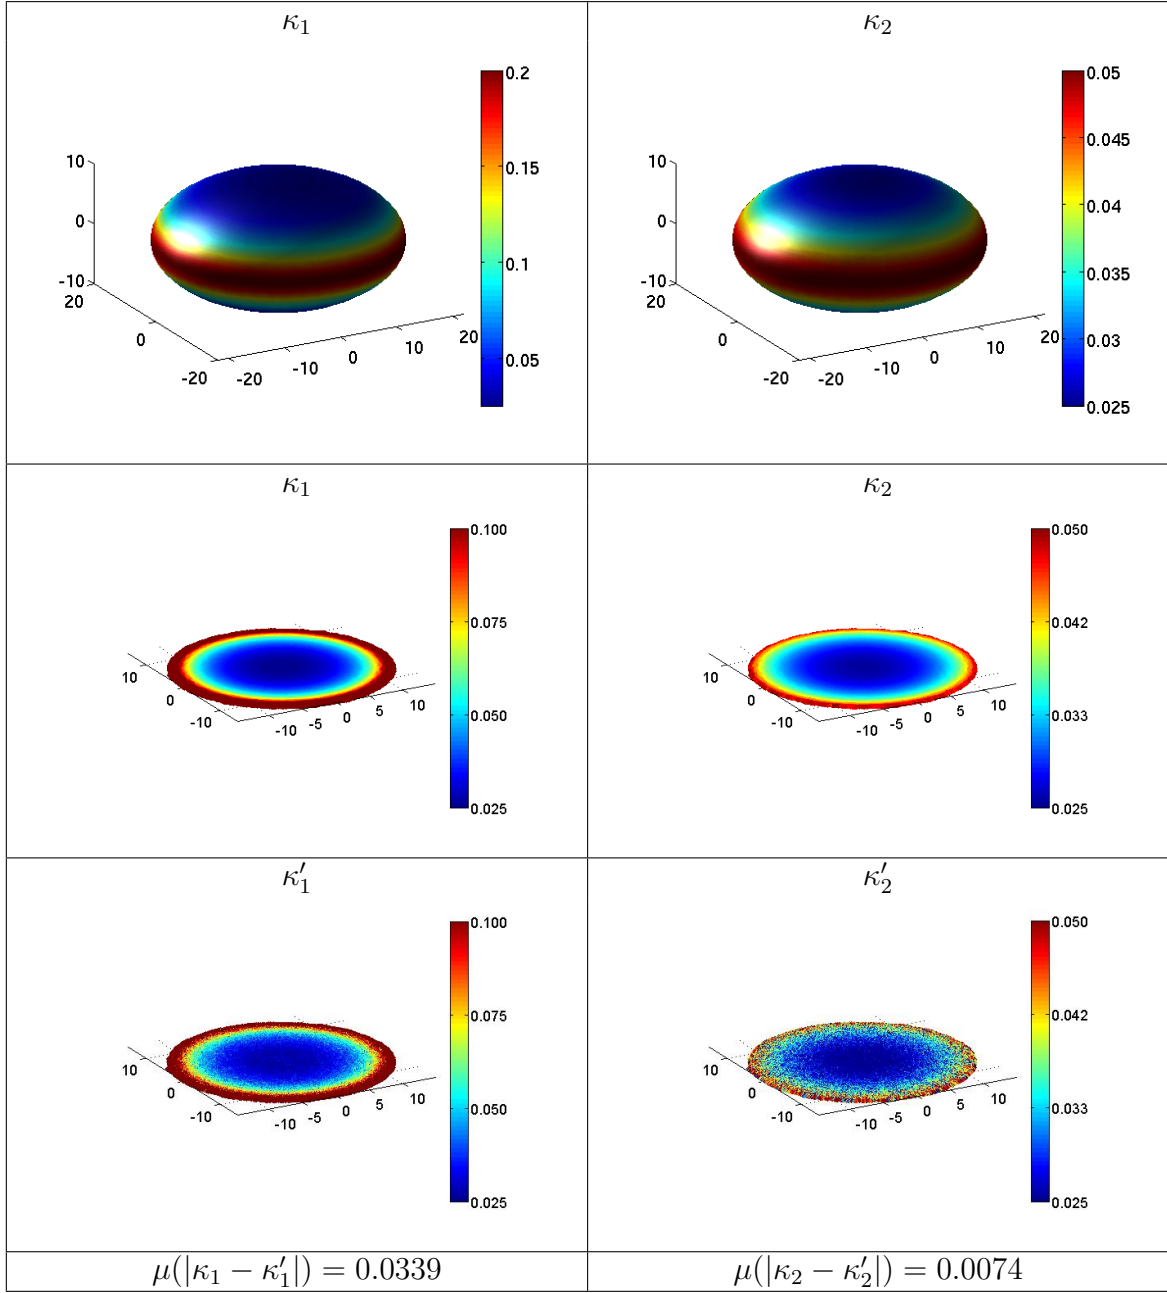


Figure 5–6: Top: An oblate spheroid coloured by principal boundary curvature  $\kappa_1$  (left) and  $\kappa_2$  (right). Middle: Points near the medial surface of the oblate spheroid coloured by principal boundary curvature. Bottom: Respective principal curvature estimates  $\kappa'_1$  and  $\kappa'_2$  shown on the medial surface. The mean absolute error of the estimation of  $\kappa_1$  and  $\kappa_2$  is shown below each column.

## Chapter 6

# A Significance Measure for Medial Sheets

As we saw in the previous chapter, the set of well-spaced medial spheres shape representation of a solid allows one to estimate principal curvatures and principal curvature directions on the boundary of the solid. In this chapter, we will show how a significance measure can be assigned to the individual parts of the union of medial spheres, namely, those spheres that have been identified as belonging to distinct medial sheets. We propose a measure of significance based on the volume of spheres restricted to their power cells. By implementing the simplification according to the significance measure proposed, we find that spheres belonging to a small number of medial sheets often reconstruct the majority of the volume of the original union of spheres. We will show that when the shape of a solid is represented using a union of medial spheres, the individual parts of the representation can be easily ordered by significance; further, this significance measure can be used to reduce the number of parts in the representation.

### 6.1 Related Work

Several algorithms assign medial sheets a significance measure in order to either remove noisy, unstable sheets of the medial surface or to guide matching. We will give an overview of methods that use a volumetric significance measure.

For applications of medial surface simplification, Styner *et al.* weigh sheets using the ratio of the volume of the solid with the sheet removed and the volume of the

solid [109]. However, this operation is reported to be expensive [109]. Tam and Heidrich [112] weigh sheets using the volume of the removed feature if the sheet is pruned. Chang and Kimia [26] measure the change in volume associated with a sheet removal divided by the average medial sphere radius. These significance measures are appropriate for pruning outer sheets for the purpose of preserving homotopy equivalence between the solid and its medial surface. However, overlap between spheres in the removable sheets can present problems. Consider, for example, Figure 6–1(a). If the pink and blue sheets are removable, then the purple volume (intersection of pink and blue spheres, but outside the orange spheres) is not taken into account by such significance measures. However, if we do choose to remove both of these sheets, the purple area will be removed. We propose an alternative volumetric significance measure that can take this overlapping region into account when computing a significance measure of spheres on a medial sheet.

For the application of shape matching, Chang and Kimia [27] consider the sum of the radii of spheres centred on medial sheet junctions as an approximate volumetric measure to guide matching. Siddiqi *et al.* [107] consider the significance of a medial sheet to be the fraction of the total number of voxels reconstructed uniquely by a medial sheet. Our significance measure should also prove useful for shape matching.

## 6.2 Significance Measure

A measure of significance of a medial sheet should reflect the role the sheet plays in the representation of the shape with respect to the other sheets. Also, this measure may need to be independent of the order in which sheets are processed. We develop such a measure based on the volume contribution of the medial balls belonging to a medial sheet to the volume of the solid represented as a union of medial balls. The volume contribution

of a sheet is also related to the object angle of medial points on that sheet, another sheet simplification measure [110].

Let  $B$  be a finite subset of the medial balls of a 3D solid  $\Omega$ . For a ball  $b_i \in B$ , consider its power cell  $P(b_i)$  in the power diagram of  $B$  (see Section 1.2.2). Then the significance of a ball  $b_i$  in the union the balls in  $B$ ,  $\lambda(b_i)$ , is given by the volume contribution of  $b_i$  restricted to its power cell, relative to the total volume of the union of balls:

$$\lambda(b_i) = \frac{\text{vol}(b_i|P(b_i))}{\text{vol}(\bigcup_i b_i)}. \quad (6.1)$$

The volume of a ball  $b_i$  restricted to its power cell can be computed using short inclusion-exclusion formulae by analyzing the simplices of the dual complex of the set of balls [47]. These formulas have been carefully implemented in the software package *AlphaBall*.<sup>1</sup> An alternative formulation of the volume of the restriction of a ball to its power cell based on volumes of generalized pyramids is given in [24].<sup>2</sup>

Let  $T \subseteq B$  be a set of medial balls whose centres lie on the same medial sheet. The significance of  $T = \{t_i \in B, i = 1, \dots, |T|\}$ ,  $\lambda(T)$ , is given by

$$\lambda(T) = \sum_{i=0}^{|T|} \lambda(t_i). \quad (6.2)$$

---

<sup>1</sup> This software was generously shared with us by Professor Patrice Koehl of UC Davis and is not in the public domain.

<sup>2</sup> An implementation is available at: <http://cgal.inria.fr/abs/Vorlume/>




Figure 6–1: (a) A union of the set of balls corresponding to three medial sheets decomposed using the power diagram; (b) the dual of the decomposition of the set of balls.

Figure 6–1 shows 7 balls belonging to three medial sheets and their dual complex (spheres in different sheets are shown in different colours). The significance measure in Equation 6.2 assigns each sheet a measure corresponding to its contribution to the union of balls.

### 6.3 Experimental Results and Discussion

Given this significance measure, we consider the subsets of balls in  $B$ , belonging to distinct medial sheets, in descending order of significance. As more sheets are added to the approximation of the shape of a solid, we can compute the degree to which  $B' \subset B$  approximates the union of the balls in  $B$  by evaluating the ratio  $\text{vol}(B')/\text{vol}(B)$ . Figures 6–2 and 6–3 show an ordering of the medial sheets of a pear and a head model and shows the percentage of the volume of the union of balls represented by subsets of balls considered. We observe that it is possible to reconstruct the majority of the volume of the original union of balls using a small fraction of the original medial sheets.




Figure 6-2: A pear model approximated using a progressively larger number of sheets of medial balls. Most of the shape is covered using 4 sheets only.

The computation of the volume of a ball restricted to its power cell is performed very quickly using the highly accurate software *AlphaBall*. The software takes approximately 5 seconds to compute this quantity for 20,000 spheres for a single 3.6 GHz Pentium IV processor with 3 GB of RAM.

Because the medial surface usually consists of a large number of sheets, the use of the full medial surface is challenging for applications that consider the individual medial sheets, such as shape segmentation, animating deformations, and shape matching. Our approach to simplifying the medial representation takes advantage of the combinatorial structure of a union of balls to quickly produce a more manageable representation in terms of the number of parts.

## 6.4 Chapter Summary and Future Work

In this chapter, we have explained how shape representations based on a union of medial balls can be simplified by assigning a significance measure to each set of balls that belongs to a distinct medial sheet. The significance measure we have proposed is




Figure 6–3: A head model approximated using a progressively larger number of sheets of medial balls. Most of the shape is covered using 6 sheets only.

independent of the ordering of sheets and can be evaluated accurately and quickly. We now present several avenues for future work.

#### 6.4.1 Homotopy-preserving Medial Surface Simplification

Since the simplification scheme we have described may remove internal sheets prior to external ones, it does not necessarily guarantee homotopy equivalence between the simplified set of spheres and the original set of spheres. By identifying sheets whose removal preserves the homotopy equivalence, the order of removal of these sheets can be based on the volumetric significance measures we have proposed.

Consider the following research direction:

*Design an algorithm for the simplification of a set of medial spheres such that the simplified union of spheres is homotopy equivalent to the union of the original set of spheres, where the choice of sheet to be removed is guided by significance measures presented in Section 6.2.*

### 6.4.2 Shape Indexing and Retrieval using Medial Surfaces

The medial surface transform has been proposed as a powerful tool for shape matching [106, 107, 27]. The significance measure presented here is an attractive choice for comparing medial sheets. The decomposition of the union of medial spheres into regions corresponding to individual sheets using the power diagram is especially appropriate for matching, as the contribution of each sheet should be independent of the order in which sheets are processed.

A matching algorithm based on the medial surface may use the graph-structure corresponding to the adjacency information of medial sheets to guide matching. However, the adjacency of medial sheets in the medial surface is inherently unstable as it changes under small boundary deformations [59, 21]. It is known that even smoothing of object boundaries may introduce new medial sheets [103]. As an example of addressing this instability, Chang and Kimia [26] replace medial sheets with the nearest degenerate sheet configuration.

In order for medial surfaces to be useful for matching, a regularization step is called for such that “similar, within category 3D shapes yield similar [medial surfaces] that are distinct from the non-category shapes” [26].

Consider the following research direction:

*Capitalize on advances in discrete and computational geometry, that facilitate efficient processing of unions of spheres, to design a matching algorithm for unions of medial spheres, where the adjacency of medial sheets is regularized.*



## **Part III**

# **Fast and Tight Shape Approximation**

## Chapter 7

# Fast and Tight Spheres

In Part I, we explained how, given a boundary representation of a solid, an alternative representation of the solid using a set of well-spaced spheres can be generated. Part III will examine the utility of this set of spheres in offering a tight approximation to a polyhedral solid. Sphere-based shape representations have numerous applications to problems in computer graphics and solid modeling, such as shape matching, generating mesh deformations, shape morphing, shape segmentation, approximating Minkowski sums, point location, proximity queries, and shadow rendering. In this chapter, we investigate if our shape representation is useful for the fast and tight approximation of a polyhedron with a small number of spheres, as is desired in the above applications, and show how our spheres may be updated as the polyhedron deforms.

Finding the minimum number of spheres that cover a given set of points on a solid's boundary is NP-hard by reduction from Set Cover [4], which is shown to be NP-hard in [70]. Finding the minimum number of spheres that provides a good-quality volumetric fit to a solid is no easier, and heuristic sphere distributions are used in practice.

We compare our method of approximating a polyhedron with spheres proposed in Chapter 2 to a state-of-the-art method for approximating solids with spheres. In Section 7.3, we show that the heuristic for generating spheres described in Chapter 2 is significantly faster and provides a tighter fit in terms of volumetric error than this state-of-the-art

method. When a polyhedron undergoes local feature size preserving deformation, we show how the sphere approximation can be quickly updated and show how the volumetric error of the new sphere sets can be evaluated in Section 7.4.

## 7.1 Generating Spheres: Previous Work

When approximating a solid with spheres, a popular strategy is to start with a subset of the spheres of the medial surface transform of the solid. A useful tool for generating such a set of spheres is the Voronoi Diagram of a point set, described in Section 1.2.2. Recall that vertices of the Voronoi diagram, called *Voronoi vertices*, are centres of empty *Voronoi spheres*, which pass through (generically) 4 points of the point set. When points  $P$  are sampled on the boundary  $\mathcal{B}$  of a solid  $\Omega$ , recall that Amenta et al. [7] show that a subset of the Voronoi vertices of  $P$ , the *poles*, converges to the medial surface of  $\Omega$  as the sampling density of  $P$  approaches infinity, and hence, the envelope of a subset of the Voronoi spheres converges to  $\mathcal{B}$ .

Among those methods that do not aim to provide a tight approximation to the solid, the following are worth noting and recalling. An early method for the approximation of a solid with a triangle mesh boundary with spheres, [96], bounds each mesh triangle with a sphere but does not necessarily offer a tight fit. The method of [4, 5] considers medial balls grown by a certain amount and seeks a minimum cardinality subset of these grown balls that covers boundary points by applying heuristics to solve this set covering problem. The method of [60] reduces the number of medial balls by growing the balls, removing covered balls, pruning the medial surface of the union of the grown balls, and shrinking the balls back.

We now overview methods that seek a sphere-based approximation having a tight fit to the object boundary. There are two methods which approximate a solid  $\Omega$  with tight-fitting spheres, starting with a set of Voronoi spheres. Hubbard [67] greedily selects adjacent Voronoi spheres for merging when their bounding sphere has the best tightness of fit. Bradshaw and O’Sullivan [23] improve the tightness of fit of Hubbard’s algorithm by using an adaptive greedy strategy that adds sample points to the boundary of  $\Omega$  to generate new Voronoi spheres as needed. Tightness of fit is evaluated as the maximum distance between each sphere and  $\Omega$ . As this quantity is difficult to compute exactly, an approximation is used. Wang *et al.*[120] propose a variational optimization method that improves on the performance of [23] and is feasible for approximations having up to several hundred spheres. The error measure used is an approximation to total sphere volume outside the object  $\Omega$ . Approximate volumetric error is also used in [54] to compare the performance of the methods of [67, 23], and an octree-based method for approximating deforming objects with spheres. In the present work, we also use a volume-based error measure.

## 7.2 Our Sphere-Based Shape Approximation

As we saw in the previous section, a popular strategy for approximating a polyhedron with a small number of well-fitting spheres is to start by computing the Voronoi spheres of a set of boundary points. The centres of these spheres, the Voronoi vertices, are inherently clustered near rounded corners of polyhedra because there are many Voronoi vertices equidistant from 4 boundary points at these locations (*e.g.* Figure 7–1(Centre)). Further, each Voronoi sphere does not necessarily describe a new salient feature of the polyhedron. Subsequently, to generate a small number of spheres offering a tight fit to the polyhedron,

such methods require an optimization step to remove and redistribute spheres, making them computationally expensive.

In Chapter 2, we described a method to approximate the salient portions of the medial surface of a polyhedron using a collection of points whose density is a function of the voxel sampling. Figure 7–1(Right) shows the sphere centres generated with our method. Compared to the set of sphere centres in Figure 7–1(Centre) generated with a Voronoi-based method, our sphere centres are much more evenly distributed on the medial surface. The reason for this is as follows. Voronoi-based methods approximate the medial surface using a set of boundary samples as input and the density of the medial surface approximation depends on the location of the boundary points. Contrarily, the distribution of medial points produced by the distance-based method we have described in Chapter 2 depends on the resolution of the subdivision of space in the interior of the object, and the density of the sphere centres produced can be controlled by varying the resolution of this subdivision. Further, as explained in Chapter 2, the radii of our spheres are distances from the approximate medial points to their nearest boundary points, ensuring that the computed medial spheres are internal and tangent to the boundary  $\mathcal{B}$ . Thus, our spheres are guaranteed to be internal and tangent to the surface, and are centred near salient parts of the medial surface. In this chapter, we will see that this set of spheres can be used directly, without an expensive optimization step, to provide a fast and tight approximation to a polyhedron and can be updated quickly as the polyhedron deforms.

As we want to generate a small number of spheres and we want the spheres to capture the shape of the object being approximated, we will use Algorithm 2 on all voxels interior to or intersected by the polyhedron’s boundary  $\mathcal{B}$ , as described in Section 2.4.3. Recall that




Figure 7-1: A polyhedron (Left) and two distributions of sphere centres for a sphere-based approximation this object. (Centre) The subset of the Voronoi vertices of points sampled on the polyhedron boundary that lies near its medial surface. (Right) Sphere centres computed with our method, developed in Chapter 2.

for each voxel with side length  $\sigma$ , interior to or intersected by  $\Omega$ , at most one point within a user-chosen tolerance of the medial surface of  $\Omega$  and that satisfies user-chosen object angle and radius conditions, is returned. Therefore, the voxel size  $\sigma$ , as well as the object angle and radius thresholds used, affect the number of medial points that are generated.

An important consequence of this sphere generation strategy is that the number of medial spheres to be generated is not known in advance. Decreasing the parameter  $\sigma$  may not increase the number of spheres generated. In fact, it is possible that the number of spheres generated decreases as  $\sigma$  is decreased. Consider for example, Figure 7-2, showing the medial surface of a nearly spherical solid that is small relative to the solid. In this example, the medial surface intersects two adjacent blue voxels. If these two blue voxels are replaced with three green voxels by decreasing  $\sigma$ , only the middle voxel of the three green voxels is intersected by the medial surface. Thus, decreasing  $\sigma$  does not necessarily increase the number of medial spheres generated. However, if the ratio between the new voxel size and the old voxel size is an integer, then the number of medial spheres generated does not decrease.




Figure 7–2: Consider the solid that is the union of the two grey balls, whose medial surface is the dashed red line. Increasing the voxel resolution (from blue to green) decreases the number of medial spheres generated for the solid.

Although fast approximate GPU-based methods for distance field computation exist [111, 99], we will use exact point-to-mesh distance to generate spheres that are guaranteed to be interior and tangent to the solid boundary. This property is essential for the error analysis in Section 7.3.1, the updates under deformations in Section 7.4.1, and the sphere dilation process described in the following chapter in Section 8.2.

### 7.3 Volumetric Error for Unions of Spheres

In this section, we evaluate the global quality of our sphere-based approximation to a polyhedron using a volumetric measure, and present comparative results against a leading method [23] that approximates polyhedra with spheres.

#### 7.3.1 Volumetric Error: Exact and Lower Bound

We evaluate tightness of fit of an approximation of  $\Omega$  with a union of balls  $U$  as the volume of  $U$  outside  $\Omega$  plus the volume of  $\Omega$  outside  $U$ . Let  $\text{vol}(\cdot)$  denote volume, and let

$\bar{A}$  be the complement of a set  $A$ . Then

$$\text{Err}_\Omega(U) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(\Omega \cap \bar{U})$$

is the volumetric error of the approximation of  $\Omega$  with  $U$ . We state the following lemma.

**Lemma 7.1.**  $\max\{\text{vol}(U) - \text{vol}(\Omega), \text{vol}(\Omega) - \text{vol}(U)\} \leq \text{Err}_\Omega(U)$  and when  $\text{vol}(U \cap \bar{\Omega}) = 0$ ,  $\text{vol}(\Omega) - \text{vol}(U) = \text{Err}_\Omega(U)$ .

*Proof.* Observe that  $\text{vol}(U) = \text{vol}(U \cap \Omega) + \text{vol}(U \cap \bar{\Omega})$  and  $\text{vol}(\Omega) = \text{vol}(\Omega \cap U) + \text{vol}(\Omega \cap \bar{U})$ . It follows that  $\text{vol}(U) - \text{vol}(\Omega) = \text{vol}(U \cap \bar{\Omega}) - \text{vol}(\Omega \cap \bar{U}) \leq \text{Err}_\Omega(U)$ . Likewise,  $\text{vol}(\Omega) - \text{vol}(U) = \text{vol}(\Omega \cap \bar{U}) - \text{vol}(U \cap \bar{\Omega}) \leq \text{Err}_\Omega(U)$ . If  $\text{vol}(U \cap \bar{\Omega}) = 0$ ,  $\text{vol}(\Omega) - \text{vol}(U) = \text{Err}_\Omega(U)$ .  $\square$

Thus, in order to compute either the exact volumetric error, or a lower bound on volumetric error, one needs to compute the volume of the polyhedron  $\Omega$ ,  $\text{vol}(\Omega)$  and the volume of the union of balls  $\text{vol}(U)$ ,  $\text{vol}(\mathcal{B})$ , exactly. In Section 1.2.2, we have described how the power diagram of a union of balls can be used to guide the efficient computation of the volume of a union of balls. The volume of a polyhedron can be computed using a standard technique in computational geometry, described in [89].

### 7.3.2 Experimental Results and Discussion

In this section, we compare the sphere-based approximations computed with our method and those computed using a leading method in terms of the volumetric error of the approximation. Approximate volumetric error has been used in the literature ([120, 54]) to evaluate the quality of sphere approximations and provides a global measure of fit.















												
$ S_V  =$	512	506	474	497	512	496	510	444	482	506	499	495
$ S_D  =$	495	475	449	483	472	481	487	422	469	500	487	480
$\text{nerr}(S_V) \geq$	0.035	0.248	0.017	0.054	0.141	0.266	0.055	0.091	0.015	0.101	0.480	0.050
$\text{nerr}(S_D) =$	0.034	0.191	0.011	0.033	0.073	0.156	0.038	0.062	0.011	0.060	0.157	0.029
$\text{Time}(S_V) =$	2,342	3,212	4,049	2,672	4,936	1,930	22,845	2,469	7,428	2,870	3,739	9,479
$\text{Time}(S_D) =$	103	138	284	99	171	94	429	278	356	114	159	243
$\# \text{ triangles} =$	6,144	5,104	6,924	5,672	11,116	5,288	46,202	5,110	18,432	5,114	13,320	19,276

Table 7-1: Sphere sets  $S_D$  are generated with our method; sphere sets  $S_V$  are generated with the method of [23]. Timings are in seconds.

The method of Bradshaw and O’Sullivan for approximating solids with spheres [23], which we will refer to as AMAA (Adaptive Medial Axis Approximation), is the state-of-the-art method for approximating objects tightly using several hundred spheres or more. The sphere approximations generated in Wang *et al.*[120] provide a tighter fit than those computed by Bradshaw and O’Sullivan [23], but the method is only feasible for generating small sphere sets (less than or equal to approximately 128 spheres). As our method is able to generate a large number of spheres quickly, we compare it to that of [23].

When approximating a solid  $\Omega$  with spheres, let  $S_D$  be the set of spheres computed with our distance-based method and let  $S_V$  be the set of spheres computed with the Voronoi-based AMAA method of [23]. Let  $U_D$  and  $U_V$  be the associated unions of balls. As explained in Section 7.2, the balls  $U_D$  are completely contained inside  $\Omega$ , *i.e.*,  $\text{vol}(U_D \cap \overline{\Omega}) = 0$ . By Lemma 7.1,  $\text{vol}(\Omega) - \text{vol}(U_D)$  gives the exact error of the approximation of  $\Omega$  with  $U_D$ , while  $|\text{vol}(U_V) - \text{vol}(\Omega)|$  is a lower bound on the error of the approximation of  $\Omega$  with  $U_V$ . Since in our experiments we examine a variety of objects of different sizes, we will define a normalized error  $\text{nerr}(U) = \frac{\text{Err}_\Omega(U)}{\text{vol}(\Omega)}$ . Thus,

$$\begin{aligned} \text{nerr}(S_D) &= \frac{\text{vol}(\Omega) - \text{vol}(U_D)}{\text{vol}(\Omega)}, \text{ and} \\ \text{nerr}(S_V) &\geq \frac{|\text{vol}(U_V) - \text{vol}(\Omega)|}{\text{vol}(\Omega)}. \end{aligned}$$

We tabulate the exact error of  $S_D$  given by  $\text{nerr}(S_D)$ , and a lower bound on the error of  $S_V$   $\text{nerr}(S_V)$  for 12 models of varying geometric complexity in Table 7–1. Figure 7–3 shows error and computation time as a function of the number of spheres generated. Timings are shown on a 3.6 GHz Pentium IV processor with 3 GB of RAM.

The AMAA construction proceeds top-down by building a hierarchy of a fixed branching factor and depth. In Table 7–1 and Figure 7–3,  $S_V$  are the leaves of an 8-ary hierarchy of depth 4 (*i.e.*, maximum number of leaf spheres generated is 512). We found that construction of a binary hierarchy with an equal number of leaf spheres using the AMAA method was significantly slower than the construction of an 8-ary hierarchy using this method.

In finding the appropriate set of spheres  $S_D$ , we found the largest voxel side length  $\sigma$  such that  $|S_D| < |S_V|$ . First,  $\sigma$  was increased until  $|S_D| \geq |S_V|$ . Then binary search on  $\sigma$  values was used to determine the largest voxel length for which  $|S_D| < |S_V|$ . Timings for the construction of the  $S_D$  sets in Table 7–1 measure multiple invocations of our method until the appropriate sphere set is found. Fast construction of a tight binary hierarchy for the spheres  $S_D$  is discussed in Section 8.4.

As can be seen from Table 7–1 and Figure 7–3, our method generates a set of tighter fitting spheres significantly faster than AMAA. In generating our sphere sets, we used a threshold on object angle of 0.6 radians.

Our distribution of sphere centres, where at most one sphere centre is produced per voxel, is typically not an optimal distribution for minimizing the volumetric error. In the case of the peanut and tooth models (columns 3 and 9 of Table 7–1) our heuristic necessarily generates a suboptimal sphere centre distribution, as an optimal distribution of spheres minimizing volumetric error would include more large radius medial spheres than small radius medial spheres. However, for these models as well, we observe that our method’s performance in terms of volumetric error is superior to that of AMAA.

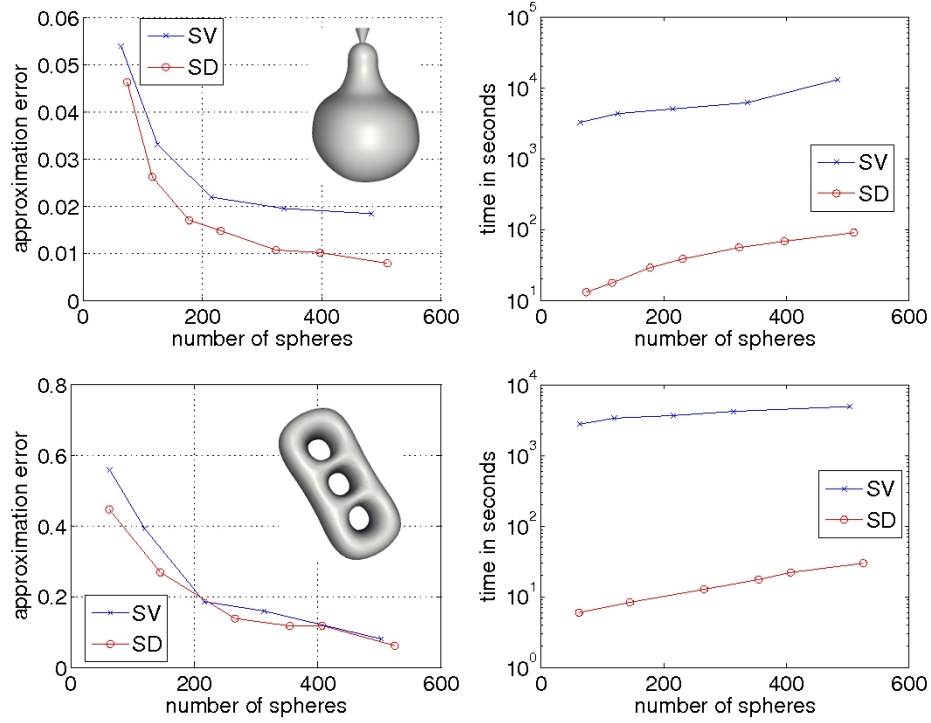


Figure 7–3: Error and timing results for the pear (5154 triangles) and triple torus (16000 triangles) models for our method (*SD*) and the AMAA method (*SV*). Note that true *SV* error may be higher than plotted.

## 7.4 Fast Updates under Deformation

In this section, we devise an algorithm for quickly updating our sphere set when the polyhedron undergoes deformation. We also compute bounds on the volumetric error of the new approximation. For the class of deformations that preserves local feature size and does not stretch the boundary, we observe that our update rule is fast and accurate.

### 7.4.1 Update Method

Let  $\mathcal{B}$  be the triangle mesh boundary of polyhedron  $\Omega$ . Consider a deformed version of  $\mathcal{B}$ ,  $\mathcal{B}'$ , bounding  $\Omega'$ , such that there exists a one-to-one mapping of vertices in  $\mathcal{B}$  to those in  $\mathcal{B}'$ ,  $F : \mathcal{B} \rightarrow \mathcal{B}'$ . We will show how to update the sphere set approximating  $\Omega$  to build an approximation to  $\Omega'$ .

For a sphere  $s$  interior to  $\Omega$  and tangent to its boundary,  $s = (c, r)$ , where  $c$  is the position and  $r$  is the radius, the position and radius of the corresponding sphere  $s' = (c', r')$  in a sphere-based approximation to  $\Omega'$  may be estimated as follows.

Let  $A$  be a nearest point to  $c$  on  $\mathcal{B}$ . Let  $N_{\mathcal{B}}(A)$  be the inward normal to  $\mathcal{B}$  at  $A$ . Since  $s$  is tangent to  $\mathcal{B}$ ,

$$c = A + rN_{\mathcal{B}}(A).$$

Suppose that  $A$  lies on  $\triangle v_i, v_j, v_k$ . Let  $\lambda_i, \lambda_j, \lambda_k$  be the barycentric coordinates of  $A$  in  $\triangle v_i v_j v_k$ , such that  $A = \lambda_i v_i + \lambda_j v_j + \lambda_k v_k$ . Then if  $v'_i, v'_j, v'_k$  are the vertices of  $\mathcal{B}'$  given by  $F(v_i), F(v_j), F(v_k)$ , let sphere  $s'$  be tangent to  $\mathcal{B}'$  at  $A' = \lambda_i v'_i + \lambda_j v'_j + \lambda_k v'_k$ . If  $s'$  is a medial sphere, then

$$c' = A' + \gamma N_{\mathcal{B}'}(A'),$$

for some constant  $\gamma$ . Given vertex normals for  $\mathcal{B}'$ , we estimate  $N_{\mathcal{B}'}(A') = \lambda_i N_{\mathcal{B}'}(v'_i) + \lambda_j N_{\mathcal{B}'}(v'_j) + \lambda_k N_{\mathcal{B}'}(v'_k)$ . We obtain an initial estimate for  $c'$  by letting  $\gamma = r$ .

Incorrect normal estimates or local shrinking of  $\mathcal{B}'$  can cause sphere  $s' = (c', r)$  to protrude outside  $\mathcal{B}'$ . This case is determined by checking the distance from  $c'$  to its closest point on  $\mathcal{B}'$ . When this distance is less than  $\gamma$ , the value for  $\gamma$  can be reduced until the sphere protrudes a user-chosen tolerable amount outside  $\mathcal{B}'$ . Algorithm 5 summarizes the procedure for shrinking spheres.

---

**Algorithm 5** SHRINK( $s = (c, r), \mathcal{B}, p, \xi$ )

---

**Input:** Mesh boundary  $\mathcal{B}$ , point  $p$  of contact of sphere  $s$  with  $\mathcal{B}$ , threshold  $\xi > 0$ .

**Output:** A radius  $d$  such that the sphere  $(p + d(c - p)/\|c - p\|, d)$  protrudes at most  $\xi$  outside  $\mathcal{B}$ .

- 1: Let  $d$  be the distance from  $c$  to  $\mathcal{B}$ .
  - 2: **while**  $r - d \geq \xi$  **do**
  - 3:   Let  $r$  be  $(r + d)/2$ .
  - 4:   Let  $c$  be  $p + r(c - p)/\|c - p\|$ .
  - 5:   Let  $d$  be the distance from  $c$  to  $\mathcal{B}$ .
  - 6: **end while**
  - 7: Return  $d$
- 

We now show that the spheres output by Algorithm 5 are either maximal or continue to protrude outside  $\mathcal{B}'$ .

**Lemma 7.2.** *Consider a sphere  $s_0 = (c_0, r_0)$  such that  $p$  is a point of contact of  $s_0$  with the boundary  $\mathcal{B}'$ . Let  $b \neq p$  be the nearest boundary point on  $\mathcal{B}'$  to  $c_0$ . Let  $d = \|b - c_0\|$ . Consider sphere  $s_1 = (c_1, r_1)$ , where  $r_1 = \frac{r_0 + d}{2}$ , and  $c_1 = p + r_1(c - p)/\|c - p\|_2$ . Then  $s_1$  is either maximal in  $\mathcal{B}'$  or  $s_1$  contains points of  $\mathcal{B}'$  in its interior.*

*Proof.* By the triangle inequality,

$$\|c_1 - b\| \leq \|c_1 - c_0\| + \|c_0 - b\| = \|c_1 - c_0\| + d. \quad (7.1)$$

Also,

$$\|c_1 - c_0\| = |r_0 - r_1| = \left| r_0 - \frac{r_0 + d}{2} \right| = \left| \frac{r_0 - d}{2} \right| = \frac{r_0 - d}{2}, \quad (7.2)$$

since  $r_0 > d$ . It follows from Equations 7.1, 7.2 that

$$\|c_1 - b\| \leq \frac{r_0 - d}{2} + d = \frac{r_0 + d}{2} = r_1. \quad (7.3)$$

Therefore,  $\|c_1 - b\| \leq r_1$ , and point  $b$  is inside or on  $s_1$ .  $\square$

By Lemma 7.2, applying Algorithm 5 to spheres that protrude outside  $\mathcal{B}$  produces spheres that continue to protrude outside  $\mathcal{B}$  or are maximal. Since  $\xi > 0$  and  $r$  is decreasing monotonically, Algorithm 5 eventually terminates.

O'Rourke and Badler [90] propose a method for decomposing a solid into spheres by shrinking large spheres that touch its surface, so that their centres remain perpendicular to the surface, until the spheres are wholly inside the solid. Ma *et al.* [81] also estimate positions of medial spheres tangent to given boundary points using normals to the given boundary points by shrinking the spheres. Our method is similar as we also shrink spheres while maintaining the property that a sphere centre moves along a perpendicular to the surface.

#### 7.4.2 Volumetric Error: Upper Bound

Given that Algorithm 5 ensures that no deformed sphere protrudes more than  $\xi$  outside the mesh boundary  $\mathcal{B}$ , we can find an upper bound for the volumetric error of the new sphere set.

**Lemma 7.3.** *Suppose that for each sphere  $s_i = (c_i, r_i)$ ,  $s_i \in S$ ,  $d_E(c_i, \mathcal{B}) > r_i - \xi$ . Let the union of the balls corresponding to the spheres in  $S$  be  $U$ . Then  $\text{Err}_\Omega(U) <$*

$\text{vol}(\Omega) - \text{vol}(U) + 2(\text{vol}(U) - \text{vol}(U^\xi))$ , where  $U^\xi$  is the union of the balls corresponding to the spheres  $S^\xi$ ,  $S^\xi = \{(c_i, r_i - \xi) | (c_i, r_i) \in S\}$ .

*Proof.* As shown in the proof of Lemma 7.1,  $\text{vol}(\Omega) - \text{vol}(U) = \text{vol}(\Omega \cap \bar{U}) - \text{vol}(U \cap \bar{\Omega})$ . Further, recall that  $\text{Err}_\Omega(U) = \text{vol}(\Omega \cap \bar{U}) + \text{vol}(U \cap \bar{\Omega})$ . It follows that

$$\text{Err}_\Omega(U) = \text{vol}(\Omega) - \text{vol}(U) + 2\text{vol}(U \cap \bar{\Omega}). \quad (7.4)$$

Let us express  $\text{vol}(U) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(U \cap \Omega)$  and  $\text{vol}(U^\xi) = \text{vol}(U^\xi \cap \bar{\Omega}) + \text{vol}(U^\xi \cap \Omega)$ . Because  $\text{vol}(U^\xi \cap \bar{\Omega}) = 0$ ,  $\text{vol}(U) - \text{vol}(U^\xi) = \text{vol}(U \cap \bar{\Omega}) + \text{vol}(U \cap \Omega) - \text{vol}(U^\xi \cap \Omega)$ . Since  $\text{vol}(U \cap \Omega) - \text{vol}(U^\xi \cap \Omega) > 0$ ,  $\text{vol}(U \cap \bar{\Omega}) < \text{vol}(U) - \text{vol}(U^\xi)$ . Combining this fact with Equation 7.4, the lemma follows.  $\square$

We now use this bound to evaluate the quality of the sphere-based approximation to the deformed polyhedra.

### 7.4.3 Experimental Results and Discussion

Figure 7–4 shows times to compute and update a reference set of spheres for three deformation sequences. The parameter  $\xi$  was set to be 0.25% of the maximum dimension of the reference polyhedron bounding box. The reference set of spheres took 146.8, 21.9, and 226.6 seconds to compute for the horse, knot and octopus reference poses, respectively on a machine with a single 3.6 GHz Pentium IV processor with 3 GB of RAM. Note that this fast update under deformations is possible because we use spheres that are interior and tangent to the solid. Using Lemmas 7.1 and 7.3, we compute lower and upper bounds on the volumetric error of our approximations.



Following [6], the *local feature size* of a point  $A$  on the boundary of a solid  $\Omega$  is the distance from  $A$  to the medial surface of  $\Omega$ . For deformations where the local feature size changes little, our strategy for updating the positions and radii of spheres whose centres lie near the medial surface should be able to quickly generate tight sphere-based approximations to deformed polyhedra. Although this approach may not produce a tight set of spheres for deformations that add ripples to the surface of an object, or if the object greatly increases in width, *i.e.*, deformations not preserving local feature size, as well as deformations that stretch the boundary, we have observed that the method does work well for deformations arising from part articulation (such as the horse sequence), as well as volume preserving elastic deformations (such as the knot and octopus sequences).

In the current approach, the spheres approximating a solid serve as the geometric primitives that represent that solid’s shape. This is in contrast with approaches that construct a set of spheres to contain either the triangles on the solid’s boundary, or a set of points sampled on the solid’s boundary. We now explain how such sphere sets are updated in the literature as the solid deforms. James and Pai [68] propose a fast conservative update rule to a sphere hierarchy when the deformation can be described as a linear superposition of displacement fields. In the proposed update, the sphere radii always increase. Wang *et al.* [120] update sphere sets when the boundary of the object deforms by setting the new sphere centres to be a linear combination of the updated positions of those boundary sample points that lie inside the original spheres, while keeping the radii fixed. Updating spheres based on the displaced positions of vertices contained inside the sphere instead of the approach we have proposed may not always produce a tight set of approximating spheres. This is because when vertices move further apart, their bounding sphere provides

a poorer estimate to a medial sphere. Such spheres can protrude greatly outside the polyhedron's boundary and since their centres may be far from the medial surface, shrinking the spheres may not produce tight fits. When the solid is a deforming chain of balls, Guibas *et al.* [65] study the construction and maintenance of its sphere-based hierarchy.

## 7.5 Chapter Summary

In chapter 2, we have shown how a sphere-based shape representation of a polyhedron can be computed. In this chapter, we have shown that this representation provides a tight volumetric approximation to a polyhedron and can be generated quickly, as compared to the leading method for sphere-based shape approximation [23]. The set of spheres in our approximation can be quickly updated as the polyhedron deforms, given a correspondence between triangles of the original polyhedron boundary and its deformed version. Union of spheres shape approximations have numerous applications, including shape matching, generation of mesh deformations, shape morphing, shape segmentation, approximate Minkowski sums, point location, proximity queries, and shadow rendering. As our representation offers a fast and tight fit to polyhedra and may be updated quickly as the polyhedra deform, it should prove valuable for these applications. The next chapter considers the application for our union of spheres to performing efficient and accurate proximity queries between solids.



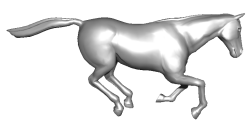

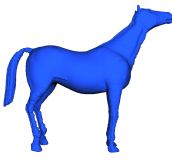
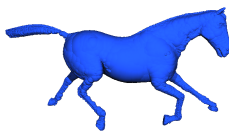
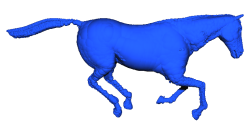







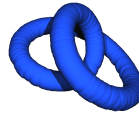
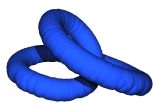

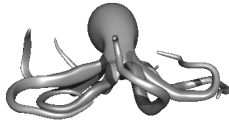



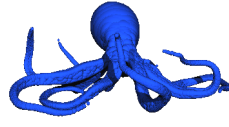


16857 triangles				
4307 spheres				
rel. speed	100%	4.04%	4.10%	4.08%
vol. error	[0.0110, 0.0110]	[0.0331, 0.0473]	[0.0362, 0.0509]	[0.0370, 0.0516]
1920 triangles				
426 spheres				
rel. speed	100%	1.16%	1.58%	1.02%
vol. error	[0.0607, 0.0607]	[0.0842, 0.0978]	[0.0928, 0.1065]	[0.0829, 0.0965]
39632 triangles				
9801 spheres				
rel. speed	100%	19.5%	19.9%	19.6%
vol. error	[0.0211, 0.0211]	[0.0256, 0.0585]	[0.0253, 0.0583]	[0.0252, 0.0579]
	(a)	(b)	(c)	(d)

Figure 7–4: (a) Reference polyhedron (top) approximated using internal spheres tangent to the polyhedron (bottom). (b)–(d) This sphere set is updated to approximate deformed versions of the reference polyhedron. The fraction of the time to compute the sphere set compared to computing the reference sphere set is shown below each pose. Also shown is the range of the volumetric error normalized by polyhedron volume.

## Chapter 8

# Application to Proximity Queries

In this chapter, we will use the set of well-spaced medial spheres shape representation of a polyhedron, in place of a boundary shape representation, to perform proximity queries between polyhedra. To allow this application, we propose a method to improve the coverage of the polyhedron boundary and describe how an efficient bounding volume hierarchy of the sphere sets can be built to accelerate proximity queries. We introduce a class of solids, called  $(\sigma, \theta)$ -fat solids. We then demonstrate experimentally that our sphere-based shape representation performs faster and with smaller error than the shape representation computed with the state-of-the-art method of [23] for approximate separation distance computations for  $(\sigma, \theta)$ -fat solids.

### 8.1 Background

For those solids which can be tightly approximated using a small number of spheres relative to the number of surface triangles, proximity queries, such as separation distance tests and collision detection, can be accelerated by working with the sphere representation instead of the boundary mesh representation. Pairwise distance tests are significantly faster between spheres than between triangles. We propose to *replace the boundary mesh representation with its approximation using spheres* and perform proximity queries using the sphere approximations.

In collision detection, one looks to either determine exactly whether two objects collide in a given instance in time, or, if an approximate solution is sought, then one can seek a *conservative* estimate of collision: one should not fail to detect any collisions, while possibly also detecting some false collisions. This is especially true when collision detection is performed in a simulation where failing to detect a collision can be highly undesirable, such as in a surgical simulation or robot motion planning. For other applications, such as physics-based animations computed for visualization purposes only, failing to detect some collisions may be acceptable. In fact, trading accuracy for speed is common when generating animations on a limited time budget [91].

We will not seek to cover the boundary of the polyhedron being approximated with spheres. For this reason, we will not be able to guarantee that our spheres can be used for conservative collision detection. We address a different, but related, problem: that of approximate separation distance computation. The *separation distance* between two solids  $M_1$  and  $M_2$  is defined as

$$\inf_{p_1 \in M_1, p_2 \in M_2} \|p_1 - p_2\|_2, \quad (8.1)$$

that is, it is the shortest distance between a pair of points, one on each solid.

In collision detection, we need to determine if the separation distance between two solids is zero. Once a collision is detected, or the lack of a collision is ascertained, a collision detection system terminates. However, in computing separation distance, the search continues until we have determined a pair of points on both models whose distance minimizes the separation distance. Thus, approximate separation distance is a different

computational problem from collision detection, but one whose solution provides a solution for collision detection. Because distances computed using our spheres are not guaranteed to be smaller or equal to the true separation distance, this method should not be used to determine collisions where conservative estimates are sought. However, when it is acceptable to trade-off accuracy of proximity queries for efficiency, we will see in this chapter that using our sphere-based representation is an attractive alternative to using the boundary representation of solids for performing proximity queries. The recent trend towards GPU-accelerated proximity queries [111, 99], including collision detection, further suggests that a certain lack of conservativeness in such computations is acceptable in the graphics community.

## 8.2 Improving Boundary Coverage by Conservative Dilation




Figure 8–1: The envelope of the dark circles  $s_1$  and  $s_2$  is non-differentiable at point  $v$ , whose nearest point on the boundary is  $m$ . The radii of both internal circles are increased to create circles that pass through  $m$ .

It is desirable that the envelope of our set of spheres provides a tight approximation to the boundary of a polyhedron  $\Omega$ . It is no longer essential that the spheres be interior to  $\Omega$ . We introduce a simple heuristic for conservatively growing the set of internal spheres that allows us to achieve a more favourable approximation to the solid boundary.

Consider the set of balls  $B$  corresponding to the set of spheres  $S$ . Recall from Section 1.2.2 that ball adjacency in the union of the balls in  $B$  is described by the dual complex of the balls,  $DC(B)$ . Edges in  $DC(B)$  correspond to pairs of balls intersecting along circular arcs on the envelope of the balls, while triangles correspond to triples of balls intersecting at vertices on the envelope of the balls.

The envelope of a set of balls is a non-differentiable object. We can solve for the locations of the vertices of this envelope, called *v-points*, by considering all triangles  $(b_i, b_j, b_k)$  in  $DC(B)$ , finding the 2 intersection points of spheres  $s_i, s_j$  and  $s_k$  corresponding to  $b_i, b_j, b_k$  and ignoring those intersection points that lie inside some sphere of  $S$ . We note which triple of spheres contributed to the creation of a v-point. For a v-point  $v$ , we consider the nearest point on the boundary of  $\Omega$  to  $v$ . The radius of each sphere  $s$  is increased sufficiently so as to cover the nearest boundary point to each of the v-points that is created by  $s$ . Figure 8–1 presents a 2D example. For those singleton edges of  $DC(B)$  identifying pairs of spheres adjacent in  $PD(B)$  that are not part of any triangles of  $DC(B)$ , we sample 2 opposite points at random on the circle of intersection of the two corresponding spheres and proceed to cover their nearest boundary points by growing the spheres.

Let  $S_D^+$  be the sphere-based approximation obtained by growing internal spheres  $S_D$  as described. Let  $S_V$  be the approximation produced by the AMAA method of [23]. We compare how well each of these sphere sets approximates the boundary of a polyhedron by evaluating the signed distance from 1) points on the envelope of the spheres to the polyhedron boundary, and 2) points on the polyhedron boundary to the envelope of the spheres. Figure 8–2 shows histograms of signed distances for several polyhedral models, where positive distance means that the point on the sphere envelope giving the distance

measurement is outside the polyhedron, and negative otherwise. We use a polygonal approximation to the envelope of spheres computed using CGAL's [1] implementation of the method in [75]. Our sample points are barycenters of boundary mesh triangles and each distance value is weighted by the area of the triangle. The maximum of the extreme distance values in the two histograms produced for each model approximates the Hausdorff distance between the boundary mesh and the envelope of the union of spheres. As we can see from the plots, the heuristic for growing internal spheres  $S_D$  generates spheres  $S_D^+$  which cover a significant portion of the boundary of the polyhedra, without generating spheres that protrude a great amount outside the polyhedra, as is often the case with  $S_V$  spheres. The mean absolute error for  $S_D^+$  is smaller than that for  $S_V$ , except for the last model, where both means are small. Note that neither the envelope of our sphere sets  $S_D^+$ , nor that of  $S_V$ , contain the polyhedra.

Because the spheres  $S_D$  are internal and tangent to the solid  $\Omega$ , points on the envelope of  $S_D$  are inside  $\Omega$  or on the boundary of  $\Omega$ . Using the method we have just described, boundary coverage is improved in the vicinity of points on the envelope of  $S_D$ . This heuristic does not consider the geometry of  $\Omega$  explicitly. Whenever the (one-sided) Hausdorff distance from the boundary of  $\Omega$  to the set of spheres is small, this strategy works well to improve boundary coverage. This strategy also offers the advantage of improving the fit of the union of spheres without an expensive optimization procedure.

### 8.3 $(\sigma, \theta)$ -fat Solids

In this section, we characterize the class of solids whose sphere-based approximations computed using the method we have proposed are connected. Consider the subset of the medial surface containing only points with object angle greater or equal to  $\theta$ ,  $\mathcal{MS}_\theta$ . Let  $\sigma$



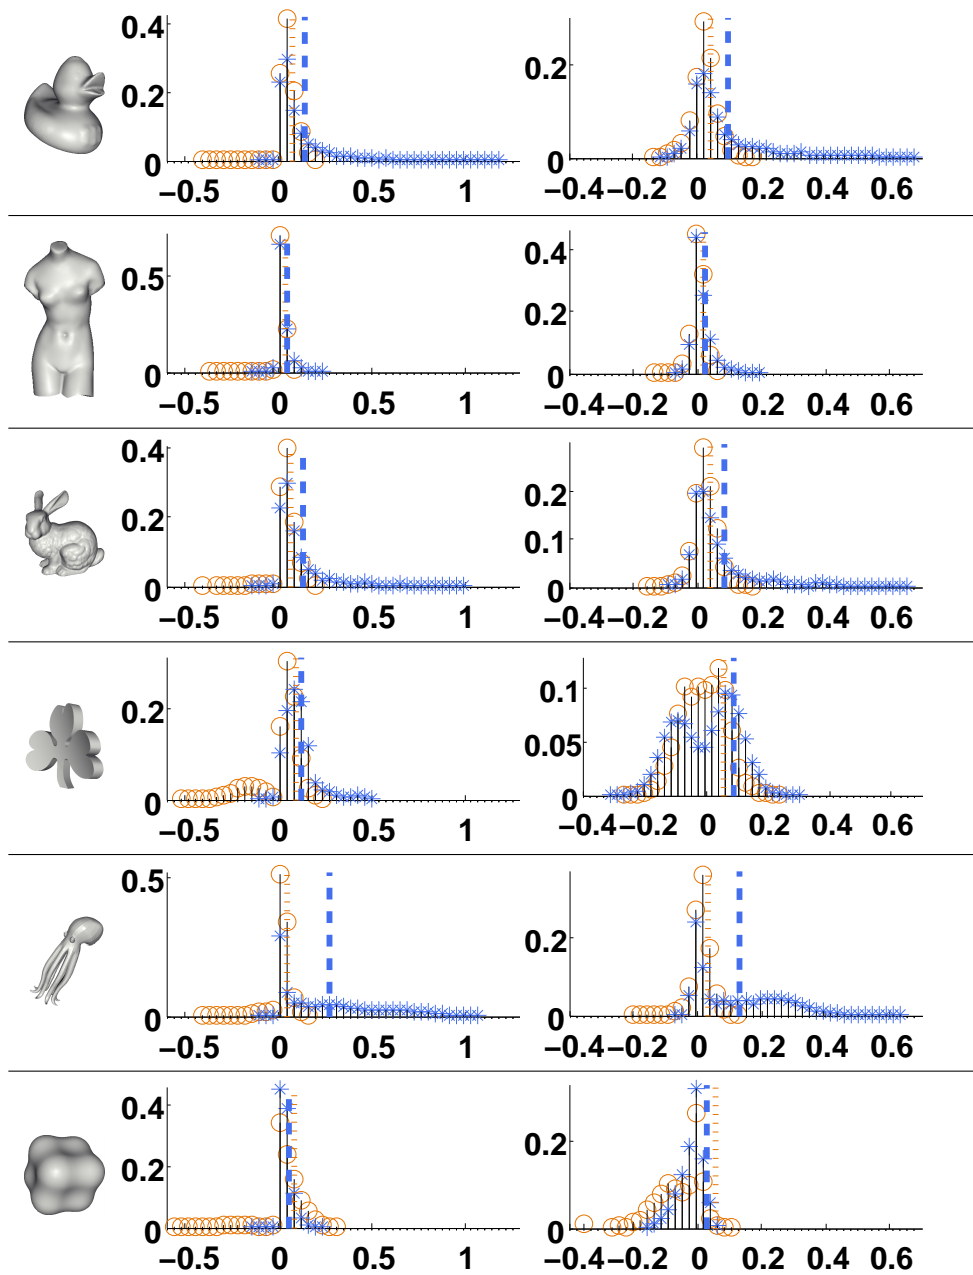


Figure 8–2: Comparison of the  $S_D^+$  and  $S_V$  approximations for various models. (Centre) Histograms of signed distances from points sampled on the boundary of polyhedra to points on the envelopes of  $S_D^+$  and  $S_V$ . (Right) Histograms of signed distances from points sampled on meshes approximating the envelopes of  $S_D^+$  and  $S_V$  to the boundary of polyhedra. Stars denote  $S_V$  spheres, while circles denote  $S_D^+$  spheres. Means of the absolute distance are shown as a dashed line for  $S_V$  and a dotted line for  $S_D^+$ .

be the side length of the voxels used to generate the approximation to the medial surface. Let  $S$  be the subset of spheres of the medial surface transform of  $\Omega$ , such that each sphere has object angle greater than  $\theta$ , and such that each voxel with side length  $\sigma$  containing a medial point with object angle greater or equal to  $\theta$  contributes a medial sphere to  $S$ .

Consider the following definition:

**Definition 8.1.** *A solid  $\Omega$  is  $(\sigma, \theta)$ -fat with respect to given values of  $\sigma$  and  $\theta$  if  $\sigma \leq \frac{r_0}{\sqrt{3}}$  for  $r_0$  the smallest medial sphere radius for medial points in  $\mathcal{MS}_\theta$ .*

Solids that are  $(\sigma, \theta)$ -fat have the following property:

**Lemma 8.1.** *Suppose that  $\mathcal{MS}_\theta$  of  $\Omega$  is connected. If  $\Omega$  is  $(\sigma, \theta)$ -fat, then the union of the spheres in  $S$  is connected.*

*Proof.* Since  $\mathcal{MS}_\theta$  of  $\Omega$  is connected, the voxels intersected by  $\mathcal{MS}_\theta$  of  $\Omega$  are connected. Consider a pair of adjacent voxels with side length  $\sigma$ , each containing medial points in  $\theta$ - $\mathcal{MS}$ . Each of these voxels contributes one sphere to  $S$ . The maximum distance between any pair of points in these voxels is  $2\sqrt{3}\sigma$ . If  $r \geq \sqrt{3}\sigma$ , for all radii  $r$  of medial points in the two voxels considered, then the two spheres touch. Since  $\Omega$  is  $(\sigma, \theta)$ -fat, medial spheres in adjacent voxels always touch. Since the set of voxels considered is connected, the union of the spheres in  $S$  is connected.  $\square$

Solids that are  $(\sigma, \theta)$ -fat do not have sharp corners or narrow parts with respect to the voxel resolution. The boundaries of those solids that are  $(\sigma, \theta)$ -fat can be better approximated using our sphere sets than other solids. Various definitions of *fatness* have been proposed in the computational geometry literature to parameterize the shape of inputs when it affects algorithm performance, as discussed in [40].




Figure 8–3: A rectangle swept sphere.

## 8.4 Hierarchy Construction using Rectangle Swept Spheres

In order to use our set of spheres  $S$  to perform fast proximity queries, we fit a bounding volume hierarchy to  $S$ . We observe that the medial surface simplified by object angle is often composed of relatively flat sheets, along which the radius of the medial spheres varies smoothly. A Rectangle Swept Sphere (RSS) is therefore a suitable bounding volume for medial spheres (see Figure 8–3). Such a bounding volume was introduced for bounding mesh triangles in the exact proximity query package PQP [76] and is popular in the design of efficient geometric algorithms [114]. The hierarchy is constructed top-down, each parent has 2 children, and each RSS is fitted to the set of leaf spheres it bounds. To find the orientation and position of the RSS rectangle and the orientation and position of the partitioning plane, we compute the centroid and the top two eigenvectors of the covariance matrix of the spheres using the CGAL library [1]. The radius is found by using a procedure similar to that in [76], modified to bound spheres instead of triangles.

Alternatively, one may compute a bounding volume hierarchy using spheres as the bounding volume. Given our sphere set approximation, we compute the hierarchy using the same top-down procedure as when using RSSs as the bounding volume. To compute the minimum bounding sphere of a set of spheres, we use CGAL’s implementation of the algorithm in [57].

## 8.5 Experimental Results

We evaluate the usefulness of our sphere approximation for measuring approximate separation distance in terms of the computation time and the accuracy of the results. In our experiments, we perform a physics-based simulation (one that simulates Newtonian mechanics) by allowing pairs of the same polyhedron to collide in a closed box. For each frame, we assign gravity a random direction and measure separation distance using both the sphere sets and the polyhedra. This setup was chosen as it can run for a long period of time with no user interaction. The box is made small enough so that the polyhedra eventually bounce off each other or the box. This allows us to compute pairwise distances for a variety of relative orientations and separation distances, without allowing for significant overlap between sphere sets (where separation distance is zero).

We compare the performance of separation distance tests for the set of spheres computed with our method ( $S_D$ ) and that computed with the AMAA method of [23] ( $S_V$ ). Given a set of voxels interior to or intersected by the boundary  $\mathcal{B}$ , we find one medial point per interior voxel and subdivide those voxels that are intersected by the boundary of  $\Omega$  into 8 voxels, thus computing at most 8 spheres for voxels intersected by the boundary of the polyhedron. We only retain those approximate medial points whose object angle estimate exceeds 0.6 radians. We grow our initial sphere set  $S_D$  using the strategy outlined in Section 8.2 to create a new sphere set  $S_D^+$ . This step takes 1-5 seconds for the models in Table 8–1. We compute both an RSS hierarchy and a sphere hierarchy of our set of spheres using the method described in Section 8.4. In proximity query experiments, we consider hierarchies with branching factor of two. References [50, 73, 94] provide arguments for the choice of branching factor of two in bounding volume hierarchies.

The timings for the AMAA sphere construction in Table 7–1 are for a top-down construction of a sphere hierarchy with a branching factor of 8 (where the leaves are the spheres  $S_V$ ). Computation of a binary AMAA hierarchy for the models shown in Table 8–1 takes significantly longer: from 3 hours (for the eight model) to 14 hours (for the octopus model). In addition to a binary sphere hierarchy, we also construct a binary RSS hierarchy of the AMAA spheres. Constructing an RSS hierarchy for the sphere sets considered takes a fraction of a second on average.

We evaluate the separation distance error as the difference between the exact distance between pairs of polyhedra and the distance computed using the sphere set approximation. As each simulation considers pairs of non-intersecting objects at a variety of separation distances and relative orientations, our error provides a meaningful measure of the quality of the sphere set approximation for the task of computing separation distance. Table 8–1 presents error statistics for the different sphere sets. Our sphere sets display a smaller average error and generally smaller maximum error than those of AMAA. The error in approximate separation distance has been used to evaluate the quality of the approximation of an object with spheres in the literature [120].

The models in Table 7–1 that are not represented in Table 8–1 are the horse, hippo, dragon and cow models. These models are not  $(\sigma, \theta)$ -fat with respect to the voxel resolution  $\sigma$  and the object angle  $\theta$  considered, because they have very narrow regions, such as their ears and tails. The simulation ran is particularly demanding of the quality of the approximation of such extremal regions as they minimize the separation distance more often than other regions. To approximate these models well using spheres generated using

our method, a small value of  $\sigma$  needs to be used, resulting in a large number of spheres. It was not possible to generate such large AMAA sphere sets to draw comparative results.

We compare average per-frame times for performing approximate separation distance tests using  $S_D^+$  and  $S_V$ , where both sphere hierarchies (SH) and RSS hierarchies (RSSH) are considered. These results are summarized in Table 8–1. Timings are shown for a 3.6 GHz Pentium 4 CPU with 3 GB of RAM. In our experiments, we find that building an RSS hierarchy of spheres significantly reduces distance query time compared to using a sphere hierarchy for both our and AMAA sphere sets, as the RSS offers greater tightness of fit. For the models in Table 8–1, we see that the fastest performance is achieved by using our sphere sets and the RSS hierarchy, even when  $|S_D^+| > |S_V|$ . The tight fit of RSS bounding volumes to our sphere sets can be attributed to the fact that our spheres are lie near relatively flat medial sheets and have a small local variation in radius.

## 8.6 Discussion

Collision detection and separation distance computations are inherently very expensive operations, especially when a large number of objects described by a large number of primitives is involved. Methods to accelerate these computations have been studied in detail by the computer graphics community. Proximity queries between pairs of triangle meshes are accelerated by the use of bounding volume hierarchies, where different choices of bounding volume have been proposed – from swept sphere volumes [76], to oriented bounding boxes [62], axis-aligned bounding boxes [118], convex surface decomposition [49],  $k$ -discrete orientation polytopes [73], and spheres [67]. Our approach is different in that we propose to *replace* the triangle mesh shape representation with a set of


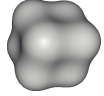






			Error		Timings	
Size			Ave.	Max.	SH	RSSH
	$S_D^+$	520	0.012	0.074	0.524	0.085
	$S_V$	498	0.024	0.063	0.268	0.088
	$S_D^+$	558	0.074	0.228	0.467	0.131
	$S_V$	397	0.094	0.331	0.627	0.132
	$S_D^+$	1052	0.026	0.102	0.672	0.134
	$S_V$	831	0.054	0.144	0.206	0.140
	$S_D^+$	296	0.009	0.024	0.385	0.066
	$S_V$	439	0.020	0.075	0.206	0.123
	$S_D^+$	695	0.116	0.297	0.709	0.143
	$S_V$	697	0.118	0.599	0.239	0.145
	$S_D^+$	1389	0.081	0.272	0.373	0.074
	$S_V$	745	0.109	0.4086	0.130	0.075
	$S_D^+$	548	0.055	0.173	1.458	0.108
	$S_V$	579	0.110	0.597	0.250	0.114
	$S_D^+$	469	0.076	0.325	0.360	0.068
	$S_V$	502	0.167	0.425	0.123	0.076

Table 8–1: Timing and error results for separation distance tests for our grown spheres  $S_D^+$  and spheres  $S_V$  of [23] for a variety of polyhedral models. Timings are shown in milliseconds. Spheres are used as bounding volumes in the  $SH$  column and  $RSS$ s are used in the  $RSSH$  column. The maximum dimension of bounding boxes for all models is 10.

spheres. However, we must note that such representations are only suitable for accelerating proximity queries for those solids that can be represented tightly using a small number of spheres relative to the number of triangles. For those solids that can be represented tightly using a small number of spheres, sphere representations can offer significant speed up for proximity queries without sacrificing quality.

With regard to collision detection, as our sphere sets are not guaranteed to cover the polyhedral boundaries they are approximating, using these spheres for computing collision and response may produce visualizations where solids interpenetrate. Further, because we do not compute the exact contact locations on the boundaries of polyhedra, incorrect contact forces can be computed. In our physics-based simulation, we used our spheres to detect collisions and compute contact forces, while displaying the effect of applying these forces on the original polyhedra. It is interesting to note that although the forces computed were not exact, and the meshes interpenetrated a certain amount in some collisions, the simulation appeared quite natural. According to a study of human perception of simulated collisions in [91], “... it is sometimes possible to produce a random collision response that is as believable as the more accurate ones, thus further masking collision anomalies.” It is also noted that people’s judgements of dynamics are often inaccurate. Also noted in [91] is that a noticeable gap between objects that bounce off each other, as happens when using conservative approximations to the object boundaries, produces unconvincing visualizations. The effect of object interpenetration on the viewer’s perception of simulation realism has not been studied, to our knowledge. For these reasons, using our non-conservative sphere sets as a proxy for meshes consisting of a large number of triangles, in order to generate natural-looking physics-based animations quickly, is a viable alternative to using



conservative approximations. This is especially important as tight conservative approximations are very costly to compute and cannot take advantage of our fast update rule as the object deforms.

## 8.7 Future Work

We now describe four possible directions for further study of sphere-based shape representations, with the application to proximity queries.

### 8.7.1 Faster Bottom-Up Sphere Hierarchies

In the timing results of Table 8–1, although the RSS hierarchy (RSSH) for our  $S_D^+$  spheres performed faster than the RSS hierarchy for  $S_V$  spheres, the sphere hierarchy (SH) of  $S_D^+$  spheres generally performed much slower than that of  $S_V$ . Our sphere-based hierarchy is constructed bottom-up, given a set of leaf spheres. The AMAA sphere-based hierarchy is built top-down, with each sphere of the hierarchy designed to offer a tight fit to the polyhedron. The disparity in timing results can be attributed to the fact that the intermediate spheres in the hierarchy that we construct do not offer a similarly tight fit to the leaf spheres as the spheres of the AMAA method.

Hierarchies with spherical bounding elements are popular because they may be quickly updated when the object deforms [68] and because considering a parent sphere instead of a leaf sphere can accelerate such applications as shadow generation [98]. Therefore, it is desirable to generate a tight-fitting hierarchy with sphere bounding volumes to our leaf spheres. There are several methods for the construction of sphere hierarchies given a set of leaf spheres described in the literature. Wang *et al.* construct a sphere hierarchy given a set of leaf spheres in a bottom-up fashion using Lloyd clustering [120]. Shamir *et al.* construct a multi-resolution sphere-based representation bottom-up by collapsing edges in the

dual complex of the associated balls (defined in Section 1.2.2) in priority order [104, 14]. Ranjan and Fournier use the ratio of the radius of the largest sphere in a cluster to the minimum enclosing sphere of a cluster to decide which clusters to replace with their enclosing sphere in a top-down construction of a sphere hierarchy [97]. Garcia *et al.* [55] build a minimum spanning tree of spheres, with distance given by the ratio of the volume of their bounding sphere to a measure evaluating the attraction force between the two spheres. Edges in the minimum spanning tree are collapsed in ascending order of weight to build a sphere hierarchy.

Consider the following research direction: *Generate a bounding volume hierarchy with sphere bounding primitives in a bottom-up fashion for our set of spheres that performs faster than the proposed hierarchy of spheres in proximity query tasks.*

### **8.7.2 Interruptible Collision Detection**

When performing collision detection on a time budget, traversal of bounding volume hierarchies (BVH) for a pair of objects is halted when the system runs out of time [91]. If any of the primitives tested are noted to collide, a collision is detected. In order for a BVH to be useable in this context, each level of the BVH must be designed to provide the tightest possible fit to the object.

The sphere sets generated by the AMAA algorithm were designed for interruptible collision detection: individual levels of the sphere hierarchy generated by this method were selected to offer a tight fit to the solid boundary. The intermediate RSS nodes of the BVH we have constructed offer a tight fit to the set of spheres we have generated. We have not investigated the suitability of this BVH for interruptible collision detection.

Consider the following research direction:

*Construct a bounding volume hierarchy of our sphere approximation of an object that offers a tight fit to the solid at each level of the hierarchy and enables one to perform interruptible collision detection of a comparable quality to leading methods.*

### **8.7.3 Sphere-Based Shape Primitives**

We have presented a method to approximate a solid with sphere primitives. However, spheres are not always the optimal choice for a primitive. Solids containing long and narrow features, such as limbs, need many spheres to be approximated well. However, if we use alternative primitives that are also long and narrow to approximate the solid, fewer primitives may need to be used.

Starting with a sphere-based shape representation, some spheres can be aggregated into more complex primitives. For example, primitives such as line-swept-spheres (capped cylinders) and rectangle-swept-spheres (see Figure 8–3), studied in [76], are a valuable alternative to sphere primitives, as they have more degrees of freedom. Likewise, cone-spheres, which are convex hulls of two spheres, have been suggested as shape modeling primitives [46, 82]. Finally, the convex hull of three spheres, which we call a *fat triangle*, is a new shape primitive suggested to us by Professor Wenping Wang in our discussions on the topic of medial representations (see Figure 8–4).

Given our set of spheres, we propose to simplify the sphere representation by aggregating certain spheres and replacing them with primitives such as fat triangles, cone-spheres, or line and rectangle swept spheres. Starting with a set of overlapping balls, their medial surface can provide valuable adjacency information that can guide this aggregation process. An algorithm for computing the medial surface of a set of balls is presented in [9].




Figure 8–4: A fat triangle, or the convex hull of 3 spheres. The spheres contribute red patches to the hull; pairwise sphere interpolants are cylindrical sections shown in blue; the linear section of the fat triangle is shown in yellow.

Since we are able to quickly compute a large number of spheres approximating a solid, it is appropriate to start with this fine sphere representation and simplify it by replacing spheres with alternative primitives, until a desired approximation quality is reached. Level of detail simplifications are of great value in computer graphics, as they allow to trade off accuracy and efficiency where appropriate. Whereas simplification of boundary mesh representations has received a considerable amount of attention, *cf.*, the survey in [56], level-of-detail simplification of volumetric primitives has not been as extensively explored. A representation of an object having a small number of volumetric primitives can facilitate geometric operators, as well as storage and transmission. Further, level of detail volumetric representations can be of value for interruptible collision detection. For the application to collision detection, we must make sure that overlap queries for the primitives used can be performed efficiently.

Consider the following research direction:

*Given a set of spheres, generate an approximation of the set of spheres using alternative sphere-based primitives that provides a good approximation to the original set of spheres, while reducing the number of primitives used.*

#### 8.7.4 Probabilistic Hierarchy Construction

When computing separation distance, certain spheres in a sphere-based shape representation serve as minimizers of separation distance more often than others.

Consider Equation 8.1. Suppose that the separation distance between solids  $M_1$  and  $M_2$  is  $d$  and is realized for points  $q_1 \in M_1$  and  $q_2 \in M_2$ . Suppose that  $M_1$  and  $M_2$  are connected unions of spheres. Consider the Voronoi diagram of the spheres in  $M_1$  and  $M_2$ . If  $q_1$  lies on sphere  $s$  of  $M_1$ , then  $q_2$  is in the Voronoi region of  $s$ ,  $V(s)$ . This relationship is discussed in [80].

Let  $s \oplus d$  be the sphere  $s$  with radius increased by  $d$ . Then  $q_2$  lies on  $s \oplus d$ . Consider the set of spheres in  $M_1$  where each sphere is dilated by  $d$ ,  $M_1 \oplus d$ . Then the likelihood that  $q_2$  lies on  $s \oplus d$  is

$$P(q_2 \in s \oplus d) = \frac{\text{surface area}(s \oplus d | \text{Vor}(s))}{\text{surface area}(M_1 \oplus d)}. \quad (8.2)$$

The Voronoi diagram of spheres and the power diagram of spheres' interiors coincide at points of sphere intersection. For this reason, Equation 8.2 is equivalent to

$$P(q_2 \in s \oplus d) = \frac{\text{surface area}(s \oplus d | P(b))}{\text{surface area}(M_1 \oplus d)}, \quad (8.3)$$

where  $P(b)$  is the power cell of ball  $b$  corresponding to sphere  $s$  in the power diagram of the balls of  $M_1$ . By exploiting this relationship between the Voronoi diagram of spheres

and the power diagram of balls, we avoid the expensive step of an explicit construction of the Voronoi diagram of spheres [71].

Given the fact that the likelihood of a sphere being a minimizer of separation distance can be easily computed when separation distance is known (using Equation 8.3), it is interesting to investigate if this information can be usefully incorporated to accelerate separation distance queries.

Consider the following research direction:

*Given an estimate of separation distance between two connected unions of spheres, build a bounding volume hierarchy for the spheres, where the length of the path from the root to a leaf is proportional to the likelihood of that leaf being a minimizer of distance in a separation distance query, such that distance queries are accelerated.*

## **Chapter 9**

# **Thesis Summary and Conclusion**

This thesis has proposed a novel shape representation for a 3D solid: the set of well-spaced medial spheres; such that at most one sphere centre lies in one voxel. We have examined the quality of this shape representation from three perspectives: 1) the ease of its generation; 2) the ability of the representation to capture valuable shape information; 3) the efficiency and accuracy of geometric operations using this representation.

Given the ability to compute nearest points on a solid's boundary to query points, we have presented algorithms to generate points within a user-chosen tolerance of the medial surface of the solid, and spaced such that at most one medial point is generated per cubic unit of space inside the solid. For each such approximate medial point, we also record its two estimated spoke vectors. Given this representation, we can estimate surface normals to medial sheets for each approximate medial point. If the approximate medial points are sufficiently dense, we can group neighbouring medial points that belong to the same smooth medial sheet using voxel neighbourhoods. The set of approximate medial points is simplified by thresholding by object angle. The result is a dense point-based approximation to the medial surface simplified by object angle, partitioned into distinct sheets.

Our algorithm for medial surface approximation considers the nearest boundary points to points sampled on spheres in order to determine whether the spheres contain points on

the medial surface. For a finite sampling rate on the sphere, our method does not guarantee that all spheres intersected by the medial surface will be detected. We have identified several properties of medial points that are not detected by our algorithm in the 3D case, and have sketched valuable components of an argument to show the relationship between the density of the samples considered on a sphere and the quality of the missed medial points in the sphere, proposed as future work. In the significantly easier 2D case, we have established such a relationship, showing that the quality of the missed medial points decreases as the density of samples on a circle increases, using several accepted criteria for quality.

In evaluating the fidelity of our shape representation of a smooth solid, we have shown how the set of approximate medial points along with their spoke vector estimates, computed with respect to a polygon mesh boundary representation of the solid, can be used to deduce principal curvatures and principal curvature directions on the smooth solid's boundary. These estimates are shown to be qualitatively consistent with those computed by a leading method that works on the triangle mesh boundary of the solid. We have also demonstrated that our shape representation can offer a part-based description of the solid, such that each part can be assigned a significance measure. This measure can be used to order medial sheets by significance and then to reduce the number of sheets required to approximate a shape.

We have shown that our shape representation is useful for generating fast and tight sphere-based approximations to a polyhedron. Unlike alternative methods for sphere-based shape approximation, our method for well-spaced medial sphere generation does not require expensive sphere redistribution or pruning steps to generate tight approximations to a solid. As a result, our method is significantly faster than existing methods and can



be used to generate sphere set approximations with a larger number of spheres than previously possible. In comparison with the state-of-the-art AMAA method, we have shown experimentally that our method generates fewer spheres having a smaller volumetric approximation error significantly faster than AMAA. Because our spheres are internal and tangent to the solid's boundary, it is possible to compare the volumetric error of sphere-based approximations produced by our method and another method, to quickly update our sphere set when the solid deforms, and to quickly dilate the sphere set to improve boundary coverage. We have demonstrated the benefit of using the rectangle swept sphere bounding volume for building a hierarchy of medial spheres. For  $(\sigma, \theta)$ -fat solids, we have presented experimental results showing that our sphere sets perform faster and more accurately than those of the AMAA method for approximate separation distance computation.

Our findings indicate that the union of well-spaced medial spheres is a valuable shape representation with respect to the three criteria mentioned above. Additional progress in terms of algorithm design and analysis, as well as extensions to the various applications and uses of medial spheres considered, can be made in order to continue developments in this fruitful research area. Some suggestions for future work are provided at the end of this chapter. Given the growing number of 3D data sets available, such research work would be particularly timely.

## Listing of Future Work

Throughout this thesis, a number of avenues for future work are described. Below is a listing of future research directions, and the chapter within the thesis where these are discussed.

- Show that if  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’, the quality of medial points present in sphere  $S$  decreases as the density of  $\Phi$  increases. (Ch. 3)
- Suppose that  $\text{DECIDEMS}(\mathcal{B}, S, \Phi)$  returns ‘Undetermined’. For each sphere in  $\mathcal{F}_f$ , we add its centre to the set of query points  $\Phi$  to create a set of query points  $\Phi'$ . Next, execute  $\text{DECIDEMS}(\mathcal{B}, S, \Phi')$ . By carrying out this operation repeatedly, what can be said about the quality of the missed medial points as more iterations are considered? (Ch. 3)
- Design other rules based on the analysis of a finite number of nearest boundary points to query points on a sphere that enable detection of medial points inside the sphere, such that the quality of the missed medial points can be shown to decrease by increasing the density of query points used, for a fixed-size query region  $S$ . (Ch. 3)
- Improve the tightness of the bounds of Theorem 4.1. (Ch. 4)
- Given a union of medial spheres, compute an estimate for the spoke vectors for each medial sphere. Use these estimates to, in turn, provide an estimate of the boundary differential geometry of the solid represented using this union of medial spheres, using tools developed in this chapter. (Ch. 5)
- Design an algorithm for the simplification of a set of medial spheres such that the simplified union of spheres is homotopy equivalent to the union of the original set of

spheres, where the choice of sheet to be removed is guided by significance measures presented in Section 6.2.(Ch. 6)

- Capitalize on advances in discrete and computational geometry, that facilitate efficient processing of unions of spheres, to design a matching algorithm for unions of medial spheres, where the adjacency of medial sheets is regularized. (Ch. 6)
- Generate a bounding volume hierarchy with sphere bounding primitives in a bottom-up fashion for our set of spheres that performs faster than the proposed hierarchy of spheres in proximity query tasks. (Ch. 8)
- Construct a bounding volume hierarchy of our sphere approximation of an object that offers a tight fit to the solid at each level of the hierarchy and enables to perform interruptible collision detection of a comparable quality as leading methods. (Ch. 8)
- Given a set of spheres, generate an approximation of the set of spheres using alternative sphere-based primitives that provides a good approximation to the original set of spheres, while reducing the number of primitives used. (Ch. 8)
- Given an estimate of separation distance between two connected unions of spheres, build a bounding volume hierarchy for the spheres, where the length of the path from the root to a leaf is proportional to the likelihood of that leaf being a minimizer of distance in a separation distance query, such that distance queries are accelerated. (Ch. 8)

## REFERENCES

- [1] CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
- [2] M.K. Agoston. *Computer Graphics and Geometric Modeling: Implementation & Algorithms*. Springer-Verlag New York, Inc., 2004.
- [3] O. Aichholzer, F. Aurenhammer, T. Hackl, B. Kornberger, M. Peternell, and H. Pottmann. Approximating boundary-triangulated objects with balls. In *European Workshop on Computational Geometry*, pages 130–133.
- [4] O. Aichholzer, F. Aurenhammer, T. Hackl, B. Kornberger, M. Peternell, and H. Pottmann. Approximating boundary-triangulated objects with balls. In *European Workshop on Computational Geometry*, pages 130–133. TU Graz, 2007.
- [5] O. Aichholzer, F. Aurenhammer, B. Kornberger, S. Plantinga, G. Rote, A. Sturm, and G. Vegter. Recovering structure from  $r$ -sampled objects. *Computer Graphics Forum*, 28(5):1349–1360, 2009.
- [6] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1999.
- [7] N. Amenta, S. Choi, and R. Kolluri. The Power Crust, Unions of Balls, and the Medial Axis Transform. *Computational Geometry: Theory and Applications*, 19(2-3):127–153, 2001.
- [8] N. Amenta, S. Choi, and R.K. Kolluri. The power crust. *Symposium on Solid Modeling and Applications*, 2001.
- [9] N. Amenta and R.K. Kolluri. The medial axis of a union of balls. *Computational Geometry: Theory and Applications*, 20:25–37, October 2001.
- [10] D. Attali and J.D. Boissonnat. Complexity of the Delaunay triangulation of points on polyhedral surfaces. *Discrete and Computational Geometry*, 30(3):437–452, 2003.

- [11] D. Attali, J.D. Boissonnat, and H. Edelsbrunner. Stability and computation of medial axes — a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, pages 109–125, 2009.
- [12] D. Attali and J.-O. Lachaud. Delaunay conforming iso-surface, skeleton extraction and noise removal. *Computational Geometry: Theory and Applications*, 19(2-3):175–189, 2001.
- [13] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding*, 67(3):261–273, 1997.
- [14] C.L. Bajaj, V. Pascucci, A. Shamir, R.J. Holt, and A. N. Netravali. Multiresolution molecular shapes. Technical report, in TICAM Technical Report, 1999.
- [15] J.F. Blinn. A generalization of algebraic surface drawing. *ACM Transactions on Graphics*, 1(3):235–256, 1982.
- [16] J. Bloomenthal and K. Shoemake. Convolution surfaces. In *SIGGRAPH*, volume 25, pages 251–256. ACM, 1991.
- [17] H. Blum. A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form*, 19(5):362–380, 1967.
- [18] H. Blum. Biological Shape and Visual Science. *Journal of Theoretical Biology*, 38:205–287, 1973.
- [19] I. Boada, N. Coll, N. Madern, and J. Antoni Sellarès. Approximations of 2d and 3d generalized Voronoi diagrams. *International Journal of Computer Mathematics*, 85(7):1003–1022, 2008.
- [20] A. Bobenko, P. Schröder, J. Sullivan, and G. Ziegler, editors. *Discrete Differential Geometry*, volume 38 of Oberwolfach Seminars. Birkhäuser, 2008.
- [21] I. Bogaevsky. Perestroikas of shocks and singularities of minimum functions. *Physica D*, 173, February 2002.
- [22] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Lévy. *Polygon mesh processing*. A. K. Peters Ltd., 2010.
- [23] G. Bradshaw and C. O’Sullivan. Adaptive medial-axis approximation for sphere-tree construction. *Transactions on Graphics*, 23(1):1–26, 2004.

- [24] F. Cazals, H. Kanhere, and S. Lorient. Computing the volume of a union of balls: a certified algorithm. *ACM Transactions on Mathematical Software*, 38(1), 2011.
- [25] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Symposium on Geometry processing*, pages 177–187, 2003.
- [26] M.-C. Chang and B.B. Kimia. Regularizing 3D medial axis using medial scaffold transforms. In *Computer Vision and Pattern Recognition*, 2008.
- [27] M.-C. Chang and B.B. Kimia. Measuring 3D Shape Similarity by Matching the Medial Scaffolds. In *ICCV Workshop on 3D Imaging and Modeling*, 2009.
- [28] M.-C. Chang, F.F. Leymarie, and B.B. Kimia. 3d shape registration using regularized medial scaffolds. In *3D Data Processing Visualization and Transmission*, pages 987–994, 2004.
- [29] F. Chazal and A. Lieutier. Stability and homotopy of a subset of the medial axis. In *Symposium on Solid Modeling and Applications*, pages 243–248, 2004.
- [30] F. Chazal and A. Lieutier. The  $\lambda$ -medial axis. *Graphical Models*, 67(4):304–331, 2005.
- [31] F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems*, 10(2):149–170, 2004.
- [32] H.-L. Cheng, H. Edelsbrunner, and P. Fu. Shape space from deformation. In *Pacific Conference on Computer Graphics and Applications*, pages 104–1012, 1998.
- [33] F. Chin, J. Snoeyink, and C.A. Wang. Finding the medial axis of a simple polygon in linear time. In *International Symposium on Algorithms and Computation*, pages 382–391, 1995.
- [34] D. Cohen-Steiner and J.-M. Morvan. Restricted delaunay triangulations and normal cycle. In *Symposium on Computational Geometry*, pages 312–321, 2003.
- [35] F.H. Croom. *Basic concepts of algebraic topology*. Springer, 1978.
- [36] T. Culver, J. Keyser, and D. Manocha. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design*, 21(1), 2004.
- [37] J. Damon. Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision*, 63(1):45–64, 2005.

- [38] J. Damon. Global geometry of regions and boundaries via skeletal and medial integrals. *Communications in Analysis and Geometry*, 15(2), 2007.
- [39] J. Damon. Tree structure for contractible regions in  $\mathbb{R}^3$ . *International Journal of Computer Vision*, 74(2):103–116, 2007.
- [40] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 2008.
- [41] D. DeCarlo, A. Finkelstein, S. Rusinkiewicz, and A. Santella. Suggestive contours for conveying shape. *Transactions on Graphics*, 22(3):848–855, 2003.
- [42] T.K. Dey and J. Sun. Normal and feature approximations from noisy point clouds. In *Foundations of Software Technology and Theoretical Computer Science*, pages 21–32, 2006.
- [43] T.K. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38:387–398, 2004.
- [44] P. Dimitrov. Flux invariants for shape. Master’s thesis, McGill University, 2003.
- [45] P. Dimitrov, J. N. Damon, and K. Siddiqi. Flux invariants for shape. In *Computer Vision and Pattern Recognition*, pages 835–841, 2003.
- [46] F. Dupont, B. Gilles, and A. Baskurt. Lossless and scalable 3d object coding method based on medial axis transformation. In *Picture Coding Symposium*, pages 381–386, 2003.
- [47] H. Edelsbrunner. The union of balls and its dual shape. In *Symposium on Computational geometry*, pages 218–231, 1993.
- [48] H. Edelsbrunner. Deformable smooth surface design. *Discrete & Computational Geometry*, 21(1):87–115, 1999.
- [49] S.A. Ehmann and M.C. Lin. Accurate and fast proximity queries between polyhedra using convex surface decomposition. *Computer Graphics Forum*, 20(3), 2001.
- [50] C. Ericson. *Real-Time Collision Detection*. Morgan Kaufmann, 2005.
- [51] M. Etzion and A. Rappoport. Computing Voronoi skeletons of a 3D polyhedron by space subdivision. *Computational Geometry: Theory and Applications*, 21, 2002.

- [52] H. Federer. Curvature measures. *Transactions of the American Mathematics Society*, 93:418–491, 1959.
- [53] M. Foskey, M.C. Lin, and D. Manocha. Efficient computation of a simplified medial axis. In *Solid Modeling and Applications*, pages 96–107, 2003.
- [54] M. Garcia, S. Bayona, P. Toharia, and C. Mendoza. Comparing sphere-tree generators and hierarchy updates for deformable objects collision detection. In *International Symposium on Visual Computing*, pages 167–174, 2005.
- [55] M.A. García, A.D. Sappa, and L. Basanez. Efficient generation of object hierarchies from 3D scenes. In *International Conference on Robotics and Automation*, pages 1359–1364, 1999.
- [56] M. Garland. State of the art report. multiresolution modeling: Survey and future opportunities. In *Eurographics*, 1999.
- [57] B. Gärtner. Fast and robust smallest enclosing balls. In *European Symposium on Algorithms*, pages 325–338, 1999.
- [58] P.J. Giblin and B.B. Kimia. A formal classification of 3D medial axis points and their local geometry. *Pattern Analysis and Machine Intelligence*, 26(2):238–251, February 2004.
- [59] P.J. Giblin, B.B. Kimia, and A.J. Pollitt. Transitions of the 3d medial axis under a one-parameter family of deformations. *Pattern Analysis and Machine Intelligence*, 31(5):900–918, 2009.
- [60] J. Giesen, B. Miklos, M. Pauly, and C. Wormser. The scale axis transform. In *Symposium on Computational Geometry*, pages 106–115, 2009.
- [61] J. Giesen, E.A. Ramos, B. Sadri, N. Amenta, and O. Cheong. Medial axis approximation and unstable flow complex. *Computational Geometry and Applications*, 18(6):533–565, 2008.
- [62] S. Gottschalk, M.C. Lin, and D. Manocha. Obbtrees: a hierarchical structure for rapid interference detection. In *SIGGRAPH*, pages 171–180. ACM, 1996.
- [63] M. Greenspan and N. Burtnyk. Obstacle count independent real-time collision avoidance. In *International Conference on Robotics and Automation*, pages 1073–1080. IEEE, 1996.



- [64] M. Gross and H. Pfister. *Point-Based Graphics*. Morgan Kaufmann, 2007.
- [65] L.J. Guibas, A. Nguyen, D. Russel, and L. Zhang. Collision detection for deforming necklaces. In *Symposium on Computational Geometry*, pages 33–42, 2002.
- [66] E. Hameiri and I. Shimshoni. Estimating the principal curvatures and the darboux frame from real 3D range data. In *3D Data Processing, Visualization, and Transmission*, pages 258–267. IEEE Computer Society, 2002.
- [67] P.M. Hubbard. Approximating polyhedra with spheres for time-critical collision detection. *Transactions on Graphics*, 15(3):179–210, 1996.
- [68] D.L. James and D.K. Pai. BD-Tree: Output-sensitive collision detection for reduced deformable models. *Transactions on Graphics*, 23(3):393–398, 2004.
- [69] M.I. Karavelas. A robust and efficient implementation for the segment Voronoi diagram. In *International Symposium on Voronoi Diagrams*, pages 51–62, 2004.
- [70] R.M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [71] D. Kim and D.-S. Kim. Region-expansion for the Voronoi diagram of 3d spheres. *Computer Aided Design*, 38, May 2006.
- [72] B.B. Kimia, A. Tannenbaum, and S.W. Zucker. On the evolution of curves via a function of curvature. I: The classical case. *Journal of Mathematical Analysis and Applications*, 163:438–458, 1992.
- [73] J.T. Klosowski, M. Held, J.S.B. Mitchell, H. Sowizral, and K. Zikan. Efficient collision detection using bounding volume hierarchies of k-dops. *Visualization and Computer Graphics*, 4:21–36, January 1998.
- [74] S.G. Krantz and H.R. Parks. Distance to  $C^k$  hypersurfaces. *Journal of Differential Equations*, 40(1):116 – 120, 1981.
- [75] N. Kruithof and G. Vegter. Envelope surfaces. In *Symposium on Computational geometry*, pages 411–420, 2006.
- [76] E. Larsen, S. Gottschalk, M.C. Lin, and D. Manocha. Fast proximity queries with swept sphere volumes. Technical report, University of North Carolina at Chapel Hill, 1999.

- [77] G. Lavoué, F. Dupont, and A. Baskurt. A new CAD mesh segmentation method, based on curvature tensor analysis. *Computer Aided Design*, 37:975–987, September 2005.
- [78] F.F. Leymarie and B.B. Kimia. The medial scaffold of 3D unorganized point clouds. *Pattern Analysis and Machine Intelligence*, 29(2), 2007.
- [79] A. Lieutier. Any open bounded subset of  $\mathbb{R}^n$  has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.
- [80] M.C. Lin. *Efficient Collision Detection for Animation and Robotics*. PhD thesis, University of California, Berkeley, 1993.
- [81] J. Ma, S.W. Bae, and S. Choi. 3d medial axis point approximation using nearest neighbors and the normal field. *The Visual Computer*, pages 1–13, 2011.
- [82] Y. Ma, C. Tu, and W. Wang. Computing the distance between canal surfaces. In *Geometric Modeling and Processing*, pages 88–103, 2010.
- [83] M. Meyer, M. Desbrun, P. Schröder, and A.H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In Hans-Christian Hege and Konrad Polthier, editors, *Visualization and Mathematics III*, pages 35–57. Springer-Verlag, Heidelberg, 2003.
- [84] B. Miklos. MESECINA. <http://www.balintmiklos.com/mesecina/>.
- [85] B. Miklos, J. Giesen, and M. Pauly. Discrete scale axis representations for 3d geometry. In *SIGGRAPH*, 2010.
- [86] L.R. Nackman and S. Pizer. Three dimensional shape description using the symmetric axis transform I: Theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(2):187–202, 1985.
- [87] B. O’Neill. *Elementary Differential Geometry*. Academic Press, 2 edition, 1997.
- [88] A. Opalach and S. Maddock. An overview of implicit surfaces. In *In Introduction to Modelling and Animation Using Implicit Surfaces*, 1995.
- [89] J. O’Rourke. *Computational Geometry in C*. Cambridge University Press, 1998.
- [90] J. O’Rourke and N. Badler. Decomposition of three-dimensional objects into spheres. *Pattern Analysis and Machine Intelligence*, 1(3):295–305, 1979.

- [91] C. O’Sullivan and J. Dingliana. Collisions and perception. *Transactions on Graphics*, 20, July 2001.
- [92] I.J. Palmer and R.L. Grimsdale. Collision detection for animation using sphere-trees. *Computer Graphics Forum*, 14(2):105–116, 1995.
- [93] S. Petitjean. A survey of methods for recovering quadrics in triangle meshes. *ACM Computing Surveys*, 34(2):211–262, 2002.
- [94] J. Pitt-Francis and R. Feathertone. Automatic generation of sphere hierarchies from CAD data. In *International Conference on Robotics and Automation*, volume 1, pages 324–329, 1998.
- [95] S.M. Pizer, P.T. Fletcher, S. Joshi, A. Thall, J.Z. Chen, Y. Fridman, D.S. Fritsch, A.G. Gash, J.M. Glotzer, M.R. Jiroutek, et al. Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision*, 55(2):85–106, 2003.
- [96] S. Quinlan. Efficient distance computations between non-convex objects. In *International Conference on Robotics and Automation*, pages 3324–3329, 1994.
- [97] V. Ranjan and A. Fournier. Volume models for volumetric data. *Computer*, 27(7):28–36, 1994.
- [98] Z. Ren, R. Wang, J. Snyder, K. Zhou, X. Liu, B. Sun, P.-P. Sloan, H. Bao, Q. Peng, and B. Guo. Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation. *Transactions on Graphics*, 25:977–986, July 2006.
- [99] G. Rong and T. Tan. Jump flooding in GPU with applications to Voronoi diagram and distance transform. In *Symposium on Interactive 3D Graphics and Games*, pages 109–116, 2006.
- [100] S. Rusinkiewicz. Estimating curvatures and their derivatives on triangle meshes. In *3D Data Processing, Visualization, and Transmission*, pages 486–493, 2004.
- [101] S. Rusinkiewicz and M. Levoy. QSplat: A multiresolution point rendering system for large meshes. In *SIGGRAPH*, pages 343–352, July 2000.
- [102] E. Saff and A. Kuijlaars. Distributing many points on a sphere. *The Mathematical Intelligencer*, 19:5–11, 1997.

- [103] D. Shaked and A.M. Bruckstein. Pruning medial axes. *Computer Vision and Image Understanding*, 69, February 1998.
- [104] A. Shamir, A. Sotzio, and D. Cohen-Or. Enhanced hierarchical shape matching for shape transformation. *International Journal for Shape Modeling*, 9(2):203–222, 2003.
- [105] K. Siddiqi, S. Bouix, A.R. Tannenbaum, and S.W. Zucker. Hamilton-Jacobi skeletons. *International Journal of Computer Vision*, 48(3):215–231, 2002.
- [106] K. Siddiqi and S. Pizer, editors. *Medial representations: mathematics, algorithms and applications*. Springer, 2008.
- [107] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. J. Dickinson. Retrieving articulated 3D models using medial surfaces. *Machine Vision and Applications*, 19(4):261–274, 2008.
- [108] S. Stolpner. Coarse-to-fine medial surfaces. Master’s thesis, McGill University, 2006.
- [109] M. Styner, G. Gerig, S. C. Joshi, and S. M. Pizer. Automatic and robust computation of 3D medial models incorporating object variability. *International Journal of Computer Vision*, 55(2-3):107–122, 2003.
- [110] A. Sud, M. Foskey, and D. Manocha. Homotopy-preserving medial axis simplification. In *Symposium on Solid and Physical Modeling*, pages 39–50, 2005.
- [111] A. Sud, N.K. Govindaraju, R. Gayle, and D. Manocha. Interactive 3D distance field computation using linear factorization. In *Symposium on Interactive 3D Graphics and Games*, pages 117–124, 2006.
- [112] R. Tam and W. Heidrich. Shape simplification based on the medial axis transform. In *Visualization*, page 63, 2003.
- [113] R. Tam and W. Heidrich. Computing polygonal surfaces from unions of balls. In *Computer Graphics International*, 2004.
- [114] M. Tang, M. Lee, and Y.J. Kim. Interactive Hausdorff distance computation for general polygonal models. *Transactions on Graphics*, 28(3):1–9, 2009.
- [115] G. Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. In *International Conference on Computer Vision*, page 902, 1995.

- [116] M. Teichmann and S. Teller. Polygonal approximation of Voronoi diagrams of triangles in three dimensions. Technical Report 766, Computer Graphics Group, Laboratory for Computer Science, MIT, Cambridge, MA, U.S.A., November 1997.
- [117] G.M. Turkiyyah, D.W. Storti, M. Ganter, H. Chen, and M. Vimawala. An accelerated triangulation method for computing the skeletons of free-form solid models. *Computer-Aided Design*, 29(1):5–19, 1997.
- [118] G. van den Bergen. Efficient collision detection of complex deformable models using AABB trees. *Journal of Graphics Tools*, 2:1–13, January 1998.
- [119] J. M. Vleugels and M. H. Overmars. Approximating generalized Voronoi diagrams in any dimension. Technical Report UU-CS-1995-14, Department of Information and Computing Sciences, Utrecht University, 1995.
- [120] R. Wang, K. Zhou, J. Snyder, X. Liu, H. Bao, Q. Peng, and B. Guo. Variational sphere set approximation for solid objects. *Visual Computer*, 22(9):612–621, 2006.
- [121] R. Weller and G. Zachmann. Inner sphere trees for proximity and penetration queries. In *Robotics: Science and Systems*, 2009.
- [122] S.A. Wilmarth, N.M. Amato, and P.F. Stiller. Motion planning for a rigid body using random networks on the medial axis of the free space. In *Symposium on Computational Geometry*, pages 173–180, 1999.
- [123] G. Wyvill, C. McPheeters, and B. Wyvill. Data structure for soft objects. *The Visual Computer*, 2(4):227–234, 1986.
- [124] Y. Yang, Y. Lai, S. Hu, and H. Pottmann. Robust principal curvatures on multiple scales. In *Symposium on Geometry processing*, pages 223–226, 2006.
- [125] H.T. Yau, L.S. Tsou, and H.M. Tseng. Automatic registration using virtual polar ball. *Computer-Aided Design and Applications*, 4:427–436, 2007.
- [126] S. Yoshizawa, A. Belyaev, and H.-P. Seidel. Skeleton-based variational mesh deformations. *Computer Graphics Forum*, 26(3):255–264, 2007.
- [127] P.A. Yushkevich, H. Zhang, and J. C. Gee. Continuous medial representation for anatomical structures. *Transactions on Medical Imaging*, 25(12):1547–64, 2006.