

© 2006 by Bardia Sadri. All rights reserved.

SURFACE AND MEDIAL AXIS TOPOLOGY THROUGH DISTANCE FLOWS
INDUCED BY DISCRETE SAMPLES

BY

BARDIA SADRI

B.S., Sharif University of Technology, 1999
M.S. University of Illinois at Urbana-Champaign, 2004

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2006

Urbana, Illinois

Abstract

The distance function induced by a surface in \mathbb{R}^n is known to carry a great deal of topological information about the surface and its embedding in space. It is an important question, from both theoretical and practical standpoints, whether such information about a surface can be extracted from the distance function induced by a discrete sample of it. Distance functions induced by discrete samples of surfaces and their associated mathematical structures are the main focus of this dissertation. These functions lead to continuous flow maps that turn out to be powerful topological tools. Based on these flow maps we design and analyze a number of simple and natural shape and medial axis reconstruction algorithms for which we can guarantee the topological type of output.

We prove that critical points of distance functions induced dense enough (relative) ε -samples of surfaces are concentrated around the surface and its medial axis. These two types of critical points can be distinguished from each other algorithmically. This “separation” of critical points is crucial to the design and analysis of of the above-mentioned algorithms.

Specifically, we present an algorithm for homeomorphic reconstruction of surfaces in 3D. This algorithm generalizes to higher dimensions with a slight change in the type of the provided topological guarantee. We also present an algorithm for medial axis approximation that computes a piece-wise linear *core* for the given sample. This core is guaranteed to be homotopy equivalent to the medial axis of the shape enclosed by the original surface. We then show that the core can be enhanced by any geometric medial axis approximation scheme without compromising the topological equivalence of the output and the medial axis being approximated. Finally, we present an analysis of Herbert Edelsbrunner’s well-known WRAP reconstruction algorithm and show that under relative ε -sampling in 3D, the output of WRAP is homotopy equivalent to the original shape.

*To my parents Mahrokh and Morteza,
for their endless patience,*

and

*to Azadeh,
for making it all worthwhile.*

Acknowledgments

It is a pleasure to be able to record my personal debt to my adviser, Edgar Ramos, without whose sincere care and guidance the very possibility of my ever succeeding to finish this work would be a matter of great doubt. He is a coauthor in all my published work surveyed in this manuscript and his technical maturity and knowledge of the subject has been a great source of inspiration at every step. But above all, I wish to express my utmost admiration for his humility, decency, and sense of responsibility, especially in a professional atmosphere in which the ruthlessness of competition has turned such virtues into unaffordable luxuries. I must also thank my other coauthors Tamal Dey and Joachim Giesen for their invaluable contributions to this work.

I am also deeply indebted to Sarel Har-Peled for co-advising me, for coauthoring my first computational geometry paper, and most importantly for being a much bigger friend any co-adviser has ever been to a co-advisee! His unorthodox sense of humor has doubtless made many of my Urbana moments unforgettable.

Many thanks are due to Jeff Erickson for being “the” person to turn to for those questions, technical and other, for which no one else knew the answer! My many discussions with him has greatly improved the quality of this work. Same gratitude is due to my other thesis committee members Rob Ghrist and John Hart for their constructive comments and ideas. I also wish to take the opportunity to thank Lenny Pitt for inspiring me, through his brilliant teaching, to take the plunge and become a theory student at the first place.

At the end, I wish to express my deepest and most sincere love and gratitude to my dearest Azadeh. Without her, life in Champaign-Urbana, surrounded by acres and acres of painfully boring corn fields would unquestionably be an utter misery. Her unparalleled grasp of reality always saved me from losing mine and kept me from living for tomorrow.

Table of Contents

List of Figures	x
Synopsis	1
Chapter 1 Background and Preliminaries	8
1.1 Convexity and General Position	8
1.1.1 Linear and Affine Subspaces	8
1.1.2 Convexity	10
1.2 Complexes	11
1.2.1 Simplicial and Abstract Simplicial Complexes	11
1.2.2 Cell Complexes	13
1.3 Distance Functions	13
1.3.1 Squared Distance Functions	15
1.3.2 Weighted Points	16
1.4 Voronoi and Delaunay Complexes	17
1.4.1 The Voronoi Complex	17
1.4.2 Weighted Points and Power Diagrams	19
1.4.3 The Delaunay Complex	20

1.4.4	Polarity and Voronoi and Delaunay Polytopes	21
1.5	An Introduction to Topology	26
1.5.1	Topological Spaces	26
1.5.2	Topology of \mathbb{R}^n	30
1.5.3	Manifolds and Submanifolds	31
1.6	Surfaces and their Medial Axes	32
1.6.1	Medial Balls and Local Feature Size	33
1.6.2	Criteria for Sampling Surfaces	37
Chapter 2	Distance Flows	43
2.1	Generalized Gradients	43
2.2	Dynamical Systems and Flows	47
2.3	Integrating the Generalized Gradient	49
2.3.1	Continuity of the Flow Map	50
2.4	Stable and Unstable Manifolds	56
2.4.1	Computing Stable and Unstable Manifolds	58
2.4.2	Stable and Unstable Flow Complexes	66
Chapter 3	Separation of Critical Points	73
3.1	Surface versus Medial Axis Critical Points	73
3.2	Distinguishing Surface and Medial Axis Critical Points	78
3.3	Tubular Neighborhoods and Reduced Shapes	83
3.3.1	The Boundary of Σ_δ	84
3.3.2	Flow on Tubular Neighborhoods	86

3.4	Distinguishing Inner and Outer Medial Axis Critical Points	90
Chapter 4 Reconstruction as a Union of Stable Manifolds		92
4.1	A Simple 3D Reconstruction Algorithm	93
4.1.1	Geometric Closeness	94
4.1.2	Convergence of Normals	99
4.1.3	Orientation of surface patches	102
4.1.4	$\tilde{\Sigma}$ and Σ are Homeomorphic	109
4.2	Extension to Higher Dimensions	110
4.2.1	Homotopy Equivalence via Distance Flow Maps	111
4.2.2	\tilde{S} and S are Homotopy Equivalent	112
Chapter 5 A Core for Approximating the Medial Axis		119
5.1	Definition of Core	119
5.2	Homotopy Type of the Core	121
5.2.1	Homotopy Equivalence of Core and Medial Axis	121
5.2.2	Extending the Core	123
5.3	Geometric Approximation	124
5.4	Experiments	132
Chapter 6 The Wrap Reconstruction Algorithm		134
6.1	Description of WRAP	135
6.2	Simulating the Minimum at Infinity	138
6.3	Flow on Tubular Neighborhoods	140
6.4	Geometric Quality	145

6.5 Topological Correctness	152
References	155
Author's Biography	160

List of Figures

1	Example of a point cloud sampled from a physical object (left) and a piecewise linear reconstruction of the surface of the object in question (right).	2
2	The knotted torus (left) and the plain (unknotted) torus (right) are homeomorphic. However they can be told apart by comparing their (outer) medial axes topologically.	4
3	Instability of medial axis under perturbations of the boundary. Left: a shape and its medial axis. Right: a small change in shape substantially changes its medial axis.	6
1.1	Left: a simplicial complex. Right: examples of collections of simplices that are not complexes.	12
1.2	Level sets of squared distance function induced by a set of (unweighted) points in the plane (left). Three dimensional epigraph of the the same distance function (right).	15
1.3	Levels sets of the squared distance function to a weighted set of points (left). Three dimensional epigraph of the the same distance function (right). Unlike the unweighted case, paraboloids are placed at different heights.	17
1.4	The Voronoi complex (diagram) of a set of points in \mathbb{R}^2	18
1.5	The Delaunay complex (triangulation) of a set of points in \mathbb{R}^2 . Empty balls <i>witness</i> the presence of simplices they circumscribe.	20
1.6	Voronoi vertices are centers of maximal Delaunay balls.	22
1.7	The projection of the intersection of Π_q and the lower envelope of Π_Q (the contribution of Π_q to the lower envelope) is the Voronoi cell V_q of q	23

1.8	The Voronoi polytope of set of weighted points in \mathbb{R}^n . The Voronoi complex of the point-set is the the projection of the the boundary complex of the Voronoi polytope back to \mathbb{R}^n	24
1.9	The Delaunay polytope of set of weighted points in \mathbb{R}^n . The Delaunay complex of the point-set is the projection of the boundary complex of the polytope.	25
1.10	The functions h_Q (left) and h_{Q^+} (right).	26
1.11	The medial balls ($B_\mu^+(x)$ and $B_\mu^-(x)$) and local feature size balls ($B_f^+(x)$ and $B_f^-(x)$) tangent to Σ at x	34
1.12	Left: a surface Σ and its medial axis M (only the inner medial axis is shown). Middle: the δ -tubular neighborhood Σ_δ of the surface. Right: the δ -tubular neighborhood M_δ of the medial axis (only the inner component shown).	37
2.1	Flow lines of the distance flow induced by a set of points. The red points are index 1 critical points.	51
2.2	Recursive computation of the stable manifold of a critical point. The algorithm $\text{INFLOW}(P, V)$ takes a Voronoi face V and a convex subset P of V and computes the inflow of P that reaches P through V	60
2.3	Computing the stable manifold of an index-2 saddle point in 3D.	63
3.1	Illustration of the proof of Lemma 6.7	75
3.2	A two-dimensional example of the separation of critical points (Theorem 3.3). Sample points (index-0 critical points) are shown in solid black. Other surface critical points are shown in blue and medial axis critical points are shown in red. Index-1 saddle points are depicted hollowed.	79
3.3	The ball of diameter pp^+ contains no critical point on the boundary of V_p	81
3.4	The inner shape S and outer shape S^* defined by a surface Σ (left). The inner δ -reduced shape S_δ and the outer δ -reduced shape S_δ^* determined by the δ -tubular neighborhood Σ_δ of Σ (right).	84
3.5	Illustration of the proof of Lemma 3.15.	87
3.6	Illustration of the proof of Lemma 3.16.	88

4.1	Algorithm RECONSTRUCT, a surface-reconstruction algorithm.	93
4.2	The output of the algorithm RECONSTRUCT: stable manifolds of medial axis maxima are the two-dimensional gray cells with blue boundaries. The common boundary of any two pair of these two-dimensional cells is itself the stable manifold of an index-1 medial axis critical point. The output of the algorithm RECONSTRUCT is the boundary of the grayed set.	94
4.3	A sink cone (left) and enclosing of the surface with sink cone envelopes (right).	95
4.4	Sink cone for a sample point p	96
4.5	A generic patch triangle on the stable manifold of a surface 2-saddle.	100
4.6	Illustration of the proof of Lemma 4.16, cases 1 (left) and 2 (right).	105
4.7	Illustration of the proof of Lemma 4.16, cases 3 (left) and 5 (right).	106
4.8	A two-dimensional demonstration of the stages in the proof of Lemma 4.22. Top-left: the source set U is shown in gray and the sink set S_{ε^2} is displayed in white. Top-right: the source set $U = U_0$ and all the critical points of the sample. Medial axis critical points are shown in red. Surface 1-saddles are shown in blue and samples in black. Middle-left: neighborhoods from index-0 critical points are removed from U_0 to result $U'_0 \simeq U_0$. Middle-right: U_1 is the union of stable manifolds of index-1 critical points and S_{ε^2} intersected with \tilde{S} . Bottom-left: Neighborhoods of index-1 saddle points are removed from U_1 . The resulting set is $U'_1 \simeq U_1$. Bottom-right: remaining spikes are pushed into S_{ε^2}	115
4.9	Left: A sampled curve. Right: Delaunay triangulation of the sample. A piecewise linear approximation of the curve can be found as a sub-complex of this triangulation.	117
5.1	An example of the core in 2D. Inner medial axis critical points are shown in red. Hollow points represent 1-saddles and solid ones are used for maxima. The inner core is shown in blue.	120
5.2	Left: a set W of Voronoi faces (in blue) that approximate parts of the inner medial axis are added to the core (also in blue). Right: flow closures $\phi_P(W)$ of W is added.	123
5.3	Proof of Theorem 5.10. The gray cap represents $D(x)$	130

5.4	Left: Core computed for the 3-holes model. The red lines are either unstable manifolds of index-2 saddle points or the one dimensional parts (hairs) of index-1 saddle points. Middle: Filtered Voronoi facets based on a condition similar to the angle condition of Dey and Zhao. Right: Extended core, i.e., the core, plus the flow closures of the facets from the middle picture.	132
6.1	Precedence relation between Delaunay simplices is shown by arrows. An arrow crossing an edge signifies that the flow on that edge is transversal while arrows ending on an edge show that the flow on the edge is tangential.	136
6.2	The original WRAP algorithm.	136
6.3	An two-dimensional example of the execution of WRAP. The removed Delaunay simplices are shown in red.	137
6.4	Extension of the sample points for simulation of the a critical point at infinity. Left: Delaunay triangulation of a set of points. Right: the point-set on the left is enclosed in a large enough ball and the boundary of the ball is sampled. The Delaunay triangulation of the original point-set is a subcomplex of the Delaunay triangulation of the extended one.	138
6.5	The modified WRAP algorithm.	140
6.6	Illustration of the proof of Lemma 6.3.	141
6.7	Illustration of the proof of Lemma 6.5.	144
6.8	Illustration of the proof of Lemma 6.10 for the case where $\phi \geq 180^\circ$	150
6.9	Illustration of the proof of Lemma 6.10 for the case where $\phi < 180^\circ$	151

Synopsis

Much of the computational power of our computers is spent on simulating or analyzing physical phenomena which involve various kinds of physical objects. Mathematical modeling of these physical objects is naturally a first step in performing any kind of computation on them. Consider for example a single snapshot of one of the now highly widespread computer animated movies and think of all the physical objects that appear in that scene. Although some, or perhaps all, of these objects do exist in physical form, no camera is pointed at any of them in the process of making of most of these movies. Instead, objects are modeled mathematically and are virtually arranged in the scene. Virtual lighting sources and textures are taken into account, and an image is rendered for the desired virtual camera location, by virtually tracing the path of light from the light sources, through reflections between the objects, all the way to the virtual camera lens. There are considerably many more applications to mathematical modeling of physical objects than rendering pictures of them. Such models allow simulation of many physical phenomena, out of which tracing the path of light is only one example.

In many applications, to model a physical object, it is enough to somehow represent the surface of the object in a mathematical form. Such a surface of course can be very complicated both geometrically and topologically; it may have very fine features and may sport several holes and handles. It is nonetheless crucial to capture such topological and geometric details accurately in a mathematical model before any kind of secondary processing can be carried out on them.

There are many ways to represent surfaces mathematically. For example, a natural way is to break a surface into several pieces and represent each piece as a graph of a function defined on a subset of the Euclidean plane \mathbb{R}^2 . Another example is to represent a surface implicitly as a level set of a continuous function defined over the space. Some surfaces can be broken into pieces where each piece has a simple representation. For example, a piece-wise linear surface can be broken into a number of planar polygonal pieces that are glued to each other along their boundaries. Other surfaces can be described as the boundary of a union of simple three dimensional objects, such a balls, cubes, cylinders, etc.

Regardless of the preferred method of representation for a particular application, it is a painstaking process to manually calculate a mathematical representation of even the simplest of physical surfaces. It is therefore rather desirable, if not completely crucial, to automatize




Figure 1: Example of a point cloud sampled from a physical object (left) and a piece-wise linear reconstruction of the surface of the object in question (right).

this process. Ideally, we wish to place a physical object in front of some kind of input device, such as a 3D scanner or digitizer, and immediately get a digital model of the object. However, scanning devices see the object in much the same as we do. They finely samples the surface of the object by tracking the reflections of some kind of light (natural light, laser, radar, etc) from the surface of object and in effect provide us with a discrete collection of sample points from the surface of the object. In fact, even if the surface of interest does not exist physically, it is sometimes possible (and much easier) to produce a sample of the surface as a first step in constructing a mathematical model of the surface.

However, a discrete set of points can barely qualify as an accurate representation of a target surface as it is very different from it topologically even though it may approximate it well geometrically. Thus we need to perform some kind of secondary processing on this point cloud and produce a model of a surface in the desired format that is approximately identical to the original surface and is topologically equivalent to it. Of course, for some applications, one may wish to bypass this secondary phase of surface reconstruction and do a different kind of post-processing directly on the input point cloud.

Inferring the topology of surfaces and their related structures from sample point clouds is at the core of every problem studied in this thesis. To make things easier, we focus only on objects (or as we call them *shapes*) that have smooth single-component boundaries. Smoothness avoids some serious complications that arise in modeling of surfaces that have sharp edges and corners. A single-component surface separate the “inside” of the shape it bounds from its “outside” and in this sense partitions its complement (the space minus the

surface) into an inner and an outer components. This assumption simplifies the discussed algorithms and their analysis but can be relaxed with little work.

The focus of this thesis is on properties of distance functions that discrete samples of surfaces induce. It turns out that such distance functions encode a rich collection of topological information about the sampled surface. To motivate the study of these distance functions, let us first look at the distance function a smooth target surface itself induces. Thus consider a compact surface Σ smoothly embedded in the n -dimensional Euclidean space. The surface Σ induces a distance function $s : \mathbb{R}^n \rightarrow \mathbb{R}$ that assigns to each point of the space the distance to its closest point on Σ :

$$s(x) = \inf_{y \in \Sigma} \|x - y\|.$$

The surface Σ itself can be trivially determined as a level set of this function; $\Sigma = s^{-1}(0)$. But more is encoded in $s(\cdot)$. Associated to every surface (in fact any compact set) embedded in \mathbb{R}^n is a structure called the *medial axis* of the set. The medial axis of a surface happens to encode information about the embedding of the surface in space. The medial axis $M(\Sigma)$ of Σ is defined as the set of points in space with more than a single closest point in Σ . Some of these points are contained in the inner open shape S enclosed by Σ and some fall into the outer one S^* . The former set makes the medial axis of the inner shape and the latter set makes the medial axis of the outer shape. In other words, the medial axis of Σ is the union of the medial axes of the two shapes S and S^* . The topological significance of medial axes becomes clear from a recent result of Lieutier [56] which states that every bounded open subset of \mathbb{R}^n has the same homotopy type as its medial axis. This means that the topological embedding of Σ into \mathbb{R}^n is essentially encoded in $M(\Sigma)$. To get a sense of this, compare the outer component of the medial axis of a knotted torus to that of an ordinary one¹ (see Figure 2). It turns out, that the medial axis of a smooth surface Σ is precisely the set of points in the complement of the surface at which the distance function s is not differentiable.

When the target surface Σ is understood by only a dense sample, the distance function s induced by Σ can be approximated by the distance function induced by the sample. It is natural to ask how much of the information in s can be retrieved from this approximation of it. The work presented in this monograph can be regarded as an attempt to answer this question.

The main two problems studied in this thesis are as follows: Given a dense sample P from an unknown smooth surface Σ ,

- is it possible to build a surface $\tilde{\Sigma}$ that can be guaranteed to be topologically equivalent to Σ and approximates it closely in a geometric sense?
- is it possible to approximate the medial axis (of the bounded component) of Σ with a (piece-wise linear) set that is provably topologically equivalent to it.

¹Of course Lieutier's results is valid for bounded open subsets of \mathbb{R}^n while in the given example the outer component is unbounded. This of course can be worked out enclosing the tori in a sphere and then study the medial axes of the resulting 2-component surfaces.




Figure 2: The knotted torus (left) and the plain (unknotted) torus (right) are homeomorphic. However they can be told apart by comparing their (outer) medial axes topologically.

Due to their practical importance, the problems studied in this manuscript have been subject to careful scrutiny of many researchers in various disciplines and there is a host of results that address many of these problems from both theoretical and practical standpoints. The significance of the proposed solutions in this work is in their use of a thorough and well-developed set of techniques and machineries that are based on the more heavily studied notions of continuous mathematics. The goal here is to approach such problems as surface and shape reconstruction, medial axis approximation, and in general problems that involve topological interpretation of discrete samples, in a more uniform and systematic way. We hope the presented arguments succeed to demonstrate the wealth and power of these techniques and structures.

The presented new results have appeared in [32, 44, 61] and journal versions [33, 45]. In the rest of this introduction our contribution to each of these problems is provided in higher detail. A survey of related results to each problem can be found at the end of its corresponding chapter.

Surface Reconstruction. The sampling model primarily used in this thesis is the ε -sampling of Amenta, Bern, and Eppstein [3]. This model of sampling uses the *local feature size* function, which assigns to each point of a surface its distance to the medial axis of the surface, as a local measure of sampling density. Specifically, an ε -sample of a smooth surface Σ must include a sample point within ε times local feature size of every point of surface from that point. See Section 1.6.2 for a detailed description this framework. The input to the surface reconstruction problem is a point set P that is assumed to be an ε -sample of a target surface Σ .

All our algorithms use the critical points of the (squared) distance function h_P induced by P . By critical points we mean local extrema and saddle points. Luckily, these points are easily computable from the Delaunay triangulation of P . Although the distance function h_P is not smooth everywhere, there is a well developed theory of critical points for such functions. Chapter 2 is dedicated to the discussion of distance functions, the generalization of the gradient vector field for such functions, and the integration of this generalized gradient vector field into a continuous flow that plays a central role in many of our results.

A fundamental contribution of this thesis is relating the critical points of h_P to both Σ itself and its medial axis $M(\Sigma)$. We show that if the surface sample is dense enough, all critical points of h_P either reside close to Σ or close to $M(\Sigma)$. This implies that one can label the critical points of h_P as either *surface* or *medial axis* critical points depending on to which one of surface or medial axis each point is close. Interestingly, all types of critical points, including local maxima, can be located close to Σ . Chapter 3 states this result on separation of critical points and offers an algorithm for distinguishing the critical points of the two types from each other.

The separation of the critical points of h_P can be turned into an algorithm that reconstructs Σ from P in 3D. This is the second main contribution of this work. We can show that the reconstructed surface is homeomorphic to Σ and approximates it closely in terms of distance relative to the local feature size. Furthermore, we show that the reconstructed surface has normals close to the original one provided the input is a *tight* ε -sample of Σ . Tight sampling, also known as (ε, δ) -sampling adds an extra condition to the original relative ε -sampling in order to avoid local over-samplings. We use the tight sampling condition in proving the topological correctness of our algorithm.

Although the design of the reconstruction algorithm of Section 4.1 is based on the properties of the critical points of h_P and their associated structures, little of this insight is used in proving the provided guarantees for the output of the algorithm. In Section 4.2 we analyze the same algorithm (modulo a slight change in the choice of final output) using techniques and machinery we develop in Chapter 2. Using these tools we can prove, in a much more elegant manner, a result that is comparable to that of Section 4.1 but generalizes to arbitrary dimensions.

The full potential of the machinery of Chapters 2 and 3 is put to test in Chapter 6 where we engage in the analysis of a classical flow-based shape reconstruction algorithm called WRAP. This algorithm, proposed by Edelsbrunner, although comes originally with no particular guarantee, has appealing properties that have made it into a successful commercial software. We analyze this algorithm for the first time after more than a decade from its original introduction and show that it captures the topology of the sampled surface and nicely approximates it geometrically if the sampling density is high enough.




Figure 3: Instability of medial axis under perturbations of the boundary. Left: a shape and its medial axis. Right: a small change in shape substantially changes its medial axis.

Approximation of the Medial Axis. As mentioned above, we can define the medial axis of a surface Σ in terms of the components of its complement. The medial axis of an open set $S \subset \mathbb{R}^n$ is the set of points in S that have at least two closest points in the boundary of S . In the sequel we sometimes refer to open subsets of \mathbb{R}^n as *shapes*. The medial axis $M(\Sigma)$ of Σ , consists of the union of the medial axes of the two shapes left in \mathbb{R}^n when Σ is removed from it.

As mentioned above, the medial axis of a bounded shape has the same homotopy type as the shape itself and therefore can be used to answer topological queries about the shape. This fact has important applications in shape analysis and motion planning among the other things.

Medial axes are *unstable* geometric objects. A tiny change on the boundary of a shape can dramatically change its medial axis (see Figure 3). Because of this, exact calculation of medial axis is seldom tried. In fact, numerical instability affects high accuracy calculations of medial axis negatively, often causing results that are less pleasing than a filtered version of themselves. For a survey of results on medial axis approximation refer to the notes at the end of Chapter 5.

The approach that we pursue in this work also falls into the Amenta-Bern framework and is presented in Chapter 5. We first compute a piecewise linear structure we call the (inner) *core* for the given sample P that is contained in the 2-skeleton (or the $(n - 1)$ -skeleton when working in \mathbb{R}^n) of the Voronoi complex of P . Not surprisingly, the idea and algorithm for computing the core come from the separation of the critical points of h_P and involves the flow map induced by the point set P .

Our main result in this direction states that the core and the medial axis are homotopy equivalent for dense enough samples in arbitrary dimensions. It turns out then that not only the core carries the topological type of the medial axis but it also allows a natural addition of further geometrically interesting components, using ideas from induced flows, without damaging its topological type. This allows us to turn any medial axis approximation

algorithm into one that in addition captures the topology of the medial axis as well.

Chapter 1

Background and Preliminaries

This first chapter is aimed at briefly summarizing the relevant previous knowledge. Accordingly, much of this material is not new but is included in the interest of orienting the reader. We devote some effort in keeping this monograph self-contained. However, an in-depth account of the introduced topics deserves many more pages than we can spare. The interested reader is thus encouraged to consult the more comprehensive sources on these subjects such as those cited at the end of this chapter.

1.1 Convexity and General Position

Basic geometric notions such as hyperplanes and affine subspaces of \mathbb{R}^n will play a central role in this manuscript. We therefore spend some time to introduce notions of general position and some fundamental properties of convex sets.

1.1.1 Linear and Affine Subspaces

Through \mathbb{R}^n denotes the n -dimensional Euclidean space. A point x in \mathbb{R}^n is an n -tuple of real numbers, $x = (x_1, \dots, x_n)$. The space \mathbb{R}^n is a vector space and one can speak of its linear subspaces, linear dependence of points, linear span of a set, and so forth. A *linear subspace* of \mathbb{R}^n is a subset closed under addition of vectors and under multiplication by real numbers. Geometric meaning of linear subspaces becomes clear by considering some examples. In \mathbb{R}^2 , linear subspaces are the origin alone, any line through the origin, and the entire space. In \mathbb{R}^3 , we have the origin, all lines and planes through the origin, and all of \mathbb{R}^3 . Given vectors $v_1, \dots, v_k \in \mathbb{R}^n$, a *linear combination* of v_1, \dots, v_k is any vector of the form

$$\alpha_1 v_1 + \dots + \alpha_k v_k,$$

where $\alpha_1, \dots, \alpha_k$ are real numbers. The subspace spanned by a set of vectors v_1, \dots, v_k , is the set of all linear combinations of these vectors.

Every linear subspace of \mathbb{R}^n must pass through the origin. *Affine subspaces* are translations

of linear subspaces; they are of the form $L + x$ where L is a linear subspace and x is a vector in \mathbb{R}^n . To understand the meaning of an *affine combination* of a set of points $a_1, \dots, a_k \in \mathbb{R}^n$ we translate these points by, say, $-a_k$ so that a_k becomes the origin and then talk about the linear combinations of the translated points; an affine combination of points a_1, \dots, a_k in \mathbb{R}^n , is any point

$$\alpha_1 a_1 + \dots + \alpha_k a_k,$$

such that $\alpha_1, \dots, \alpha_k$ are real numbers and

$$\alpha_1 + \dots + \alpha_k = 1.$$

The *affine hull* of a set $X \subset \mathbb{R}^n$ of points, denoted $\text{aff } X$, is defined as the intersection of all affine subspaces of \mathbb{R}^n that contains X . It is not hard to see that the affine hull of X is the set of all affine combinations of the points in X .

A set of vectors v_1, \dots, v_k in \mathbb{R}^n are *linearly dependent* if there are coefficients $\alpha_1, \dots, \alpha_k$, not all zero, such that

$$\alpha_1 v_1 + \dots + \alpha_k v_k = 0, \tag{1.1}$$

and *linearly independent* if no such coefficients exist, i.e. Equation (1.1) implies that $\alpha_1 = \dots = \alpha_k = 0$. Again with a translation of origin *Affine dependence* and *affine independence* are defined accordingly; points a_1, \dots, a_k are affine dependent if one can find coefficients $\alpha_1, \dots, \alpha_k$ such that (1.1) holds and in addition $\alpha_1 + \dots + \alpha_k = 0$ (note that for affine combinations this sum must be 1 but for affine dependence it must be 0).

Affine subspaces of \mathbb{R}^n of certain dimensions have specific names. An $(n - 1)$ -dimensional affine subspace is called a *hyperplane* while the word *plane* is a 2-dimensional affine subspace of \mathbb{R}^n for any n . A 1-dimensional subspace is a *line* and a k -dimensional subspace is called a *k-flat*. A hyperplane in general is specified by a linear equation of the form

$$a_1 x_1 + \dots + a_n x_n = b,$$

or simply as $\langle a, x \rangle = b$ where $\langle a, x \rangle$ defined as

$$\langle a, x \rangle = \sum_{i=1}^n a_i x_i$$

denotes the *scalar*, *inner*, or *dot product* of the vectors $a = (a_1, \dots, a_n)$ and $x = (x_1, \dots, x_n)$. So, a hyperplane can be expressed as the set

$$\{x \in \mathbb{R}^n : \langle a, x \rangle = b\}, \tag{1.2}$$

where $a \in \mathbb{R}^n \setminus \{0\}$ and $b \in \mathbb{R}$. A (closed) half-space of \mathbb{R}^n is a set of the form

$$\{x \in \mathbb{R}^n : \langle a, x \rangle \geq b\},$$

for which the hyperplane in (1.2) is the *boundary*.

General k -flats can be given as intersection of k hyperplanes. Equivalently, they can be regarded as solutions of systems of linear equations of the form

$$Ax = b,$$

where $x \in \mathbb{R}^n$ is regarded as column vector and A is a $k \times n$ matrix and $b \in \mathbb{R}^k$.

General Position. Many of the proofs in this thesis make an assumption of “general position” for point-sets (or other objects) they deal with. Often, this assumption is made without being stated. Intuitively, general position means that no “unlikely coincidences” happen in the point-set. For example if three points are chosen on the plane, with no special intentions, say “randomly”, then they are unlikely to lie on a common line. For points in a plane, we always require that no three points be on the same line. In general, for points in \mathbb{R}^n we assume that no “unnecessary” affine dependence exists.

There is no fully standard definition of general position. Depending on the context the above definition may be augmented or even replaced by other assumptions on non-degeneracy. For example, one may require for points on the plane that no four points lie on the same circle or that no two points have the same x -coordinate. Of course, not every condition passes as a general position condition. A determinant factor in accepting an assumption on the input configuration as a general position assumption is that the general position configurations can always be found in abundance arbitrarily close to every possible configuration. In other words, almost all small movements (*perturbations*) of any given configuration must result a configuration in general position. More formally, the set of general position configurations in any neighborhood of every possible configuration must have measure one.

1.1.2 Convexity

A set $C \subseteq \mathbb{R}^n$ is called *convex* if it has no “dips” or “cavities” or more accurately if for every pair of points $x, y \in C$, the line segment \overline{xy} connecting x and y , i.e. the set

$$\{\alpha x + (1 - \alpha)y : \alpha \in [0, 1]\}$$

is contained in C .

The intersection of any arbitrary family of convex sets is clearly convex. The *convex hull* of a set $X \subseteq \mathbb{R}^n$, denoted $\text{conv } X$ is the intersection of all convex subsets of \mathbb{R}^n that contain X .

A *convex combination* of points $p_1, \dots, p_k \in \mathbb{R}^n$ is an affine combination of p_1, \dots, p_k in which all the assigned coefficients are non-negative. In other words, a point q is a convex

combination of p_1, \dots, p_k if q can be written as

$$a = \alpha_1 p_1 + \dots + \alpha_k p_k$$

with $\alpha_i \geq 0$ for all $1 \leq i \leq k$, and furthermore $\alpha_1 + \dots + \alpha_k = 1$. It is an easy exercise to verify that the convex hull of a set of points X is exactly the set of all convex combinations of the points in X .

A basic yet important results about convex sets is the separability of disjoint convex sets by hyperplanes.

Theorem 1.1 (Separation Theorem) *Let $C, D \subseteq \mathbb{R}^n$ be convex sets with $C \cap D = \emptyset$. Then there is a hyperplane h of \mathbb{R}^n such that C is contained in one of the half-spaces determined by h and D is in the opposite half-space.*

An important Corollary of the Separation Theorem that we use in this thesis is the following.

Corollary 1.2 *Let $C \subseteq \mathbb{R}^n$ be a closed convex set and let $x \notin C$ be a point. Let $y \in C$ be a closest point of C to x . Then the hyperplane h through x and orthogonal to $x - y$ separates x from C .*

Proof. The proof is a simple application of the Separation Theorem on C and the open ball $B = B(x, \|x - y\|)$. The proof follows from the observation that the only hyperplane through y that has B completely on one side is the one tangent to B at y . ■

1.2 Complexes

Geometric objects are often made of simple pieces according to certain combinatorial rules. As such, they can be described as *complexes* with their constituent *cells*, which are usually polytopes or simplices. The incidence relation between cells of a complex may be governed by many constraints of topological and combinatorial nature. In this section, we introduce a few classes of complexes that we will find useful in the following chapters.

1.2.1 Simplicial and Abstract Simplicial Complexes

Simplicial complexes generalize and formalize the geometric notion of a triangulation in the plane. A building blocks of a simplicial complex is called a *simplex*. Because of their combinatorial nature, simplicial complexes have been used broadly to analyze geometric and topological properties of shapes and spaces.




Figure 1.1: Left: a simplicial complex. Right: examples of collections of simplices that are not complexes.

Simplices. A finite set of points in \mathbb{R}^n is *affinely independent* if no i -dimensional affine subspace of \mathbb{R}^n contains more than $i + 1$ of them (and this is true for every i). A k -*simplex* σ is the convex hull of a set S of $k + 1$ affinely independent points. The *dimension* of σ is $\dim \sigma = k$. Every subset of \mathbb{R}^n of size $n + 2$ or more is affinely dependent. Therefore, the maximum dimension of a simplex in \mathbb{R}^n is n . It is customary to define the empty set as a simplex of dimension -1 , or a (-1) -simplex.

Faces and Cofaces. If $\sigma = \text{conv } S$ is a simplex, the convex hull of any subset $T \subset S$ is also a simplex τ . It is a subset of σ and is called a *face* of σ . The latter relation between σ and τ is denoted $\tau \leq \sigma$. If $\dim \tau = \ell$, then τ is an ℓ -*face* of σ . If $\tau \leq \sigma$ but $\tau \neq \sigma$, we write $\tau < \sigma$ and call τ a *proper face* of σ . If τ is a face of σ then σ is called a *coface* of τ .

Simplicial Complexes. A *simplicial complex* is a collection of faces of a finite number of simplices, any two of which are either disjoint or meet in a common face. Formally, a collection K of simplices is a simplicial complex if

1. whenever $\sigma \in K$ and $\tau \leq \sigma$ then $\tau \in K$, and
2. for any two simplices τ and σ in K , $\tau \cap \sigma \leq \sigma, \tau$.

Abstract Simplicial Complexes. The face-coface relation between simplices in a simplicial complex can be interpreted as a subset-superset relation between the vertices of those simplices. If we rid ourselves of the concern over geometric realizability of a simplicial complex and only look at each simplex as a set consisting exactly of its vertices, we get an abstract version of the simplicial complexes. Therefore, a finite system A of finite sets is an abstract simplicial complex if $\alpha \in A$ and $\beta \subseteq \alpha$ implies $\beta \in A$. In accordance to embedded simplicial complexes, the dimension of an abstract simplex α is defined to be $\dim \alpha = |\alpha| - 1$ where $|\alpha|$ represents the cardinality of the set α .

Incidence Poset and Graph. The simplices of a simplicial or abstract simplicial complex make a partial order under inclusion. We sometimes call this partial order the *incidence poset*, or the *incidence graph* of the complex. The latter terminology stems in the idea that a poset can be treated like an acyclic directed graph in which the vertices corresponding to simplices τ and σ are joined by an arc from τ to σ if $\tau < \sigma$.

1.2.2 Cell Complexes

Cell complexes generalize simplicial complexes by replacing simplices with topological balls called *cells*. A *k-cell* (or a *k-face*) in a cell complex is a closed *k*-dimensional ball. In a cell complex K , the intersection of two cells is itself a cell of K . The relative interiors of any two cells in K are disjoint and the boundary of every cell is a union of other cells in K .

We sometimes refer to a cell complex matching the above definition as a *closed cell complex* because its cells are closed balls. Sometimes cell complexes are defined with *open* cells, meaning that the cells are taken as open balls. The cells of an open cell complex K' are pairwise disjoint. The boundary of a cell is defined as the closure of the cell minus the cell itself. The requirement that boundary of a cell must be a union of other cells in the complex remains intact.

The *underlying space* of a complex K , i.e. the union of all cells in K , is denoted $|K|$ ¹.

There is a natural correspondence between open and closed cell complex in which every closed cell complex K gives rise to an open cell complex K' satisfying $|K| = |K'|$ and vice versa. The cells of K' are exactly the relative interiors of the cells of K .

1.3 Distance Functions

Throughout this thesis we focus solely on the *n*-dimensional Euclidean space \mathbb{R}^n endowed with the ℓ_2 norm. In other words, every point or vector $x \in \mathbb{R}^n$ can be written as $x = (x_1, \dots, x_n)$ where the *coordinates* x_1, \dots, x_n of x are real numbers. We do not distinguish points and vectors from one another; a point x is identically a vector $x-0$ where $0 = (0, \dots, 0)$ is the origin.

The *length* of a vector x (or equivalently, the distance from the origin to a point x) is given

¹The notation $|\cdot|$ is used in this text to represent both the cardinality of a set and the underlying space of a complex. Although complexes can be regarded as sets, we never use this notation to refer to the number of simplices in a complex.

by the ℓ_2 norm $\|x\|$ of x which is defined as

$$\|x\| = \left(\sum_{i=1}^n x_i^2 \right)^{1/2}.$$

The *distance* between two points $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ in \mathbb{R}^n , denoted $\|x - y\|$, must always be interpreted as the ℓ_2 normal of the vector $x - y$ which is

$$\|x - y\| = \left(\sum_{i=1}^n (x_i - y_i)^2 \right)^{1/2}.$$

By $B(c, R)$ where $c \in \mathbb{R}^n$ and $R \in \mathbb{R}$, we denote the *open* ball with center c and radius R .

$$B(c, R) = \{x \in \mathbb{R}^n : \|x - c\| < R\}.$$

The closure of $B(c, R)$ is denoted by $\overline{B}(c, R)$. Thus

$$\overline{B}(c, R) = \{x \in \mathbb{R}^n : \|x - c\| \leq R\}.$$

The distance between a point x and a set of points $S \subset \mathbb{R}^n$ is defined as

$$\text{dist}(x, S) = \inf_{y \in S} \|x - y\|.$$

When S is compact, the map

$$y \mapsto \|x - y\|$$

from S to \mathbb{R} is bounded and attains its minimum. Thus for every $x \in \mathbb{R}^n$, there is a point $y \in S$ satisfying $\|x - y\| = \text{dist}(x, S)$. As such, for a compact set S , the “infimum” in the definition of $\text{dist}(x, S)$ can be replaced with “minimum”.

We can also measure the distance between two sets A and B as the smallest distance between a pair of points of which one is in A and the other is in B . Equivalently, the distance between the two sets can be defined as the infimum of the distance to one of the sets over the points in the other.

$$\text{dist}(A, B) = \inf_{x \in A} \text{dist}(x, B) = \inf_{x \in A, y \in B} \|x - y\|.$$

As before, if both of the sets A and B are compact there is a pair (x, y) that realizes this infimum.

Sometimes we need to measure the difference between two sets A and B of points. Such a difference can be interpreted as the distance between the sets A and B . The *one-way*




Figure 1.2: Level sets of squared distance function induced by a set of (unweighted) points in the plane (left). Three dimensional epigraph of the the same distance function (right).

Hausdorff distance between A and B is defined as

$$\vec{d}_H(A, B) = \sup_{x \in A} \text{dist}(x, B).$$

Notice that under this definition $\vec{d}_H(A, B)$ and $\vec{d}_H(B, A)$ are not necessarily equal.

The *Hausdorff* distance between two subsets A and B of \mathbb{R}^n is defined as the maximum of the two one-way Hausdorff distances between A and B .

$$d_H(A, B) = \max \left\{ \vec{d}_H(A, B), \vec{d}_H(B, A) \right\}.$$

1.3.1 Squared Distance Functions

For reasons that become clear later, we would rather deal with the square of the distance to a point and not the distance itself. Any set of points S in \mathbb{R}^n induces a *squared distance function* $h_S : \mathbb{R}^n \rightarrow \mathbb{R}$ which assigns to every point of the spaces the square of its distance to S ,

$$h_S(x) = \inf_{y \in S} \|x - y\|^2 = \text{dist}(x, S)^2.$$

As with ordinary distances, when S is compact we can safely replace the “infimum” with “minimum” in the above definition. In this thesis, we will only deal with squared distance functions induced by finite point-sets which are of course compact.

For a point $x \in \mathbb{R}^n$, the set of points in S that realize $h_S(x)$ is denoted $A_S(x)$, i.e.

$$A_S(x) = \{y \in S : \|x - y\| = h_S(x)\}.$$

1.3.2 Weighted Points

Chapter 6 takes advantage of squared distance functions induced by points with real weights assigned to them. The weight of a point manipulates the distance (or rather the squared distance) to the the weighted point. Formally, a *weighted* point q is an ordinary point in \mathbb{R}^n furnished with a *weight* $w_q \in \mathbb{R}$. It is perhaps more rigorous mathematically to think of a weighted point q as a pair $(\tilde{q}, w_q) \in \mathbb{R}^n \times \mathbb{R}$. But with an abuse of terminology, we will refer to $q = (\tilde{q}, w_q)$ as a *weighted point in* \mathbb{R}^n . An *unweighted* point is assumed to have weight zero. The *squared distance* to a weighted point q is given by the function

$$\pi_q : \mathbb{R}^n \rightarrow \mathbb{R}, \quad x \mapsto \|x - \tilde{q}\|^2 - w_q.$$

The function π_q is sometimes called the *power* of q . It is convenient to interpret a weighted point q as a ball $B_q = B(\tilde{q}, \sqrt{w_q})$ centered at q and with radius $\sqrt{w_q}$ (at least when $w_q \geq 0$). The squared distance of a point x to q is negative inside this ball, zero on its boundary and, positive outside. Conversely, we sometimes talk about the power of a ball $B = (c, R)$, which one can interpret as the power of a weighted point $q = (\tilde{q}, w_q)$ where $\tilde{q} = c$ and $w_q = R^2$.

There are also geometric interpretations of the squared distance to q . For example, when x is outside the ball $B_q = B(q, \sqrt{w_q})$, $\pi_q(x)$ is the square of the distance between x and any point of tangency between a hyper-plane through x and tangent to B_q . When x is inside the ball B_q , $-\pi_q(x)$ is the square of the radius of the ball of intersection between B_q and the hyper-plane through x and orthogonal to $\tilde{q} - x$.

Given a discrete set Q of weighted points, the weighted squared distance to the set Q (See Figure 1.3) is defined by the map

$$h_Q : \mathbb{R}^n \rightarrow \mathbb{R}, \quad x \mapsto \min_{q \in Q} \pi_q(x) = \min_{q \in Q} (\|x - q\|^2 - w_q).$$

Given a set of weighted points $Q \subset \mathbb{R}^n$, and a point $x \in \mathbb{R}^n$, $A_Q(x)$ denotes the preimage of x under h_Q , i.e.

$$A_Q(x) = \{q \in Q : h_Q(x) = \pi_q(x)\}.$$

For a set Q of weighted points in \mathbb{R}^n , the underlying set of unweighted points is denoted by \tilde{Q} . In formula

$$\tilde{Q} = \{\tilde{q} : (\tilde{q}, w_q) \in Q\}.$$

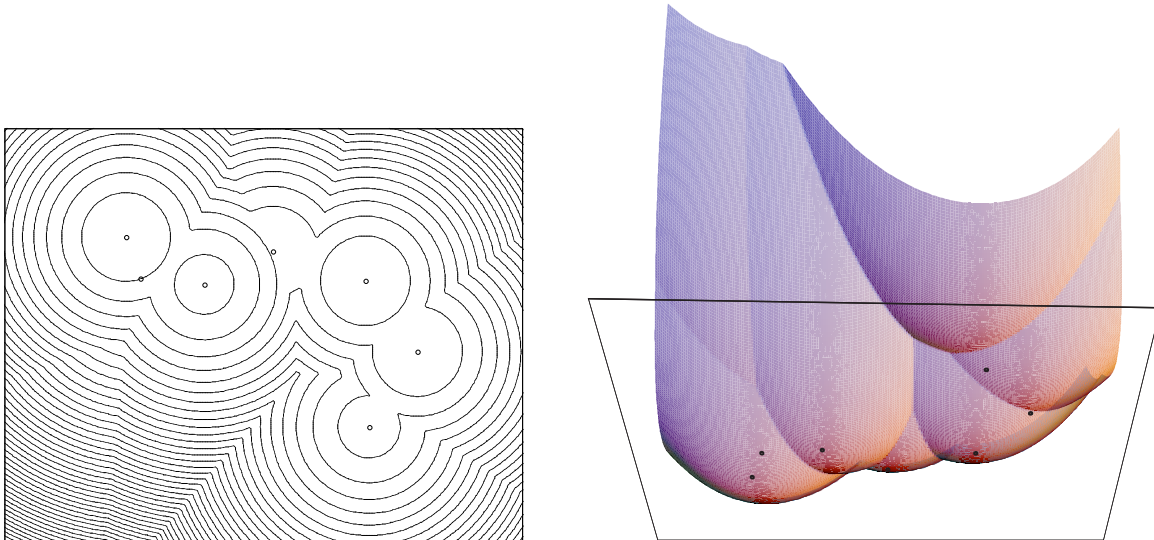


Figure 1.3: Levels sets of the squared distance function to a weighted set of points (left). Three dimensional epigraph of the the same distance function (right). Unlike the unweighted case, paraboloids are placed at different heights.

1.4 Voronoi and Delaunay Complexes

Consider a set P of points (also called *sites* in the context of Voronoi complexes) in the plane. Roughly speaking, the Voronoi diagram of the set P , subdivides the plane by determining for each point in P its *region of influence*. The region of influence of a point p (or the *Voronoi cell* of p as we shall call it), consists of all points in the plane to which p is a nearest neighbor in P . As we will see later, the Voronoi cell of every site $p \in P$ is a possibly unbounded convex polygonal region. Together, these regions cover the entire plane (See Figure 1.4).

The present section is a concise introduction to Voronoi diagrams and some of their important properties. We will see how these diagrams generalize to allow weighted points in Euclidean spaces of arbitrary dimensions. We will also introduce Delaunay complexes and investigate their duality with Voronoi diagrams.

1.4.1 The Voronoi Complex

Let P be a set of m points (*sites*) in \mathbb{R}^n . We assign to every point $p \in P$, a Voronoi cell V_p consisting of the set of points in \mathbb{R}^n that are no closer to any other point $q \in P$ than to p .

$$V_p = \{x \in \mathbb{R}^n : \|x - p\|^2 \leq \|x - q\|^2, \forall q \in P\}.$$




Figure 1.4: The Voronoi complex (diagram) of a set of points in \mathbb{R}^2 .

The above definition of V_p is equivalent to

$$\begin{aligned} V_p &= \bigcap_{q \in P} \{x \in \mathbb{R}^n : \|x - p\|^2 \leq \|x - q\|^2\} \\ &= \bigcap_{q \in P} \{x \in \mathbb{R}^n : 2\langle q - p, x \rangle - (\|q\|^2 - \|p\|^2) \leq 0\}. \end{aligned}$$

Thus, V_p is the intersection of at most m half-spaces and is therefore an n -dimensional (potentially unbounded) convex polytope.

The boundary of the Voronoi cell V_p of every site $p \in P$ is made of lower dimensional convex polytopes. A k -dimensional polytope on the boundary of V_p is called a k -cell. The Voronoi diagram (or *Voronoi complex*) of P , denoted $\text{Vor } P$, is the collection of Voronoi cells V_p of every $p \in P$ along with all their lower dimensional k -cells for $k = 0, \dots, n-1$. Every Voronoi cell V_p is a (non-proper) face of itself and because of that is called an n -cell. A 0-cell of $\text{Vor } P$ is a vertex of some V_p and is called a *Voronoi vertex*. A 1-cell is called a *Voronoi edge* and an $(n-1)$ -cell is referred to as a *Voronoi facet*.

For simplicity from now on, whenever we speak of the Voronoi complex of an unweighted point-set P , we assume that P is in *general position* in the following sense. It can be shown that an arbitrarily small random perturbation makes any point-set non-degenerate with probability one. In essence, this means that the a general position assumption is in a sense “reasonable” and applies to almost all situations.

Assumption 1.3 (Voronoi general position) A set $P \subset \mathbb{R}^n$ of unweighted points is in *Voronoi general position* if for all $0 \leq i \leq n-1$, no $i+3$ points in P lie on an i -sphere. For the rest of this text, we assume, whenever dealing with Voronoi complex of an unweighted point-set P , that P is in Voronoi general position.

Let p be any site in P . For every point x on the boundary of V_p at least two of the m inequalities

$$2\langle q - p, x \rangle - (\|q\|^2 - \|p\|^2) \leq 0, \quad q \in P \tag{1.3}$$

become equalities (including the trivial one from $q = p$). In other words, every point on the boundary of V_p has at least two closest points in P , p being one of them. It turns out that a k -cell of $\text{Vor } P$ on the boundary of V_p for some $p \in P$ is the locus of all points in the space for which exactly $n - k + 1$ particular inequalities in (1.3) become equalities. This entails that $\text{Vor } P$ is indeed a cell complex; the face-coface relationship is determined by the sets of inequalities that become equalities for individual cells.

For a different perspective on Voronoi complexes, consider the equivalence relation “ \sim ” between points in \mathbb{R}^n given by

$$x \sim y \quad \Leftrightarrow \quad A_P(x) = A_P(y).$$

Every equivalence class of “ \sim ” correspond to a subsets T of P ; for every x in the equivalence class, $A_P(x) = T$. Such an equivalence class consists of all those points x in the space (if any) that are at equal distance $h_P(x)$ from every point in T and are greater than $h_P(x)$ apart from every point in $P \setminus T$. Therefore, when $T = \{p\}$ is a singleton, the equivalence class of “ \sim ” corresponding to T is exactly the interior of V_p . In general, every nonempty class corresponding to a subset T of P of size k is the relative interior of an $(n - k + 1)$ -cell in $\text{Vor } P$. Thus the cells of $\text{Vor } P$ are closures of equivalence classes of “ \sim ”.

1.4.2 Weighted Points and Power Diagrams

Although the definition of Voronoi diagram involves distances between points of space and Voronoi sites, it only *compares*, for every point x , the distances between x and individual sites. Because of this, the definition remains sound even if the distances are measured in a different metric. It is therefore natural to consider other variants of Voronoi diagrams resulting from various ways of measuring distances. For instance, one can define ℓ_p -norm Voronoi diagrams for any $p \in \{1, 3, 4, \dots, \infty\}$ in addition to the traditional ℓ_2 version.

In fact, it is even possible to have different distance measures for individual sites and one particularly interesting instance of such a setting is when the sites are *weighted points*. Let $Q \subset \mathbb{R}^n$ be a set of weighted points². The Voronoi cell V_q of a point $q \in Q$ is defined in a fashion similar to the unweighted case.

$$V_q = \{x \in \mathbb{R}^n : \pi_q(x) \leq \pi_p(x), \forall p \in Q\}.$$

²To emphasize the distinction between weighted point-sets and unweighted ones, we try to use the symbol Q for weighted point-sets as opposed to P for unweighted ones.




Figure 1.5: The Delaunay complex (triangulation) of a set of points in \mathbb{R}^2 . Empty balls *witness* the presence of simplices they circumscribe.

This is equivalent to

$$\begin{aligned} V_q &= \bigcap_{p \in P} \{x \in \mathbb{R}^n : \|x - q\|^2 - w_q \leq \|x - p\|^2 - w_p\} \\ &= \bigcap_{p \in P} \{x \in \mathbb{R}^n : 2\langle p - q, x \rangle - (\|q\|^2 - \|p\|^2) - (w_p - w_q) \leq 0\}. \end{aligned}$$

As in the unweighted case, V_q is the intersection of at most $m = |Q|$ half-spaces and is therefore a (potentially unbounded) convex polytope. However, unlike the unweighted case, it is possible that the Voronoi cell V_q of a point $q \in Q$ is empty or does not contain q itself. The Voronoi complex $\text{Vor } P$ is defined in exactly the same way as in the unweighted case and has the same properties. Notice that if all the points in Q have the same weight c , then $\text{Vor } Q = \text{Vor } \tilde{Q}$ where \tilde{Q} is the unweighted version of Q . More generally, if we add the same amount c to the weight of every point $q \in Q$, the Voronoi cell V_q of every point remains the same. As a result, one can always add a large enough constant c to all the weights so that the resulting weights are all non-negative. Therefore, from this point on, we only consider sets of weighted points in which all the weights are non-negative.

1.4.3 The Delaunay Complex

The *Delaunay complex* $\text{Del } Q$ of a set of weighted points Q is defined as a *dual* to the Voronoi complex $\text{Vor } Q$. It is defined as

$$\text{Del } Q = \{\text{conv } S : S \subseteq Q, \bigcap_{q \in S} V_q \neq \emptyset\}.$$

The cells of Q are therefore convex hulls of subsets of Q . It can be verified that $\text{Del } Q$ is indeed a cell complex. The underlying space $|\text{Del } Q|$ of $\text{Del } Q$ is contained in $\text{conv } \tilde{Q}$. In fact, this containment can be shown to be equality.

An important observation on the relationship between Voronoi faces and their dual Delaunay faces makes the statement of the following proposition.

Proposition 1.4 *Let Q be a set of (possibly weighted) points in \mathbb{R}^n . Let $V \in \text{Vor } Q$ be a Voronoi face and let $D \in \text{Del } Q$ be the Delaunay face dual to V . Then the affine hulls of V and D are orthogonal.*

Proof. Let $\{q_0, q_1, \dots, q_i\} \subseteq Q$ be the set of vertices of D . Without loss of generality assume that q_0 is the origin and therefore the affine hull of D is a linear subspace of \mathbb{R}^n . For every other vertex q_i of D , V is contained in the (weighted) bisector of the segment $\overline{q_0 q_i}$ which is a hyperplane of \mathbb{R}^n orthogonal to the segment in question. The set D is convex and this entails that the set of vectors $\{q_i - q_0 : i \geq 1\}$ span the affine hull of D . All of these vectors are orthogonal to $\text{aff } V$ and this implies that the affine hulls of V and D are orthogonal. ■

The duality of the Delaunay complex and Voronoi complex of a set P of unweighted points gives an alternate definition for the complex $\text{Del } P$.

Empty Balls Property. Let P be a set of weighted points in the \mathbb{R}^n . Then for any subset $T \subset P$, the convex hull of T is included in $\text{Del } P$, if and only if, there is a closed ball containing the point in T in its boundary and no points of P in its interior. (See Figure 1.5). Such a ball is called the *Delaunay ball* witnessing T (or $\text{conv } T$). A *maximal* Delaunay ball is one that witnesses a set T such that no proper superset of T is witnessed by any other Delaunay balls. It is worth noticing that the circumcenters of maximal Delaunay balls are exactly Voronoi vertices (See Figure 1.6).

Another important property of the Delaunay complex of an unweighted set of points is that it is a simplicial complex. This is a direct consequence of our Voronoi general position assumption.

1.4.4 Polarity and Voronoi and Delaunay Polytopes

In the sequel, let us imagine the space \mathbb{R}^n as the hyperplane $x_{n+1} = 0$ in \mathbb{R}^{n+1} . We will call the hyperplane $x_{n+1} = 0$ our *workspace* and casually refer to the direction x_{n+1} as *vertical*. Thus we say a point (x_1, \dots, x_n, y) is *above* (*below*) another point (x_1, \dots, x_n, z) if $y > z$ ($y < z$).

The *epigraph* of a real-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is the set of points in \mathbb{R}^{n+1} that lie on




Figure 1.6: Voronoi vertices are centers of maximal Delaunay balls.

the graph of the function f . More formally,

$$\text{epi } f = \{(x_1, \dots, x_n, f(x_1, \dots, x_n)) : (x_1, \dots, x_n) \in \mathbb{R}^n\}.$$

The $(n + 1)$ -tuple $(x_1, \dots, x_n, f(x_1, \dots, x_n))$ is also denoted by the shorter notation of $(x, f(x))$.

The epigraph of the squared distance to the origin of \mathbb{R}^n is the paraboloid

$$\Pi_0 = \{(x, \|x\|^2) : x \in \mathbb{R}^n\}.$$

We call this paraboloid the *unit paraboloid*. The epigraph of the squared distance to a weighted point $q = (\tilde{q}, w_q)$ is

$$\Pi_q = \{(x, \pi_q(x)) : x \in \mathbb{R}^n\}.$$

It is easy to see that Π_q is the translation of Π_0 by the vector $(\tilde{q}, -w_q)$, i.e.

$$\Pi_q = \Pi_0 + (\tilde{q}, -w_q).$$

When Q is a set of weighted points in \mathbb{R}^n , the epigraph of h_Q is the *lower envelope* (in the direction x_{n+1}) of the collection

$$\Pi_Q = \{\Pi_q : q \in Q\},$$

of translations of the unit paraboloid for individual points in Q . Here, the lower envelope of a collection $\{E_1, \dots, E_k\}$ in which E_i is the epigraph of a function f_i , $i = 1, \dots, k$, is defined as $\text{epi } \min_{i \in \{1, \dots, k\}} f_i$.




Figure 1.7: The projection of the intersection of Π_q and the lower envelope of Π_Q (the contribution of Π_q to the lower envelope) is the Voronoi cell V_q of q .

In particular, this trivially shows that h_Q is a continuous function for any set Q of weighted points. The Voronoi cell V_q of every point $q \in Q$ is precisely the projection of the intersection of Π_q and the lower envelope of Π_Q (that is, the contribution of Π_q to this lower envelope) along x_{n+1} into the workspace \mathbb{R}^n . (See Figure 1.7).

Let us now consider a transformation on \mathbb{R}^{n+1} and given by

$$(x_1, \dots, x_n, x_{n+1}) \mapsto \left(x_1, \dots, x_n, x_{n+1} - \sum_{i=1}^n x_i^2 \right).$$

In other words, the above transformation drops every point of space vertically by the square of its distance to the vertical axis. The image of the workspace $x_{n+1} = 0$ under this transformation is the vertical mirror image of Π_p , i.e. $-\Pi_p$. More interestingly, the image of a paraboloid Π_q is a hyperplane H_q which is the epigraph of a linear map $\eta_q : \mathbb{R}^n \rightarrow \mathbb{R}$ where

$$\begin{aligned} \eta_q(x) &= \pi_q(x) - \|x\|^2 \\ &= \|x - \tilde{q}\|^2 - w_q - \|x\|^2 \\ &= -2\langle x, \tilde{q} \rangle + \|\tilde{q}\|^2 - w_q. \end{aligned}$$

It is not hard to verify that for any point $q \in Q$ (or equivalently, for any translation Π_q of the unit paraboloid Π_0) H_q is the hyperplane tangent to Π_q at $x = 0$. We can also apply the transformation to all paraboloids in Π_Q simultaneously. Paraboloids go to hyperplanes,

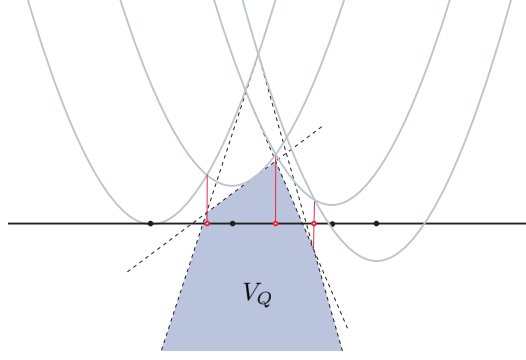


Figure 1.8: The Voronoi polytope of set of weighted points in \mathbb{R}^n . The Voronoi complex of the point-set is the the projection of the the boundary complex of the Voronoi polytope back to \mathbb{R}^n .

intersection of paraboloids got to intersections of hyperplanes, and the lower envelope of the paraboloids goes to lower envelope of the hyperplane since this transformation clearly maintains the vertical order or points.

If we replace each hyperplane H_q with the halfspace bounded by H_q from above, the intersection of all these halfspaces is a convex polyhedron V in \mathbb{R}^n . Since our transformation moves points vertically, the projection in to the workspace of the lower envelope of paraboloids and the lower envelope of hyperplanes are the same. In particular, the projection of k -dimensional face of the polytope V_Q is a k -cell of $\text{Vor } P$. We call the polytope V_Q , the *Voronoi polytope* of Q (See Figure 1.8).

As we saw the above transformation maps every weighted point q (or equivalently paraboloid Π_q) into a hyperplane H_q that is the epigraph of the function $\eta_q : \mathbb{R}^n \rightarrow \mathbb{R}$ given by

$$\eta_q(x) = -2\langle x, \tilde{q} \rangle + \|\tilde{q}\|^2 - w_q.$$

Every non-vertical hyperplane can be written as the epigraph of a function of the above form for some point $q = (\tilde{q}, w_q)$. The *polar point* of a non-vertical hyperplane H_q is the point

$$H_q^* = (\tilde{q}, -\|\tilde{q}\|^2 + w_q)$$

in \mathbb{R}^{n+1} . Therefore, *polarity* establishes one-to-one correspondence between points and non-vertical hyperplanes in \mathbb{R}^{n+1} . An important observation about this relation is the *order reversal property* which states that if the polar point η_1^* of a hyperplane H_1 is above, on, or below a hyperplane H_2 , then the polar point of H_2 , i.e. H_2^* is above, on, or below H_1 .

Now, let Q be any set of (possibly) weighted points in \mathbb{R}^n . For each $q \in Q$, let $q^* \in \mathbb{R}^n$ be the polar point of the hyperplane H_q , i.e. $q^* = H_q^*$ and let $Q^* = \{q^* : q \in Q\}$. We define the *Delaunay polytope* D_Q of Q as the polytope that is the intersection of all halfspaces of \mathbb{R}^{n+1} which non-vertical boundary that include Q^* . The boundary of D_Q extends to infinity




Figure 1.9: The Delaunay polytope of set of weighted points in \mathbb{R} . The Delaunay complex of the point-set is the projection of the boundary complex of the polytope.

in the $-x_{n+1}$ direction. A hyperplane H is said to *support* a polytope if H intersects the boundary of the polytope but is disjoint from its interior. The following proposition follows from the definitions of polytopes V_Q and D_Q (See Figure 1.9).

Proposition 1.5 *A hyperplane H supports D_Q if and only if the polar point H^* of H is on the boundary of V_Q .*

If we move a hyperplane H all the while supporting the Delaunay polytope D_Q , the polar point H^* of H traces the boundary of V_Q . It can be observed that a k -dimensional face of V_Q corresponds to an $(n - k)$ -dimensional face of D_Q . There is however the exception of the vertical faces of D_Q that correspond to no faces of V_Q . It can be shown that the projection into the workspace of the boundary complex of D_Q is exactly the Delaunay complex of Q . The vertical faces of D_Q project to the boundary of the convex hull of Q .

The important observation to make after the above discussion is that, essentially, the projection into \mathbb{R}^n of the boundary complex of every convex polytope in \mathbb{R}^{n+1} (that is unbounded from below and has no vertical faces) is the Voronoi complex of some set of weighted points. Similarly, the projection into \mathbb{R}^n of the boundary complex of every polytope in \mathbb{R}^{n+1} (with only vertical unbounded faces) is the Delaunay complex of some set of weighted points in \mathbb{R}^n . Consequently, the Voronoi complex of any set of points, is “similar to” the Delaunay complex of some other set of points and vice versa; if we replace the cells of the Voronoi complex of the former set with their intersections with its convex hull, then the equality becomes exact.

A special case of this relationship between pairs of point-sets will be of particular interest to us in Chapter 6. Namely, we explore for an unweighted point-set P , the point-set Q whose Voronoi complex is similar to the Delaunay complex of P . It is not hard to observe that the point-set Q must be the set of Voronoi vertices of $\text{Del } P$. The weight of any point $q \in Q$ is




Figure 1.10: The functions h_Q (left) and h_{Q^+} (right).

the square of its distance to P , i.e. the circumradius of the Delaunay ball centered at q . In other words if we treat the weighted points in Q as balls, the ball corresponding to a Voronoi vertex q is exactly the Delaunay ball centered at q . The bisector of two weighted points $q_1, q_2 \in Q$ that are centers of intersecting Delaunay balls, is the hyperplane that contains the intersection of the two balls. This is exactly the hyperplane of the Delaunay face dual to the edge q_1q_2 of $\text{Vor } P$. Figure 1.10 (left) shows a 1-dimensional example of this setting. The black curve in the figure is the graph of the squared distance function to an unweighted point-set P . The Voronoi vertices of $\text{Vor } P$ are shown with red hollow bullets. We weight each of these vertices by their vertical distance to the black curve (their squared distance to P) and call the resulting set of weighted points Q . The red curve is the graph of the squared distance to the set Q . The projection of the red curve to the real line results three intervals. The middle interval connects the middle two points of P . It is indeed a 1-cell in $\text{Del } P$. However, the other two intervals are not exactly 1-cells of $\text{Del } P$ but become so if we clip (intersect) them with the convex hull of P . Alternatively one can imagine to add to Q a symbolic weighted point at infinity with infinite weight. The resulting set is denoted Q^+ . It is then possible to assume that the parabola placed infinitely low at infinity reaches the real line at the extreme points of P , altering the graph of $\text{Vor } Q$ only outside $\text{conv } P$. The graph of Q^+ can be seen in Figure 1.10 (right). The vertical dotted lines are pieces of the parabola at infinity.

1.5 An Introduction to Topology

1.5.1 Topological Spaces

A *topology* on a set X is a collection $T \subseteq 2^X$ of subsets of X that satisfies three conditions:

1. The emptyset \emptyset and X are both in T .
2. The intersection of $S_1 \cap S_2$ of any pair of sets S_1 and S_2 in T is also in T .
3. The union of any collection of elements in T is also in T .

The elements of T are called the *open sets* of X . The pair (X, T) is called a *topological space*. When the topology T of a topological space (X, T) is understood, we simply refer to X as a topological space.

Thus the definition implies that finite intersections and arbitrary unions of open sets are also open. The complement $X \setminus S$ of an open set S is a *closed set*. Note that the definition allows for a subset of X to be open, closed, open and closed, or neither.

Subspace Topology. Every subset of a topological space can be given the *subspace topology* in which the open sets are the intersections of the open sets of the larger space with the subset. More formally, if (X, T) is a topological space and $X' \subset X$, then (X', T') where

$$T' = \{S \cap X' : S \in T\},$$

is also a topological space with subspace topology inherited from (X, T) .

Interior, Closure, Boundary, and Neighborhood. The *interior* $\text{int } S$ of a set $S \subseteq X$ is the union of all open sets contained in S . Since arbitrary union of open sets is open, the interior of any set is an open set. The *closure* $\text{cl } S$ of S is the intersection of all closed sets that contain S . Closed sets are complements of open sets. The intersection of a collection of closed sets is the complement of union of their (open) complements which is itself open. Thus the closure of every set is closed. The *boundary* of a set S is $\partial S = \text{cl } S \setminus \text{int } S$. A *neighborhood* of a point x is any subset A of X that contains x in its interior $\text{int } A$.

Continuous Functions and Homeomorphisms. A function between topological spaces is said to be *continuous* if the inverse image of every open set is open. This is an attempt to capture the intuition that there are no “breaks” or “separations” in the function. A homeomorphism is a bijection f that is continuous and whose inverse f^{-1} is also continuous. Two spaces are said to be *homeomorphic* if there exists a homeomorphism between them. From the standpoint of topology, homeomorphic spaces are essentially identical.

Compact Sets. Given a topological space (X, T) , an *open cover* of a set $S \subseteq X$ is any collection A of open sets whose union is S , i.e.

$$\bigcup_{U \in A} U = S.$$

A *sub-cover* of an open cover A for a set S is any subset of A that is also a cover for S . A set S is called *compact* for a topology (X, T) if every open cover of S has a *finite* sub-cover. The compactness of a set is generally difficult to establish for arbitrary topological spaces. In the topology of \mathbb{R}^n (see below) compact sets are shown to be exactly the closed bounded sets. A

fundamental property of subsets of \mathbb{R}^n that are closed and bounded is that the the following classical theorem of calculus, known as the *Extreme Value Theorem*, holds for them.

Theorem 1.6 (Extreme Value Theorem) *Any continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ attains its maximum and minimum on any closed and bounded subset of \mathbb{R}^n .*

The term “bounded” is of course is not meaningful for subsets of abstract topological spaces. But the essential property of closed and bounded subsets of \mathbb{R}^n that makes the Extreme Value Theorem hold for them is captured by the notion of compactness for arbitrary topological spaces. The above theorem generalizes to the following version that is meaningful any topological space.

Theorem 1.7 *Let X and Y be topological spaces and let $f : X \rightarrow Y$ be a continuous map. If X is compact, then $f(X) = \{f(x) : x \in X\}$ is also compact.*

Homotopy Equivalence. Homeomorphisms are the most prominent forms of topological equivalence; topology is concerned with equivalences of spaces modulo continuous transformations and homeomorphisms are exactly capturing this notion. There are however alternate notions of topological equivalence between spaces. Homotopy equivalence is one such notion which is weaker than homeomorphism. Intuitively, two spaces X and Y are homotopy equivalent if they can be transformed into one another by bending, shrinking, and expanding operations. For example, a solid disk or solid ball is homotopy equivalent to a point, and $\mathbb{R}^2 \setminus \{(0, 0)\}$ is homotopy equivalent to the unit circle S^1 .

Two continuous functions from a topological space to another are called *homotopic* if one can be *continuously deformed* into the other. Such a continuous deformation deformation is called a *homotopy*. Formally, a *homotopy* between continuous functions $f, g : X \rightarrow Y$, where X and Y are topological spaces, is a continuous map

$$H : [0, 1] \times X \rightarrow Y,$$

such that for all points $x \in X$, $H(0, x) = f(x)$ and $H(1, x) = g(x)$. If there is a homotopy H between function f and g , then we write $H : f \simeq g$ or simply $f \simeq g$ when the homotopy H itself is not of interest and say f and g are *homotopic*. It can be easily verified that homotopy equivalence is an equivalence relation on the set of continuous functions with the same domain and range.

A simple type of homotopy that is particularly used in this thesis. Let $B \subset \mathbb{R}^n$ and let X be any topological space and let $f, g : X \rightarrow B$ be any two continuous maps with the property that for all $x \in X$, the line segment connecting $f(x)$ and $g(x)$ is contained in B . The homotopy $H : [0, 1] \times X \rightarrow B$ defined as

$$H(t, x) = (1 - t)f(x) + tg(x).$$

is called the *straight-line-homotopy* between f and g .

Two topological spaces X and Y are *homotopy equivalent* or have the same homotopy type (written $X \simeq Y$) if there are continuous maps $f : X \rightarrow Y$ and $g : Y \rightarrow X$, such that $g \circ f$ is homotopic to the identity map on X and $f \circ g$ is homotopic to the identity map on Y . The maps f and g are called homotopy equivalences.

Clearly, every homeomorphism is a homotopy equivalence as well but the converse is not true. Homotopy equivalence between spaces X and Y implies the X and Y share most of their topological characteristics; they have matching connected components, holes, tunnels, and cavities. Roughly speaking, the difference between homotopy equivalence and homeomorphism for topological spaces that are subsets of Euclidean spaces is that the dimension of the spaces are preserved under homeomorphism but can change under homotopy equivalence. For example, a ball and a single point are all homotopy equivalent even though they are not homeomorphic. A topological space that has the same homotopy type of as a single point is called *contractible*.

A Criterion for Homotopy Equivalence. One particular way of establishing the homotopy equivalence of two topological spaces X and Y is through using the following proposition.

Proposition 1.8 *Let X and $Y \subseteq X$ be arbitrary sets and let $H : [0, 1] \times X \rightarrow X$ be a continuous function on both variables satisfying the following three conditions.*

1. $\forall x \in X, H(0, x) = x,$
2. $\forall x \in X, H(1, x) \in Y,$ and
3. $\forall y \in Y, \forall t \in [0, 1], H(t, y) \in Y.$

Then X and Y have the same homotopy type.

Intuitively, the first argument of the map H can be interpreted as “time”. Using a simple reparameterization of the first argument, we can replace the interval $[0, 1]$ in the above Proposition with any interval $[0, T]$ where $T > 0$ is a real number. The closedness and in particular finiteness of the considered time interval $[0, T]$ is crucial to the validity of Proposition 1.8. The map H can therefore be thought to continuously move every point in X into Y as time changes from 0 to T . At time 0, the map $H(0, \cdot)$ is the identity of X . By time T , all the points of X arrive in Y . Finally, the points in Y stay in Y at all times.

Deformation Retraction. There is a special kind of homotopy equivalence that is relatively easy to visualize.

Let X be a topological space and let A be a subset of X . A continuous map $r : X \rightarrow A$ is called a *retraction* if the restriction of r to A is the identity map of A . Note that the retraction r maps X *onto* A . In this case, A is called a *retract* of X .

For a topological space X , a subset $A \subset X$ is a *deformation retract* of X if there is a retraction $r : X \rightarrow A$ such that the identity map of X is homotopic to r . The homotopy map H realizing this is called a *deformation retraction*. Intuitively, this means that X can be continuously deformed into A in a way that the points in A end up where they started. We say that A is a *strong deformation retract* of X if in addition, the points of A , $H[t, x] = x$ for every point $x \in A$ and $t \in [0, 1]$.

1.5.2 Topology of \mathbb{R}^n

The abstract definitions given above allow boundless examples of topological spaces. In this thesis we will only work with one particular example in this bunch as well as some of its subsets endowed with subspace topology (sometimes called the *relative topology of \mathbb{R}^n*). The domain of this topological space is the n dimensional Euclidean space \mathbb{R}^n . Its topology is the collection of all open sets in the sense of real analysis: open subsets of \mathbb{R}^n are those sets that contain an open ball neighborhood of every point they include. Here an open ball neighborhood of a point x is any ball $B(x, r)$ with $r > 0$. Thus subset $X \subseteq \mathbb{R}^n$ is *open* if for every point $x \in X$ there exists a radius $r_x > 0$, such that $B(x, r_x)$ is contained in X .

Interior. The geometric structure of \mathbb{R}^n leads to geometric interpretations of otherwise abstract notions such interior, closure, or boundary of sets. An *interior point* of a set $S \subseteq \mathbb{R}^n$ is a point that has an open ball neighborhood contained in S . The interior of a set $S \subseteq \mathbb{R}^n$ is then the set of all interior points of S . Similarly, a point $x \in \mathbb{R}^n$ is a *point of closure* for a set $S \subseteq \mathbb{R}^n$ if every open ball neighborhood of x intersects S (possibly in x itself). The *closure* of a set $S \subseteq \mathbb{R}^n$ is the set of all points of closure.

Boundary. The boundary points of a set $S \subseteq \mathbb{R}^n$ are points of closure of S that are not interior points. Therefore, every neighborhood $B(x, r)$ of a point $x \in \partial S$ must intersect S but must not be entirely contained in S . In other words, a boundary point of S is a point for which every neighborhood contains a point in S and a point outside S .

Relative Interior and Boundary. In many cases we deal with a subsets S of \mathbb{R}^n for which $\dim \text{aff } S < n$. In other words, the set S is contained in some lower-dimensional affine subspace of \mathbb{R}^n . The *relative interior* and *relative boundary* of a set S are respectively the interior and boundary of S relative to $\text{aff } S$. For example all the points of a line segment in

\mathbb{R}^3 are boundary points for it. But only the two endpoints of the segment make its relative boundary and the rest of its points form its relative interior.

Compact sets in \mathbb{R}^n find a simple characterization.

Theorem 1.9 (Heine-Borel Theorem) *A subset S of \mathbb{R}^n is compact if and only if S is bounded and closed.*

Finally, proving two subset A and B of \mathbb{R}^n homeomorphic can be simpler than finding a continuous map from A to B with continuous inverse.

Theorem 1.10 *If A is a compact subset of \mathbb{R}^n and there exists a continuous bijection between A and $B \subseteq \mathbb{R}^n$, then A and B are homeomorphic.*

In other words, to prove that two compact subsets A and B of \mathbb{R}^n are homeomorphic, it is enough to find a continuous map between A and B that is one-to-one and onto.

1.5.3 Manifolds and Submanifolds

Intuitively speaking, a *manifold* is a topological space that locally “looks” like \mathbb{R}^n . To make this notion more clear, try to imagine the perspective of a very small ant walking on the surface of a fairly large doughnut! At any point, the ant is unable to realize, only by looking around, that it is not walking on a plane.

More formally, a topological space M is called *locally Euclidean* of dimension n if every point of M has a neighborhood homeomorphic to an open ball in \mathbb{R}^n . An *n -dimensional manifold* or an *n -manifold* is essentially a locally Euclidean topological space of dimension n that satisfies two extra technical conditions (it is Hausdorff and second-countable) the descriptions of which we skip in this manuscript. We are particularly interested in *smooth manifolds*. The rigorous definition of *smooth* or more accurately *differentiable manifolds* is involved. However, we only deal with very a special type of smooth manifolds which are relatively easy to define properly, namely smooth manifolds that are *submanifolds* of \mathbb{R}^n .

In our familiar \mathbb{R}^3 , with coordinates (x, y, z) , a locus $z = F(x, y)$, where $F : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a continuous function, describes a (2-dimensional) surface, where as a locus of the form $y = G(x)$, with $G : \mathbb{R} \rightarrow \mathbb{R}$, describes a (1-dimensional) curve. These notions have higher dimensional analogues. A subset $M \subset \mathbb{R}^{n+r}$ is said to be an n -dimensional *submanifold* of \mathbb{R}^{n+r} if *locally* M can be described by giving r of the coordinates *differentiably* in terms of the n remaining ones. More precisely, give an point $p \in M$, a neighborhood of p on M can be described in some coordinate system $(x, y) = (x_1, \dots, x_n, y_1, \dots, y_r)$ of \mathbb{R}^{n+r} by r differentiable functions

$$y_i = f_i(x_1, \dots, x_n), \quad i = 1, \dots, r.$$

Let us clear what the terms *differentiable* and *smooth* mean. A function can fall in any of a number of differentiability classes. The class C^0 comprises all continuous functions. The class C^k comprises functions all whose partial derivatives are C^{k-1} functions. The class C^∞ also known as the *smooth functions* is that class of all functions that have continuous derivatives of all orders. There are classes of even smoother functions. For example C^ω or the class of *analytic functions* comprises all functions convergent Taylor series, i.e. for any point a and any radius r , the Taylor expansion of the function around any point a within any radius r converges to the function itself. When the differentiability of a manifold M is interpreted in class C^k (for $k \in \{0, 1, \dots, \infty, \omega\}$) M is called a C^k manifold.

1.6 Surfaces and their Medial Axes

A submanifold of \mathbb{R}^n of dimension $n - 1$, i.e. of *codimension 1*, is called a *hypersurface* or simply a *surface* in \mathbb{R}^n , although the turn “surface” is most commonly used for a hypersurface in \mathbb{R}^3 .

The set of points of a surface in \mathbb{R}^n that have an open neighborhood homeomorphic to \mathbb{R}^{n-1} is called the *interior* of the surface which is always non-empty. The complement of the interior is called the *boundary*. The boundary of surfaces in \mathbb{R}^3 is a union of closed curves. The simplest example of a surface with boundary is the closed disk in \mathbb{R}^2 ; its boundary is a circle. A surface with an empty boundary is called *boundaryless*. A *closed* surface is one that is boundaryless and compact. The 2-dimensional sphere or the 2-dimensional torus are examples of closed surfaces.

As mentioned before, in this thesis we shall only deal with smooth surfaces. Throughout, the term “smooth” must be interpreted as C^2 or smoother³. All the considered problems deal with an unknown *target* surface that is known only through a dense sample (See the following section). The *target* surface in question will always be denoted by Σ . For simplicity we assume throughout that the target surface Σ has only one connected component. However, this is not a restrictive assumption and surfaces with multiple components can always be easily handled with little or no extra effort (See section 3.4).

To a surface Σ one can associate the two *open* components of its complement $\mathbb{R}^n \setminus \Sigma$ which we refer to as the bounded or *inner* component or *shape*, and the unbounded or *outer* component or *shape*. Whenever the target surface is denoted by Σ , the inner shape associated to Σ is denoted by S and the outer one by S^* .

The distance function induced by the target surface Σ is involved in much of the calculations in our work. We denote this function by $s(\cdot)$ when the target surface Σ is understood from

³In fact for all applications in this text, a $C^{1,1}$ -smooth surface suffices. The class $C^{1,1}$ consists of surfaces that are C^1 continuous and in addition the normals to surface satisfy a Lipschitz condition.

the context. In formula for every $x \in \mathbb{R}^n$

$$s(x) = \text{dist}(x, \Sigma) = \min_{y \in \Sigma} \|x - y\|.$$

Notice that in contrast to unlike the way we treated discrete point-sets, we prefer to work with the actual distance to Σ and not the square of it.

Medial Axis. The *medial axis* $M(S)$ of the open set S is the set of all points in S that have at least two closest points in Σ , i.e.,

$$M(S) = \{x \in S : |A(x)| > 1\},$$

where

$$A(x) = \{y \in \Sigma : \|x - y\| = s(x)\}$$

is the set of closest points to x in Σ . Note that since Σ is compact, the Extreme Value Theorem implies that $A(x)$ is well-defined and non-empty for every $x \in \mathbb{R}^n$. The medial axis M of a surface Σ is the union of the medial axes of the inner and outer components S and S^* associated to Σ , i.e.,

$$M(\Sigma) = M(S) \cup M(S^*).$$

When Σ is understood as the target surface, $M(\Sigma)$ is simply put as M . We also call $M(S)$ the *inner* medial axis and $M(S^*)$ the *outer* medial axis of Σ , respectively. Thus, M consists of all points in \mathbb{R}^n that have at least two closest points in Σ .

The medial axis is of particular interest because of its topological properties.

Theorem 1.11 (Medial Axis Homotopy) *Every bounded open subset of \mathbb{R}^n is homotopy equivalent to its medial axis.*

In particular, for a surface Σ , the bounded shape S enclosed by Σ has the same homotopy type as its medial axis $M(S)$.

1.6.1 Medial Balls and Local Feature Size

It is possible for a general surface to approach infinitesimally close to its medial axis or even intersect with it. However, it can be shown that for closed C^2 (or smoother) surfaces, the Hausdorff distance between the surface and its medial axis is strictly positive.

A closed ball B is called *empty* with respect to a surface Σ if $\text{int } B \cap \Sigma = \emptyset$. An empty ball is *maximal* if it is not strictly contained in any other empty balls. Any point x of a smooth surface (C^2 -smooth or smoother) is on the boundary of exactly two maximal empty balls. These balls which we call *medial balls* are tangent to Σ at x . the interior of one of these balls




Figure 1.11: The medial balls ($B_\mu^+(x)$ and $B_\mu^-(x)$) and local feature size balls ($B_f^+(x)$ and $B_f^-(x)$) tangent to Σ at x .

is contained in S and the other in S^* . We call the former the *inner medial ball* and the latter the *outer medial ball* of x and denote them respectively by $B_\mu^-(x)$ and $B_\mu^+(x)$ and their radii by $\mu^-(x)$ and $\mu^+(x)$ which are respectively called the *inner* and *outer medial feature size* at x .

For a smooth surface Σ , there is a unique tangent hyperplane to every point $x \in \Sigma$. The unit vector normal to this hyperplane in the direction pointing toward S^* is denoted by n_x^+ and its negation (pointing toward S) by n_x^- .

A point on the line $\{x + tn_x^+ : t \in \mathbb{R}\}$ is called a *focal point* if $t = \kappa_i^{-1}(x)$ where $\kappa_i(x)$ is one of the principal curvatures of Σ at x . The centers of the ball $B_\mu^+(x)$ and $B_\mu^-(x)$ are on the segment connecting the closest focal points of x on the two sides of Σ . Furthermore,

$$\min \{\mu^+(x), \mu^-(x)\} \leq \min_i \left| \frac{1}{\kappa_i(x)} \right|.$$

In other words, the smaller of the two medial feature size is a lower bound for the smallest curvature radius at x . When the center of a medial ball tangent to Σ at a point x is not a focal point of x , that medial ball is tangent to Σ in at least one more point $y \neq x$ in which case the center of that medial ball is a medial axis point. Intuitively, a very small ball tangent to Σ at x , say at the exterior of Σ , can be enlarged, all the while kept tangent to Σ at x , by moving its center away from x in the direction n_x^+ . The growth of the ball continues until the moving center reaches a point where moving further forward causes the interior of the growing ball to intersect with Σ . This happens if either (1) the growing ball has become tangent to Σ in a point $y \neq x$, or (2) the center of the ball has reached a focal point of x .

and the growing ball is coinciding with a curvature ball of x .

The above observation motivates the definition of the *skeleton* of a surface Σ as the set of centers of all maximal empty balls with respect to Σ and is denoted $\text{sk } \Sigma$. Clearly, every medial axis point is also in the skeleton. However, we saw above that the converse is not true. Thus the medial axis is contained in the skeleton. It can be shown that skeleton itself is contained in the closure of the medial axis which is sometimes called the *cut locus* of Σ . Thus we have

$$M \subseteq \text{sk } \Sigma \subseteq \text{cl } M.$$

It is not hard to build specific example in which the above containments are strict. The topological properties of the medial axis are not necessarily shared with the skeleton or the cut locus. For example, there are piece-wise analytic surfaces for which the skeleton of the bounded enclosed shape does not have the same homotopy type as the shape itself.

Geometrically however, the medial axis and the skeleton differ only slightly. It is therefore convenient, whenever only the geometry of the medial axis is concerned, to think of the medial axis and the skeleton as the same object. Specifically, it is helpful to consider the the centers of all medial balls, including those that are focal points of some surface points, as medial axis points.

Local Feature Size. The function

$$f : \Sigma \rightarrow \mathbb{R}, \quad x \mapsto \text{dist}(x, M)$$

is called the *local feature size*. The centers of the medial balls tangent to Σ at a point $x \in \Sigma$ are in the medial axis M of Σ (unless they are focal points of x in which case they are contained in the closures of M and therefore every neighborhood of them intersects M). This means that there is a medial axis point within distance $\mu^+(x)$ from x and likewise, there is a medial axis point within distance $\mu^-(x)$ from x . From this we have

$$f(x) \leq \min \{ \mu^+(x), \mu^-(x) \}.$$

Thus if one places two closed balls $B_f^+(x)$ and $B_f^-(x)$, both of radius $f(x)$ and one in each side of Σ such that the boundaries of these two balls are tangent to Σ at x , then the interiors of both of these balls avoid Σ . More precisely, we have

$$B_f^+(x) \subseteq B_\mu^+(x) \quad \text{and} \quad B_f^-(x) \subseteq B_\mu^-(x).$$

Recall that each of $\mu^+(x)$ and $\mu^-(x)$ are smaller than the radius of the smallest curvature ball and $f(x)$ being smaller than both of these implies that

$$\frac{1}{f(x)} \geq \max_i |\kappa_i(x)|.$$

Because of this, local feature size is used as a conservative measure of local fineness or coarseness of surface features and can suggest the local density of a good sample around a surface point x ; the smaller $f(x)$, the higher the density of a sample near x should be in order to be useful in capturing of topological and geometric details of the surface.

Since underlying role of local feature size is to provide an approximation to the radius of the smallest curvature ball one should expect find it related to the local rate of change in surface normals. This relationship is provided by the following lemma.

Lemma 1.12 [2] *Let x and y be points on Σ with $\|x - y\| \leq \xi f(x)$ for $\xi \leq 1/3$. Then*

$$\angle(n_x^+, n_y^+) = \angle(n_x^-, n_y^-) \leq \frac{\xi}{1 - 3\xi}.$$

A function $g : A \rightarrow B$ with $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$ is k -Lipschitz if for all $x, y \in A$

$$\|g(x) - g(y)\| \leq k \cdot \|x - y\|.$$

Every k -Lipschitz function is uniformly continuous. In general for any set $S \subseteq \mathbb{R}^n$ the function $d(x, S)$ is 1-Lipschitz as a result of the triangle inequality:

$$\text{dist}(x, S) \leq \|x - y\| + \text{dist}(y, S).$$

In particular the local feature size function $f(\cdot) = \text{dist}(\cdot, M)$ is 1-Lipschitz. Because of this, $f(x)$ is often preferred over $\mu^+(x)$ or $\mu^-(x)$, or even $\min\{\mu^+(x), \mu^-(x)\}$ since none of μ^+ , μ^- , or their minimum are even continuous.

Now, let us look at an arbitrary point $x \in \mathbb{R}^n$. If $x \notin M$, then it has a *unique* closest point in Σ . This closest point to x is denoted by \hat{x} . For a point $x \in \Sigma$, $\hat{x} = x$. For any point $x \notin (\Sigma \cup M)$, the center of the medial ball tangent to Σ at \hat{x} and at the same side of Σ as x is represented by \check{x} . Although for a medial axis point \hat{x} is not well-defined, we take for such points \check{x} to be the same as x . Notice that \check{x} can be at infinity. This happens exactly when the hyperplane tangent to Σ at \hat{x} separates x and Σ .

For a point $x \notin \Sigma$, we defined the function $\mu(x)$ as the distance between \hat{x} and \check{x} . More precisely,

$$\mu(x) = \begin{cases} \|\hat{x} - \check{x}\| & x \notin M, \\ s(x) & x \in M. \end{cases}$$

Finally, for every point $x \notin \Sigma$ the function $m(x)$ is defined as the distance between x and \check{x} , i.e.

$$m(x) = \|x - \check{x}\|.$$

Thus for every $x \in \mathbb{R}^n \setminus \Sigma$, we have

$$\mu(x) = s(x) + m(x).$$




Figure 1.12: Left: a surface Σ and its medial axis M (only the inner medial axis is shown). Middle: the δ -tubular neighborhood Σ_δ of the surface. Right: the δ -tubular neighborhood M_δ of the medial axis (only the inner component shown).

δ -tubular neighborhoods. For a constant $\delta \leq 1$, define the δ -neighborhood of the medial axis, denoted M_δ as the set

$$M_\delta = \{x \in \mathbb{R}^n \setminus \Sigma : m(x) < \delta\mu(x)\}.$$

Similarly we define the δ -tubular neighborhood of Σ (See Figure 1.12) as a (non-uniform) thickening of Σ as

$$\Sigma_\delta = \{x \in \mathbb{R}^n \setminus M : s(x) \leq \delta f(\hat{x})\}.$$

The δ -reduced shapes S_δ and S_δ^* are defined as

$$S_\delta = S \setminus \Sigma_\delta \quad \text{and} \quad S_\delta^* = S^* \setminus \Sigma_\delta.$$

1.6.2 Criteria for Sampling Surfaces

A *Piece-wise linear* surface is a surface that is made of patches that are each contained in a hyperplane in \mathbb{R}^n . These patches are glued along their boundaries to make the surface. Naturally, intersections of such patches are contained in lower-dimensional affine subspaces of \mathbb{R}^n . A piece-wise linear surface can therefore be thought of as the underlying space of a cell complex in which every cells is contained in some affine subspace of \mathbb{R}^n .

All the problems considered in this thesis deal with a smooth *input* or *target* surface. This input surface is always understood through a finite sample. Clearly, as a discrete set of points, the sample does not share the topology of the surface. Indeed, it is the goal of surface reconstruction algorithms to construct a piece-wise linear surface that agrees with the target surface topologically.

It is clear that no nontrivial statement about the output of an algorithms, that takes for input a sample of a surface, can be made if the sample is defined as just a finite subset of the surface with no further requirements; a single-point taken from the surface fits this description of a sample but is trivially insufficient for producing any meaningful output.

Uniform ε -Sample. To allow provable guarantees, a sample needs to be *dense* enough meaning that it should not leave large areas of the surface unsampled. The most natural definition of a dense sample is that of a *uniform* or *non-relative* ε -sample. Given a surface Σ , a sample $P \subset \Sigma$ is called an uniform ε -sample of Σ , if for every $x \in \Sigma$, there is a sample point p in the ε -neighborhood of x . In other words, P is an ε -sample of Σ if

$$\forall x \in \Sigma : B(x, \varepsilon) \cap P \neq \emptyset.$$

Relative ε -Sample. A uniform ε -sample is sufficient for capturing the topology of a surface only if ε is chosen proportional to the size of the smallest features of the surface. For some intuition, consider a very small handle attached to a very large sphere. Regardless of how large the sphere is, it is the size of the handle that determines the density of a useful sample. Roughly speaking, if ε is not small enough, a uniform ε -sample of a sphere alone would pass as a valid ε -sample of the sphere with handle. If the handle is substantially small compared to the sphere, a good ε -sample has to cover the entire sphere very densely and this amounts to using a huge number of sample points. This motivates sampling of surfaces with varying densities; finer features are to be sampled more densely and coarser features can be sampled more sparsely.

Various measures of fineness of a surface features can be defined. The *local feature size* $f(x)$ of a surface point x defined as the the distance from x to the medial axis of the surface has proved to be a reliable measure. A relative sample requires every point x of a surface to be represented by a sample point that is not too far away from it compared to $f(x)$. More precisely, a sample $P \subset \Sigma$ is called a (relative) ε -sample if for all $x \in \Sigma$, there exists a sample point $p \in P$ within distance $\varepsilon f(x)$ from x , i.e.

$$\forall x \in \Sigma : B(x, \varepsilon f(x)) \cap P \neq \emptyset.$$

Tight (ε, δ) -Sample. Both uniform and relative ε -sampling criteria only put upper bounds on the distance between surface points and their closest sample points. In other words, they have no problem with *over-sampling*. In certain cases, this over-sampling can complicate the proofs substantially or even make some of the claims that are valid for “reasonable sample” invalid for an unevenly dense sample.

There are several ways of enforcing “evenness” of sampling. When dealing with uniform samples, one common extra condition that prevents over sampling is to require that no two sample points are too close to each other. Alternatively, it can be required that the ball $B(x, \varepsilon)$ around every surface point x (which is required to include a sample) has at most k sample points in it.

Likewise, in relative sampling, one can ask the ball $B(x, \varepsilon f(x))$ around every surface point x

to have at most k sample points in it. Alternatively, the distance between two sample points can be bounded from below. This latter condition is the one we use in this thesis. Formally, a variable (ε, δ) -sample of a surface Σ is a finite subset P of Σ if

1. P is an ε -sample of Σ , and
2. For any $p \in P$, there are no points of P within distance $\delta f(p)$ from p , i.e.

$$B(p, \delta f(p)) \cap P = \{p\}.$$

Bibliography and Remarks. Linear spaces (also called vector spaces) along with linear transformations and systems of linear equations make the main subjects of linear algebra. Vector spaces are a central theme in modern mathematics and linear algebra is widely used in abstract algebra and functional analysis and has a concrete representation in analytic geometry. It has extensive applications in engineering, natural sciences as nonlinear models are often approximated by linear ones. For a comprehensive introduction to linear algebra refer to [30] or [52]. In geometry, affine spaces are seen in the study of *affine transformations*. An affine transformation is a linear transformation followed by a translation. In a geometric setting, affine transformations are precisely the functions that map straight lines to straight lines. In general, *affine geometry* is the geometry that avoids the involvement of any notions for origin, length, or angle, but relies on the notion of subtraction of points giving a vector. The notion of convexity and basics of convex geometry have been around since antiquity. Nonetheless, convex geometry and convex optimization are active field of research to this date. See the text of Barvinok [9] for a good introduction to this subject.

Simplicial complexes generalize planar triangulations which on their own have numerous applications. Triangulating a domain or more generally subdividing it into a simplicial complex is a standard discretization step in finite element methods for approximate solution of partial differential equations or integral equations. Modeling surfaces as 2-dimensional simplicial complexes is a practical way of representing them and is broadly used in computer graphics, simulations, games, and many other places. The output of surface reconstruction algorithms studied in this thesis are also simplicial complexes.

Abstract simplicial complexes are prominent in algebraic topology since as topological spaces these are found to be the easiest to deal with, at least in terms of concrete calculations. See the book of Hatcher [48] for an excellent treatment of abstract simplicial complexes and their application in simplicial homology.

Cell complexes are studied in topology under the name of CW complexes. A CW complex is a type of topological space introduced by J.H.C. Whitehead [68] to meet the needs of homotopy theory. The idea is to have a class of spaces that is broader than simplicial complexes but still retains a combinatorial nature, so that computational considerations are not ignored. The name CW stands for

closure-finite weak topology. The building block of a CW complex is called a *cell*. A closed k -cell is a topological space homeomorphic to a simplex, or equally the closed unit k -ball B^k defined as

$$B^k = \{x \in \mathbb{R}^k : \|x\| \leq 1\}.$$

If σ is a k -cell we denote by $\partial\sigma$, the subset of σ corresponding to the unit $(k-1)$ -sphere S^{k-1} which is

$$S^{k-1} = \{x \in \mathbb{R}^k : \|x\| = 1\}.$$

CW complexes are built using a basic operation called *attaching a cell*. Let X be a topological space, σ a k -cell and $f : \partial\sigma \rightarrow X$ a continuous map. We let $X \cup_f \sigma$ denote the disjoint union of X and σ quotiented out by the equivalence relation that each point of $x \in \partial\sigma$ is identified with $f(x) \in X$. The map f is called the *attaching map*. It should be emphasized that the attaching map must be defined on all of $\partial\sigma$ meaning that the entire boundary of sigma must be glued to X when attaching σ to X . A *finite CW complex* is a topological space X for which there is a finite nested sequence

$$\emptyset \subset X_0 \subset \cdots \subset X_n = X,$$

such that for each $i = 0, 1, \dots, n$, X_i is the result of attaching a cell to X_{i-1} (note that by this definition, X_0 has to be a 0-cell). The above nested sequence is called a *CW decomposition* of space X .

An important property of CW complexes is that attaching the same cell σ to a space X using two homotopic attaching maps f_1 and f_2 results two homotopy equivalent spaces $X \cup_{f_1} \sigma$ and $X \cup_{f_2} \sigma$. The connection between CW complexes and cell complexes as we defined in this chapter becomes clear from the following result.

Theorem 1.13 *Let*

$$\emptyset \subset X_0 \subset \cdots \subset X_n = X, \tag{1.4}$$

be a CW decomposition of a finite CW complex X . Then X is homotopy equivalent to a CW complex with a CW decomposition with precisely the same number of cells of each dimension as (1.4) and with the cells attached so that their dimensions form a nondecreasing sequence.

Voronoi diagrams are named after the Ukrainian mathematician George Voronoi because of his seminal work on the subject at the beginning of the twentieth century [67] although P. G. L. Dirichlet had similar results half a century earlier. There are also unpublished notes from René Decartes from early 17th century suggesting he was using these diagrams. Delaunay triangulations are named after Russian mathematician Boris Delaunay (or Delone) [31]. An in depth survey of Voronoi diagrams can be found in [8].

For thorough handling of the introduced topics in topology, the reader is encour-

aged to consult standard text books such as [48] or [57]. For an introduction to topological manifolds consult [54] and for differentiable manifolds see [40] or [15].

Medial axis of a surface (or a shape) was introduced by Blum [11] and has been subject to extensive studies. An associated concept is the *medial axis transform* (MAT) which consists of all medial axis points and the distance to the boundary of the shape from each point. Medial axis and MAT have applications in image analysis and computer vision [60], solid modeling [10], mesh generation and finite element analysis [64, 65], shape simplification [66], motion planning [39], etc.

In [25] Choi, Choi, and Moon, provided detailed description of the medial axis of an open subset of \mathbb{R}^2 that has a boundary consisting of a piece-wise analytic curve. Stability and finiteness properties of the medial axis of subanalytic open sets was explored by Chazal and Soufflet [24].

The homotopy equivalence of bounded open sets of \mathbb{R}^n and their medial axes (Theorem 1.11) was proven by Lieutier [56]. A special case of this theorem for open subsets of \mathbb{R}^n with C^2 -smooth boundaries, or open subsets of \mathbb{R}^3 with piece-wise C^2 -smooth boundaries was proven by Wolter [69].

The definition of local feature size function as the distance to the medial axis and the relative ε -sampling based on it is due to Amenta, Bern, and Eppstein [3]. The ε -sampling framework has received widespread attention and there are numerous results in shape, surface, and medial axis analysis and approximation that are based on this framework.

Among other sampling criteria, one must mention those that allow the presence of *noise* in the sample. By a *noisy sample* we mean a sample that is not a subset of the target surface and is only required to be suitably close to it. An ε -noisy sample of a surface Σ in the uniform sense can be simply described as any finite set P for which $d_H(\Sigma, P) \leq \varepsilon$. There are various versions of relative noisy sampling. For example one can define such a sample for a surface Σ as a set P such that

1. for every $p \in P$, $\|p - \hat{p}\| \leq \delta f(\hat{p})$, and
2. for every $x \in \Sigma$, there exists a $p \in P$ such that $\|x - p\| \leq \varepsilon f(x)$.

For simplicity, throughout this thesis we only work with noise-free samples, i.e. we always assume that the supplied sample is a subset of the target surface. However many of the presented results can be extended to allow noisy sample at the cost of requiring denser sampling.

Finally, we must mention a few recent approaches in sampling of non-smooth surfaces. As mentioned before, smooth surfaces have been the focus of most shape and surface reconstruction algorithms, simply because they are easier to deal with; the long-developed machinery of calculus loses a great deal of power and becomes a great deal more complicated for dealing with surfaces if the smoothness

assumption is relaxed in any way. However, many of the surfaces of practical interest are not smooth.

Boissonnat and Oudot [14] introduced the family of Lipschitz surfaces which properly include smooth surfaces but include also some surfaces with sharp edge and vertices. The k -Lipschitz radius $\text{lr}_k x$ of a point x on a (not necessarily smooth) surface Σ is defined as the maximum radius of a ball $B = B(x, R)$ for which the intersection $B \cap \Sigma$ of the surface and the ball is the epigraph of a k -Lipschitz function. A surface Σ is a k -Lipschitz if $\inf_{x \in \Sigma} \text{lr}_k x > 0$. An ε -sample of a k -Lipschitz surface Σ is a subset $P \subset \Sigma$ such that every point $x \in \Sigma$ has a point of $p \in P$ within distance $\varepsilon \text{lr}_k x$. In [14] the authors show that a sample with this specifications from a k -Lipschitz surface is almost as good for reconstruction of the surface, as an ε -sample is for smooth surfaces.

Yet another approach for sampling of non-smooth surfaces is introduced by Chazal, Cohen-Steiner, and Lieutier [21]. Their approach uses a parameterization of the notion of feature size which interpolates between local feature size at one end and *weak feature size* at the other. The weak feature size of a surface Σ (or rather of the shape S enclosed by Σ) is the smallest distance to Σ of a point $x \in M(S)$ that is contained in $\text{conv } A(x)$. This notion, introduced by Chazal and Lieutier [22] is shown to play a central role in topologically correct approximation of shapes. The parameterized feature size function, called the μ -reach is defined for a compact set K as the minimum distance between a point of K and the μ -medial axis of $\mathbb{R}^n \setminus K$ which is itself a filtered version of the medial axis. The sampling is uniform (but noise is allowed) and can guarantee homotopic approximation of a large class of compact sets in \mathbb{R}^n . We skip further details of this approach and encourage the interested reader to consult [21].

Chapter 2

Distance Flows

Distance functions, especially those induced by discrete samples of surfaces, are at the heart of this dissertation. However, we do not use such functions directly. Rather, we exploit the continuity of “flow maps” that result from integrating certain “generalized gradient” vector fields associated to these functions. In general, the gradient of a distance function is not defined everywhere. Nevertheless, except for a finite number of “critical points”, the distance function in question has a unique direction of steepest ascent everywhere. It turns out that a generalized gradient vector field can be characterized in terms of the Voronoi and Delaunay complexes of the inducing point-set. Such vector fields can then be integrated to result continuous flow maps. These flow maps and their associated structures are central to the algorithms that are presented in the rest of this thesis. A study of the properties of such flow maps allows us to analyze these algorithms in a somewhat uniform manner.

2.1 Generalized Gradients

We observed in Section 1.4.4 that the squared distance map h_Q induced by any set Q of weighted points in \mathbb{R}^n is continuous. However, h_Q is not a smooth function and in particular, its partial derivatives are not defined everywhere. We are particularly interested in the *gradient* ∇h_Q of h_Q . Recall that for a real-valued function f defined on \mathbb{R}^n , the gradient of f at a point x is the vector

$$\nabla f(x) = \left(\frac{\partial f(x)}{\partial x_1}, \dots, \frac{\partial f(x)}{\partial x_n} \right).$$

It is well-known that at a point x where the partial derivatives $\partial f(x)/\partial x_i$ all exist (and therefore $\nabla f(x)$ is defined) this vector points in the *direction of steepest ascent* of f at x ; this is the direction toward which the f increases the fastest. In other words the quantity

$$D_v f(x) = \lim_{t \rightarrow 0} \frac{f(x + tv) - f(x)}{t},$$

known as the *directional derivative* of f in direction v at x , is maximized over all *unit* vectors v when v is parallel to $\nabla f(x)$. Thus for any vector v ,

$$D_v f = \langle \nabla f, v \rangle.$$

The points x for which $\nabla f(x) = 0$ are called the *critical points* of function f ; the function f in such points has more than one directions of steepest ascent (local minima or saddle points) or has no such directions at all (local maxima).

In the interior of any full-dimensional Voronoi cell V_q , the function h_Q coincides with the power of q , π_q . The latter function is smooth and its partial derivatives of all orders exist and are continuous. In particular, the gradient $\nabla \pi_q$ is defined everywhere. Therefore, ∇h_Q is defined in the interior of every Voronoi n -cell V_q and is given by

$$\nabla h_Q(x) = 2(x - q).$$

However, it is *only* in the interiors of full-dimensional Voronoi cells that this gradient is defined. In other words, $\nabla h_Q(x)$ is defined if and only if $|A_Q(x)| = 1$.

Interestingly, there are many points $x \in \mathbb{R}^n$ with $|A_Q(x)| > 1$ at which there is a unique *direction of steepest ascent*. In fact, such a direction exists everywhere (although it can be zero at certain points which we shall call “critical” in analogy to the smooth setting). In fact, it is possible to characterize a vector field $v_Q : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that

- (1) agrees with ∇h_Q wherever this gradient is defined,
- (2) is parallel to the direction of steepest ascent, wherever this direction uniquely exists (even if ∇h_Q is undefined), and
- (3) is 0 otherwise.

In order to explicitly describe $v_Q(x)$ for a point $x \in \mathbb{R}^n$, we need to distinguish the Voronoi face to which x belongs and the Delaunay face dual to it. Specifically, for any point $x \in \mathbb{R}^n$, let $V_Q(x)$ denote the *lowest dimensional* Voronoi face of $\text{Vor } Q$ that contains x and let $D_Q(x)$ be the Delaunay face dual to $V_Q(x)$ in $\text{Del } Q$.

The *generalized gradient* vector field v_Q can then be characterized as follows. For every point $x \in \mathbb{R}^n$, let $d_Q(x)$, called the *driver* of x , be the closest point to x in $D_Q(x) = \text{conv } A_Q(x)$. We define the vector field $v_Q : \mathbb{R}^n \rightarrow \mathbb{R}^n$ at a point x as

$$v_Q(x) = 2(x - d_Q(x)). \tag{2.1}$$

Observe that for any point x in the interior of a Voronoi cell V_q , $A_Q(x) = \{q\}$ is a singleton. Therefore $D_Q(x) = \{q\}$ and from that $d_Q(x) = q$. Thus we get

$$v_Q(x) = 2(x - d_Q(x)) = 2(x - q) = \nabla h_Q(x),$$

confirming that v_Q agrees with ∇h_Q wherever the latter is defined.

The fact that $A_Q(x)$ is the same for every point x in the relative interior of every Voronoi face V and the fact that affine hulls of $D = D_Q(x)$ and $V_Q(x)$ are orthogonal (Proposition 1.4) implies that every point on $\text{aff } V$ has the same closest point on the $\text{aff } D$. This in turn leads to a the following crucial observation.

Proposition 2.1 *Given a set Q of (possibly weighted) points in \mathbb{R}^n , let x and y be two points in the relative interior of the same face of $\text{Vor } Q$, i.e. $V_Q(x) = V_Q(y)$. Then, $d_Q(x) = d_Q(y)$.*

We therefore use the notations $A_Q(V)$ and $d_Q(V)$ respectively for the set of the closest points and the driver common to all points in the relative interior of a Voronoi face V .

A point x for which $v_Q(x) = 0$, or identically a point that satisfies $x = d_Q(x)$, is called a *critical points* of h_Q . All other (non-critical) points are *regular points*. Recall that the affine hulls of a Voronoi face and its dual Delaunay face are orthogonal (Proposition 1.4). Consequently, these affine hulls intersect in at most one point. In particular, a Voronoi face and its dual Delaunay face can intersect in at most one point. Since $D_Q(x) = \text{conv } A_Q(x)$, a point x is critical if and only if it is “the” intersection point of a Voronoi face and its dual Delaunay face.

Generically, if a Voronoi face V intersects the affine hull of its dual Delaunay face D , the intersection point (which is unique by Proposition 1.4) is a point of $\text{rel int } V$. This is easier to observe for a set P of unweighted points but is also valid for weighted sets through similar reasoning. Consider a delaunay face D with vertex set $P_D = \{p_1, \dots, p_i\} \subset P$. Let V be the Voronoi face dual to D . The set $\text{aff } V$ is the locus of all points in \mathbb{R}^n that are at equal distance from all points in P_D . Essentially, the set P_D by itself only determine $\text{aff } V$. The polytope V is then determined in conjunction with the points in $P \setminus P_D$. Let z be the intersection point of affine hulls of D and V . The affine hull of each facet V' of V is the locus of all points that are at equal distance from sites in P_D and one site $p' \in P \setminus P_D$. In order for this locus to include z , p' must be chosen from a sphere centered at z and of radius equal to $\text{dist}(z, P_D)$. Such a sphere is an $n - 1$ dimensional surface and has Lebesgue measure zero in \mathbb{R}^n . In other words, only if p' is chosen from a set of Lebesgue measure zero, the point z can possibly be on the boundary of V . We therefore define a point-set in general position with respect to critical points of the distance function induced by them as follows.

Assumption 2.2 (critical general position) A set Q of (possibly weighted) points is in *critical general position with respect to the critical points of h_Q* if every critical point c of h_Q is contained in the relative interior of the lowest dimensional Voronoi face that includes c , i.e. $c \in \text{rel int } V_Q(c)$. For the rest of this monograph, whenever the critical points of a distance function h_Q induced by a point-set Q are considered, the set Q is assumed to be in critical general position with respect to the critical points of h_Q .

As mentioned before the vector field v_Q generalizes the gradient of h_Q ; we saw above that

$v_Q(x) = \nabla h_Q(x)$ wherever the gradient is defined. The characterization of v_Q as the “direction of steepest ascent” at regular points (where this direction uniquely exists) is given by the following Theorem from [51]. We reproduce the proof for the purpose of completeness.

Theorem 2.3 [51] *Let Q be a set of (possibly weighted) points in \mathbb{R}^n and let $x \in \mathbb{R}^n$ be a point with $v_Q(x) \neq 0$. Let $v = v_Q(x)/\|v_Q(x)\|$ be the unit vector in the direction of $v_Q(x)$. Then there is a small enough $\varepsilon_0 > 0$ such that for any $\varepsilon \leq \varepsilon_0$,*

$$h_Q(x + \varepsilon v) > h_Q(x + \varepsilon u)$$

for any unit vector $u \neq v$.

Proof. The squared distance of a point $x \in \mathbb{R}^n$ to a (weighted) point q is $\pi_q(x) = \|x - \tilde{q}\|^2 - w_q$, and its squared distance to the set Q is

$$h_Q(x) = \min_{q \in Q} \pi_q(x).$$

Recall that $A_Q(x)$ denotes the set of points in Q closest to x (at weighted distance $h_Q(x)$ from x). We want to determine the direction at x that results in a maximum rate of increment for $h_Q(x)$ — the steepest ascent direction of h_Q at x . Let $u \in \mathbb{S}^{n-1}$ be a direction in \mathbb{R}^n . To determine the rate of increment of $h_Q(x)$ in the direction of u , consider the hyperplane H_u consisting of the set of points z satisfying

$$\langle u, z - x \rangle = \rho$$

where ρ is smallest value for which $\langle u, \tilde{q} - x \rangle \leq \rho$ for every $q \in A_Q(x)$; in words, H_u is orthogonal to u and shifted in the direction opposite to u as far as possible all the while keeping the points in $A_Q(x)$ in the halfspace H_u^- that extends in the direction opposite to u and is bounded by H_u . Let q be any point in $A_Q(x) \cap H_u$, and let c be the point in H_u closest to x . Then

$$\pi_q(x) = \|x - \tilde{q}\|^2 - w_q = \|x - c\|^2 + \|c - \tilde{q}\|^2 - w_q.$$

It is easily observed that for $x' = x + \varepsilon u$, with ε infinitesimally small, $A_Q(x') = A_Q(x) \cap H_u$ (look at the balls $B(x, \pi_q(x))$ and $B(x', \pi_q(x'))$); their intersection with H_u is the same — this is most easily seen for zero weights; though it is true for weighted points as well). Also

$$\pi_q(x') = \|x' - \tilde{q}\|^2 - w_q = \|x' - c\|^2 + \|c - \tilde{q}\|^2 - w_q.$$

So it follows that the rate of change of $h_Q(x)$ in the direction of u is equal to the rate of change of $\|x - c\|^2 + \|c - p\|^2$ which is $2\|x - c\|$. This is largest when c is farthest from x while $x - c$ is in the direction of u . A well-known geometric fact implies that the rate of change of $h_Q(x)$ is maximized when u is chosen in such a way that c is the closest point to x in $\text{conv } A_Q(x)$. ■

Another important fact about the vector field v_Q relates the drivers of incident faces of

Vor Q and is given in the following proposition. Intuitively, it says that for Voronoi face V with “opaque” relative interior, all the proper faces of V that are “visible” from the driver $d = d_Q(V)$ of V have also driven by d .

Proposition 2.4 [51] *Let Q be a set of (possibly weighted) point in \mathbb{R}^n . Let V be a face of Vor Q and let $V' < V$ be a proper face of V . Let $d = d_Q(V)$ and $d' = d_Q(V')$. Then $d = d'$ if and only if for all $x \in \text{rel int } V'$, the segment $\overline{d'x}$ is disjoint from the relative interior of V , i.e. $\overline{d'x} \cap \text{rel int } V = \emptyset$.*

2.2 Dynamical Systems and Flows

A *dynamical system* on an open set $S \subseteq \mathbb{R}^n$ is a C^1 map $\phi : \mathbb{R} \times S \rightarrow S$ that writing $\phi(t, x) = \phi_t(x)$ satisfies

- (a) $\phi_0 : S \rightarrow S$ is identity, i.e. $\phi(0, x) = x$ for every $x \in S$;
- (b) For any $s, t \in \mathbb{R}$, the composition $\phi_t \circ \phi_s$ equals ϕ_{t+s} , i.e.

$$\phi(t, \phi(s, x)) = \phi(t + s, x)$$

for all $x \in S$.

Note that by definition $\phi_t : S \rightarrow S$ is C^1 for each t and has a C^1 inverse ϕ_{-t} . Any dynamical system ϕ on S in general gives rise to a differential equation on S by defining a vector field $v : S \rightarrow \mathbb{R}^n$. Given ϕ_t , define v by

$$v(x) = \left. \frac{d}{dt} \phi_t(x) \right|_{t=0}. \quad (2.2)$$

We think of the vector $v(x)$ as the tangent vector to the curve $t \mapsto \phi_t(x)$ at $t = 0$. It is possible to rewrite this in more conventional terms. For any $x \in S$, let $x(t) = \phi_t(x)$ and let v be the vector field defined above. Then Equation (2.2) can be written as

$$x' = v(x), \quad (2.3)$$

where x' denotes dx/dx . Thus $x(t)$ or $\phi_t(x)$ is the solution curve of the above differential equation for the initial condition $x(0) = x$.

Nontrivially, the above process be reversed. Any vector field v on \mathbb{R}^n naturally induces a first order differential equation identical to Equation (2.3):

$$x' = v(x).$$

If vector field v is smooth enough, this equation locally has a unique solution.

Theorem 2.5 (Picard-Lindelöf Theorem) Let $S \subseteq \mathbb{R}^n$ be open¹ and let $v : S \rightarrow \mathbb{R}^n$ be a C^1 (continuously differentiable) vector map. For any $x_0 \in S$, there is an open interval $J(x_0) = (\alpha(x_0), \beta(x_0))$ and a unique solution

$$\phi : J(x_0) \rightarrow S$$

of the differential equation

$$x' = v(x)$$

satisfying the initial condition

$$\phi(0) = x_0.$$

Moreover, the the solution $\phi(t)$ depends continuously on the initial condition $\phi(0)$.

The solution curve $\phi : J(x_0) \rightarrow S$ corresponding to each $x_0 \in S$ is called a *flow curve* or *integral curve*. The reason the flow curve corresponding to the initial condition $\phi(0) = x_0$ is only defined locally on some interval $(\alpha(x_0), \beta(x_0))$ and cannot be generally extended to the entire real line is that this flow curve can reach the boundary of S . It is however possible to extend the proof of the above theorem so that the interval $J(x_0)$ becomes maximal in the sense that that each solution curve at each end of its corresponding interval either reaches the boundary of S or goes to infinity (or both). To indicate the dependence of a flow curve $\phi(t)$ resulting from the initial condition $\phi(0) = x_0$ to x_0 , we write

$$\phi(t) = \phi(t, x_0).$$

Thus $\phi(0, x_0) = x_0$. Let $\Omega \subset \mathbb{R} \times S$ be

$$\Omega = \{(t, y) \in \mathbb{R} \times S : t \in J(y)\}.$$

The map $(t, y) \mapsto \phi(t, y)$ is then a function

$$\phi : \Omega \rightarrow S,$$

which we call the *flow* of Equation (2.3). It turns out that the flow map ϕ is C^1 and has the properties of a dynamical system.

Our interest in flow maps lies on their continuity on both parameters. If for for some $S' \subset S$ and $t_0 > 0$,

$$[0, t_0] \times S' \subseteq \Omega,$$

then ϕ defines a homotopy between the two sets S' and $\phi(t_0, S')$. Such homotopies can be used to determine topological types of certain subsets of S .

The main idea behind most topological proofs in this thesis is to use of the generalized gradient vector field v_Q induced by a set of (possible weighted) points to define a differential

¹The theorem can in fact be stated for an open subset S of any normed vector space.

equation

$$x' = v_Q(x).$$

Unfortunately, v_Q is not C^1 and therefore the results mentioned above do not apply to it. It is nonetheless possible to integrate v_Q and produce a flow map ϕ_Q which although not C^1 is still continuous.

2.3 Integrating the Generalized Gradient

Although the vector field v_Q induced by a set Q of weighted points is not smooth, it is still possible to integrate v_Q in a manner somewhat similar to what we saw in the previous section. Restricted to the relative interior of a full-dimensional Voronoi cell V_q of $\text{Vor } Q$ associated to a point $q \in Q$, v_Q agrees with $\nabla h_Q = \nabla \pi_q$ and is therefore smooth and can be integrated using Theorem 2.5. The resulting flow map consists of straight integral lines that meet at q . On each of these integral lines, the flow moves away from q at a speed that increases linearly with distance to q . A similar statement is true about lower dimensional Voronoi faces. By Proposition 2.1, every point x in the relative interior of any Voronoi face V in $\text{Vor } Q$ has the same driver $d = d_Q(V)$ and the vector $v_Q(x)$ is therefore given by

$$v_Q(x) = 2(x - d_Q(V)) = 2(x - d).$$

Thus, at every point in the relative interior of V , moving in the direction of steepest ascent corresponds to moving straight away from d . If $d \in \text{aff } V$, the vector $v_Q(x)$ is parallel to $\text{aff } V$ and therefore escaping d results integral lines that are contained in the relative interior of V . In the case where $d \notin \text{aff } V$ the driver $d = d_Q(V)$ of V is the same as the driver $d_Q(V')$ of the coface V' of V into which $v_Q(x)$ points (Proposition 2.4). In this case, one can extend the flow lines in the relative interior of V' to V as well.

Thus the space \mathbb{R}^n can be partitioned into a number of *regions* in such a way that the vector field v_Q can be integrated restricted to each region. Specifically, each region is made of the relative interior of a Voronoi face V plus the relative interiors of all those proper faces V' of V that share the same driver as V , i.e. $d_Q(V) = d_Q(V')$. The integral lines on incident regions can then be glued to each other to results a flow map that is defined everywhere in the space. Intuitively, a particle can follow a flow line restricted to one region (moving opposite to the driver of that region on a straight line) until it reaches another region. At that point, it is driven by the driver of the latter region until it hits a third region, and so forth. The trajectory of the particle can be interpreted as a “combined” integral line. This results a map ϕ_Q that is defined all over \mathbb{R}^n . The resulting *integral lines* or *flow paths* of ϕ_Q are in general piece-wise linear curves that may turn only at points where they reach the relative interior of a new Voronoi face. All the points in a linear section of a flow path have the same driver d . Thus the magnitude $\|v_Q(x)\|$ of $v_Q(x)$ which is given by

$$\|v_Q(x)\| = 2\|x - d_Q(x)\| = 2\|x - d\|$$

grows linearly with distance to d on that section. The speed of the particle at any point x on a flow path is given by the vector $v_Q(x)$. Thus we get a flow map

$$\phi_Q : (\mathbb{R}^+ \cup \{+\infty\}) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$$

that assign to each pair (t, x) with $t \in \mathbb{R}^+ \cup \{+\infty\}$ and $x \in \mathbb{R}^n$, the point $y = \phi_Q(t, x)$ that is reached at time t by moving along the flow line that starts at x in such a way that the speed of movement at any point $z = \phi(\tau, x)$, $\tau \leq t$, is $v_Q(z)$, where \mathbb{R}^+ denotes the set of *non-negative* real numbers.

We allow the time parameter t in the definition of ϕ_Q to take a symbolic value of $+\infty$ in order to simplify our notation when referring to the limit points of the flow lines of ϕ_Q which are precisely the critical points of h_Q . It is of course possible for a flow line to go to infinity. Because of this we augment the space \mathbb{R}^n with a *symbolic critical point at infinity*. Since the distance to the point-set is infinite at it, this symbolic critical point is regarded as a maximum. The map ϕ_Q has the standard properties of a flow map, i.e.

$$\phi_Q(0, x) = x; \tag{2.4}$$

$$\phi_Q(s + t, x) = \phi_Q(s, \phi_Q(t, x)). \tag{2.5}$$

Moreover, at every point $\phi_Q(t, x)$ of a flow line $\phi_Q(x)$, except its turning points,

$$\frac{d}{dt}\phi_Q(t, x) = v_Q(\phi_Q(t, x)).$$

Recall that each integration region consists of the relative interior of a Voronoi faces V along with the relative interior of every proper face $V' < V$ that shares the same driver as V . Consequently, the intersection of a flow line and such a region is a half-open line segment (or a half-open half-line) which includes its extreme point at its end where it enters the region. This implies that the map $t \mapsto v_Q(\phi_Q(t, x))$ is in fact the *right-derivative* of $t \mapsto \phi_Q(t, x)$. In particular, since there are only a countable (in fact finite) number of turns in a flow path, one can express the latter map by integrating v_Q :

$$\phi_Q(t, x) = x + \int_0^t v_Q(\phi_Q(\tau, x))d\tau. \tag{2.6}$$

2.3.1 Continuity of the Flow Map

Clearly, the map ϕ_Q changes continuously with changing of its first parameter; the image of $\phi_Q(\cdot, x)$ is a continuous curve.

We shall first characterize the rate of flow along each linear piece of a flow path and then prove that ϕ_Q is indeed continuous on its second parameter as well.




Figure 2.1: Flow lines of the distance flow induced by a set of points. The red points are index 1 critical points.

Lemma 2.6 *Let $x \in \mathbb{R}^n$ and $T \in \mathbb{R}^+$ be such that the flow path $\phi_Q([0, T], x)$ consists of a single line segment. Then for all $t \in [0, T]$,*

$$\phi_Q(t, x) = x + \frac{1}{2}v_Q(x) (e^{2t} - 1).$$

Proof. Notice that the $\phi_Q(t, x)$ is a point on the line segment connecting x and $\phi_Q(T, x)$. This direction of this segment is indicated by $v_Q(x)$. Thus, for simplicity we align the real line with $v_Q(x)$ in such a way that $d_Q(x)$ becomes the origin and x lies on the positive side of the real line. We indicate the distance to origin on this real line of the point $\phi_Q(t, x)$ by $\psi(t)$. Thus $\phi_Q(t, x)$ is related to $\psi(t)$ by the equation

$$\phi_Q(t, x) = d_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \psi(t). \quad (2.7)$$

We have

$$\begin{aligned} \psi(0) &= \|x - d_Q(x)\| \\ \dot{\psi}(t) &= 2\psi(t), \end{aligned}$$

where $\dot{\psi}$ denotes $d\psi(t)/dt$. This gives the differential equation

$$\frac{d\psi}{\psi} = 2dt,$$

which has the solution

$$\psi(t) = \psi(0)e^{2t}.$$

Replacing in (2.7) gives us

$$\begin{aligned}
\phi_Q(t, x) &= d_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \cdot \|x - d_Q(x)\| e^{2t} \\
&= x - \frac{1}{2}v_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \cdot \frac{1}{2}v_Q(x)e^{2t} \\
&= x + \frac{1}{2}v_Q(x) (e^{2t} - 1).
\end{aligned}$$

■

A vector field v defined on \mathbb{R}^n is k -semi-Lipschitz if for every pair of vectors x and y

$$\langle x - y, v(x) - v(y) \rangle \leq k \cdot \|x - y\|^2.$$

Lemma 2.7 *Let Q be a set of weighted points in \mathbb{R}^n . Then the vector field v_Q is 2-semi-Lipschitz, i.e. for any $x, y \in \mathbb{R}^n$,*

$$\langle x - y, v_Q(x) - v_Q(y) \rangle \leq 2\|x - y\|^2.$$

Proof. Let p be a point in $A_Q(x)$ and $q \in A_Q(y)$. We show that

$$\langle x - y, p - q \rangle \geq 0. \tag{2.8}$$

Since $p \in A_Q(x)$, $\pi_p(x) \leq \pi_q(x)$ or equivalently

$$\|x - p\|^2 - \|x - q\|^2 \leq w_p - w_q.$$

Similarly since $q \in A_Q(y)$, $\pi_p(y) \geq \pi_q(y)$ and therefore

$$w_p - w_q \leq \|y - p\|^2 - \|y - q\|^2.$$

Eliminating $w_p - w_q$ from the above two inequality we obtain

$$\langle x - p, x - p \rangle - \langle x - q, x - q \rangle \leq \langle y - p, y - p \rangle - \langle y - q, y - q \rangle.$$

Expanding we get

$$\langle x, q \rangle - \langle x, p \rangle \geq \langle y, q \rangle - \langle y, p \rangle,$$

which gives (2.8) by rearranging.

Inequality (2.8) can be written also as

$$\langle x - y, p - y \rangle \geq \langle x - y, q - y \rangle. \tag{2.9}$$

The above inequality in particular holds for the point $p_0 \in A_Q(x)$ minimizing $\langle x - y, p_0 - y \rangle$ and the point $q_0 \in A_Q(y)$ maximizing $\langle x - y, q_0 - y \rangle$. Therefore the hyperplane $\Pi = \{z \in$

$\mathbb{R}^n : \langle x - y, z - y \rangle = \gamma\}$, where $\gamma = \frac{1}{2}(\langle x - y, p_0 - y \rangle + \langle x - y, q_0 - y \rangle)$, separates the two sets $A_Q(x)$ and $A_Q(y)$ as well as their convex hulls and in particular $d_Q(x) \in \text{conv}(A_Q(x))$ and $d_Q(y) \in \text{conv}(A_Q(y))$. Thus we get

$$\langle x - y, d_Q(x) - d_Q(y) \rangle \geq 0. \quad (2.10)$$

Now to prove the statement of the lemma we write

$$\begin{aligned} \langle x - y, v_Q(x) - v_Q(y) \rangle &= \langle x - y, 2(x - d_Q(x)) - 2(y - d_Q(y)) \rangle \\ &= 2\langle x - y, (x - y) - (d_Q(x) - d_Q(y)) \rangle \\ &= 2\|x - y\|^2 - 2\langle x - y, d_Q(x) - d_Q(y) \rangle \\ &\leq 2\|x - y\|^2, \end{aligned}$$

where the final inequality follows (2.10). ■

Lemma 2.8 *Let Δ be an upper bound for the diameter of $\text{conv} Q$. Then for any point x ,*

$$\|v_Q(x)\| \leq 2(\Delta + \text{dist}(x, \tilde{Q})).$$

Proof. By definition $v_Q(x) = 2(x - d_Q(x))$. Let p be a point in $\text{conv} \tilde{Q}$. By triangle inequality $\|x - d_Q(x)\| \leq \|x - p\| + \|p - d_Q(x)\|$. Since $d_Q(x) \in \text{conv} A_Q(x) \subset \text{conv} \tilde{Q}$, $\|p - d_Q(x)\| \leq \Delta$. Thus $\|x - d_Q(x)\| \leq \Delta + \text{dist}(x, \tilde{Q})$ and this implies the statement of the lemma. ■

Our main goal in this section is to prove that the flow ϕ_Q is continuous on its second variable. This, along with its continuity on its first variable, guarantees that the map $\phi_Q : \mathbb{R}^+ \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a continuous map and can therefore be used for the role of map H in Proposition 1.8 for establishing homotopy equivalences.

Since the flow path $\phi_Q([0, T], x)$ is a piece-wise linear curve, we can study the flow in individual linear pieces of a flow path.

Lemma 2.9 *The map $t \mapsto \|v_Q(\phi(t, x))\|^2$ is the right-derivative of $t \mapsto h_Q(\phi_Q(t, x))$. In other words, for any x and $\varepsilon > 0$, there exists a $t_0 > 0$, such that for any $t \in [0, t_0]$,*

$$\left| h_Q(\phi_Q(t, x)) - (h_Q(x) - t\|v_Q(\phi_Q(t, x))\|^2) \right| < t\varepsilon. \quad (2.11)$$

Proof. Again it is enough to prove the lemma for the first linear piece of $\phi_Q(x)$. Let $y = \phi_Q(t, x)$ and let $\lambda = \|x - y\|$. Let p be a point in $A_Q(x)$ and in the hyper-plane Π containing $d_Q(x)$ and orthogonal to $v_Q(x)$. Notice that such a point must exist since by definition $d_Q(x)$ is the closest point to x on $\text{conv} A_Q(x)$ and Π is a supporting hyper-plane orthogonal to $x - d_Q(x)$. Thus $\text{conv} A_Q(x)$ must have a vertex on Π . Let $\alpha = \angle(x - d_Q(x), x - p)$. Since y is in the same linear pieces of the flow line as x , $A_Q(x) = A_Q(y)$ and therefore $p \in A_Q(y)$.

Using the cosine law on triangle xyp we get

$$\|y - p\|^2 = \|x - p\|^2 + \lambda^2 + 2\lambda\|x - p\| \cdot \cos \alpha.$$

Subtracting w_p from both sides gives

$$h_Q(y) = h_Q(x) + \lambda^2 + 2\lambda\|x - p\| \cdot \cos \alpha.$$

From Lemma 2.6, $\lambda = \frac{1}{2}\|v_Q(x)\|\gamma$ where $\gamma = e^{2t-1}$. Thus we have for the left hand side of the inequality (2.11)

$$\begin{aligned} & |h_Q(y) - h_Q(x) + t\|v_Q(y)\|^2| \\ &= |\lambda^2 + 2\lambda\|x - p\| \cdot \cos \alpha - t\|v_Q(y)\|^2| \\ &= |\lambda^2 + 2\lambda\|x - p\| \cdot \cos \alpha - t(\|v_Q(x)\| + 2\lambda)^2| \\ &= |\lambda^2 + 2\lambda\|x - p\| \cdot \cos \alpha - 4t\lambda^2 - t\|v_Q(x)\|^2 - 4t\lambda\|v_Q(x)\|| \\ &= \left| \|v_Q(x)\|^2 \left(\frac{1}{4}\gamma^2 - t\gamma^2 - 2t\gamma - t \right) + \|v_Q(x)\|\gamma\|x - p\| \cos \alpha \right|. \end{aligned}$$

From our choice of p and the discussion above,

$$\|x - p\| \cdot \cos \alpha = \|x - d_Q(x)\| = \frac{1}{2}\|v_Q(x)\|.$$

Thus we have

$$|h_Q(y) - h_Q(x) + t\|v_Q(y)\|^2| = \|v_Q(x)\|^2 \left| \left(\frac{1}{4} - t \right) \gamma^2 - 2t\gamma - t + \frac{1}{2}\gamma \right|$$

Therefore, for (2.11) to hold we must have for any $\varepsilon > 0$,

$$\|v_Q(x)\|^2 \left| \left(\frac{1}{4} - t \right) \gamma^2 - 2t\gamma - t + \frac{1}{2}\gamma \right| < t\varepsilon,$$

when t is sufficiently small. Substituting $e^{2t} - 1$ for γ the above inequality becomes

$$\|v_Q(x)\|^2 \left| \left(\frac{1}{4} - t \right) (e^{2t} - 1)^2 - 2t(e^{2t} - 1) - t + \frac{1}{2}(e^{2t} - 1) \right| < t\varepsilon,$$

which leads to the following inequality after simplification and rearrangement

$$\frac{1}{t} \left| \left(\frac{1}{4} - t \right) e^{4t} - \frac{1}{4} \right| < \frac{\varepsilon}{\|v_Q(x)\|^2}.$$

The left hand side of the above inequality is an increasing continuous function on $t \in (0, +\infty)$ which converges to 0 when $t \rightarrow^+ 0$ and this completes the proof of the Lemma.

■

The above Lemma immediately implies the following Corollary.

Corollary 2.10 *For any $x \in \mathbb{R}^n$ and any $t \in [0, +\infty)$,*

$$h_Q(\phi_Q(t, x)) = h_Q(x) + \int_0^t \|v_Q(\phi_Q(\tau, x))\|^2 d\tau.$$

In particular, the map $t \mapsto h_Q(\phi_Q(t, x))$ is strictly increasing.

Theorem 2.11 *The flow map ϕ_Q is continuous on its second variable. In other words for all $x \in \mathbb{R}^n$ and for all $t \geq 0$ and for all $\varepsilon > 0$, there exists a $\delta > 0$, such that for every $y \in \mathbb{R}^n$ satisfying $\|x - y\| < \delta$, $\|\phi_Q(t, x) - \phi_Q(t, y)\| < \varepsilon$.*

Proof. Consider two points x and y and real number t and let for each $0 \leq \tau \leq t$, $x(\tau) = \phi_Q(\tau, x)$ and $y(\tau) = \phi_Q(\tau, y)$. The two flow lines $\tau \mapsto x(\tau)$ and $\tau \mapsto y(\tau)$ for $\tau \in [0, t]$ are piecewise linear curves each with a finite number of turns. Let $0 = t_1 < t_2 < \dots < t_r = t$ be the set of times at which at least one of these two curves makes a turn (enters a new Voronoi cell and switches drivers). We show that for each of the intervals $[t_i, t_{i+1}]$, $1 \leq i \leq r - 1$,

$$\|x(t_{i+1}) - y(t_{i+1})\| \leq \|x(t_i) - y(t_i)\| \cdot e^{t_{i+1} - t_i}.$$

This will imply that

$$\|x(t_r) - y(t_r)\| \leq \|x(0) - y(0)\| \cdot \prod_{i=1}^{r-1} e^{t_{i+1} - t_i} = e^t \cdot \|x(0) - y(0)\|.$$

In particular, for any $\varepsilon > 0$, if $y \in B(x, \delta)$ where $\delta \leq \varepsilon/e^t$, then

$$\phi_Q(t, y) \in B(\phi_Q(t, x), \varepsilon).$$

It suffices to prove the above claim only for the first interval $[t_1, t_2]$; the claim will then hold for the subsequent intervals by a reparameterization of time. Equivalently, it will be enough to show that the claim is valid for the case where each of the two flow paths $x([0, t])$ and $y([0, t])$ is a single line segment.

For $0 \leq \tau \leq t$, we define the function $\lambda(\tau)$ as the square of the distance between $x(\tau)$ and $y(\tau)$, i.e.

$$\lambda(\tau) = \|x(\tau) - y(\tau)\|^2.$$

Then we have

$$\begin{aligned}
d\lambda(\tau) &= \lambda(\tau + d\tau) - \lambda(\tau) \\
&= \|x(\tau + d\tau) - y(\tau + d\tau)\|^2 - \|x(\tau) - y(\tau)\|^2 \\
&= \|x(\tau) + v_Q(x(\tau))d\tau - y(\tau) - v_Q(y(\tau))d\tau\|^2 - \|x(\tau) - y(\tau)\|^2 \\
&= 2 \langle x(\tau) - y(\tau), v_Q(x(\tau)) - v_Q(y(\tau)) \rangle d\tau + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 (d\tau)^2 \\
&\leq 2\|x(\tau) - y(\tau)\|^2 d\tau + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 (d\tau)^2,
\end{aligned}$$

in which the last inequality follows from Lemma 2.7. This results

$$\frac{d\lambda(\tau)}{d\tau} \leq 2\lambda(\tau) + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 d\tau$$

Lemmas 2.8 and 2.9 bound $\|v_Q(x(\tau)) - v_Q(y(\tau))\|^2$ as a function of τ . Since $d\tau$ is infinitesimally small, we get the differential inequality

$$\frac{d\lambda(\tau)}{d\tau} \leq 2\lambda(\tau),$$

which for $0 \leq \tau \leq t$, yields to the solution

$$\|x(\tau) - y(\tau)\|^2 = \lambda(\tau) \leq \lambda(0) \cdot e^{2\tau}.$$

In particular for $\tau = t$,

$$\|x(t) - y(t)\| \leq e^t \cdot \|x(0) - y(0)\|.$$

■

Corollary 2.12 *For any finite set Q of weighted points in \mathbb{R}^n , the induced flow map ϕ_Q is continuous on $[0, +\infty) \times \mathbb{R}^n$.*

2.4 Stable and Unstable Manifolds

Flow Orbit. For a given flow map ϕ_Q , the *flow orbit* of a regular point x , denoted $\phi_Q(x)$ is the set of all points $\phi_Q(t, x)$ for all $t \geq 0$.

$$\phi_Q(x) = \phi_Q([0, +\infty), x) = \{y : y = \phi_Q(t, x) \text{ for some } t \in \mathbb{R}^+ \cup \{+\infty\}\}.$$

The flow orbit of a critical point is the critical point itself. Thus critical points are the *fixed-points* of flow curves $\phi_Q(x)(\cdot) = \phi_Q(\cdot, x)$.

The flow orbit or *flow closure* of $\phi_Q(T)$ of a set $T \subseteq \mathbb{R}^n$ is the union of flow orbits of all

points in T , i.e.

$$\phi_Q(T) = \bigcup_{x \in T} \phi_Q(x).$$

Notice that by this definition

$$T \subseteq \phi_Q(T).$$

Inflow. One can also consider the set of points that flow into any point in the space. The *inflow* of a point x is the set

$$\phi_Q^{-1}(x) = \{y : x \in \phi_Q(y)\}.$$

In other words, the inflow of x is the union of all flow paths that reach x . This definition can also be extended to sets of points.

$$\phi_Q^{-1}(T) = \bigcup_{x \in T} \phi_Q^{-1}(x).$$

for a subset T of \mathbb{R}^n .

The flow orbit $\phi_Q(x)$ of every point x converges to a critical point c of h_Q or goes to infinity as $t \rightarrow +\infty$. Notice that we consider c to be also in the flow orbit of x . In particular we consider the symbolic maximum at infinity to be in the flow orbit of every point whose orbit flows to infinity.

Stable Manifold. For a critical point c of h_Q , the set of all points x whose flow orbits converges to c is called the *stable manifold* of c and is denoted by $\text{Sm}(c)$. In other words,

$$\text{Sm}(c) = \{x : \phi_Q(+\infty, x) = c\} = \phi_Q^{-1}(c).$$

In other words $\text{Sm}(c)$ is simply the inflow of the critical point c .

Unstable Manifold. Although there is no flow out of a critical point c , it is interesting to know where the points very close to c flow. Some of these points flow into c while other flow away from it. We define the *unstable manifold* $\text{Um}(c)$ of a critical point c , as the set of all points into which points arbitrarily close to c flow. Formally,

$$\text{Um}(c) = \bigcap_{\varepsilon > 0} \phi(B(c, \varepsilon)).$$

In other words, the unstable manifold of c consists of c and all the integral lines that start infinitesimally close to c . With an abuse of terminology, we may say that $\text{Um}(c)$ are the flow lines that flow “out of” c .

Flow-Tight and Flow-Repellant Sets. We are particularly interested in sets that include their flow orbit. A set T is said to be *flow-tight* if $\phi_Q(T) = T$. We call a set T *flow-repellant* if $T \cap \phi_Q(T^c) = \emptyset$, i.e. no flow lines of ϕ_Q enters T .

Stable and unstable manifolds of critical points are perhaps the most natural examples of flow-tight sets. The unstable manifold $\text{Um}(c)$ of a critical point c is defined as the flow closure of an infinitesimally small neighborhood of the critical point c . Dealing with infinity may cast some doubt on the validity of the claim that $\phi_Q(\text{Um}(c)) = \text{Um}(c)$. However, as we shall see in the following section (Proposition 2.19), the unstable manifold of every critical point c can be described as the flow closure of a fixed closed set that contains cc and therefore is flow-tight. Flow-tightness of stable manifolds is even easier to establish; if $x \in \text{Sm}(c)$ then $c \in \phi_Q(x)$. But then by Equation (2.5) $c \in \phi_Q(y)$ for every $y \in \phi_Q(x)$ as well. Thus $y \in \text{Sm}(c)$.

If T_1 and T_2 are flow-tight sets for a flow ϕ_Q , their union is clearly flow-tight as well. It can be easily observed that the intersection $T_1 \cap T_2$ of these flow-tight sets is also flow-tight. This is because for every $x \in T_1 \cap T_2$, $\phi_Q(x) \subseteq T_1$ and $\phi_Q(x) \subseteq T_2$ and therefore $\phi_Q(x) \subseteq T_1 \cap T_2$.

Proposition 2.13 *Given a distance flow map ϕ_Q , induced by a (possibly weighted) point-set Q , if two sets T_1 and T_2 are flow tight for ϕ_Q , then their union and their intersection are also flow-tight for ϕ_Q . In particular, unions and intersections of stable and unstable manifolds of critical points of h_Q are flow-tight for ϕ_Q .*

2.4.1 Computing Stable and Unstable Manifolds

In the following Chapters of this dissertation, we present algorithms that explicitly utilize stable or unstable manifolds of critical points of certain distance functions. Computation of these sets is therefore an integral part of these algorithms. In this section, we present algorithms for computing these objects in arbitrary dimensions. We also explore a number of their fundamental properties.

The general idea behind computing these objects for a flow map ϕ_Q , induced by a set of weighted points Q , is to exploit the simple structure of ϕ_Q within individual faces of $\text{Vor } Q$ and trace the flow lines of ϕ_Q “in bulk”, one Voronoi face at a time. For computing stable manifolds, we trace these flow lines backwards, and for unstable manifolds, forwards.

Computation of Stable Manifolds. Consider a set Q of weighted points in \mathbb{R}^n and a critical point c of h_Q . Before giving a formal description of an algorithm for computing the stable manifold of c , let us develop some intuition by examining the local structure of the inflow of c in the Voronoi faces that include c . By the general position assumption 2.2, c is contained in the relative interior of the Voronoi face $V = V_Q(c)$. Clearly, c is the driver of V , i.e. $c = d_Q(V)$. Therefore, the flow lines of ϕ_Q on V escape c . In other words, no point of

V (except c itself) flows into c . Consequently, all flow lines of ϕ_Q that end in c must reach it through a proper coface of V . Therefore, to compute the stable manifold of c (the inflow of c) we compute for each coface V' of V the part of the inflow of c that reaches c through V' .

Thus, let V' be a proper coface of V . Naturally, the first step is to find the points in V' itself that flow into c directly through V' . This subset of V' is easily determined from the location of the driver $d' = d_Q(V')$ of V' . If $d' \notin \text{aff } V'$, then no point of V' can flow into c through V' because $c \in \text{aff } V'$ but flow lines starting at points in V' leave $\text{aff } V'$ and enter a coface of it. However, if $d' \in \text{aff } V'$, then any point in the set $S = V' \cap \text{conv}\{d', c\}$, except perhaps for d' itself, flows into c . Notice that if $d' \in S$ and therefore in V' , then it is a critical point and does not flow to c (if d' is critical then $\phi_Q(d') = \{d'\}$). No point of V' other than those in S reach c through V' .

Any other point that flows into c through V' has to be in the inflow of the set S . These points can reach S either through a proper face of V' or through a proper coface of it. The proper faces of V' through which flow lines can reach S are those that intersect S . Proper cofaces of V' all intersect (in fact contain) S . If d' is a critical point, no proper face of V' (other than V) intersect S and all the flow that reaches S has to come from proper cofaces of V' . The above process can then be repeated recursively (in a slightly more general form) for each of these faces. This generalization results the algorithm of Figure 2.2.

The algorithm `INFLOW` takes two input arguments: a polytope P and a Voronoi face V . The polytope P is assumed to be contained in V . Informally, P is the part of the stable manifold of c that falls in the relative interior of V . It is assumed that P gets no inflow through V from proper faces of V unless those already contained in P , i.e.,

$$\phi_Q^{-1}(\text{rel int } P) \cap V = \text{rel int } P.$$

Upon completion, `INFLOW(P, V)` returns the closure I of the inflow

$$\phi_Q^{-1}(\text{rel int } P)$$

of the relative interior of P . To compute the stable manifold $\text{Sm}(c)$ of a critical point c , we run the algorithm `INFLOW` with parameters $P = \{c\}$ and $V = V_Q(c)$; the relative interior of $\{c\}$ is the point c itself and the inflow $\phi_Q^{-1}(c)$ is precisely $\text{Sm}(c)$.

The algorithm iterates over every proper coface V' of V and determines the set of points in the relative interior of V' that flow into P through V' . Similar to the special case considered above where $P = \{c\}$, this is exactly the intersection of the set

$$S = \text{rel int}(\text{conv}(P \cup \{d'\})),$$

and the Voronoi face V' . As before, $d' = d_Q(V')$ is the driver of V' . Once S is computed, we recursively compute the inflow of $S \cap V'$. This is done by looking at every Voronoi face V'' of V' (including V' itself) and computing `INFLOW($\text{cl } S \cap V'', V''$)`.

Algorithm INFLOW(convex polytope P , Voronoi face V)

```

1   $I \leftarrow P$ 
2  for each Voronoi face  $V' > V$ 
3       $d' \leftarrow d_Q(V')$ 
4       $S \leftarrow \text{rel int}(\text{conv}(P \cup \{d'\}))$ 
5      if  $S \cap V' \not\subseteq P$ 
6          for each Voronoi face  $V'' \leq V'$ 
7              if  $(S \cap \text{rel int } V'') \not\subseteq V$ 
8                   $I \leftarrow I \cup \text{INFLOW}(\text{cl } S \cap V'', V'')$ 
9  return  $I$ 

```

Figure 2.2: Recursive computation of the stable manifold of a critical point. The algorithm $\text{INFLOW}(P, V)$ takes a Voronoi face V and a convex subset P of V and computes the inflow of P that reaches P through V .

Proposition 2.14 *Let Q be a set of (possibly weighted) points in \mathbb{R}^n . For any critical point c of h_Q , the algorithm INFLOW , when run with parameters $P = \{c\}$ and $V = V_Q(c)$, returns the closure of $\text{Sm}(c)$.*

Let us pause here for a bit to study the INFLOW algorithm a little more carefully for the special case of flow maps induced by unweighted point-sets in three dimensions. In addition to offering intuition on how INFLOW behaves, this allows us to characterize the structure of stable manifolds of critical points of such distance functions. Flow maps induced by unweighted point-sets play a central role in the subsequent chapters. In particular, the analysis of the 3D surface reconstruction algorithm given in Section 4.1, makes explicit use of the resulting characterizations. Thus, let $P \subset \mathbb{R}^3$ be a set of unweighted points. Every critical points of h_P is of one of the following four types.

1. Index-0 critical points are the minima of h_P and are precisely the points in P .
2. Index-1 critical points (or *1-saddles*) and are intersection points of Voronoi facets and their dual Delaunay edges.
3. Index-2 critical points (also called *2-saddles*) and are intersection points of Voronoi edges and their dual Delaunay triangles.
4. Index-3 critical points are the local maxima of h_P . They are Voronoi vertices that are contained in their dual Delaunay tetrahedra.

The stable manifold of a minimum is just the minimum itself; nothing flows into a minimum (Corollary 2.10). A 1-saddle c is the midpoint of a Delaunay edge l that intersects its dual

Voronoi facet f . These Delaunay edges are sometimes called *Gabriel* edges. Let $p, q \in P$ be the sites at the ends of l . By our general position assumption for critical points, l intersects f in its relative interior. The algorithm INFLOW puts the two segments \overline{pc} and \overline{qc} in $\text{Sm}(c)$. But since both p and q are critical, the algorithm stops at both of these points and nothing else is added to $\text{Sm}(c)$. Consequently, the stable manifold of an index-1 critical point c is precisely the relative interior of the Delaunay (Gabriel) edge that contains c .

The study of the stable manifold of a 2-saddle is particularly illuminating. Figure 2.3(a) shows an index-2 critical point c which is the intersection point of a Delaunay triangle t_0 and its dual Voronoi edge e_0 . The Voronoi cells of the sites that make the vertices of t all have e_0 (and of course c) on their boundary and are therefore the 3-dimensional cofaces of e_0 . Each of these sites is the driver of its corresponding Voronoi cell. This means that the segment connecting each of these sites to c is contained in $\text{Sm}(c)$ (Figure 2.3(b); only two of these segments are shown). The 2-dimensional cofaces of e_0 are Voronoi facets that have e_0 as an edge. There are three such facets each dual to one of the edges of t . Figure 2.3(c) shows one of these facets. Let us call this facet f_0 and its dual Delaunay edge l_0 . The driver of a 2-dimensional facet is the midpoint of its dual Delaunay edge. Suppose that f_0 does not contain its driver which we call d_0 and let z_0 be the point of intersection between the segment $\overline{d_0c}$ and the boundary of f_0 that is closest to d_0 . All the points on the segment $\overline{z_0c}$ flow into c through f_0 . If we now consider any of the cofaces of f_0 , which are the Voronoi cells corresponding to the endpoints of l_0 , the inflow of $\overline{z_0c}$ from each of these cells is a triangle that has $\overline{z_0c}$ as an edge and the site of the Voronoi cell in question as the third vertex (Figure 2.3(d)). Any other point in the space that flows to c through f_0 can only do so by reaching z_0 .

We therefore compute the inflow of the point z_0 . The point z_0 is the point of the boundary of f_0 that intersects the segment $\overline{d_0c}$. In general, since f_0 is a two dimensional Voronoi face, its boundary is made of Voronoi edges and vertices and therefore z_0 is either on the relative interior of a Voronoi edge or coincides with a Voronoi vertex. It is not hard to convince oneself that the latter case is rather unlikely to happen for, say, a random point-set. As we shall see below, this is indeed a degenerate case that only happens for a measure zero collection of point-sets. We therefore assume (see the general position assumption below) that z_0 is in the relative interior of a Voronoi edge e_1 incident to the facet f_0 . Generically, there are two other Voronoi facets other than f_0 that are incident to e_1 . Let f_1 be one of these and let f'_1 be the other. We only follow the algorithm's progress on f_1 . The work done on f'_1 is similar and comprises another branch in recursion tree that is followed by the INFLOW algorithm. Let l_1 be the Delaunay edge dual to f_1 . The end-points of l_1 are sites in P whose corresponding Voronoi cells have e_1 on their boundary. As in the previous case, the line segments connecting each of these sites to z_0 flow into z_0 and should therefore be included in $\text{Sm}(c)$ (Figure 2.3(e)). To compute the set of points in f_1 that flow into z_0 , we connect the driver d_1 of f_1 (which is the midpoint of l_1) to z_0 and let z_1 be the intersection point of the segment $\overline{d_1z_0}$ with the boundary of f_1 that is closer to d_1 . All the points on the segment $\overline{z_1z_0}$ flow into z_0 through f_1 . We then look at the cofaces of f_1 which are the Voronoi cells corresponding to the endpoints of l_1 . Each of these cells contributes a triangle to $\text{Sm}(c)$ that has $\overline{z_1z_0}$ as an edge and its corresponding site as the third vertex (Figure 2.3(f)). Of

course, the inflow of z_1 is computed similarly and is appended to the output. The process stops when the driver of a Voronoi facet is contained in the facet and is therefore a critical point c' (Figure 2.3(g)). Note that this critical point will be the midpoint of a Delaunay edge that intersects its dual Voronoi facet and therefore has index 1. This case also adds to $\text{Sm}(c)$ a segment attached to two triangles (Figure 2.3(h)). However, we no longer trace the inflow of c' since c' itself is not in $\text{Sm}(c)$ even though it is in its closure.

It is not hard to see from the above discussion that the stable manifold of an index-2 critical point of the distance function h_P induced by an unweighted set of points P is a piece-wise linear surface that is made of a finite number of triangles which we call *patch triangles*. Each patch triangle has exactly one vertex in the point-set P .

A closer inspection of the construction of $\text{Sm}(c)$ of a critical point c as given by the algorithm INFLOW (Figure 2.2) can reveal more details about the structure of this object.

Lemma 2.15 *Let Q be a set of (possibly weighted) points in \mathbb{R}^n and let c be a critical point of h_Q of index k . Then every critical point c' in the boundary of $\text{Sm}(c)$ has index less than k , provided that $\text{Sm}(c)$ does not intersect the $(n - k - 1)$ -skeleton of $\text{Vor } Q$.*

Proof. Observe that every critical point c' on the boundary of $\text{Sm}(c)$ is a drivers d' encountered in line 3 of the algorithm that is the intersection point of a Voronoi face and its dual Delaunay face in $\text{Vor } Q$ and $\text{Del } Q$, respectively. Also, notice that by construction, every neighborhood of d' intersects $\text{Sm}(c)$. In particular, the lowest dimensional Voronoi face $V_Q(c')$ that contains c' intersects $\text{Sm}(c)$. By the assumption of the Lemma,

$$\dim V_Q(c') \geq n - k.$$

implying that

$$\dim D_Q(c') = \text{ind } c' \leq k.$$

This is weaker than the claim of part (1); we must show that the above inequality is strict. To this end, observe that the algorithm INFLOW starts on an $n - k$ dimensional Voronoi face $V = V_Q(c)$. It proceeds by only considering the *proper* cofaces of V . All these cofaces have dimensions strictly greater than $n - k$. Tracing the inflow that comes from proper cofaces takes the algorithm to processing of higher dimensional Voronoi faces. Thus the only way the algorithm can trace the inflow of c to another Voronoi face of dimension $n - k$ is by reaching this face as a proper face V'' of a higher dimensional Voronoi face V' in line 6. Observe that in this case, the flow from V'' enters V' transversally and therefore $d_Q(V'') = d_Q(V')$ by Proposition 2.4. Thus if $V'' \cap D_Q(c') \neq \emptyset$, then c' is the intersection point of V' and its dual Delaunay face as well. But c' being on the boundary of V' contradicts our critical general position assumption (Assumption 2.2). This completes the proof of the Lemma. ■

The requirement that $\text{Sm}(c)$ must stay clear from faces of $\text{Vor } Q$ of dimension $n - k - 1$ or smaller may seem too strong at first glance. However we shall argue below that this condition is met by the stable manifolds of all critical points of almost every set of points Q .

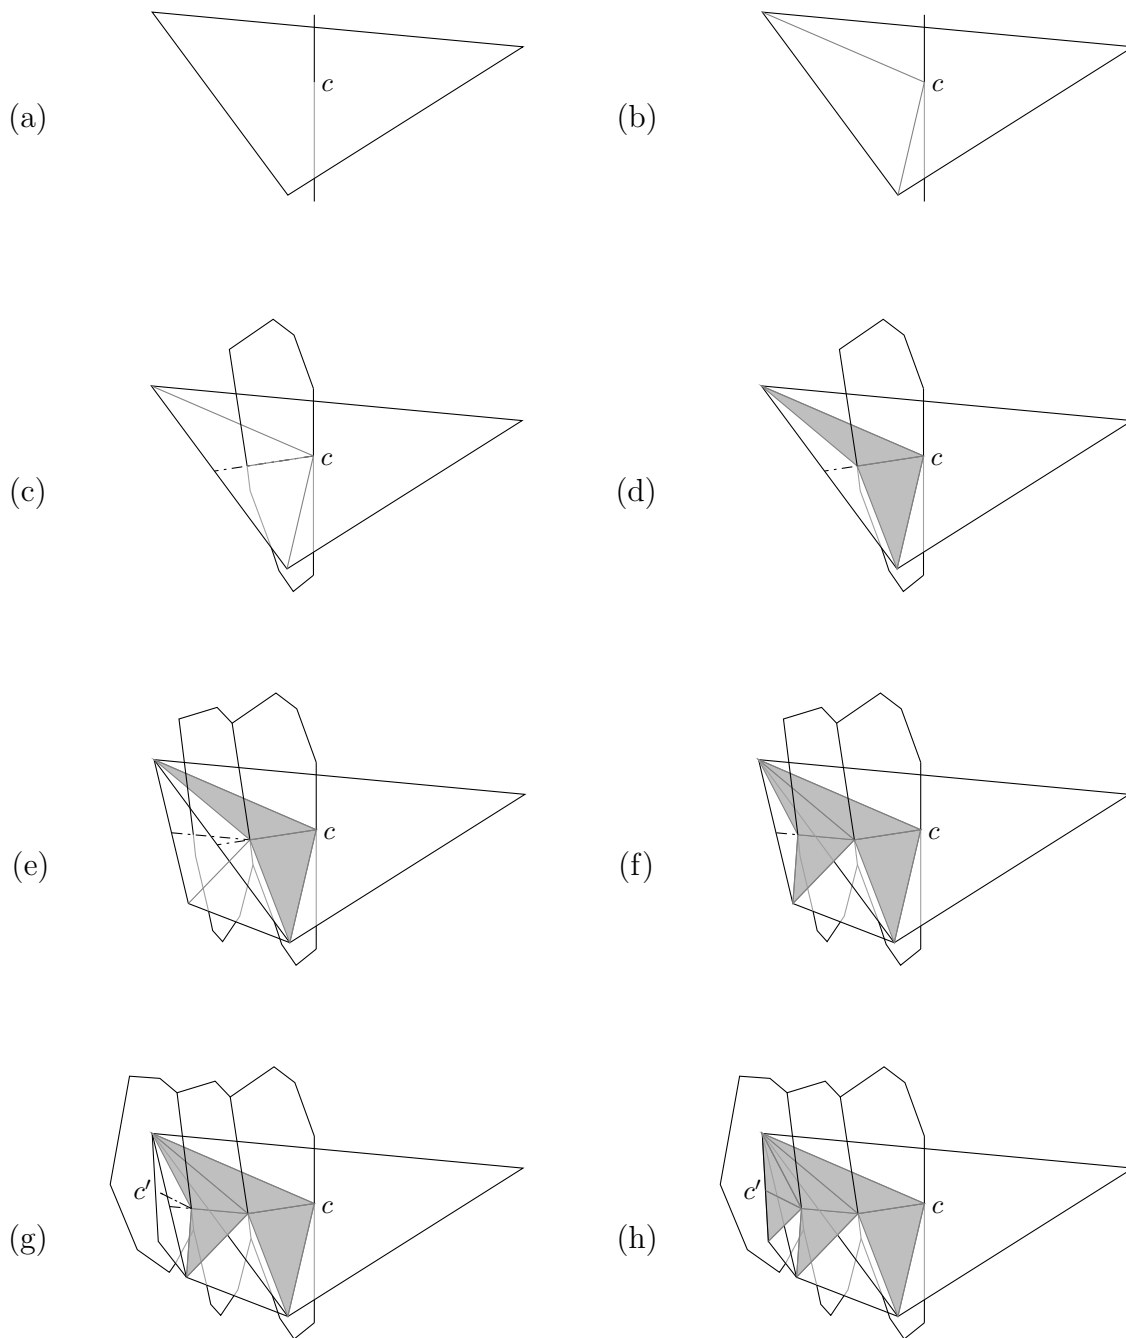


Figure 2.3: Computing the stable manifold of an index-2 saddle point in 3D.

Observe that when $n = 3$ and $k = 2$, this condition requires that stable manifolds of index-2 saddle points stay clear of Voronoi vertices; this is exactly what was assumed in the above examining of such stable manifolds for unweighted point-sets in \mathbb{R}^3 . We will use the following simple exercise from linear algebra to show that this is indeed a sensible assumption.

Lemma 2.16 *Let A be a fixed i -flat in \mathbb{R}^n and let B be a generic j -flat of \mathbb{R}^n . Then $A \cap B$ is an $(i + j - n)$ -flat. In particular, if $i + j < n$, $A \cap B = \emptyset$.*

Consider the invocation of INFLOW with parameters (P, V) where V is an i -dimensional Voronoi face and P is an ℓ -dimensional convex polytope in V and consider the stage of the algorithm in which INFLOW examines an r -dimensional coface V' of V . The affine hull of the set $S = \text{rel int}(\text{conv}(P \cup \{d'\}))$ where $d' = d_Q(V')$ is $\ell + 1$ dimensional. This is because the point $d' \notin \text{aff } V$ since otherwise the flow on V would be tangential and the Voronoi face V' would be ignored. The affine hull of an i -dimensional Voronoi face is determined by $n - i$ sites. Now, let V'' be a j -dimensional face of V' as considered by the algorithm. We are interested in knowing the dimension of the the intersection of affine hulls of V'' and S . Notice that the location of the driver d' of V' is determined only by the sites that define $\text{aff } V'$. Thus let us treat these sites, and therefore $\text{aff } S$, as fixed and the face V'' as variable. Since $V'' \leq V'$, the defining sites of V'' contain those of V' . But being j -dimensional, it has $n - j - (n - r) = r - j$ more points. These $r - j$ points can determine freely any j dimensional affine subspace of the r dimensional space $\text{aff } V'$. Thus by Lemma 2.16, with $\text{aff } S$ being fixed, the dimension of $\text{aff } S \cap \text{aff } V''$ is generically $j + \ell + 1 - r$ and in particular if $j + \ell + 1 < r$, $\text{aff } S$ and $\text{aff } V'$ do not intersect.

Now let us remember that for computing the stable manifold of an index k -critical point c , INFLOW starts with parameters $P = \{c\}$ and $V = V_Q(c)$. Therefore, $\ell = \dim P = 0$ and $i = \dim V = n - k$. Proper cofaces of V have dimensions $r > n - k$ and therefore the Voronoi faces V'' that can serve as second parameters for the subsequent recursive calls can have dimensions $n - k$ or higher. The dimension of the resulting intersection between $\text{aff } S$ and $\text{aff } V''$ for a j -dimensional V'' with $j \geq n - k$ will be $j + 1 - r$. Using this, one can show inductively that in subsequent recursive calls as well, the stable manifold being computed never intersects a Voronoi face of $n - k - 1$ or lower dimensions.

We formalize the following observations into a general position assumption.

Assumption 2.17 (stable manifold general position) A set of (possibly weighted) points $Q \subset \mathbb{R}^n$ is said to be in *stable manifold general position* if for every index- k critical point c of h_Q , $0 \leq k \leq n$, $\text{Sm}(c)$ does not intersect any face of $\text{Vor } Q$ of dimension $n - k - 1$ or lower. From this point on, we only study stable manifolds of critical points of distance to point-sets in stable manifold general position.

Lemma 2.18 *Let x be a point on the boundary of $\text{Sm}(c)$ for a critical point c of the squared distance function h_Q induced by a set Q of (possibly weighted) points in \mathbb{R}^n . If $x \in \text{Sm}(c')$ for a critical point c' other than the the maximum at infinity, then $c' \in \partial \text{Sm}(c)$.*

Proof. The proof is a simple consequence of the continuity of the flow map ϕ_Q . Since $x \in \partial \text{Sm}(c)$, every neighborhood of x contains a point in $\text{Sm}(c)$. Consider the flow line $\phi_Q(x)$ between x and c' . Since c' is not at infinity, this flow line has bounded length and is a piece-wise linear curve with a finite number of turning points. The speed of the flow at every linear segment of this curve is strictly increasing. Therefore, the flow from x reaches c' in finite time. In other words, there is a $t \geq 0$ such that $c' = \phi_Q(t, x)$. Now by the continuity of ϕ_Q , for every $\varepsilon > 0$, there is a $\delta > 0$, such that for every $y \in B(x, \delta)$, $\phi_Q(t, y) \in B(c', \varepsilon)$. If y is chosen from $\text{Sm}(c)$, then $\phi_Q(t, y) \in \text{Sm}(c)$ as well. This implies that there is a point of $\text{Sm}(c)$ at the ε -neighborhood of c' for every $\varepsilon > 0$. Therefore $c' \in \partial \text{Sm}(c)$ as desired. ■

Computation of Unstable Manifolds. Now let us see how we can compute the unstable manifold of any critical point of a distance function. Let c be a critical point of h_Q induced by a set of weighted points $Q \subset \mathbb{R}^n$ and let $V = V_Q(c)$ be the lowest dimensional Voronoi face of $\text{Vor } Q$ that contains c . The critical point c is the driver of the points in the relative interior of V . Also, since V is convex, the boundary of V is visible from c . Consequently for any $r > 0$,

$$V \subset \phi_Q(B(c, r) \cap V),$$

and from this $V \subseteq \text{Um}(c)$. Since by definition $\text{Um}(c)$ is flow-tight for ϕ_Q , we must have

$$\phi_Q(V) \subseteq \text{Um}(c).$$

In fact it can be proven that the above containment is an equality.

Proposition 2.19 *Let c be a critical point of h_Q for a set of weighted points Q . Then*

$$\text{Um}(c) = \phi_Q(V),$$

where $V = V_Q(c)$. Moreover, $\text{Um}(c)$ is closed in the topology of \mathbb{R}^n .

Proof. Let W denote the union of V and all its cofaces. We show that for any point x not in $\phi_Q(V)$, there is an $r > 0$, such that no flow line starting in the r -neighborhood of c reaches x , i.e. $x \notin \phi_Q(y)$ for any $y \in B(c, r)$. It is enough to show this only for points $x \in W \setminus V$ (i.e. an arbitrary point on a proper coface of V) because this implies that for some neighborhood $U = \text{int } W$ of c , $\text{Um}(c) \cap U \subset V$ which implies $\text{Um}(c) = \phi_Q(U \cap V) = \phi_Q(V)$.

As a first step, we show that for every proper coface V' of V , $d_Q(V') \notin \text{aff } V$. In particular, this means that $d_Q(V') \neq c$. To see this, notice that if D and D' respectively denote the duals to V and V' in $\text{Del } Q$, from $V < V'$ we have $D' < D$. Since c is a critical point, $\text{aff } V \cap D = \{c\}$. By our critical general position assumption (Assumption 2.2), $c \in \text{rel int } D$. However, $d_Q(V') \in D'$ and $D' \subseteq \partial D$.

To prove the proposition, we trace the inflow of x “backwards” in a manner similar to the algorithm INFLOW and show that the inflow of x is a strictly positive distance away from c . Thus, let x be any point in a proper coface V_0 of V . Let $d_0 = d_Q(V_0)$ be the driver of

V_0 . The flow lines of ϕ_Q on V_0 are line segments whose supporting lines concur in d_0 . We showed above that $d_0 \notin \text{aff } V$. Since $x \notin \text{aff } V$, from the convexity of the relative interior of V_0 , there is a neighborhood $B(c, r_0)$ of c that avoids the compact segment $S_0 = \overline{d_0 x}$.

Now let S'_0 be the intersection of the segment S_0 with V_0 and let V_1 be a coface of V_0 (and therefore a coface of V). The driver $d_1 = d(V_1)$ is not in $\text{aff } V$ as argued above. The set $(V_1 \cup V_0) \setminus V$ is convex and excludes c but includes the segment S_0 and point d_1 and consequently $S_1 = \text{conv}(S'_0 \cup \{d_1\})$. The set S_1 is compact and does not include the point c . Therefore $\text{dist}(c, S_1) > 0$, let $S'_1 = S_1 \cap V_1$. The same argument can be repeated for cofaces of V_1 and so on.

The closed-ness of $\text{Um}(c)$ follows easily from the above characterization. We need the following observation. Let V be a Voronoi face of $\text{Vor } Q$ and let $S \subset \partial V$ be a closed set. Let V' be the subset of ∂V that is visible from $d = d_Q(v)$. Then the two sets $A = \phi_Q(S) \cap V$ and $B = \phi_Q(S) \cap \partial V$ are both closed. To see this, first remember that by Proposition 2.4 all the points in V' have the same driver, i.e. d . The reason is that $\phi_Q(S) \cap V = S \cup (\phi_Q(S') \cap V)$ where $S' = S \cap V'$ is the subset of S that is visible from d . The set V' is a closed set and therefore its intersection with S , i.e. S' is also closed. The flow lines of ϕ_Q in V concur in d and can be used to define a continuous maps $\alpha : V \rightarrow V'$ that maps every point $x \in V$ to the point $\alpha(x) = \overline{dx} \cap V'$. The set A is the inverse image of the closed set S' under the map α and is therefore closed. If we define the map $\beta : \partial V \rightarrow V'$ as the restriction of α to ∂V , then the set B is the inverse image of S' under β and is therefore closed as well.

Using this observation, the unstable manifold $\text{Um}(c)$ of a critical point c can be broken into a finite number of closed sets the union of which has to be closed as well. ■

By Proposition 2.19 computing the unstable manifold of a critical point c consists of computing the flow closure of the Voronoi face $V_Q(c)$.

2.4.2 Stable and Unstable Flow Complexes

Consider a set of (possibly weighted) points Q and let C_Q be the set of critical points of the squared distance function h_Q induced by Q (including the critical point at infinity). The *stable flow complex* of Q , denoted $\text{Sfc } Q$ is the collection of stable manifolds of all critical points in C .

$$\text{Sfc } Q = \{\text{Sm}(c) : c \in C_Q\}.$$

The following proposition follows from Lemmas 2.15 and 2.18 and states some fundamental properties of the stable flow complex induced by a set Q of weighted points.

Proposition 2.20 *The stable flow complex $\text{Sfc } Q$ of a weighted set of points in \mathbb{R}^n has the following properties.*

1. $Sfc Q$ is an open cell complex whose underlying space $|Sfc Q|$ is the entire space \mathbb{R}^n . In other words, the boundary of the stable manifolds of every critical point is made of stable manifolds of other critical points.
2. If $Sm(c) < Sm(c')$ where “ $<$ ” denotes the face/coface relation in $Sfc Q$, then $ind c < ind c'$. In other words, the boundary of stable manifold of a critical point of index i is made of stable manifolds of critical points of index less than i .

Lemma 2.21 *Let c and c' be two critical points of the distance function h_Q induced by a weighted set of points Q . Then in $Sfc Q$, $Sm(c) < Sm(c')$ if and only if $c' \in Um(c)$.*

Proof. If $c' \in Um(c)$, by definition, every neighborhood of c has a point that flows into c' . But this is by definition identical to $c \in \partial Sm(c')$ which by Proposition 2.20 means that $Sm(c) < Sm(c')$. ■

Corollary 2.22 *Let c be a critical point of the distance function h_Q induced by a set of (possibly weighted) points Q . Then there exists a constant $r_c > 0$ such that every point x in the open ball $B(c, r_c)$ is either on the unstable manifold $Um(c)$ of c , or if $x \in Um(c')$ for some $c' \neq c$, then $ind c' < ind c$.*

Proof. Suppose to the contrary that there is a critical point c' with $ind c' > ind c$, such that every neighborhood of c has a point of $Um(c')$. This will then imply that c is a boundary point of $Um(c')$. However, by Proposition 2.19, $Um(c)$ is a closed set in \mathbb{R}^n and therefore includes its boundary. Therefore $c \in Um(c')$. But then by Lemma 2.21, $Sm(c) > Sm(c')$ which by Proposition 2.20 implies $ind c > ind c'$, a contradiction. ■

Like stable flow complex, the unstable flow complex of a set of (possibly weighted) points $Q \subset \mathbb{R}^n$ is defined as

$$Ufc Q = \{Um(c) : c \in C_Q\}.$$

Stable manifolds $Sm(c)$ and $Sm(c')$ of any two distinct critical points c and c' are disjoint. But this is not in general true about unstable manifolds. In particular unstable manifold of a critical point includes its boundary which is made of unstable manifolds of higher index critical points. This relationship between index of incident unstable manifolds follows from Lemma 2.21. Specifically, we get for any pair of points c and c' that

$$Sm(c) < Sm(c') \quad \text{if and only if} \quad Um(c) > Um(c').$$

From the implementation point of view, computing the unstable flow complex of a set of points can be achieved through computating the unstable manifold of every critical point separately. The fact that unstable manifolds of two critical points can intersect poses a problem on efficient computation of the unstable flow complex: the computed polytopes

may be repeated or may highly overlap. In other words, a naive computation of individual unstable manifolds gives us a polytope soup that doesn't have the structure of an embedded cell-complex and is often highly redundant. To this end we enhance the definition of the unstable flow complex for practical purposes through the following idea.

We can define an embedded cell complex decomposition of space by grouping together, as cells, the points of the space that are contained in the unstable manifolds of the the same set of critical points. We therefore define a relation " \sim_u " on pairs of points in \mathbb{R}^n by declaring $x \sim_u y$ if and only if x is contained in the unstable manifolds of exactly the same set of critical points as y . Obviously " \sim_u " is an equivalence relation. The *unstable flow complex* induced by a point-set P , denoted $\text{Ufc } P$ is the cell complex whose cells are connected components of the subdivision of space into equivalence classes of the " \sim " relation. Each unstable manifold in $\text{Ufc } P$ can be represented as a subcomplex of this cell complex. Namely, $\text{Um}(c)$ consists of precisely all those faces of this complex that have c in their corresponding sets.

We are particularly interested in computation of unstable flow complexes of unweighted point-sets. For an unweighted point-sets P , it is easy to observe that the no flow enters the relative interior of a full-dimensional Voronoi cell. In other words, every full dimensional Voronoi cell V_p of a site $p \in P$ is contained in the unstable manifold $\text{Um}(p)$ of p . Thus all the computation is performed on the $(n - 1)$ -skeleton of $\text{Vor } P$. In fact it turns out that the $(n - 1)$ -skeleton of $\text{Ufc } P$ is a refinement of the $(n - 1)$ -skeleton of $\text{Vor } P$. For point-sets in \mathbb{R}^3 , this refinement of the Voronoi 2-skeleton can be computed by computing of the flow orbit of every Voronoi vertex and subdividing every Voronoi facet along those flow orbits that cross it.

Bibliography and Remarks.

Generalized gradients are generally defined and rigorously studied in *non-smooth analysis* [26, 27]. The generalized gradient vector field v_Q defined in this chapter and its critical points agree with the corresponding notions in the setting of non-smooth analysis. Distance functions to closed sets and their critical point theories have been subject of intensive studies (see [41, 46]). The notion of critical points for distance functions to a point in Riemannian geometry has been introduced by Grove and Shiohoma [47].

Flows and orbits are usually defined in the context of dynamical systems. For a proof of Theorem 2.5 and related results see [49]. In differentiable dynamical systems, time inversion results another dynamical systems. In the distance flow introduced in this chapter however, flow lines can join and because of this, reversing the time can lead to non-deterministic flows; a reverse flow line can reach a multi-way fork.

The notion of critical points is also closely related to *Morse Theory*. In differential topology, the techniques of Morse theory give a very direct way of analyzing the topology of a manifold by studying critical points of differentiable functions on that manifold. Consider a smooth function $f : M \rightarrow \mathbb{R}$ defined on a smooth compact d -manifold M without boundary. A point $c \in M$ is critical if

$$\left(\frac{\partial f}{\partial x_1}(c), \dots, \frac{\partial f}{\partial x_d}(c) \right) = (0, \dots, 0).$$

At regular (non-critical) points of f , the gradient of f is non-zero and f has a direction of steepest ascent. A *non-degenerate* critical point of f is one for which the determinant of Hessian of f at c is non-zero, i.e. the matrix $(H_{ij})_{i,j=1}^n$, where

$$H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$$

is full-rank. If all critical points of f are non-degenerate, f is called a *Morse function*. The values $f(c)$ where c is a critical point of f are called *critical values* of f .

The central Morse lemma states that when f is a Morse function, at every critical point c of f the local coordinates can be chosen in such a way that in a neighborhood of c ,

$$f(x_1, \dots, x_d) = \pm x_1^2 \pm \dots \pm x_d^2 + \text{higher order terms.}$$

The number of minuses in the above representation is called the *index* of c .

Let f be a Morse function defined on a manifold M and let M_α denote the subset of M on which the value of the function f is less than or equal to α . The fundamental theorem of Morse theory states that if f has not critical points in the interval $[\alpha, \beta]$, then M_α is a deformation retract of M_β . In other words the

topology of M_α changes only at those values of α which are critical values of f . Moreover, the type of change in the topology of the M_α at a critical value α of f depends on the index of the critical point c for which $f(c) = \alpha$. More precisely if α is a critical value of f corresponding to a critical point of index k and f has no other critical values in the interval $[\alpha - \varepsilon, \alpha + \varepsilon]$, then $M_{\alpha+\varepsilon}$ is homotopy equivalent to $M_{\alpha-\varepsilon}$ with a k -cell attached (see the notes on CW complexes at the end of Chapter 1).

Integration of generalized distance gradient fields induced by weighted points was introduced by Giesen and John [43, 51]. It was they who coined the term “driver” and showed that for any set of weighted points in \mathbb{R}^n and any point x , the vector $x - d_Q(x)$ gives the direction of steepest ascent of h_Q at x (Theorem 2.3). They also characterized and gave algorithms for computing stable flow complex of weighted sets of points in two and three dimensions. The characterization of the stable flow complex in 2 and 3 dimensions is due to John [51]. In 3D, he presents an algorithm that exactly computes the stable flow complex of every index 2 critical point. Doing so automatically leads to the computation of stable manifolds of index-1 critical points (as boundary faces of index-2 stable manifolds). The algorithm INFLOW presented in this Chapter is a generalization of the algorithm of John which is due to Buchin and Giesen [18]. Proposition 2.20 is a consequence of their algorithm and the resulting characterization.

It must be mentioned however that distance flow ideas existed before Giesen and John although not fully formalized. For example, it was based on a distance flow that Edelsbrunner introduced his WRAP shape reconstruction algorithm (which is the subject of our study in Chapter 6).

In continuous setting a prolific line of work on topology and reconstruction of compact sets and their medial axes based on distance functions was started by Lieutier’s seminal paper on the homotopy of the medial axis [56]. In this paper, the author studies the distance function induced by the boundary Σ of an arbitrary bounded open set S in \mathbb{R}^n . He defines for a point $x \in S$ the distance function

$$s(x) = \text{dist}(x, \partial S) = \text{dist}(x, \Sigma),$$

and defines for any point $x \in S$, the set $A(x)$ as the subset of Σ that realize $s(x)$, i.e.

$$A(x) = \{y \in \Sigma : \|x - y\| = s(x)\}.$$

He then defines the closed ball $\overline{B}(d(x), r(x))$ as the smallest ball that encloses the set $A(x)$. Notice that by definition, $r(x) \leq s(x)$ for all $x \in S$. The center of this ball is the *driver* of x . The generalized gradient vector $v(x)$ for points $x \in S$ is defined as

$$v(x) = \frac{x - d(x)}{s(x)}.$$

It turns out that for points $x \notin M(S)$ the gradient $\nabla s(x)$ is defined and equals $v(x)$. It is then shown that the Euler schemes of the vector field v on S uniformly converge to a flow map $\phi : \mathbb{R} \times S \rightarrow S$ that is continuous on both variables.

It can then be shown that for any $x \in S$, the map $t \mapsto v(\phi_Q(t, x))$ is the right-derivative of the map $x \mapsto \phi(t, x)$ and that the map $t \mapsto \|v(\phi(t, x))\|^2$ is the right-derivative of the map $t \mapsto s(\phi(t, x))$. As a result one gets the integral equations

$$\phi(t, x) = \int_0^t v(\phi(\tau, x)) d\tau,$$

and

$$s(\phi(t, x)) = \int_0^t \|v(\phi(\tau, x))\|^2 d\tau. \quad (2.12)$$

Notice the similarity of the above equation and the one proved in Corollary 2.10 for the squared distance flow in this chapter. In particular, this implies that the map $t \mapsto s(\phi(t, x))$ is increasing.

Finally, Lieutier proved in [56] that the map

$$t \mapsto r(\phi(t, x))$$

is lower-semi-continuous and non-decreasing. A consequence of this is that the medial axis $M(S)$ of S is flow tight for ϕ . This is because a point x is in $M(S)$ if and only if $r(x) > 0$. Since moving along a flow path doesn't decrease r , the flow orbit of any point of $M(S)$ is contained in $M(S)$.

After developing the above tools, it becomes easy to show that the open set S and its medial axis are homotopy equivalent. Let Δ be an upper bound on the diameter of S and let x be an arbitrary point in S . Since $M(S)$ is flow-tight for ϕ , if $\phi(\Delta, x) \notin M(S)$, it must be that $\phi(t, x) \notin M(S)$ for all $t \in [0, \Delta]$. Since for every point x outside $M(S)$, $\|v(x)\| = 1$, if $\phi(\Delta, x) \notin M(S)$ by Equation (2.12) we get

$$\begin{aligned} s(\phi_S(\Delta, x)) &= s(x) + \int_0^\Delta \|v(\phi_S(\tau, x))\|^2 d\tau \\ &\geq s(x) + \int_0^\Delta d\tau = s(x) + \Delta. \end{aligned}$$

However, $s(x) < \Delta$ for every point $x \in S$ and this implies that $\phi(\Delta, x) \in M(S)$ for all $x \in S$. This along with the fact that $M(S)$ is flow tight fulfills all the requirements of Proposition 1.8 and therefore S and $M(S)$ are homotopy equivalent.

The application of the continuity of the flow map ϕ in the above argument is the key to several of our proof ideas in the following chapters. The above-mentioned results of Lieutier (continuity of the map ϕ , integral equation 2.12, etc) can perfectly be adopted for discrete sets of unweighted points. If we enclose the set P of points in question with a very large open ball B and take the bounded open set S to be $B \setminus P$, then near the points of P , the vector fields v and v_P agree (except in magnitude) and the flow lines of the maps ϕ and ϕ_P coincide. This is enough to prove all the results we intend to prove in the following chapters as

long as the points sets involved are not weighted (all chapters except Chapters 6).

However, the proofs of Lieutier do not apply to weighted point-sets. In fact, the Voronoi diagrams (power diagrams) of weighted points sets have to be defined for squared distance functions. This becomes more clear when we remember the definition of Voronoi diagrams as a lower envelope of parabolas. Changing square distances to ordinary distances correspond to replacing parabolas with cones. The problem with doing this is that in general the projection of the intersection of two cones back to the n -dimensional working space will not be an affine subspace and is instead a hyperboloid of some dimension. This entirely compromises the definition of Voronoi diagrams and the results of Lieutier in particular.

Chapter 3

Separation of Critical Points

Our main objective in this monograph is to study the flow maps induced by samples of surfaces and their related structures and in order to investigate their capacity for topological analysis of the sampled surfaces and their medial axes. The present chapter is devoted primarily to a fundamental result of this thesis (Theorem 3.3) that restricts the whereabouts of the critical points of the squared distance function h_P induced by an ε -sample P of a surface Σ as function of the sampling density ε . Loosely put, it states that every critical point of h_P lies either very close to the surface itself or rather close to its medial axis and the extent of this “closeness” depends on ε . In Section 3.1 we explicate and prove this “separation” of critical points. Section 3.2 presents an algorithm for distinguishing critical points of the two families from each other. Section 3.3 introduces a topological thickening of the surface (as well as corresponding shrunk versions of the inner and outer shapes) that lay the foundation for feature flow-based analyses. In particular, in Section 3.3.2 we see that these shrunk shapes are flow-tight for ϕ_P . This observation is central to several of the proofs presented later. Finally, in Section 3.4 we show how the critical points near the medial axis can further be subdivided algorithmically to those near the inner versus those near the outer medial axis.

In the following P is always an ε -sample (with ε to be specified) of a smooth closed surface Σ embedded in \mathbb{R}^n and is treated as a set of unweighted points. The medial axis of Σ is denoted by M , the bounded volume enclosed by Σ by S , and the complement of $S \cup \Sigma$ by S^* . As usual, h_P is the squared distance function to P and ϕ_P is the flow map from integrating the vector field v_P that generalizes the gradient of h_P .

3.1 Surface versus Medial Axis Critical Points

For any point $x \in \mathbb{R}^n \setminus (\Sigma \cup M)$ the points \hat{x} and \tilde{x} are uniquely defined and the point x lies on the segment $\overline{\hat{x}\tilde{x}}$ (See Section 1.6.1). Recall that $\mu(x) = \|\tilde{x} - \hat{x}\|$. Of course \tilde{x} can be at infinity in which case the segment in question becomes a half-line. If $\mu(x) \neq \infty$, the ratio

$$0 < \frac{\|x - \hat{x}\|}{\|\tilde{x} - \hat{x}\|} = \frac{s(x)}{\mu(x)} < 1,$$

is a relative measure of how close to the surface or the medial axis the point x is; the closer this ratio is to zero, the closer x is to $\hat{x} \in \Sigma$ and the closer the ratio is to 1, the closer x is to $\check{x} \in M$. The “separation” result of this chapter uses this measure of proximity to states that the critical points of h_P are concentrated close to either the surface Σ or its medial axis M . More precisely, every critical point c of h_P is either only a small fraction (depending on ε) of $\mu(c)$ apart from \hat{c} or is within a similar distance from \check{c} . Of course a constant fraction of $\mu(c)$ when $\mu(c) = \infty$ is meaningless. The following lemma shows that this should not be a concern.

Lemma 3.1 *Let x be a point in $\mathbb{R}^n \setminus (\Sigma \cup M)$ with $\mu(x) = \infty$. Then, x is not a critical point of h_P and the angle between the vectors $\check{x} - x$ and $v_P(x)$ is strictly less than $\pi/2$.*

Proof. When $\mu(x) = \infty$, \check{x} is at infinity and the plane H tangent to Σ at \hat{x} separates x from Σ and in particular from $A_P(x) \subset \Sigma$ and therefore from $\text{conv } A_P(x)$. Consequently, x cannot be a critical point of h and for every point $y \neq \hat{x}$ on the ray shot from \hat{x} in direction $x - \hat{x}$, the angle between the vectors $v_P(y)$ and $\hat{y} - y$ is strictly less than $\pi/2$. ■

In general, if $\mu(x) = \infty$ the plane tangent to Σ at \hat{x} is a supporting hyperplane of Σ , meaning that \hat{x} is on the boundary of $\text{conv } \Sigma$. Thus the previous Lemma immediately implies that no critical points of h_P can be found outside $\text{conv } \Sigma$. In fact, every critical point of h_P is contained in some Delaunay simplex of $\text{Del } P$ and is therefore contained in $\text{conv } P \subset \text{conv } \Sigma$.

Recall that it is crucial to consider a critical point (a maximum) at infinity. This doesn't tamper with the validity of the proof of the previous lemma since the proof clearly examines “specific” points of the space. The critical point at infinity, however, is symbolic and is excluded from the following discussion.

Lemma 3.2 *Let $x \in \mathbb{R}^n \setminus (\Sigma \cup M)$ be a point satisfying $\|x - \hat{x}\| = \lambda f(\hat{x})$. Let B_μ be the medial ball tangent to Σ at \hat{x} and at the same side of Σ as x and let B_f be the ball of radius $f(\hat{x})$ tangent to Σ at \hat{x} but on the opposite side of Σ relative to x (and B_μ). Let $\mu_0 = \mu(x)/f(\hat{x})$. Define the set $L(x) \subset \mathbb{R}^n$ as*

$$L(x) = \overline{B}(x, \ell(\varepsilon, \lambda)f(\hat{x})) \setminus (B_\mu \cup B_f),$$

in which ε is the sampling density and

$$\ell(\varepsilon, \lambda) = \sqrt{\lambda^2 + \varepsilon^2(1 + \lambda)}.$$

Then,

1. $A_P(x) \subset L(x)$. In particular, the points in $A_P(x)$ are within distance $\ell(\varepsilon, \lambda)f(\hat{x})$ from x .




Figure 3.1: Illustration of the proof of Lemma 6.7

2. $L(x) \subset \overline{B}(x, r_{\mu_0}(\varepsilon, \lambda)f(\hat{x})) \subseteq \overline{B}(x, r(\varepsilon, \lambda)f(\hat{x}))$ where

$$r_{\mu_0}(\varepsilon, \lambda) = \varepsilon \sqrt{\frac{1 + \lambda}{1 - \lambda/\mu_0}} \quad \text{and} \quad r(\varepsilon, \lambda) = \varepsilon \sqrt{\frac{1 + \lambda}{1 - \lambda}}.$$

3. $L(x)$ is contained in a cone with apex x , axis $\hat{x} - x$, and half-angle ψ where

$$\psi = \psi(\varepsilon, \lambda) = \arcsin\left(\frac{r(\varepsilon, \lambda)}{\ell(\varepsilon, \lambda)}\right).$$

Proof. Both B_μ and B_f avoid Σ . By the ε -sampling assumption, there is a sample point in the ball $B_\varepsilon = \overline{B}(\hat{x}, \varepsilon f(\hat{x}))$. Thus the closest point in P to x is within distance ℓ from x , where ℓ is the distance between x and y , where y is a farthest point from p in the set $L_0 = B_\varepsilon \setminus (B_\mu \cup B_f)$. Notice that ℓ is an upper bound for $h_P(x)$. Figure 3.1 shows a planar section of this setting. In the figure, the region L_0 is shaded with the darkest gray. Let B_ℓ denote the ball $\overline{B}(x, \ell)$. Since P does not intersect B_μ or B_f , $A_P(x)$ must be contained in the region

$$L_1 = L(x) = B_\ell \setminus (B_\mu \cup B_f).$$

Let z be a point in this region farthest away from \hat{x} and let $r = \|\hat{x} - z\|$. Let γ be the angle between $y - \hat{x}$ and the hyper-plane tangent to Σ at \hat{x} . It can be easily seen from Figure 3.1 (right) that $\sin \gamma = \varepsilon/2$.

To simplify notation, let us take $f(\hat{x})$ as unit length. Since the angle $\angle(y - \hat{x}, x - \hat{x})$ is $\pi/2 + \gamma$ we have using the cosine rule

$$\begin{aligned}\ell^2 = \|x - z\|^2 = \|x - y\|^2 &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \cos(\pi/2 + \gamma) \\ &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \sin \gamma \\ &= \lambda^2 + \varepsilon^2(1 + \lambda).\end{aligned}$$

Now by applying the cosine rule to the triangle $xz\tilde{x}$, we have for the angle $\beta = \angle(x - \tilde{x}, z - \tilde{x})$

$$\begin{aligned}\cos \beta &= \frac{\mu_0^2 + (\mu_0 - \lambda)^2 - \ell^2}{2\mu_0(\mu_0 - \lambda)} \\ &= \frac{\mu_0^2 + (\mu_0 - \lambda)^2 - \lambda^2 - \varepsilon^2(1 + \lambda)}{2\mu_0(\mu_0 - \lambda)} \\ &= 1 - \frac{\varepsilon^2}{2\mu_0^2} \cdot \frac{1 + \lambda}{1 - \lambda/\mu_0}.\end{aligned}$$

If we rewrite the above equality as

$$\cos \beta = 1 - 2 \left(\frac{\varepsilon}{2\mu_0} \sqrt{\frac{1 + \lambda}{1 - \lambda/\mu_0}} \right)^2,$$

and observe on the figure that $\sin(\beta/2) = (r/2)/\mu_0$, we can use the identity $\cos \beta = 1 - 2\sin^2(\beta/2)$ to obtain,

$$r = \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda/\mu_0}}.$$

To complete the proof, we need to only show that the angle $\beta' = \angle(\hat{x} - x, z - x)$ is smaller than $\psi(\varepsilon, \delta)$ as given in the statement of the Lemma. From the figure $\sin \beta' = h/\ell$ where h is the distance between z and the line supporting the segment $x\hat{x}$. Since $h \leq r$, $\sin \beta' \leq r/\ell = \sin \psi$.

■

Theorem 3.3 *Let $\varepsilon < 1/\sqrt{3}$ and let x be a point in $\mathbb{R}^n \setminus (\Sigma \cup M)$ with $\mu(x) < \infty$. If*

$$\varepsilon^2 f(\hat{x}) \leq s(x) \leq \mu(x) - 2\varepsilon^2 f(\hat{x}),$$

then x is not a critical point. Furthermore, then $\angle(\tilde{x} - x, v_P(x)) < \pi/2$.

Proof. To simplify the notation, we take $f(\hat{x})$ as unit length. Let us use the shorthand λ for $s(x)$. By Lemma 3.2, $A_P(x) \subset B_\ell = B(x, \ell)$, where $\ell = \ell(\varepsilon, \lambda)$ as defined in Lemma 3.2. On the other hand $A_P(x)$ is disjoint from $B_\mu = B(\tilde{x}, \mu)$ where $\mu = \mu(x)$. Let H be the hyperplane normal to $\tilde{x} - \hat{x}$ through x and let R be the radius of the ball of intersection between H and B_μ . The plane H is at distance $\|x - \tilde{x}\| = \mu(x) - \lambda$ from \tilde{x} . By the

Pythagorean theorem

$$R^2 = \mu^2 - (\mu - \lambda)^2.$$

If the radius ℓ of B_ℓ is less than R , then $B_\ell \setminus B_\mu$ is strictly contained in the open half-planes of $\mathbb{R}^n \setminus H$ that contains \hat{x} . Since $x \in H$, this implies that $x \notin \text{conv}(B_\ell \setminus B_\mu)$ which further entails that $x \notin \text{conv} A_P(x)$. Since $d_P(x) \in \text{conv} A_P(x)$, this would imply that $\angle(\check{x} - x, v_P(x)) < \pi/2$. In particular, x cannot be a critical point if $R > \ell$ or equivalently if

$$\mu^2 - (\mu - \lambda)^2 > \lambda^2 + \varepsilon^2(1 + \lambda).$$

Rearranging the above inequality gives us

$$2\lambda^2 + (\varepsilon^2 - 2\mu)\lambda + \varepsilon^2 < 0.$$

Solving for λ , we get $\lambda_{\min} < \lambda < \lambda_{\max}$, where

$$\begin{aligned} \lambda_{\min} &= \frac{1}{2} \left(\mu - \varepsilon^2/2 - \sqrt{(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2} \right) \\ \lambda_{\max} &= \frac{1}{2} \left(\mu - \varepsilon^2/2 + \sqrt{(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2} \right). \end{aligned}$$

Since $\mu \geq 1$, $\varepsilon \leq 1/\sqrt{3}$ is sufficient to have $(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2 \geq 0$. Thus for $\varepsilon \leq 1/\sqrt{3}$, both λ_{\min} and λ_{\max} are real.

The assumption of $\varepsilon \leq 1/\sqrt{3}$ can be written as $3\varepsilon^2 \leq 1$ from which

$$2\varepsilon^2 + 1 \leq 2 - \varepsilon^2 \leq 2\mu - \varepsilon^2 = 2(\mu - \varepsilon^2/2).$$

Multiplying by $2\varepsilon^2 \geq 0$, gives us

$$2\varepsilon^2(2\varepsilon^2 + 1) \leq 4\varepsilon^2(\mu - \varepsilon^2/2).$$

By adding $(\mu - \varepsilon^2/2)^2$ to both sides and rearranging we get

$$(\mu - \varepsilon^2/2)^2 - 2\varepsilon^2 \geq ((\mu - \varepsilon^2/2) - 2\varepsilon^2)^2.$$

The smaller side being non-negative allows us to take square roots of both sides which by rearranging results $\lambda_{\min} \leq \varepsilon^2$.

As for λ_{\max} , using the inequality $\sqrt{1-t} \geq 1-t$ for $0 \leq t \leq 1$

$$\begin{aligned}
\lambda_{\max} &= \frac{1}{2}(\mu - \varepsilon^2/2) \left(1 + \sqrt{1 - \frac{2\varepsilon^2}{(\mu - \varepsilon^2/2)^2}} \right) \\
&\geq \frac{1}{2}(\mu - \varepsilon^2/2) \left(2 - \frac{2\varepsilon^2}{(\mu - \varepsilon^2/2)^2} \right) \\
&= \mu - \varepsilon^2/2 - \frac{\varepsilon^2}{\mu - \varepsilon^2/2} \\
&\geq \mu - \varepsilon^2/2 - \frac{\varepsilon^2}{1 - \varepsilon^2/2} \\
&\geq \mu - 2\varepsilon^2.
\end{aligned}$$

Thus if $\varepsilon^2 < \lambda < \mu - 2\varepsilon^2$, the point x is separated from $\text{conv } A_P(x)$ and therefore x cannot be a critical point. ■

Corollary 3.4 (Separation of Critical Points) *Let P be an ε -sample of a smooth surface Σ in \mathbb{R}^n for $\varepsilon \leq 1/\sqrt{3}$. Then for every critical point c of h_P , either $c \in \Sigma_{\varepsilon^2}$, i.e.*

$$s(c) = \|c - \hat{c}\| \leq \varepsilon^2 f(\hat{c}),$$

or

$$m(c) = \|c - \check{c}\| \leq 2\varepsilon^2 f(\check{c}).$$

The previous Corollary motivates the following definition.

Surface and Medial Axis Critical Points. Let $\varepsilon < 1/\sqrt{3}$ and let P be an ε -sample of a smooth surface Σ . We call a critical point of h_P a *surface critical point* if it is contained in Σ_{ε^2} and we call it a *medial axis critical point* if it is contained in $M_{2\varepsilon^2}$.

Figure 3.2 shows an example of separation of critical points for a sample of a closed curve in \mathbb{R}^2 .

3.2 Distinguishing Surface and Medial Axis Critical Points

So far we have shown that the set of critical points of the distance function h_P can be partitioned into two subsets: surface critical points and medial axis critical points. Given the sample P , it is easy to compute the set of all critical points of h_P ; it is enough to compute the Delaunay complex $\text{Del } P$ and find all the simplices in it that intersect their dual Voronoi




Figure 3.2: A two-dimensional example of the separation of critical points (Theorem 3.3). Sample points (index-0 critical points) are shown in solid black. Other surface critical points are shown in blue and medial axis critical points are shown in red. Index-1 saddle points are depicted hollowed.

faces. Once the critical points are computed, the next question is that of distinguishing the surface critical points from the medial axis ones.

The present section addresses this algorithmic issue. For the rest of this section, let P be an ε -sample of Σ where bounds on ε are determined as we proceed. Of course $\varepsilon \leq 1/\sqrt{3}$ is assumed to ensure that Corollary 3.4 holds and surface and medial axis critical points are well-defined.

The idea here is to look at where on the boundary of the Voronoi cell of a sample point $p \in A_P(c)$ a critical point c is located. Of course if we only regard the Voronoi cell of p as a convex polytope, there is not much to distinguish between various points of its boundary. However, for dense samples of surfaces, the Voronoi cell V_p of a sample point p is known to be *long and skinny*, extending in the direction of the normal to surface at p . In other words, V_p resembles a long and thin cylinder and therefore it makes sense to tell apart the ends of the cylinder from its middle.

This intuition can be formalized by the introduction of the *poles* of a Voronoi cell. But first that let us introduce some notation.

Notation. The angle between two vectors u and v , denoted $\angle(u, v)$ is always smaller than π . The angle between two lines (or a line and a vector) is the non-obtuse angle formed by them. The *acute angle* between vectors u and v is the smaller of the two angles made by the lines through u and v .

Poles. Recall that for a sample point $p \in P$ the Voronoi cell of p is denoted by V_p . If V_p is bounded, the *positive pole* of p , denoted p^+ , is the Voronoi vertex of V_p farthest away from p . The *positive pole vector* ν_p^+ is the vector $p^+ - p$. If V_p is unbounded, ν_p^+ is taken as the unit vector in the direction which is the average of all unbounded Voronoi edges in V_p (or just any direction in which V_p is unbounded). In the latter case we informally refer to a point at infinity in the direction ν_p^+ as the positive pole. The *negative pole* p^- of p is the farthest Voronoi vertex of V_p from p for which the smaller angle between the vectors ν_p^+ and $\nu_p^- = p^- - p$, called the *negative pole vector* at p , is greater than $\pi/2$ (or equivalently the inner product $\langle \nu_p^-, \nu_p^+ \rangle$ is negative).

The connection between the poles and the shape of Voronoi cells becomes clear by the following lemma and the corollary immediately following it.

Lemma 3.5 [2] *Let p be a sample point in an ε -sample P of a surface Σ . Let x be any point in V_p with $\|x - p\| \geq \xi f(p)$ for $\xi > 0$. Then*

$$\angle(x - p, n_p) \leq \arcsin \frac{\varepsilon}{1 - \varepsilon} + \arcsin \frac{\varepsilon}{\xi(1 - \varepsilon)}.$$

Corollary 3.6 *For any point p of an ε -sampling P of a surface Σ , the acute angle between the normal to surface at p , n_p , and either of ν_p^+ and ν_p^- is at most $2 \arcsin(\varepsilon/(1 - \varepsilon))$.*

The first step in distinguishing the surface critical points from the medial axis ones is to restrict the locations on the boundary of a Voronoi cell where a critical point may reside.

Lemma 3.7 *For any sample point $p \in P$ the ball $B = (\frac{1}{2}(p + p^+), \frac{1}{2}\|p - p^+\|)$, i.e. the ball with diameter pp^+ , does not contain any critical point c with $p \in A_P(c)$. The analogous statement holds for the negative pole p^- .*

Proof. If p^+ lies at infinity then the ball B becomes a half-space with normal ν_p^+ . The boundary of this half-space is a plane that supports the convex hull of P . Since any critical point of h_P must be contained in the convex hull of P it follows that the interior of B , i.e., the open half space, cannot contain any critical point of h_P . Thus we can assume that $\|p^+ - p\|$ is finite. Let c be a critical point for which $p \in A_P(c)$. If c is a minimum, then $c = p$ and there is nothing to prove. Otherwise, $|A_P(c)| > 1$. All points in $A_P(c)$ lie on the boundary of the ball B' of radius $\|c - p\|$ centered at c . Let B'' be the open ball of radius $\|p - p^+\|$ centered at p^+ . By construction there can be no points of P in B'' . Thus all points of $A_P(c)$ must belong to $\partial B' \setminus B''$. On the other hand, for c to be a critical point, it must be in the convex hull of $A_P(c)$. This happens only if the angle $\angle pcp^+$ is smaller than $\pi/2$. The latter condition is in turn identical to c being outside the ball B (see Figure 3.3). The proof of the analogous statement for p^- follows the same lines. ■

Lemma 3.8 *For a surface critical point c and a sample point $p \in A_P(c)$,*

$$\|c - p\| \leq \frac{\varepsilon(1 + \varepsilon)}{1 - \varepsilon(1 + 2\varepsilon)} \cdot \|p - p^-\|.$$




Figure 3.3: The ball of diameter pp^+ contains no critical point on the boundary of V_p .

Proof. Let $q \in P$ be the closest sample point to \hat{c} . By the ε -sampling condition $\|\hat{c} - q\| \leq \varepsilon f(\hat{c})$. By the triangle inequality we obtain

$$\begin{aligned}
 \|c - p\| &\leq \|c - q\| \\
 &\leq \|c - \hat{c}\| + \|\hat{c} - q\| \\
 &\leq \varepsilon^2 f(\hat{c}) + \varepsilon f(\hat{c}) \\
 &= \varepsilon(1 + \varepsilon)f(\hat{c}).
 \end{aligned} \tag{3.1}$$

Since the local feature size function $f(\cdot)$ is 1-Lipschitz we can write

$$\begin{aligned}
 f(\hat{c}) &\leq f(p) + \|\hat{c} - p\| \\
 &\leq f(p) + \|c - \hat{c}\| + \|c - p\| \\
 &\leq f(p) + \varepsilon^2 f(\hat{c}) + \varepsilon(1 + \varepsilon)f(\hat{c}) \\
 &= f(p) + \varepsilon(1 + 2\varepsilon)f(\hat{c}).
 \end{aligned}$$

By rearranging we get

$$f(\hat{c}) \leq \frac{1}{1 - \varepsilon(1 + 2\varepsilon)} f(p). \tag{3.2}$$

Combining (3.1) and (3.2) we finally get

$$\begin{aligned}
 \|c - p\| &\leq \varepsilon(1 + \varepsilon)f(\hat{c}) \\
 &\leq \frac{\varepsilon(1 + \varepsilon)}{1 - \varepsilon(1 + 2\varepsilon)} \cdot f(p) \\
 &\leq \frac{\varepsilon(1 + \varepsilon)}{1 - \varepsilon(1 + 2\varepsilon)} \cdot \|p - p^-\|.
 \end{aligned}$$

The last inequality follows from the observation that the centers of the medial balls tangent

to Σ at p at each side of Σ are contained in V_p . This is because p is a closest surface point to each of these centers and therefore has to be a closest sample as well. ■

From Lemmas 3.7 and 3.8 we derive the following corollary.

Corollary 3.9 *Let c be a surface critical point and let p be a sample point in $A_P(c)$. If $\varepsilon \leq 1/5$ then the angle between the vector $c - p$ and each of the vectors ν_p^- and ν_p^+ is at least 70 degrees.*

Proof. By Lemma 3.7 neither of the two balls with diameters pp^+ and pp^- contain any critical points c with $p \in A_P(c)$. Thus we get from using Lemma 3.8 and Thales theorem that

$$\cos \angle(c - p, \nu_p^+) \leq \frac{\varepsilon(1 + \varepsilon)}{1 - \varepsilon(1 + 2\varepsilon)}.$$

(the same bound holds for $\cos \angle(c - p, \nu_p^-)$). Thus for $\varepsilon \leq 1/5$, the resulting angles are both larger than 70 degrees. ■

Next we show that a medial axis critical point on the boundary of a Voronoi cell V_p behaves similar to a pole of p .

Lemma 3.10 *Let c be a medial axis critical point and let p be a sample point in $A_P(c)$. Then for any $\varepsilon < 1/\sqrt{3}$, $\|c - p\| \geq (1 - 2\varepsilon^2)f(p)$.*

Proof. Since c is a medial axis critical point $m(c) = \|c - \check{c}\| \leq 2\varepsilon^2\mu(c)$. Clearly, $s(c) \leq \|c - p\|$ since $s(c)$ is the distance from c to Σ and $p \in \Sigma$. Moreover, $\mu(c) < \infty$ by Lemma 3.1. Thus we have

$$m(c) \leq 2\varepsilon^2 f(\check{c}) \leq 2\varepsilon^2 \mu(c) \leq 2\varepsilon^2 (m(c) + s(c)).$$

From which we get

$$m(c) \leq \frac{2\varepsilon^2}{1 - 2\varepsilon^2} \cdot s(c) \leq \frac{2\varepsilon^2}{1 - 2\varepsilon^2} \cdot \|c - p\|,$$

which results

$$\begin{aligned} f(p) &\leq \|\check{c} - p\| \\ &\leq m(c) + \|c - p\| \\ &\leq \left(\frac{2\varepsilon^2}{1 - 2\varepsilon^2} + 1 \right) \|c - p\| \\ &= \frac{1}{1 - 2\varepsilon^2} \cdot \|c - p\| \end{aligned}$$

and thus $\|c - p\| \geq (1 - 2\varepsilon^2)f(p)$. ■

Corollary 3.11 *Let c be a medial axis critical point and let $p \in A_P(c)$ be a sample point. If $\varepsilon \leq 1/5$ then the angle between the vector $c - p$ and exactly one of the vectors ν_p^- and ν_p^+ is less than 60 degrees.*

Proof. By Lemma 3.10 $\|c - p\| \geq (1 - 2\varepsilon^2)f(p)$ and therefore using $\xi = 1 - 2\varepsilon^2$ in Lemma 3.5 implies that the angle between $c - p$ and the normal line n_p at p is at most

$$\arcsin \frac{\varepsilon}{1 - \varepsilon} + \arcsin \frac{\varepsilon}{(1 - 2\varepsilon^2)(1 - \varepsilon)}.$$

The above bound is valid for the angle that the vector $c - p$ makes with exactly one of the two vectors n_p^+ or n_p^- . We assume, without loss of generality, that it is n_p^+ . By Corollary 3.6, one of the polar vectors ν_p^+ or ν_p^- (say ν_p^+) makes an angle of at most $2 \arcsin(\varepsilon/(1 - \varepsilon))$ with n_p^+ . Thus the angle between the two vectors ν_p^+ and $c - p$ is at most

$$3 \arcsin \frac{\varepsilon}{1 - \varepsilon} + \arcsin \frac{\varepsilon}{(1 - 2\varepsilon^2)(1 - \varepsilon)}.$$

To complete the proof we simply use $\varepsilon \leq 1/5$ in the above bound. ■

Corollary 3.9 and Corollary 3.11 together imply the following result.

Theorem 3.12 *Let P be an ε -sample of a surface Σ for $\varepsilon \leq 1/5$. Let c be a critical point of h_P and let p be any point in $A_P(c)$. Then c is a medial axis critical point if and only if both of the angles*

$$\angle(c - p, \nu_p^+) \quad \text{and} \quad \angle(c - p, \nu_p^-)$$

are greater than 70 degrees.

Thus the two types of critical points can be distinguished simply by measuring the two angles described in Theorem 3.12.

3.3 Tubular Neighborhoods and Reduced Shapes

Recall from Chapter 1 that for any $0 < \delta < 1$, the δ -tubular neighborhood of a surface Σ (See Figure 3.4) is defined as the set

$$\Sigma_\delta = \{x \in \mathbb{R}^n \setminus M : \|x - \hat{x}\| < \delta f(\hat{x})\}.$$

and the δ -reduced shapes S_δ and S_δ^* are defined as

$$S_\delta = S \setminus \Sigma_\delta \quad \text{and} \quad S_\delta^* = S^* \setminus \Sigma_\delta.$$

Notice that the definition of Σ_δ puts the medial axis $M(S)$ of S in S_δ and the medial axis $M(S^*)$ of S^* in S_δ^* . As a first step in exploring the important properties of these sets we characterize the boundary points of Σ_δ (which happen to coincide with the union of boundary points of S_δ and S_δ^*).




Figure 3.4: The inner shape S and outer shape S^* defined by a surface Σ (left). The inner δ -reduced shape S_δ and the outer δ -reduced shape S_δ^* determined by the δ -tubular neighborhood Σ_δ of Σ (right).

3.3.1 The Boundary of Σ_δ

We begin by showing that every point of $M(S)$ is in fact an interior point of S_δ and the same relation is valid between $M(S^*)$ and S_δ^* .

Lemma 3.13 *For every $0 < \delta < 1$, every point of $M(S)$ is an interior point of S_δ and every point of $M(S^*)$ is an interior point of S_δ^* .*

Proof. We only prove the first claim. The proof of the second claim is identical. Take $x \in M(S)$ and let $B = B(x, \omega)$ where

$$\omega = \frac{1 - \delta}{1 + 2\delta} \cdot s(x) > 0.$$

We will show that $B \subset S_\delta$ implying that x is an interior point of S_δ . Take any point $y \in B$. If $y \in M(S)$, then $y \in S_\delta$ by definition. Thus assume that $y \notin M(S)$ and therefore has a unique closest point \hat{y} in Σ . Notice that since $B' = B(x, s(x)) \subset S$ but $\partial B'$ intersects Σ , we get

$$s(x) - \omega \leq s(y) \leq s(x) + \omega. \quad (3.3)$$

Moreover, since $x \in M(S)$, we get

$$f(\hat{y}) \leq s(y) + \|x - y\| \leq s(y) + \omega \leq 2\omega + s(x). \quad (3.4)$$

Combining Equations (3.3) and (3.4) and using the definition of ω we obtain

$$s(y) \geq s(x) - \omega = \delta s(x) + 2\delta\omega \geq \delta f(\hat{y}). \quad \blacksquare$$

Note, that the previous lemma also implies that every boundary point of S_δ has a unique closest point in Σ . The following lemma gives a complete characterization of the boundary points of S_δ .

Lemma 3.14 *The boundary of S_δ consists of exactly those points $x \in S \setminus M(S)$ that satisfy*

$$\|x - \hat{x}\| = \delta f(\hat{x}).$$

An analogous claim characterizes the boundary points of S_δ^ .*

Proof. Again we only prove the lemma for S_δ since the proof for S_δ^* is similar. Lemma 3.13 shows that no point of $M(S)$ can be a boundary point for S_δ . For every other point $x \in S$, there is a unique $\hat{x} \in \Sigma$. Consider the segment $\hat{x}\tilde{x}$. For every point in this segment the closest point in Σ is the same point \hat{x} . Thus, all points in the relative interior of the segment $\hat{x}\tilde{x}$ lie in Σ_δ and all points of the relative interior of $x\tilde{x}$ lie in S_δ . This implies that x is a boundary point for S_δ .

We now show that S_δ has no other boundary points. First we show that no point $x \in S_\delta \setminus M(S)$ with $s(x) = \lambda f(\hat{x})$ where $\lambda > \delta$ can be a boundary point of S_δ . Let $B = B(\tilde{x}, \mu f(\hat{x}))$ be the medial ball tangent to Σ at \hat{x} at the same side of Σ as x . We show that the open ball centered at x and with radius

$$\omega = \frac{(\mu - \lambda)(\lambda - \delta)}{\mu - \lambda + 2\delta} f(\hat{x}) > 0$$

is entirely contained in S_δ . By definition $B \subset S$. Let y be a point at distance less than ω from x . The angle $\alpha = \angle(x - \tilde{x}, y - \tilde{x})$ grows when y moves away from x in the direction of $y - x$. When $\|x - y\|$ is fixed, α is maximized when $\angle(\tilde{x} - y, x - y) = \pi/2$ in which case $\sin \alpha = \omega / ((\mu - \lambda)f(\hat{x}))$. The point \hat{y} can be at most as far away from y as \hat{x} is. Therefore, \hat{y} lies in the ball $B' = B(y, \|\hat{x} - y\|)$. Since B is an empty ball, this implies that the distance between \hat{y} and \hat{x} is at most $2f(\hat{x}) \sin \alpha = 2\omega / (\mu - \lambda)$ and therefore

$$f(\hat{y}) \leq f(\hat{x}) + \|\hat{x} - \hat{y}\| \leq f(\hat{x}) + \frac{2\omega}{\mu - \lambda}.$$

On the other hand,

$$s(y) \geq s(x) - \omega = \lambda f(\hat{x}) - \omega \geq \delta \left(f(\hat{x}) + \frac{2\omega}{\mu - \delta} \right) \geq \delta f(\hat{y}).$$

Since y was chosen arbitrarily from $B(x, \omega)$, it follows that $B(x, \omega) \subset S_\delta$, meaning x is an interior point of S_δ .

Finally we show that no point of Σ_δ can be a boundary point of S_δ . Let $x \in \Sigma_\delta$ be a point for which $s(x) = \lambda f(\hat{x})$ where $\lambda < \delta$. Let μ be as above. Consider the open ball of radius

$$\omega = \frac{(\delta - \lambda)(\mu - \lambda)}{\mu - \lambda + 2\delta} f(\hat{x}) > 0$$

and let y be a point in this ball. Similar to the previous case, we can show that $\|\hat{y} - \hat{x}\| \leq 2\omega/(\mu - \lambda)$ and therefore

$$f(\hat{y}) \geq f(\hat{x}) - \frac{2\omega}{\mu - \lambda}.$$

On the other hand $s(y) < s(x) + \omega$. Combining these we get

$$s(y) < s(x) + \omega = \delta \left(f(\hat{x}) - \frac{2\omega}{\mu - \lambda} \right) < \delta f(\hat{y}). \quad \blacksquare$$

3.3.2 Flow on Tubular Neighborhoods

Lemma 3.15 *Let x be a point on the boundary of S_δ . Any vector v at x that makes an angle α less than*

$$\arctan\left(\frac{1 - \delta}{2\delta}\right)$$

with the vector $\tilde{x} - x$ points into S_δ .

Proof. Let c be the point on the line segment $\hat{x}\tilde{x}$ at distance $f(\hat{x})$ from \hat{x} . Without loss of generality we assume that $v = y - x$, where $y \in S$ is chosen close enough to x so that the inner angle θ of the triangle xcy at c is less than $\pi/2 - \alpha$. Note that by definition α is the inner angle of the triangle xcy at x , see Figure 3.5. Let $B = B(c, f(\hat{x}))$ and let $B' = \overline{B}(y, \|\hat{x} - y\|)$. B does not contain any point from Σ and $\hat{y} \in B'$. Therefore $\hat{y} \in B' \setminus B$ and $\|\hat{x} - \hat{y}\| \leq 2f(\hat{x}) \sin \theta$ which in turn, together with the fact that f is 1-Lipschitz, implies that

$$f(\hat{y}) \leq f(\hat{x})(1 + 2 \sin \theta).$$

Let w be the intersection point of the boundary of B with the ray through $y - c$ and let z be the closest point to x on this ray. By construction $z \in B$. Our assumption on θ implies that that on the ray through $c - y$ the point y comes before the point z as seen from c . Putting everything together we get

$$\begin{aligned} \|y - \hat{y}\| &\geq \|y - w\| \\ &= \|y - z\| + \|z - w\| \\ &> \|y - z\| + \|x - \hat{x}\| \\ &= \|y - z\| + \delta f(\hat{x}) \\ &= \|x - z\| \cot(\alpha + \theta) + \delta f(\hat{x}) \\ &= (1 - \delta) f(\hat{x}) \sin \theta \cot(\alpha + \theta) + \delta f(\hat{x}). \end{aligned}$$




Figure 3.5: Illustration of the proof of Lemma 3.15.

Thus,

$$\begin{aligned} \frac{\|y - \hat{y}\|}{f(\hat{y})} &> \frac{(1 - \delta)f(\hat{x}) \sin \theta \cot(\alpha + \theta) + \delta f(\hat{x})}{f(\hat{x})(1 + 2 \sin \theta)} \\ &= \frac{(1 - \delta) \sin \theta \cot(\alpha + \theta) + \delta}{1 + 2 \sin \theta}. \end{aligned}$$

In order for y to be in S_δ we want this fraction to be larger than δ . This amounts to $\cot(\alpha + \theta) > 2\delta/(1 - \delta)$ or equivalently to

$$\tan(\alpha + \theta) < \frac{1 - \delta}{2\delta}.$$

Since we assumed $\tan \alpha < (1 - \delta)/2\delta$, by the continuity of the tangent function, we can find a $\theta_0 > 0$ such that $\tan(\alpha + \theta) < (1 - \delta)/(2\delta)$ for all $0 < \theta < \theta_0$. This implies the existence of $\lambda_0 > 0$ such that

$$(1 - \lambda)x + \lambda y = x + \lambda(y - x) = x + \lambda v \in S_\delta,$$

for all $0 \leq \lambda < \lambda_0$. In other words, the vector v points into S_δ at x . ■

Lemma 3.16 *Let x be a point on the boundary of S_δ . The angle α that $v_P(x)$ makes with $\tilde{x} - x$ is bounded by*

$$\arccos \left(\frac{2\delta(1 - \varepsilon - \delta) - \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)} \right),$$

provided that the argument of the arccos is between 0 and 1.




Figure 3.6: Illustration of the proof of Lemma 3.16.

Proof. Let c be the point on the line segment $\overline{\hat{x}\tilde{x}}$ at distance $f(\hat{x})$ from \hat{x} . Let $B = B(c, f(\hat{x}))$ and let $B' = \overline{B}(x, (\delta + \varepsilon)f(\hat{x}))$. The driver $d_P(x)$ of x has to be contained in the convex hull of $B' \setminus B$. Let w be a point in the intersection of ∂B and $\partial B'$. Consider the triangle cxw , see Figure 3.6. The inner angle of this triangle at x is at least $\pi - \alpha$. From the cosine rule we get

$$\begin{aligned} \cos(\pi - \alpha) &\leq \frac{(1 - \delta)^2 f(\hat{x})^2 + (\delta + \varepsilon)^2 f(\hat{x})^2 - f(\hat{x})^2}{2(1 - \delta)(\delta + \varepsilon)f(\hat{x})^2} \\ &= \frac{2\delta(\delta + \varepsilon - 1) + \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)}. \end{aligned}$$

It follows

$$\cos \alpha \geq \frac{2\delta(1 - \delta - \varepsilon) - \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)},$$

which implies the statement of the lemma. ■

Lemma 3.17 *The closure of δ -reduced shape clS_δ is flow-tight under the flow ϕ_P , for any $\varepsilon^2 \leq \delta \leq 10\varepsilon^2$ provided that $\varepsilon \leq 0.14$.*

Proof. Lemmas 3.15 and 3.16 imply that for any point x on the boundary of S_δ the vector $v_P(x)$ points into the interior of S_δ if $\varepsilon^2 \leq \delta \leq 10\varepsilon^2$ and $\varepsilon \leq 0.14$ as can be checked by plugging in the values into the bounds provided by these two lemmas. ■

Topology of the Tubular Neighborhood and Reduced Shapes. Before concluding this section let us make a quick look on the topology of the tubular neighborhood Σ_δ and the reduced shapes S_δ and S_δ^* for $0 \leq \delta < 1$.

Lemma 3.18 *For any $0 \leq \delta < 1$, $\text{cl} S_\delta$ is homotopy equivalent to S . In fact, the former is a strong deformation retract of the latter.*

Proof. Consider the retraction map $r : S \rightarrow \text{cl} S_\delta$ given by

$$r(x) = \begin{cases} \hat{x} + \delta f(\hat{x}) \cdot n_{\hat{x}}^- & x \in S \setminus \text{cl} S_\delta \\ x & x \in \text{cl} S_\delta \end{cases}$$

The map r is continuous on $S \setminus \text{cl} S_\delta$ since $n_{\hat{x}}^-$ changes continuously with \hat{x} (because the surface is smooth), the map $x \mapsto \hat{x}$ is continuous because the only points of discontinuity of this map are medial axis points of which there are none in $S \setminus \text{cl} S_\delta$, and because the local feature size function is 1-Lipschitz and therefore continuous. The continuity of r on all of its domain follows from a gluing argument using the fact that the points on the boundary of S_δ are mapped to themselves both with the identity map and with the mapping $x \mapsto \hat{x} + \delta f(\hat{x}) \cdot n_{\hat{x}}^-$.

If we now define the map $R : [0, 1] \times S \rightarrow \text{cl} S_\delta$ as

$$R(t, x) = \begin{cases} (1-t)x + tr(x) & x \in S \setminus \text{cl} S_\delta \\ x & x \in \text{cl} S_\delta, \end{cases}$$

the map R is a straight-line homotopy from the identity of S to the retraction map r . ■

Lemma 3.19 *The two sets Σ_δ and Σ are homotopy equivalent, for any $0 < \delta < 1$. In fact, the latter is a strong deformation retract of the former.*

Proof. Consider the retraction map $r : \Sigma_\delta \rightarrow \Sigma$ given by $r(x) = \hat{x}$. The map r is clearly continuous on Σ_δ because $\Sigma_\delta \cap M = \emptyset$ and M consists of the only points in space where the map $x \mapsto \hat{x}$ is not continuous (in fact undefined). Since the map r is identity on Σ , r is a retraction.

If we now define the map $R : [0, 1] \times \Sigma_\delta \rightarrow \Sigma_\delta$ as

$$R(t, x) = (1-t)x + t\hat{x}$$

the map R is a straight-line homotopy from the identity on Σ_δ to the retraction r . ■

3.4 Distinguishing Inner and Outer Medial Axis Critical Points

In Section 3.2 we showed how for a dense enough sample P of a surface Σ , the surface and medial axis critical points can be distinguished from each other. In this section we show how medial axis critical points can further be partitioned into inner and outer medial axis critical point in an algorithmic manner. This further subdivision is based on the cell complex structure of the stable flow complex induced by h_P .

Recall that the incidence poset of a cell complex defines a binary relation “ $<$ ” between cells of a cell complex in which $\tau < \sigma$ if τ is a proper face of σ , i.e. $\tau \neq \sigma$ and τ is contained in the boundary of σ .

In the case of the stable flow complex of a point-set P , each cell of the complex corresponds to a critical point of h_P . The relation “ $<$ ” can therefore be adopted for pairs of critical points by letting

$$c_1 < c_2 \quad \text{if and only if} \quad \text{Sm}(c_1) < \text{Sm}(c_2).$$

Lemma 3.20 *Let c_i and c_o respectively be inner and outer medial axis critical points of the distance function h_P induced by an ε -sample P of a smooth surface Σ for $\varepsilon \leq 0.14$. If a critical point c of h_P satisfies $c < c_i$ and $c < c_o$, then c is a surface critical point.*

Proof. Suppose to the contrary that c is a medial axis critical point. Without loss of generality assume that c is an inner medial axis critical point, i.e. $c \in S_{\varepsilon^2}$. However, since $c < c_o$, by Lemma 2.21, $c_o \in \text{Um}(c)$ and therefore there is a flow path from every neighborhood of c to c_o . But this is in contradiction with S_{ε^2} being flow-tight for ϕ_P . ■

Lemma 3.21 *Let c_1 and c_2 be two inner medial axis critical points of the distance function h_P as defined in Lemma 3.20, such that*

$$\text{cl Sm}(c_1) \cap \text{cl Sm}(c_2) \cap S_{\varepsilon^2} \neq \emptyset.$$

In other words, the boundaries of stable manifolds of c_1 and c_2 intersect in S_{ε^2} . Then there is an inner medial axis critical point c satisfying $c < c_1$ and $c < c_2$. The analogous claim with c_1 and c_2 being outer medial axis points for which the boundaries of $\text{Sm}(c_1)$ and $\text{Sm}(c_2)$ intersect in $S_{\varepsilon^2}^$ is also valid.*

Proof. Let z be a point in

$$\text{cl Sm}(c_1) \cap \text{cl Sm}(c_2) \cap S_{\varepsilon^2}.$$

There are flow paths from arbitrarily small neighborhood of z and each of c_1 and c_2 (since z is on the boundary of both stable manifolds of c_1 and c_2 and every point in each of these stable manifolds flows to the corresponding critical point). If z is a critical point, then there is nothing to prove. Otherwise $z \in \text{Sm}(c)$ for some critical point c . By the cell-complex

structure of the stable flow complex, $c < c_1$ and $c < c_2$. Also, there is a flow path between z and c since $z \in \text{Sm}(c)$. Since S_ε^2 is flow-tight, c must be also be in S_ε^2 . ■

To separate inner and outer medial axis critical points from each other, it is enough to look at the incidence poset of the critical points. First, we remove the vertices in the poset corresponding to surface critical points. This separates the inner and outer medial axis critical points by Lemma 3.20. By Lemma 3.21 the inner medial axis critical points form a connected component and so do the outer medial axis critical points. The component corresponding to the outer medial axis critical points can be detected as the one that contains the critical point at infinity.

Handling surfaces with multiple components. It must be mentioned that the all the results of this chapter are valid for surfaces with multiple components. In particular Theorems 3.3 and 3.12 make no specific assumption on the number of components of the surface in question and therefore stand valid for multi-component surfaces. A surface Σ with k components defines k corresponding components of the δ -tubular neighborhood for any $0 \leq \delta < 1$. It also partitions $\mathbb{R}^n \setminus \Sigma$ into $k + 1$ connected components (shapes) and the same number of corresponding reduced shapes. Medial axis critical points are then partitioned based on the shape in which they are contained. The flow-tightness of the closures of the reduced shapes follows from exactly the same argument as that of Lemma 3.17. The *inner* volume enclosed by a k -component surface Σ can be defined as the union of those shapes S (components of $\mathbb{R}^n \setminus \Sigma$) for which a path connecting a point $x \in S$ to a point at infinity (which only intersects Σ transitively) crosses Σ an odd number of times. Lemmas 3.20 and 3.21 continue to hold in this case and allow us to separate inner and outer medial axis critical points.

Bibliography and Remarks.

Theorem 3.3 was proven by Dey, Giesen, Ramos, and Sadri in [32] although the version of the theorem proven there is weaker than the one given here in that (1) the medial axis critical points c are shown to be within $(1 - 2\varepsilon)\mu(c)$ from \check{c} (instead of $1 - 2\varepsilon^2$ of Theorem 3.3) and (2) ε is required to be less than $1/3$ (instead of $1/\sqrt{3}$ in the version presented here). The upper bound on the allowed ε for the algorithmic separation of surface and medial axis critical points (Theorem 3.12) is also improved from $1/10$ in [32] to $1/5$.

In [23] a separation result similar to Theorem 3.3 is presented for critical points of the distance functions that induced by noisy samples.

Chapter 4

Reconstruction as a Union of Stable Manifolds

Separation of critical points suggests a comparison between stable manifolds of surface critical points versus those of the medial axis ones. Consider the stable manifold of a surface maximum. In Section 3.3.2 we saw that the flow lines of ϕ_P do not enter Σ_{ε^2} from either of S_{ε^2} or $S_{\varepsilon^2}^*$. Since surface critical points are contained in Σ_{ε^2} , their stable manifolds have to also be entirely contained in Σ_{ε^2} . Since the union of stable manifolds of all critical points must cover the entire space, the entire volume of S_{ε^2} and $S_{\varepsilon^2}^*$, must be covered by stable manifolds of medial axis critical points. Also, since Σ_{ε^2} is flow-repellent for ϕ_P , the stable manifolds of inner medial axis critical points do not extend beyond $S \cup \Sigma_{\varepsilon^2}$ and likewise those of the exterior medial axis critical points are contained in $S^* \cup \Sigma_{\varepsilon^2}$. Thus intuitively, S is covered closely by the union of medial axis critical points of the interior and S^* by the exterior analogue of that.

In Section 4.1 we describe an algorithm for reconstructing a surface Σ in \mathbb{R}^3 as a subcomplex of the stable flow complex induced by a sample P of Σ . Specifically, the reconstructed surface is the boundary of the union of stable manifolds of inner (or outer) medial axis critical points. We then show that the output produced by the algorithm is contained in a tubular neighborhood of Σ and is homeomorphic to it if the given sample is *tight* enough (Refer to Section 1.6.2 for the definition of tight sampling). The presented proofs make no use of our knowledge about the continuity of the flow map ϕ_P and the fact that certain tubular neighborhoods of the surface are flow-tight. Instead, we build a homeomorphism; in fact we show that under the considered sampling assumption, the projection of the reconstructed set into the surface Σ constitutes a homeomorphism. There are parts of this proof that are inherently three-dimensional and do not generalize to higher dimensions even though the algorithm is well-defined regardless of the dimension of the input.

Section 4.2 presents a completely different analysis of the same algorithm that takes advantage of the continuity of the flow map ϕ_P and some other tools developed in earlier chapters. The proof is valid for a hypersurfaces in Euclidean spaces of arbitrary dimension. Moreover the assumption on tightness of sampling can be dropped. As drawback, we do not establish a homeomorphism between the target and reconstructed surfaces and instead settles for a homotopy equivalence between the inner (outer) shape associated to Σ with the the union of stables manifolds of inner (outer) medial axis critical points of h_P the boundary to which is taken as the reconstruction in the algorithm of Section 4.1.

Algorithm RECONSTRUCT(sample point set P)

```

1   $C \leftarrow$  set of the critical points of  $h_P$ .
2   $(C_M^+, C_M^-, C_\Sigma) \leftarrow$  SEPARATE( $P, C$ ).
3   $\tilde{S} \leftarrow \bigcup_{c \in C_M^-} \text{Sm}(c)$ .
4   $\tilde{\Sigma} \leftarrow \partial \tilde{S}$ .
5  return  $\tilde{\Sigma}$ 

```

Figure 4.1: Algorithm RECONSTRUCT, a surface-reconstruction algorithm.

4.1 A Simple 3D Reconstruction Algorithm

In surface reconstruction, we are given an ε -sample P of a smooth surface Σ and we wish to construct for either the surface Σ or the bounded volume S enclosed by Σ , a corresponding piece-wise linear surface or volume that is topologically equivalent with the original surface or volume and closely approximates it geometrically.

The above observations suggest a very natural surface/volume reconstruction algorithm: approximate the volume S by the union of stable manifolds of the inner medial axis critical points and Σ by the boundary of this volume. Figure 4.1 gives a formal description of this algorithm. By SEPARATE(P, C) we refer to a “combined” version of the algorithms described in Chapter 3 that partitions the input set of critical points C into three sets C_M^+ of outer medial axis critical points, C_M^- of inner medial axis critical points, and C_Σ of surface critical points, and returns the tripple (C_M^+, C_M^-, C_Σ) . The algorithm shown in Figure 4.2 returns the boundary of the union of stable manifolds of inner medial axis critical points as output. All of the following results would remain valid if the inner medial axis critical points were replaced with their outer analogues. Figure 4.2 demonstrates the output of RECONSTRUCT for a two-dimensional example.

In this rest of this section, we give geometric and topological guarantees for the output of the algorithm RECONSTRUCT under (ε, δ) -sampling for surfaces in \mathbb{R}^3 . The summary of the results is given in the following Theorem.

Theorem 4.1 *For any $0 < \rho < 1$ there exists ε_0 such that given an (ε, δ) -sample P from a smooth closed surface Σ , where $\varepsilon < \varepsilon_0$ and $\delta/\varepsilon \geq \rho$, the algorithm RECONSTRUCT outputs a sub-complex $\tilde{\Sigma}$ of the flow complex of P with the following properties:*

- (1) $\tilde{\Sigma}$ is contained in the tubular neighborhood $\Sigma_{3\varepsilon^2}$.
- (2) The normal to every triangle pqr in $\tilde{\Sigma}$, with $p \in P$, forms an angle of $O(\varepsilon)$ with the normal to Σ at p , n_p .




Figure 4.2: The output of the algorithm RECONSTRUCT: stable manifolds of medial axis maxima are the two-dimensional gray cells with blue boundaries. The common boundary of any two pair of these two-dimensional cells is itself the stable manifold of an index-1 medial axis critical point. The output of the algorithm RECONSTRUCT is the boundary of the grayed set.

(3) $\tilde{\Sigma}$ is homeomorphic to Σ .

In particular these claims hold for $\rho \geq 1/3$ and $\varepsilon_0 \leq 0.01$.

4.1.1 Geometric Closeness

To analyze local geometry of the flow near the surface, we place at sample points $p \in P$ cones that open along inner and outer normal directions at p . We show that, under certain conditions, such cones are *sinks* for the flow ϕ_P , i.e., on their surfaces and sufficiently close to Σ , the flow is either tangential or points to the inside of the cones. By overlapping together these close-reaching cones (see Figure 4.3) we obtain inner and outer envelopes that enclose the surface and are in a sense “one-way” for the flow. This means that flow cannot escape from these envelopes leading to properties of the flow complex central to the analysis of the output of our algorithm.

Sink Cones. For a point p , and a direction vector n , let $\nabla = \text{cone}(p, n, \theta, r)$ be the *cone-patch* consisting of points x for which $\|x - p\| \leq r$ and $\angle(n, x - p) = \theta$. We call θ , and r , the *angle* and the *reach* of C , respectively. A cone-patch is essentially part of the surface of an infinite cone. With a slight abuse of notation we refer to a cone-patch also as a cone. The *boundary* of ∇ consists of points $x \in \nabla$ for which $\|x - p\| = r$. We say that ∇ is a *sink* if at




Figure 4.3: A sink cone (left) and enclosing of the surface with sink cone envelopes (right).

every point in the relative interior of ∇ (thus excluding the boundary of ∇) the flow is either tangential or directed toward the interior of the convex hull of ∇ .

Lemma 4.2 *For any point $p \in P$, the two cones*

$$\nabla_p = \text{cone}(p, n_p^+, \theta, r) \quad \text{and} \quad \Delta_p = \text{cone}(p, n_p^-, \theta, r)$$

are sinks for any $0 < \theta < \frac{\pi}{2}$ and $r = f(p) \cos \theta$. Furthermore, the interior of the convex hulls of ∇_p and Δ_p do not contain any critical points.

Proof. We only prove that ∇_p is a sink. The proof for Δ_p is similar. Let c be the point on the ray in direction n_p^+ at distance $f(p)$ from p . Let B' be the ball with radius $f(p) = \|p - c\|$ centered at c . Note that B' cannot contain any points from Σ in its interior since it is contained in a medial ball. Let B be the ball centered at c with radius $f(p) \sin \theta$. Then ∇_p is the cone tangent to B with apex p (see Figure 4.4). For every $x \in \nabla_p$, the set $A_P(x)$ of closest sample points to x are inside the ball B'' of radius $\|x - p\|$ centered at x but outside B' . Let H be the plane tangent to ∇_p at x . Since $r = f(p) \cos \theta$ we have that $\text{conv}(B'' \setminus B')$ is entirely on the opposite side of H with respect to B and therefore, $v_P(x)$ points toward the interior of ∇_p . Thus ∇_p is a sink and no point in the relative interior of ∇_p is a critical point. Every point in the interior of the convex hull of ∇_p is on a cone $\text{cone}(p, n_p^+, \theta', r)$ for some $\theta' < \theta$ or in the relative interior of the line segment \overline{pc} . As we have seen the points on the cones cannot be critical. But neither can be any point y in the relative interior of the line segment pc since for such a point $A_P(y) = \{p\}$ and $y \neq p$. ■

Fixing cone angles. Notice that the above lemma puts at every sample point two sink cones with the same apex, angle, reach, and axis but in opposite direction. For the rest of the section, we shall consider only such cones with a fixed cone angle that depends only on the density of sampling ε . Indeed, we fix

$$\theta = \theta(\varepsilon) = \frac{\pi}{2} - 1.1\varepsilon$$

and we respectively denote the outer and inner cones at a vertex p by

$$\Delta_p = \text{cone}(p, n_p^-, \theta, f(p) \cos \theta) \quad \text{and} \quad \nabla_p = \text{cone}(p, n_p^+, \theta, f(p) \cos \theta).$$




Figure 4.4: Sink cone for a sample point p .

Lemma 4.3 For any point $x \in V_p \cap \Sigma$, the ray shot in direction n_x^+ (n_x^-) hits ∇_p (Δ_p) provided that $\varepsilon \leq 0.05$.

Proof. Let β be the angle between n_x^+ and n_p^+ . Since $x \in V_p \cap \Sigma$, we have $\|x - p\| \leq \varepsilon f(x)$ and therefore $\beta \leq \frac{\varepsilon}{1-3\varepsilon} < \theta$. Therefore, the ray shot from x in direction of n_x^+ hits cone(p, n_p^+, θ, ∞), the infinite extension of ∇_p , at some point x' . Let $\eta = \|x' - p\|/f(p)$. If $\eta \leq \cos \theta$, then $x' \in \nabla_p$ and there is nothing to prove. Otherwise $\|x' - p\| > f(p) \cos \theta$. It can be observed that the closest point to p on the line through x' and x is at distance no less than

$$\|x' - p\| \sin(\theta - \beta) > \cos \theta \sin(\theta - \beta) f(p)$$

from p . On the other hand, since $x \in V_p \cap \Sigma$, $\|x - p\| \leq \frac{\varepsilon}{1-\varepsilon} f(p)$. This implies

$$\begin{aligned} \frac{\varepsilon}{1-\varepsilon} f(p) &\geq \|x - p\| \\ &> \cos \theta \sin(\theta - \beta) f(p) \\ &= \sin(1.1\varepsilon) \cos(1.1\varepsilon + \beta) f(p) \\ &\geq \sin(1.1\varepsilon) \cos\left(1.1\varepsilon + \frac{\varepsilon}{1-3\varepsilon}\right) f(p) \end{aligned}$$

which is a contradiction for $\varepsilon \leq 0.05$. ■

Lemma 4.4 Let x be a point on ∇_p or Δ_p . Then $\|x - \hat{x}\| \leq 3\varepsilon^2 f(\hat{x})$ when $\varepsilon \leq 0.05$.

Proof. Without loss of generality assume that $x \in \nabla_p$ (the proof for the case where $x \in \Delta_p$ is similar). Shoot a ray from x , parallel to n_p^- until it hits Δ_p at a point x' . Since each of ∇_p and Δ_p is completely contained in its corresponding medial ball tangent to Σ at p , the line segment xx' intersects Σ . Therefore

$$\|x - \hat{x}\| \leq \|x - x'\| = 2\|x - p\| \cos \theta \leq 2f(p) \cos^2 \theta. \quad (4.1)$$

On the other hand, by the triangle inequality,

$$\begin{aligned}\|\hat{x} - p\| &\leq \|x - \hat{x}\| + \|x - p\| \\ &\leq 2f(p) \cos^2 \theta + f(p) \cos \theta \\ &\leq 3f(p) \cos \theta.\end{aligned}$$

Since $f(\cdot)$ is 1-Lipschitz

$$\begin{aligned}f(\hat{x}) &\geq f(p) - \|\hat{x} - p\| \\ &\geq f(p) - 3f(p) \cos \theta,\end{aligned}$$

which gives us

$$f(p) \leq \frac{f(\hat{x})}{1 - 3 \cos \theta}.$$

Plugging this into (4.1) we get

$$\|x - \hat{x}\| \leq \frac{2 \cos^2 \theta}{1 - 3 \cos \theta} \cdot f(\hat{x}) \leq 3\varepsilon^2$$

for $\varepsilon \leq 0.05$ and $\theta = \pi/2 - 1.1\varepsilon$. ■

Cone envelopes. Let

$$\nabla = \bigcup_{p \in P} \nabla_p \quad \text{and} \quad \Delta = \bigcup_{p \in P} \Delta_p .$$

Define Σ_∇ as the set of points $x \in \nabla$ for which the segment $\overline{x\hat{x}}$ has no point of ∇ in its relative interior. We call Σ_∇ the *outer cone envelope*. The *inner cone envelope* Σ_Δ is defined analogously with respect to Δ .

Lemma 4.5 *The surface Σ is homeomorphic to both Σ_Δ and Σ_∇ . Both cone envelopes divide \mathbb{R}^3 into a bounded and an unbounded component. The bounded component of the inner cone envelope and the unbounded component of the outer cone envelope are flow-tight under the flow ϕ_P .*

Proof. We only present the proof for the outer cone envelope. The proof for the inner cone envelope follows the same lines.

Let $\pi : \mathbb{R}^3 \setminus M \rightarrow \Sigma$ project into Σ , i.e. π maps every point in $\mathbb{R}^3 \setminus M$ into its closest point \hat{x} in Σ . The map π is continuous. Let π^+ be the restriction of π to Σ_∇ . Since $\Sigma_\nabla \cap M = \emptyset$ (because $\Sigma_\nabla \subset \Sigma_{3\varepsilon^2}$ by Lemma 4.4), π^+ is defined and continuous on Σ_∇ . By definition of Σ_∇ , π^+ is injective and by Lemma 4.3, π^+ is surjective. This along with the compactness of Σ_∇ (inferred from the compactness of Σ and the continuous one to one mapping given by π^+) implies that π^+ is a homeomorphism.

Since we assumed that Σ is a manifold without boundary, so is Σ_∇ . Thus Σ_∇ divides \mathbb{R}^3 into a bounded and an unbounded component. By Lemma 4.2 the bounded component has to be flow-tight for ϕ_P . ■

Cone neighborhood. The closed volume sandwiched between Σ_Δ and Σ_∇ is called the *cone neighborhood* of Σ and is denoted by $\Sigma_{\bar{\chi}}$.

Theorem 4.6 *The output of the algorithm RECONSTRUCT lies in $\Sigma_{\bar{\chi}}$ and the latter itself is contained in $\Sigma_{3\varepsilon^2}$.*

Proof. By Lemmas 4.2 and 4.5, the stable manifold $\text{Sm}(c)$ of any surface critical point c has to be contained in $\Sigma_{\bar{\chi}}$. Thus the output of RECONSTRUCT completely lies in $\Sigma_{\bar{\chi}}$. By Lemma 4.4, Σ_∇ and Σ_Δ are contained in $\Sigma_{3\varepsilon^2}$. This implies that $\Sigma_{\bar{\chi}}$ is also contained in $\Sigma_{3\varepsilon^2}$. ■

The following Corollary will be needed later.

Corollary 4.7 *Let $\varepsilon < 1/2$. For every $x \in \Sigma_\xi$ with $0 \leq \xi < (1 - \varepsilon)/2$, i.e., for every x satisfying $\|x - \hat{x}\| \leq \xi f(\hat{x})$, and for every $p \in A_P(x)$,*

$$\|x - p\| \leq \frac{\varepsilon + \xi}{1 - \varepsilon - 2\xi} f(p).$$

In particular, when $\varepsilon \leq 0.05$, for every point $x \in \Sigma_{\bar{\chi}}$, and every $p \in A_P(x)$, $\|x - p\| \leq 1.23\varepsilon f(p)$, and for every surface critical point c , and every $p \in A_P(c)$, $\|c - p\| \leq 1.12\varepsilon f(p)$.

Proof. Let x be a point in Σ_ξ . From the definition of ε -sampling, $\|\hat{x} - q\| \leq \varepsilon f(\hat{x})$ for $q \in A_P(\hat{x})$. For any sample point $p \in A_P(x)$, by the triangle inequality

$$\|x - p\| \leq \|x - q\| \leq \|x - \hat{x}\| + \|\hat{x} - q\| \leq (\varepsilon + \xi)f(\hat{x}). \quad (4.2)$$

Thus

$$\|p - \hat{x}\| \leq \|x - p\| + \|x - \hat{x}\| \leq (\varepsilon + \xi)f(\hat{x}) + \xi f(\hat{x}) \leq (\varepsilon + 2\xi)f(\hat{x}),$$

and from this and because the local feature size is 1-Lipschitz,

$$f(p) \geq f(\hat{x}) - \|\hat{x} - p\| \geq f(\hat{x}) - (\varepsilon + 2\xi)f(\hat{x}). \quad (4.3)$$

Combining (4.2) and (4.3) we get

$$\|x - p\| \leq \frac{\varepsilon + \xi}{1 - \varepsilon - 2\xi} f(p).$$

Using $\varepsilon \leq 0.05$ along with Theorem 4.6 implies the bounds for the case where $x \in \Sigma_{\bar{\chi}}$. In the case of x being a surface critical point we invoke Corollary 3.4, instead. ■

4.1.2 Convergence of Normals

The output $\tilde{\Sigma}$ produced by the algorithm RECONSTRUCT consists of stable manifolds of index-2 saddle points that lie in a small tubular neighborhood of the surface. We refer to these stable manifolds as *surface patches*. We want to show that under (ε, δ) -sampling, with a fixed $\rho = \delta/\varepsilon$, the normal of triangles in these surface patches is within $O(\varepsilon)$ from the normal to surface at a nearby point, of course for sufficiently small ε . We use the following two lemmas.

Lemma 4.8 [2] *For any two points $p, q \in \Sigma$, the angle between segment pq and either of n_p^+ and n_p^- is greater than*

$$\frac{\pi}{2} - \arcsin \frac{\|p - q\|}{2f(p)}.$$

Lemma 4.9 [2] *For points $p, q, r \in \Sigma$, let p be a vertex of the triangle pqr with the largest angle and let r be its circumradius. If $r = \lambda f(p)$, then the acute angle between the normal to pqr and the normal to surface at p is at most $\beta(\lambda)$ where*

$$\beta(\lambda) = \arcsin(\lambda) + \arcsin \left(\frac{2}{\sqrt{3}} \sin(2 \arcsin \lambda) \right) \leq 4\lambda,$$

for $\lambda \leq \frac{1}{4}$.

The stable manifold $\text{Sm}(c)$ of every 2-saddle is a piece-wise linear surface made of a finite number of triangles, which we call *patch triangles*. Each patch triangle t has exactly one vertex in P . Note that for every point x in a patch triangle t , the vertex of t that belongs to P is a closest sample point to x (see Section 2.4.1). If x is on the boundary of t , it can have more than one closest sample point as it belongs to more than one patch triangle. The following lemma shows that under tight sampling, each patch triangle must have a normal close to surface normal at its vertex in P .

Lemma 4.10 *For any $0 < \rho < 1$, there exists ε_0 such that if P is an (ε, δ) -sample of Σ with $\varepsilon \leq \varepsilon_0$ and $\delta = \rho\varepsilon$, then for any $x \in \text{Sm}(c)$, the stable manifold of a surface 2-saddle c , the acute angle between n_p , where p is a closest sample point to x , and n_t , the normal direction of the patch triangle $t \subset \text{Sm}(c)$ that contains x and has p as a vertex, is at most*

$$\arcsin \left(\frac{\sin(1.23\varepsilon)}{2 \sin \left(\frac{1}{2} \arcsin(\rho/2.46) \right)} \right) = O(\varepsilon/\rho).$$

Proof. Let P be an ε -sample of Σ for $\varepsilon \leq \varepsilon_0$. By Corollary 4.7, $\|x - p\| \leq 1.23\varepsilon f(p)$ for every $p \in A_P(x)$. Every point x on $\text{Sm}(c)$ is on a patch triangle $t = pou$ of $\text{Sm}(c)$ with the following structure (see Figure 4.5): t has exactly one vertex p in P . The edge uo of t opposite to p is on the Voronoi facet dual to a Delaunay edge pq and ends on the dual Voronoi edge e of a Delaunay triangle pqr in which $\|r - p\| > \|q - p\|$. The mid-point d of pq is the driver of the




Figure 4.5: A generic patch triangle on the stable manifold of a surface 2-saddle.

points on uo . Furthermore, the line containing e does not intersect the triangle pqr except when o is the critical point c in which case the patch triangle t is coplanar with (and in fact contained in) the Delaunay triangle t_c containing c . We postpone the study of this special case for later. Let s be the circumcenter of pqr and let p' be a point on the circumcircle of pqr opposite to p with respect to s . Then $\angle p'qp = \pi/2$ and that $\|d - s\| = \frac{1}{2}\|q - p'\|$. Furthermore, $\|q - p'\| > \|q - r\|$. Therefore we get $\|d - s\| \geq \frac{1}{2}\delta f(r)$. On the other hand, since $r \in A_P(o)$, by Corollary 4.7, $\|s - p\| = \|s - r\| < \|r - o\| \leq 1.23\varepsilon f(r)$. Combining these we get for the angle $\alpha = \angle qps$:

$$\sin \alpha \geq \frac{\|d - s\|}{\|p - s\|} \geq \frac{\delta}{2.46\varepsilon} = \frac{\rho}{2.46}.$$

On the other hand, $\pi/2 > \angle qpo > \angle qps \geq \alpha$. Also, $o \in \text{Sm}(c)$ and $\text{Sm}(c)$ is contained in $\Sigma_{\bar{\gamma}}$ and therefore, po makes an angle of at least $\theta = \frac{\pi}{2} - 1.1\varepsilon$ with n_p . Moreover, $\|p - q\| \leq 2.46\varepsilon f(p)$ and therefore by Lemma 4.8, $q - p$ makes an angle of at least $\frac{\pi}{2} - 1.23\varepsilon < \theta$ with n_p . Thus, the three points p , q , and o , make a triangle $t' = pqo$ with an angle of at least α at vertex p satisfying $\arcsin(\rho/2.46) < \alpha < \frac{\pi}{2}$, and with both of the edges incident to p making an angle of at least $\pi/2 - 1.23\varepsilon$ with n_p . It can be shown through elementary calculations that under these conditions, $n_{t'}$, the normal to t' , and n_p , make an angle of at most

$$\arcsin\left(\frac{\cos(\pi/2 - 1.23\varepsilon)}{2 \sin(\alpha/2)}\right),$$

which matches the bound in the statement of the lemma. We now consider the special case when o coincides with c . This happens when the Voronoi edge e intersects its dual Delaunay triangle pqr at c and the patch triangle t in question becomes coplanar with the Delaunay triangle $t_c = pqr$. Notice that $A_P(c) = \{p, q, r\}$ and by Corollary 4.7, $\|p - c\| \leq 1.12\varepsilon f(p)$. Similar inequalities hold for q and r . As in the previous case, let p' be the point on the circumcircle of pqr opposite to p with respect to circumcenter c . We denote the angles $\angle qpp'$ and $\angle rpp'$ by β and γ respectively, and their sum by α . Since the angles $\angle pqp'$ and $\angle prp'$ are each 90 degrees, we have $\sin \beta = \frac{1}{2}\|q - p'\|/\|p - c\|$ and $\sin \gamma = \frac{1}{2}\|r - p'\|/\|p - c\|$. In order for c to be a critical point, all the angles of the triangle t_c must be acute. Since the

sine function is concave for acute angles we have

$$\begin{aligned}\sin \frac{\alpha}{2} &\geq \frac{\sin \beta + \sin \gamma}{2} = \frac{\|r - p'\| + \|q - p'\|}{4 \cdot \|p - c\|} \\ &\geq \frac{\|r - q\|}{4 \cdot \|p - c\|} \geq \frac{\delta f(r)}{4 \cdot 1.12\varepsilon f(r)} = \frac{\rho}{4.48}.\end{aligned}$$

On the other hand $\|p - q\| \leq \|p - c\| + \|q - c\| \leq 2.24\varepsilon f(r)$ by Lemma 4.8. Thus pq makes an angle of at most $\pi/2 - 1.12\varepsilon$ with n_p . A similar argument establishes the same bound for the angle between pr and n_p . Similar to the previous case, We have shown for the triangle pqr that the angle α at p is at least $2 \arcsin(\rho/4.48)$ and that the edges pq and pr make an angle of at least $\pi/2 - 1.12\varepsilon$ with n_p . This implies that the angle between the normal to the plane of this triangle and the normal to Σ at p is at most

$$\arcsin \left(\frac{\cos(\pi/2 - 1.12\varepsilon)}{2\rho/4.48} \right).$$

It can be verified that the this bound results a smaller angle than the one obtained above (matching the bound in the statement of the lemma) for the general case, whenever the two bounds are defined for any $0 < \varepsilon \leq 0.05$ and $0 < \rho < 1$. \blacksquare

Corollary 4.11 *For $\varepsilon \leq 0.05$ and $\rho \geq 1/2$, we have for every point x on the stable manifold of a surface 2-saddle c that the acute angle between normal n_t to any patch triangle t of $Sm(c)$ that contains x , and the normal $n_{\hat{x}}$ to Σ at \hat{x} is at most 23 degrees.*

Proof. Plugging $\varepsilon \leq 0.05$ and $\rho \geq 1/2$ in Lemma 4.10, gives an upper bound of 18 degrees for the angle between n_t and n_p , where n_p is the normal to Σ at a closest sample point p to x . Since by Theorem 4.6, $Sm(c)$ is contained in $\Sigma_{3\varepsilon^2}$, we have $\|x - \hat{x}\| \leq 3\varepsilon^2 f(\hat{x})$. Let q be a closest sample point to \hat{x} . By sampling condition, $\|\hat{x} - q\| \leq \varepsilon f(\hat{x})$ and therefore

$$\|\hat{x} - p\| \leq \|\hat{x} - x\| + \|x - p\|.$$

On the other hand $\|x - p\| \leq \|x - \hat{x}\| + \|\hat{x} - q\|$. Therefore we get

$$\|\hat{x} - p\| \leq 2\|x - \hat{x}\| + \|\hat{x} - q\| \leq (6\varepsilon^2 + \varepsilon)f(\hat{x}) \leq 1.3\varepsilon f(\hat{x}),$$

for $\varepsilon \leq 0.05$. Therefore by Lemma 1.12 the angle between n_p^+ and $n_{\hat{x}}^+$ is at most $1.3\varepsilon/(1 - 3 \cdot 1.3\varepsilon) \leq 5^\circ$. \blacksquare

The following proposition is directly based on the structure of the stable manifolds of 2-saddles (Section 2.4.1) and the fact that flow paths never turn more than 90 degrees.

Proposition 4.12 *If t_1 and t_2 are patch triangles of $Sm(c)$ for a surface 2-saddle c such that t_1 and t_2 have one edge in common, then the dihedral angle between t_1 and t_2 is no less than $\pi/2$.*

Lemma 4.13 *Let c be a surface 2-saddle. Suppose we orient the patch triangles in $\text{Sm}(c)$ arbitrarily but consistently so that for any patch triangle t , n_t^+ and n_t^- are respectively the outer and inner normal directions on t with respect to the applied orientation. Then, under the assumptions of Corollary 4.11 exactly one of the following cases holds.*

1. $\angle(n_t^+, n_x^+) \leq 23^\circ$, for every patch triangle t of $\text{Sm}(c)$ and for every $x \in t$.
2. $\angle(n_t^+, n_x^-) \leq 23^\circ$, for every patch triangle t of $\text{Sm}(c)$ and for every $x \in t$.

Proof. First notice that as was shown in the proof of Corollary 4.11, for any point $x \in t$, where t is a patch triangle of $\text{Sm}(c)$, $\angle(n_x^+, n_p^+) \leq 5^\circ$, where p is the vertex of t that is a sample point. Thus, if for the arbitrary orientation of t and for a point $x \in t$, $\angle(n_t^+, n_x^+) = \alpha$, the same holds for every other point y in t , modulo changing α by 5 degrees.

Let t_c be the Delaunay triangle that contains c . All patch triangles $t \subset t_c$ of $\text{Sm}(c)$, have the same n_t^+ which agrees with one of the two orientations of the direction normal to t_c . By Corollary 4.11, the normal direction of t_c makes an angle of at most 23° with either n_c^+ or n_c^- . Assume without loss of generality that the first case holds, i.e. $\angle(n_t^+, n_c^+) \leq 23^\circ$. We show now that this will imply that for every patch triangle t of $\text{Sm}(c)$ and every $x \in t$, $\angle(n_t^+, n_x^+) \leq 23^\circ$. We prove this by extending the result for the triangles we already have this property for to their neighboring patch triangles. Thus, assume t and t' are two patch triangles with an edge e in common. Let z be a point on e . Since t and t' are oriented consistently, the dihedral angle between t and t' is $\pi - \angle(n_t^+, n_{t'}^+)$. By Proposition 4.12 this angle is at least $\pi/2$ and therefore $\angle(n_t^+, n_{t'}^+) \leq \pi/2$. Therefore using triangle inequality for angles we get $\angle(n_z^+, n_{t'}^+) \leq \angle(n_t^+, n_z^+) + \angle(n_t^+, n_{t'}^+) \leq 90 + 23 = 113^\circ$. But by Corollary 4.11, $\angle(n_z^+, n_{t'}^+)$ is either less than 23° or more than $180^\circ - 23^\circ$ and we have just shown that the latter case does not hold. ■

4.1.3 Orientation of surface patches

The output the algorithm RECONSTRUCT is a collection of stable manifolds of surface 2-saddles (patches) that are attached to each other at Gabriel edges (stable manifolds of surface 1-saddles). In order to establish that this reconstruction has the same topology as the original surface Σ , we shall provide a homeomorphism between the two surfaces. We have shown in the previous section that, roughly speaking, each patch is almost flat and lies almost orthogonal to the normal to Σ at the 2-saddle in the patch. In order to achieve the desired homeomorphism we need to show that neighboring patches do not fold over each other. Our way of showing this can be summarized as follows: we observe that the normal to a patch induced by a 2-saddle c at c itself is close to the normal to Σ at a near surface point to c . This gives a natural orientation of the patch. We prove afterwards that the side of the patch that faces the union of stable manifolds \tilde{S} of the critical points in the component computed by the algorithm RECONSTRUCT is consistently determined by the given orientation.

We will need the following two technical lemmas.

Lemma 4.14 *Let ∇ and Δ be two infinite cones with cone angle θ , with the same apex p and same axis, extended in opposite directions. Let x be a point not in the interior of either of the convex hulls of ∇ or Δ . Consider a line ℓ passing through x , making an angle of $\alpha < \theta$ with the common axis of ∇ and Δ , and hitting ∇ and Δ in points x_1 and x_2 , respectively. Then*

$$\|x_1 - x_2\| \leq 2 \cdot \|x - p\| \cdot \frac{\cos \theta}{\sin(\theta - \alpha)}.$$

Proof. Without loss of generality assume that p is the origin and that the common axis of ∇ and Δ is the z -axis. By the assumptions of the lemma, x_1 and x_2 are in opposite sides of x on ℓ . Consider the vertical plane H containing x and the z axis. When x is fixed, if we consider an arbitrary line ℓ through x making an angle of α with the z -axis, it is easy to observe that $\|x - x_1\|$ is maximized when ℓ is contained in H , in which case by the law of sines $\|x - x_1\| = \|x - y_1\| \cdot \sin \theta / \sin(\theta - \alpha)$, where y_1 is the vertical projection of x to ∇ . Thus in general

$$\|x - x_1\| \leq \|x - y_1\| \cdot \sin \theta / \sin(\theta - \alpha). \quad (4.4)$$

Similarly we get for the distance between x and x_2

$$\|x - x_2\| \leq \|x - y_2\| \cdot \sin \theta / \sin(\theta - \alpha), \quad (4.5)$$

where y_2 is the vertical projection of x to Δ . On the hand, when $\|x - p\|$ is fixed, $\|y_1 - y_2\|$ is maximized when x is in the plane $z = 0$, in which case $\|y_1 - y_2\| = 2 \cdot \|x - p\| \cdot \cot \theta$. So, in general

$$\|y_1 - y_2\| \leq 2 \cdot \|x - p\| \cdot \cot \theta. \quad (4.6)$$

Combining (4.4), (4.5), and (4.6) we get

$$\begin{aligned} \|x_1 - x_2\| &= \|x - x_1\| + \|x - x_2\| \\ &\leq (\|x - y_1\| + \|x - y_2\|) \cdot \sin \theta / \sin(\theta - \alpha) \\ &= \|y_1 - y_2\| \cdot \sin \theta / \sin(\theta - \alpha) \\ &\leq 2 \cdot \|x - p\| \cdot \cos \theta / \sin(\theta - \alpha). \quad \blacksquare \end{aligned}$$

Lemma 4.15 *Assume $\varepsilon \leq 0.01$. Let x be a point in $\Sigma_{\bar{\chi}}$ with $\{p, q, r\} \subset A_P(x)$. Then the acute angle between normal to the Delaunay triangle pqr and each of the normals n_p , n_q , and n_r is at most 8ε .*

Proof. Assume without loss of generality that p is the vertex of pqr with the largest face angle. Since $x \in \Sigma_{\bar{\chi}} \subset \Sigma_{3\varepsilon^2}$, by Corollary 4.7,

$$\|x - p\| = \|x - q\| = \|x - r\| \leq 1.23\varepsilon f(p).$$

On the other hand, $\|x - p\|$ is an upper bound for the circumradius of pqr . Thus using Lemma 4.9, the acute angle between n_p and normal to pqr is at most $\beta(1.23\varepsilon) \leq 5\varepsilon$. On the

other hand, $\|p - q\| \leq \|p - x\| + \|x - q\| \leq 2 \cdot 1.23f(p)$. Therefore by Lemma 1.12,

$$\angle(n_p^+, n_q^+) \leq \frac{2.46\varepsilon}{1 - 3 \cdot 2.46\varepsilon} \leq 3\varepsilon$$

for $\varepsilon \leq 0.01$. The same argument can be repeated with r instead q . ■

We say that a triangle $t = pqr$ with $p, q, r \in P$ *lies flat to surface* or simply *is flat* if the normal of t is within 8ε from one of n_p, n_q , or n_r . For such a triangle, it is meaningful to distinguish between the side that faces the interior of Σ and the one that faces its exterior. We refer to these sides as *inner* and *outer* sides, respectively.

Let c be a surface 2-saddle. By definition, c is the intersection point of a Delaunay triangle t_c and its dual Voronoi edge e_c . Thus $|A_P(c)| = 3$ and by Lemma 4.15, the normal to t_c makes an angle of 8ε or less with the surface normal at any of the vertices of t_c . Thus, t_c lies flat to surface. Since t_c intersects e_c in a point of its relative interior (by our non-degeneracy assumption), we can distinguish between the two endpoints of e_c as its inner and outer vertices and refer to them as v_c^- and v_c^+ , respectively. We denote the the segment cv_c^+ excluding c by e_c^+ , and define e_c^- similarly. Notice that c is the driver for points on e_c and therefore the flow direction on $e_c \setminus \{c\}$ is toward its endpoints at each side of c . Therefore, every point of e_c between c and v_c^+ flows to the same maximum that v_c^+ flows into. A similar statement holds for the points between c and v_c^- . We define $U_c^+ = e_c^+ \cup \phi_P(v_c^+)$ and $U_c^- = e_c^- \cup \phi_P(v_c^-)$. In fact, U_c^+ and U_c^- together constitute the *unstable manifold* of c . Thus, if U_c^+ intersects Σ_∇ then the flow originated at any point of e_c^+ , arbitrarily close to c must end up in an exterior medial axis maximum m implying that $\text{Sm}(c)$ is incident to $\text{Sm}(m)$ through the outer side of t_c . Similar statements can be made by replacing U_c^+ with U_c^- and Σ_∇ with Σ_Δ .

Lemma 4.16 *For any $0 < \rho < 1$, there exists ε_0 small enough such that if P is an (ε, δ) -sampling of Σ for $\varepsilon \leq \varepsilon_0$ and $\delta = \rho\varepsilon$, then for any $x \in U_c^+ \cap \Sigma_\nabla$, $\angle(v_P(x), n_p^+) \leq 8\varepsilon$, and for every point $x \in U_c^- \cap \Sigma_\nabla$, $\angle(v_P(x), n_p^-) \leq 8\varepsilon$, where p is any point in $A_P(x)$. In particular, for $\rho = \delta/\varepsilon \geq 1/3$, $\varepsilon_0 \leq 0.01$ suffices.*

Proof. We only prove the lemma for points in $U_c^+ \cap \Sigma_\nabla$. The proof for points in $U_c^- \cap \Sigma_\nabla$ is analogous. For simplicity we enforce $\varepsilon_0 \leq 0.01$ although the statement of the lemma may hold for larger values of ε_0 . Let P be an (ε, δ) -sampling of Σ with $\varepsilon \leq \varepsilon_0 \leq 0.01$ and $\delta = \rho\varepsilon$.

Since $x \in \Sigma_\nabla \subset \Sigma_{3\varepsilon^2}$, by Corollary 4.7, $\|x - p\| \leq 1.23\varepsilon f(p)$ for every $p \in A_P(x)$. Notice that it suffices to prove that $\angle(v_P(x), n_p^+) \leq 5\varepsilon$ for only one point $p \in A_P(x)$. This is because for any other point $q \in A_P(x)$, $\|p - q\| \leq \|p - x\| + \|q - x\| \leq 2 \cdot 1.23\varepsilon \max\{f(p), f(q)\}$ and therefore by Lemma 1.12,

$$\angle(n_p^+, n_q^+) \leq \frac{2.46\varepsilon}{1 - 3 \cdot 2.46\varepsilon} \leq 3\varepsilon$$

for $\varepsilon \leq 0.01$.




Figure 4.6: Illustration of the proof of Lemma 4.16, cases 1 (left) and 2 (right).

As above, let t_c and e_c be the Delaunay triangle and its dual Voronoi edge for which $\{c\} = t_c \cap e_c$. Notice that U_c^+ is a piece-wise linear curve. Let u_0, u_1, \dots, u_k be the vertices of this curve with $u_0 = c$, $u_1 = v_c^+$, and $u_k = m$, where m , the maximum at which U_c^+ ends. Notice of course that u_0 itself does not belong to U_c^+ as $v_P(u_0) = v_P(c) = 0$. We prove the lemma inductively starting from the segment u_0u_1 and going up to $u_{i-1}u_i$ for the smallest i for which $u_{i-1} \in \Sigma_{\bar{\chi}}$ but $u_{i-1}u_i$ has a point outside $\Sigma_{\bar{\chi}}$. For this last line segment $u_{i-1}u_i$, the argument we provide will be true only for the initial part $u_{i-1}u^*$, where u^* is the first point of U_c^+ (starting from u_{i-1}) not in $\Sigma_{\bar{\chi}}$. Since no flow enters $\Sigma_{\bar{\chi}}$, no point of U_c^+ past u^* will be in $\Sigma_{\bar{\chi}}$. From the structure of flow complex it is easy to see that every vertex in $\{u_1, \dots, u_k\}$ is either a Voronoi vertex or lies on a Voronoi edge. Furthermore, the relative interior of every segments $u_{i-1}u_i$, $i = 1, \dots, k$, falls entirely inside a Voronoi edge or facet.

For the base case of our induction we observe that the lemma holds for points $x \in u_0u_1$ (excluding u_0). To see this, notice that the direction of $v_P(x)$ for such points agrees with the vector $v_c^+ - c$. Let p be the vertex of t_c with the largest angle in t_c . Using Lemma 4.15, and taking into account that x is on the outer side of t , implies that the angle between $n^+(p)$ and $v_c^+ - c$ is at most 8ε .

In fact, by Lemma 4.15, for any point $x \in U_c^+ \cap \Sigma_{\bar{\chi}}$ that flows on a Voronoi edge e , the Delaunay triangle t dual to e must lie flat to surface and thus we can distinguish between its side facing outward and the one facing inward. Informally, we will say that in such a case x is *above* t if x is on the side of t facing outward, or *below* t otherwise.

For the induction step, we assume that the statement of the lemma holds for points on a segment $u_{i-1}u_i$ of $U_c^+ \cap \Sigma_{\bar{\chi}}$ and show that this entails the same for the point on the segment u_iu_{i+1} . Let f_1 be the Voronoi face of dimension d_1 that contains the relative interior of $u_{i-1}u_i$, and let f_2 be the Voronoi face of dimension d_2 containing u_i . Finally let f_3 be the Voronoi face that contains the relative interior of u_iu_{i+1} and let d_3 be its dimension. Notice that f_1 and f_3 are cofaces of f_2 and therefore d_1 and d_3 are both greater than d_2 . We prove the induction step by going over all possible combinations of f_1 , f_2 , and f_3 .




Figure 4.7: Illustration of the proof of Lemma 4.16, cases 3 (left) and 5 (right).

1. *Edge-vertex-edge.* First we study the case in which the flow on a Voronoi edge e , reaches a Voronoi vertex v and enters another Voronoi edge e' . We assume that the statement of the lemma holds for points on e and show that it remains true as the flow moves on to e' . To see this, let $t = pqr$ be the Delaunay triangle dual to e and let $t' = qrs$ be the one dual to e' (See Figure 4.6 (left)). The Voronoi vertex v must be dual to the Delaunay tetrahedron Δ with vertex set $\{p, q, r, s\}$. Since the flow through v continues on e' , the driver of points in e' must lie in the interior of the triangle t' or in other words, the line through e' must intersect t' . As discussed above both t and t' are flat and v is above t . It is easy to see that extending the statement of the Lemma to e' , is identical to showing that v is also above e' . By Lemma 4.15, the outward normal direction for both t and t' are within 8ε from n_q^+ . If v is not above t' , it must be in Δ . But in that case v is a maximum and the flow does not leave it to enter e' .

2. *Facet-vertex-edge.* Next, we consider the case where the flow through a Voronoi facet f dual to Delaunay edge pq reaches a Voronoi vertex v dual to tetrahedron $\Delta = pqr s$ and continues on a Voronoi edge e dual to Delaunay triangle $t = qrs$ (See Figure 4.6 (right)). We assume that the lemma holds for the points x on $U_c^+ \cap f$ and show that this extends to the points on e . As in the previous case t is flat and we only need to show that v is above t . By induction hypothesis, $\angle(v - d, n_q^+) \leq 8\varepsilon$ where $d = \frac{1}{2}(p + q)$ is the driver of the points on f . It can be verified that if v is not above t , it must be that $v \in \Delta$ making v a local maximum, a contradiction.

3. *Edge-vertex-facet.* Consider now the case where the flow through a Voronoi edge e reaches a Voronoi vertex v and enters a Voronoi facet f incident to v . Let pqr be the Delaunay triangle dual to e and let rs be the Delaunay edge dual to f . Note that v is the circumcenter of the Delaunay tetrahedron $\Delta = pqr s$. We assume that the lemma holds for points on e which is identical to assuming that v is above t . Since the flow through v continues on f , the closest point of Δ to v is the midpoint $m = \frac{1}{2}(r + s)$ of the edge rs (see Figure 4.7, left). For this to happen, v must be in the wedge made by two half-planes π_1 and π_2 both having

the line through rs as boundary and respectively being orthogonal to triangles $t_1 = prs$ and $t_2 = qrs$. Since $A_P(v) = \{p, q, r, s\}$, by Lemma 4.15, the normals to both t_1 and t_2 make an angle of at most 8ε with n_r . Since v is above t but not contained in Δ , it must be above both t_1 and t_2 . If we base at m , two vectors v_1 and v_2 , respectively normal to t_1 and t_2 in their outward directions, v_1 will lie in π_1 and v_2 in π_2 . The segment vm is on the plane bisecting rs and so are v_1 and v_2 . It follows from the triangle inequality for angles that $v - m = v_P(x)$, for $x \in f \cap U_c^+$, also makes an angle of at most 8ε with n_r^+ . Notice that with exactly the same argument but with using s instead of r , we get the same bound with respect to n_s^+ .

4. *Facet-vertex-facet.* The proof of this case is a simple combination of the proofs of cases 2 and 3.

5. *Facet-edge-facet.* We show now that under tight enough sampling, i.e. by choosing ρ large enough, if the flow through a Voronoi facet f arrives at a Voronoi edge e of f , it will continue on e and does not enter another facet f' incident to e , given that the statement of lemma holds for the points of $U_c^+ \cap f$. Suppose to the contrary that this is not the case, i.e. (see Figure 4.7, right) the flow crosses e and enters another facet f' incident to e . Let rs be the Delaunay edge dual to f . The driver of the flow on f is $m = \frac{1}{2}(r + s)$. Let y be the point where the flow reaches e . The dual Delaunay triangle t to e has r and s for vertices plus another vertex s' . For the flow to cross e and enter f' , f' must be dual to the Delaunay edge ss' . Furthermore, the line of e must not intersect t . Let o be the circumcenter of t . By our assumption, the flow direction on f , which coincides with $y - m$ makes an angle of no more than 8ε with n_r^+ . On the other hand, $y \in \Sigma_{\bar{x}}$ and thus t lies flat to surface and since the largest angle in t is at r , the normal to t , i.e. direction of e , makes an acute angle of at most 5ε with n_r (See proof of Lemma 4.15). This in particular implies that y is above t . Therefore, $\angle myo$ is at most 13ε . In order for the line of e not to intersect t , it must hold that $\angle mss' < \angle mso$. The two triangles mso and myo both share the edge mo and both have a right angle on one of the end-points of this edge. We will show below that $\|m - y\| < \|m - s\|$. Since $\|m - s\| < \|o - y\|$, this will imply that $\angle mso < \angle myo$. Using exactly the same argument as in the proof of Lemma 4.10, we get $\angle mso \geq \arcsin(\rho/2.46)$, and therefore since $\angle myo \leq 13\varepsilon$, we must have

$$13\varepsilon > \arcsin\left(\frac{\rho}{2.46}\right).$$

This inequality is violated for $\varepsilon \leq \varepsilon_0$ for $\varepsilon_0 \leq \frac{1}{13} \arcsin(\rho/2.46)$ (in particular for $\varepsilon_0 \leq 0.01$ when $\rho \geq 1/3$) giving us the desired contradiction.

Now we prove that $\|m - y\| \leq \|m - s\|$. Notice that s is a closest sample point to y and by our assumption $y \in \Sigma_{\bar{x}}$. Therefore, y is between the cones ∇_s and Δ_s . On the other hand by Corollary 4.7, $\|r - s\| = 2\|m - s\| \leq 2\|y - s\| \leq 2 \cdot 1.23\varepsilon f(s)$ and therefore using Lemma 4.8, ms makes an angle of at least $\frac{\pi}{2} - 1.23\varepsilon > \theta$ with the normal to Σ at s . This implies that m is also between ∇_s and Δ_s . By our inductive hypothesis, my makes an angle of at

most 8ε with n_s^+ . Lemma 4.14 can now be used to get

$$\|m - y\| \leq 2 \cdot 1.23\varepsilon \cdot f(s) \frac{\sin(1.1\varepsilon)}{\cos(9.1\varepsilon)} < \frac{1}{2} \delta f(s) \leq \|m - s\|,$$

where the middle inequality holds for $\varepsilon \leq 0.01$ and $\rho \geq 1/3$. In fact for any constant $0 < \rho < 1$ the above inequality holds (and the desired contradiction is achieved) for any $\varepsilon \leq \varepsilon_0$ for small enough ε_0 since the left hand side has a quadratic dependence on ε .

Thus we have proved that whenever the flow on U_c^+ moves to a Voronoi facet f , it leaves f by either hitting a Voronoi edge e and continuing on e , or by hitting a vertex v . Thus we have covered all cases in the inductive step and this completes the proof of the lemma. ■

In the following Lemma we show that if $\text{Sm}(c)$ is incident to the stable manifold $\text{Sm}(m)$ of an interior (exterior) medial axis maximum m , then the part of $\text{Sm}(c)$ that is contained in t_c is incident to $\text{Sm}(m)$ at the inner (outer) side of t_c .

Lemma 4.17 *For any surface 2-saddle c , U_c^+ does not intersect Σ_Δ and U_c^- does not intersect Σ_∇ .*

Proof. We prove the claim for U_c^+ . The other claim is proved analogously. Suppose to the contrary that U_c^+ intersects Σ_Δ at x . Let v be the last turning point of U_c^+ before reaching x . Let q be a sample point for which $x \in \Delta_q$ and let p be a closest sample point to x . Then $\|x - p\| \leq \|x - q\| \leq f(q) \cos \theta = f(q) \sin(1.1\varepsilon)$. Therefore, $\|p - q\| \leq 2f(q) \sin(1.1\varepsilon)$ and therefore by Lemma 1.12,

$$\angle(n_p^+, n_q^+) \leq \frac{2 \sin(1.1\varepsilon)}{1 - 3 \cdot 2 \sin(1.1\varepsilon)} = O(\varepsilon).$$

On the other hand, by Lemma 4.16, the vector $x - v$ makes an angle of $O(\varepsilon)$ with n_p^+ . It is easy to observe that this contradicts the assumption that the flow hits Δ_q . ■

The following lemma is a direct consequence of Lemma 4.16 and Lemma 4.13.

Lemma 4.18 *Let c_1 and c_2 be two surface critical points with $\text{Sm}(c_1)$ and $\text{Sm}(c_2)$ put by RECONSTRUCT into $\tilde{\Sigma}$, such that boundaries of $\text{Sm}(c_1)$ and $\text{Sm}(c_2)$ have a Gabriel edge e in common. Let t_1 and t_2 be the patch triangles incident to e in $\text{Sm}(c_1)$ and $\text{Sm}(c_2)$, respectively. Then, the dihedral angle between t_1 and t_2 is larger than $\pi/2$.*

Proof. Orient patch triangles of $\text{Sm}(c_1)$ by taking for every patch triangle t of $\text{Sm}(c_1)$, the normal to t pointing to the side of t incident to the interior of the reconstruction. Denote this normal by n_t^- . Lemma 4.16 implies that in this case for every point $x \in t$ where t is a patch triangle of $\text{Sm}(c_1)$, $\angle(n_x^+, n_t^-) < 23^\circ$. In particular, by letting $t = t_1$ and choosing x to be a point on e , we get $\angle(n_x^+, n_{t_1}^-) < 23^\circ$

If we do a similar orientation on $\text{Sm}(c_2)$, we get $\angle(n_{\hat{x}}^+, n_{t_2}^+) < 23^\circ$. Thus, the dihedral angle between t_1 and t_2 is at least $180^\circ - 46^\circ = 134^\circ$. ■

4.1.4 $\tilde{\Sigma}$ and Σ are Homeomorphic

Theorem 4.19 *Under assumptions of Theorem 4.1, the output $\tilde{\Sigma}$ produced by the algorithm RECONSTRUCT is a 2-manifold without boundary and is homeomorphic to Σ .*

Proof. We consider in this proof the case where $\tilde{\Sigma}$ is the boundary of the union of stable manifolds of the inner medial axis critical points (the outer case is similar). We argue that $\tilde{\Sigma}$ and Σ are homeomorphic. Consider the restriction $\zeta: \tilde{\Sigma} \rightarrow \Sigma$ of the closest point map $x \mapsto \hat{x}$. We prove that ζ is a homeomorphism. Since both $\tilde{\Sigma}$ and Σ are compact, it is sufficient to show that ζ is continuous, one-to-one and onto.

First, we argue that ζ is one-to-one. Orient the normal to each patch triangle t so that it makes an angle less than $\frac{\pi}{2}$ with the oriented normal n_p^+ at the vertex p of t which is a sample point. Because of Lemma 4.13 and Lemma 4.18, the triangles of $\tilde{\Sigma}$ can be oriented consistently satisfying this condition. We denote this oriented normal for a patch triangle t by n_t . Notice that although Lemmas 4.13 and 4.18 are stated for the special case where $\rho \geq 1/3$ and $\varepsilon \leq 0.01$, they can effectively be reproduced for any smaller ρ provided that ε is chosen small enough accordingly.

By Lemma 4.10, for every point x in a patch triangle t the oriented triangle normal n_t makes an angle of $O(\varepsilon/\rho)$ with $n_{\hat{x}}^+$. In particular when $\varepsilon \leq 0.01$ and $\rho \geq 1/3$, this angle is at least 23° . Suppose ζ is not one-to-one. Then, there are two points x and x' in $\tilde{\Sigma}$ that are both mapped to the same point \hat{x} by ζ . Consider the line ℓ normal to Σ at \hat{x} . This line passes through both x and x' . Assume without loss of generality that x and x' are consecutive intersection points of ℓ and $\tilde{\Sigma}$. Then, at one of x and x' the line ℓ enters and at the other exits the interior bounded by $\tilde{\Sigma}$. In other words, if we orient ℓ along $n_{\hat{x}}^+$, it makes an angle at least $\frac{\pi}{2}$ with one of the oriented normals of $\tilde{\Sigma}$ at x or x' , an impossibility.

Next, we argue that $\tilde{\Sigma}$ is a manifold. For this we first observe that each edge in $\tilde{\Sigma}$ is incident to at least two triangles. This of course holds by definition for the interior edges of each surface patch. If a Gabriel edge on the boundary of a surface patch is incident only to that patch, the patch must be incident to the stable manifold of the same inner (or outer) medial axis maximum on both sides. This contradicts Lemma 4.17. We show now that the triangles incident to each vertex v of $\tilde{\Sigma}$ form a topological disk and hence $\tilde{\Sigma}$ is a 2-manifold. If not, there are two triangles incident to v so that a normal line stabs both of them at points arbitrarily close to v since they lie almost parallel to Σ . This is in contradiction with ζ being one-to-one.

We are left to show that ζ is continuous and onto. The continuity of ζ follows from the fact

that the original closest point function $x \mapsto \hat{x}$ is continuous everywhere except at the medial axis. To show that ζ is onto, consider $\zeta(\tilde{\Sigma}) \subseteq \Sigma$. We claim that $\zeta(\tilde{\Sigma}) = \Sigma$. Since $\tilde{\Sigma}$ is a 2-manifold without boundary and ζ maps it homeomorphically to $\zeta(\tilde{\Sigma})$, we have $\zeta(\tilde{\Sigma})$ as a compact 2-manifold without boundary and $\zeta(\tilde{\Sigma}) \subseteq \Sigma$. This is only possible if $\zeta(\tilde{\Sigma}) = \Sigma$ as both $\zeta(\tilde{\Sigma})$ and Σ are compact 2-manifolds without boundary. ■

4.2 Extension to Higher Dimensions

The proof presented in the previous section does not take full advantage of the fact that the flow ϕ_P is a continuous map. Neither does it use the important fact that ε^2 -tubular neighborhood Σ_{ε^2} of the surface Σ is flow-repellant for the flow ϕ_P . Using these extra information we can generalize the result of the previous section to higher dimensions; well, almost!

The proof to which the rest of this chapter is dedicated is not exactly the higher-dimensional equivalent of Theorem 4.1 and differs from it in the followings details.

- The type of “topological equivalence” achieved in the previous section was homeomorphism. The result of the present chapter will however be a homotopy equivalence.
- The algorithm of the previous section outputs the boundary of the union of stable manifolds of the inner (or outer) medial axis critical points as a surface (a 2-manifold) which by Theorem 4.1 captures the topology of the surface Σ . In contrast, in the present section we establish the homotopy equivalence of the union of the stable manifolds of the inner medial axis critical points (and not their boundary) with the bounded shape S enclosed by Σ (and not with Σ itself). Notice that these results are not the same; two open sets in \mathbb{R}^n can be homotopy equivalent while their boundaries are not and two surfaces can be homeomorphic while the volumes they enclose are not homotopy equivalent.
- The geometric and topological guarantees that we gave in the previous section strongly depended on the tightness of the input sample. In particular, the density ε of the sample required by Theorem 4.1 depends on the tightness of the input sample. Yet the argument we supply later in this section merely needs a simple ε -sample of the surface.

As before, the set P is an ε -sample of the target surface Σ . An upper bound for ε will be determined later. We intend to use the flow map ϕ_P for the role of the continuous map H in Proposition 1.8 to show the desired homotopy equivalence between S and the output of the algorithm which is denoted by \tilde{S} and is the union of stable manifolds of inner medial axis critical points.

4.2.1 Homotopy Equivalence via Distance Flow Maps

Let Q be a set of (possibly weighted) points in \mathbb{R}^n . If X and Y are subsets of \mathbb{R}^n with $Y \subset X$, in order to establish a homotopy equivalence between X and Y by applying Proposition 1.8 using ϕ_Q as H , one must show that

1. the flow orbit of every point in X stays in X , i.e. $\phi_Q(X) = X$ (and thus the map ϕ_Q can be restricted to X alone),
2. the flow orbit of every point in Y stays in Y , i.e. $\phi_Q(Y) = Y$, and
3. within a finite amount of time, the flow orbit of every point in X ends in Y .

Notice that the first condition of Proposition 1.8 is automatically satisfied for any flow map ϕ_Q because for every $x \in \mathbb{R}^n$, $\phi_Q(0, x) = x$. If X is bounded, the finiteness of flow time into Y can be guaranteed using Corollary 2.10 provided there is a lower bound c for $\|v_Q(x)\|$ for every $x \in X \setminus Y$: let Δ be an upper bound on the diameter of X and assume that for some constant $c > 0$, $\|v_Q(x)\| \geq c$ for every $x \in X \setminus Y$. Let $y = \phi_Q(t, x)$ be in $X \setminus Y$. Since Y is flow-tight for ϕ_Q , $y \notin Y$ implies that $\phi_Q(\tau, x) \notin Y$ for all $0 \leq \tau \leq t$. Consequently $\|v_Q(\phi_Q(\tau, x))\| \geq c$. Then by Corollary 2.10 we have for $y = \phi_Q(t, x)$

$$\begin{aligned}
 h_Q(y) &= h_Q(\phi_Q(t, x)) \\
 &= h_Q(x) + \int_0^t \|v_Q(\phi_Q(\tau, x))\|^2 d\tau \\
 &\geq h_Q(x) + \int_0^t c^2 d\tau \\
 &= h_Q(x) + tc^2.
 \end{aligned} \tag{4.7}$$

But since X is flow-tight for ϕ_Q , $y \in X$ and therefore

$$\begin{aligned}
 h_Q(y) &= \min_{q \in Q} (\|y - q\|^2 - w_q) \\
 &\leq \max_{x \in X, q \in \tilde{Q}} \|x - q\|^2 - \min_{q \in Q} w_q \\
 &\leq \left(\Delta + \text{dist}(X, \tilde{Q}) + \text{diam } \tilde{Q} \right)^2 - \min_{q \in Q} w_q.
 \end{aligned}$$

The latter quantity is bounded and this put a finite upper bound on the value of t in the inequality (4.7). Thus we have proved the following theorem.

Theorem 4.20 *Let Q be a finite set of (possibly weighted) points in \mathbb{R}^n . If for sets $Y \subset X \subset \mathbb{R}^n$, X and Y are both flow-tight for ϕ_Q , i.e. $\phi_Q(X) = X$ and $\phi_Q(Y) = Y$, and there is a constant $c > 0$ for which $\|v_Q(x)\| \geq c$ for all $x \in X \setminus Y$, then X and Y are homotopy equivalent.*

The above theorem is the key to all homotopy equivalence proofs in the rest of this monograph. To invoke the theorem we need two flow-tight sets X and Y which we sometimes call the *source set* and the *sink set*, respectively. Stable and unstable manifolds, their unions, and their intersections are flow-tight by definition. However, these won't be the only examples of flow-tight sets we will consider.

4.2.2 \tilde{S} and S are Homotopy Equivalent

We now move back to the proof of the homotopy equivalence of S and \tilde{S} . First we show that removing the Voronoi faces corresponding to surface critical points from \tilde{S} does not change its homotopy type.

Lemma 4.21 *Let c be a critical point of h_P and let $U \subseteq \mathbb{R}^n$ be a flow-tight set for ϕ_P that does not include c , i.e. $c \notin U$. Let $V = \text{rel int } V_P(c)$ be the relative interior of the lowest dimensional Voronoi face in $\text{Vor } P$ that contains c . For $r \geq 0$, let V_r be the intersection of V with the open ball $B(c, r)$, i.e. $V_r = V \cap B(c, r)$. Then for every $r \geq 0$, U and $U \setminus V_r$ have the same homotopy type. In fact $U \setminus V_r$ is a strong deformation retract of U . Moreover, if $U \cap B(c, r) \subset V$ then $U \setminus V_r$ is flow-tight for ϕ_P .*

Proof. We build a deformation retraction from U to $U \setminus V_r$. Since c is a critical point and V is the relative interior of the lowest-dimensional Voronoi face that contains c , c is the driver of the points in V . Consequently, if $x \neq c$ is a point in $V \cap U$, $d_P(x) = c$ and since U is flow-tight for ϕ_P , we have

$$\{x + t(x - c) : t \geq 0\} \cap \text{cl } V \subseteq U.$$

Now, define the map $\rho_r : V_r \rightarrow \partial V_r$ (where ∂V_r is defined relative to the affine hull of V_r) as

$$\rho_r(x) = \underset{t \geq 0}{\text{argmin}} \{y : y = x + t(x - c), y \in \partial V_r\}.$$

In other words, $\rho_r(x)$ is the point where of the ray shot from x in the direction of $x - c$ hits the boundary of V_r . Since V_r is convex (it is the intersection of V and $B(c, r)$ which are both convex), it is easy to see that the map ρ_r is continuous (it is a projection from a point in a convex set to the boundary of the convex set) and this implies that the retraction map $\rho_r^* : U \rightarrow U \setminus V_r$ defined below is also continuous.

$$\rho_r^*(x) = \begin{cases} \rho_r(x) & x \in V, \\ x & x \in U \setminus V_r. \end{cases}$$

We now define the map $R_r : [0, 1] \times U \rightarrow U$ as

$$R_r(t, x) = \begin{cases} (1-t)x + t\rho_r(x) & x \in V_r, \\ x & x \in U \setminus V_r. \end{cases}$$

which gives us a straight line homotopy between from the identity map of U to the retraction ρ_r .

If $U \cap B(c, r) \subset V$, then for any $y \in V_r \cap U$, only points in V_r can flow into y . In other words, $y = \phi_P(t, x)$ for some $t \geq 0$ and $x \in U$ implies that $x \in V_r$. Therefore, all flow lines that are affected by removal of V_r from U start in V_r . But we saw above that each such flow line loses an initial segment in $U \setminus V_r$. Thus $U \setminus V_r$ is flow tight for ϕ_P . ■

Lemma 4.22 *Let C_Σ be the set of all surface critical points of h_P where P is an ε -sample of a smooth surface Σ for $\varepsilon \leq 0.14$ and let $U_\Sigma = \bigcup_{c \in C_\Sigma} \text{Um}(c)$ be the union of unstable manifolds of all surface critical points. Then the set U defined as*

$$U = (S_{\varepsilon^2} \cup U_\Sigma) \cap \tilde{S}$$

is homotopy equivalent to S .

Proof. We show that $U \simeq S_{\varepsilon^2}$ which proves the lemma since $S_{\varepsilon^2} \simeq S$ by Lemma 3.18. First notice that U is a flow-tight set. This is because S_{ε^2} and U_S are both flow-tight and so is \tilde{S} .

Recall that the index of a critical point c , $\text{ind } c$, is $\dim D_P(c)$, i.e. the dimension of the Delaunay face dual to the lowest dimensional Voronoi face of $\text{Vor } P$ that contains c . Equivalently $\text{ind } c = |A_P(c)| - 1$. Let C_Σ^i , $i = 0, \dots, n$ denote the set of surface critical points of index i . Thus

$$C_\Sigma = C_\Sigma^0 \cup \dots \cup C_\Sigma^n.$$

We define for every $0 \leq i < n$ the set $U_i \subset \mathbb{R}^n$ as

$$U_i = \left(S_{\varepsilon^2} \cup \bigcup_{j \geq i} \bigcup_{c \in C_\Sigma^j} \text{Um}(c) \right) \cap \tilde{S}$$

In other words, U_i is the restriction to \tilde{S} of the union of the reduced shape S_{ε^2} plus the unstable manifolds of surface critical points of index i or higher. Notice that $U_i = \tilde{S}$ and $U_n = S_{\varepsilon^2}$ because surface critical points in general are not contained in \tilde{S} and the unstable manifold of a critical point of index n , i.e. a maximum is the critical point itself.

First observe that all U_i 's are flow-tight for ϕ_P by definition. To complete the proof, we show by induction on i that all U_i 's, $i = 0, \dots, n$ are homotopy equivalent and this will prove our Lemma. For a base case, we show that U_0 and U_1 are homotopy equivalent. To see this, consider an index 0 critical point of h_P , i.e. a minimum. An index 0 critical point is exactly a point $p \in P$. The unstable manifold of p is $\phi_P(V_p)$. Therefore, for a small enough $r > 0$

the ball $B(p, r)$ is contained in V_p and as a result,

$$U_0 \cap B(p, r) \subset V_p.$$

Thus by Lemma 4.21 we can remove a neighborhood of every surface critical of index 0 from U_0 to get a set U'_0 that is flow-tight for ϕ_P and is homotopy equivalent to U_0 . Next we show that $U'_0 \simeq U_1$. For this we use Theorem 4.20. Since U'_0 and U_1 are both flow-tight for ϕ_P , all we need to do is to find a lower bound for $\|v_P(x)\|$ for points $x \in U_1 \setminus U'_0$.

Now, for any point $x \in U_1 \setminus U'_0$, let $V_P(x)$ be the lowest dimensional Voronoi face of $\text{Vor } P$ that contains x and let $D_P(x)$ be the Delaunay face dual to $V_P(x)$. The driver $d_P(x)$ is contained in $D_P(x)$. There are two cases to consider; depending on whether $V_P(x)$ and $D_P(x)$ intersect or not.

If $V_P(x) \cap D_P(x) = \emptyset$ then let ζ be the minimum over all pairs of dual faces $V \in \text{Vor } P$ and $D \in \text{Del } P$ of $\text{dist}(V, D) > 0$. We then have

$$\|v_P(x)\| = 2\|x - d_P(x)\| \geq 2 \text{dist}(V_P(x), D_P(x)) \geq 2\zeta.$$

On the other hand if $V_P(x)$ and $D_P(x)$ do intersect, their intersection will (by definition) be a critical point c which coincides with $d_P(x)$. If c is a medial axis critical point, then $c \notin \Sigma_\delta$ for any $\delta \leq 1 - 2\varepsilon^2$. Thus

$$\|v_P(x)\| = 2\|x - c\| \geq 2 \text{dist}(\Sigma_{\varepsilon^2}, S_{1-2\varepsilon^2}) \geq 2\xi,$$

where it can be shown similar to the argument of Lemma 3.14 that

$$\xi \geq (1 - 3\varepsilon^2) \cdot f_0,$$

in which f_0 is the minimum over Σ of local feature size which was assumed to be strictly positive. Notice that c cannot be a medial axis critical point since the inner medial axis critical points and their unstable manifolds are contained in $S_{2\varepsilon^2}$ which is flow-tight for ϕ_P .

Thus the only remaining possibility is that c is a surface critical point. Note that since $c \in V_P(x)$, $V_P(c) = V_P(x)$ and therefore $V_P(x) \subset \text{Um}(c)$. But in that case c must have index 0 since unstable manifolds of surface critical point of index 1 and higher are included in U_1 while $c \in U'_0 \setminus U_1$. Thus c is an index 0 critical point. Therefore

$$\|v_P(x)\| = 2\|x - c\| \geq 2 \text{dist}(U'_0, C_\Sigma^0) \geq 2\eta,$$

where C_Σ^0 denotes the set of critical points of index 0 or less. Thus

$$\|v_P(x)\| \geq 2 \min\{\zeta, \xi, \eta\} > 0.$$

Theorem 4.20 now implies that $U'_0 \simeq U_1$.




Figure 4.8: A two-dimensional demonstration of the stages in the proof of Lemma 4.22. Top-left: the source set U is shown in gray and the sink set S_{ϵ^2} is displayed in white. Top-right: the source set $U = U_0$ and all the critical points of the sample. Medial axis critical points are shown in red. Surface 1-saddles are shown in blue and samples in black. Middle-left: neighborhoods from index-0 critical points are removed from U_0 to result $U'_0 \simeq U_0$. Middle-right: U_1 is the union of stable manifolds of index-1 critical points and S_{ϵ^2} intersected with \tilde{S} . Bottom-left: Neighborhoods of index-1 saddle points are removed from U_1 . The resulting set is $U'_1 \simeq U_1$. Bottom-right: remaining spikes are pushed into S_{ϵ^2} .

The proof of homotopy equivalence of U_i and U_{i+1} is done exactly the same way as we did for U_0 and U_1 . The only detail is removing of a neighborhood of index i surface critical point in such a way to be able to guarantee that the resulting set U'_i is flow-tight for ϕ_P . For this to be done using Lemma 4.21 as we did to make U'_0 , we need to make sure that if c is an index i surface critical point, then

$$U_i \cap B(c, r) \subset V_P(c)$$

for sufficiently small $r > 0$. This is the case unless the unstable manifold of some other surface critical point c' reaches arbitrarily close to c and is not contained in $\text{Um}(c)$. Corollary 2.22 now implies that in this case $\text{ind } c' < \text{ind } c$. But the unstable manifolds of surface critical points of index less than $i = \text{ind } c$ are not included in U_i . This completes the proof of the Lemma. ■

Theorem 4.23 *For $\varepsilon \leq 0.14$ the union of unstable manifolds of inner medial axis critical point, \tilde{S} is homotopy equivalent to S . Moreover, $S_{\varepsilon^2} \subset \tilde{S} \subset S \cup \Sigma_{\varepsilon^2}$.*

Proof. Using Lemma 4.22 we only need to show that $U \simeq \tilde{S}$. Again we use Theorem 4.20. Recall that both of the set U and \tilde{S} are flow-tight for ϕ_P . Therefore we only need to show that for every point $x \in \tilde{S} \setminus U$, $\|v_P(x)\|$ is bounded away from 0. As in the proof of Lemma 4.22 we distinguish between two cases of $V_P(x) \cap D_P(x) = \emptyset$ and $V_P(x) \cap D_P(x) = \{c\}$ where c is a critical. The first case is handled exactly like Lemma 4.22. For the second case we note as in the proof of the previous Lemma that $x \in \text{Um}(c)$ by definition. However, if c is a surface critical point, $\text{Um}(c) \cap \tilde{S}$ is contained in U and therefore $x \in U$. Thus c must be a medial axis critical point. Since all inner medial axis critical points are in $S_{1-2\varepsilon^2}$, c can only be an outer medial axis critical point. It is easy again now to bound the distance between x and c as in the proof of Lemma 4.22.

The geometric guarantees of the Theorem follow the flow-tightness of S_{ε^2} and $S_{\varepsilon^2}^*$. ■

Bibliography and Remarks.

The surface reconstruction problem has an extremely rich literature that spans several disciplines inside and outside the computer science community. Two-dimensional versions of this problem arise in pattern recognition, image processing, and computer vision [58]. In particular, the use of Delaunay triangulations for boundary recognition in images have been studied by Brandt and Algazi [16] and by Robinson, et al. [62].

In computer graphics community, Hoppe et al. [50] provided the first clean abstraction for the problem of surface reconstruction from unorganized point clouds. They suggested an algorithm for locally estimating a “signed distance” to the target surface using the distance to the closest sample point. The output surface is then a piece-wise linear approximation of the zero level set of this signed distance function. Reconstruction of the surface as an approximation of the zero set




Figure 4.9: Left: A sampled curve. Right: Delaunay triangulation of the sample. A piecewise linear approximation of the curve can be found as a sub-complex of this triangulation.

of a function that is interpolated from the input sample is used in many other works. Following Hoppe et al., Curless and Levoy [29] used a distance function represented on a voxel grid. Boissonnat and Cazals [13] used natural neighbor interpolation in conjunction with Delaunay triangulations and introduced a surface reconstruction algorithm that provided geometric and topological guarantees. Levin [55] introduced the moving least square surfaces (MLS), which used the least square method to fit a surface into the sampled point cloud. This method was later adopted by Alexa et al. [1] to make a surface reconstruction algorithm which was widely used subsequently for surface modeling and rendering. Theoretical results on MLS surface reconstruction, including sampling conditions and geometric and topological guarantees, were later provided by Amenta and Kill [6], Kolluri [53], Bremer and Hart [17], Dey, Goswami, and Sun [34], and Dey and Sun [36]. Recently learning theoretic techniques involving support vector machines (SVMs) are employed to produce smooth functions the zero set of which can be used for approximation of the input surface (See for example Schölkopf et al. [63], or Carr et al. [19]). Thus far, no topological guarantees have been established for these methods.

Delaunay triangulation based surface reconstruction algorithms have perhaps been the most successful in terms of theoretical guarantees they provide for the geometric and topological quality of the output. All these algorithms exploit the Delaunay complex or its dual Voronoi complex of the input sample. The Delaunay triangulation in particular works well, as observed first by Boissonnat [12], because of its capturing of the notion of *neighboring* points to any given sample point (see Figure 4.9).

Given an ε -sample of a surface Σ for a certain value of ε , Amenta and Bern [2] gave an algorithm with topological and geometric guarantees that selects a subset of Delaunay triangles as output. Their algorithm first filters a set of “candidate” Delaunay triangles and then extracts a manifold from this set. A drawback of the manifold extraction method is that it fails easily if the sampling density requirement is not met, as often happens in practice, resulting an output which is not a manifold. Amenta, Choi, Dey and Leekha [4] provided a similar algorithm but with a considerably simpler proof of the homeomorphism of the output to

the original surface. The algorithm of Boissonnat and Cazals [13] avoids manifold extraction and uses a natural neighbor interpolation technique to produce a smooth surface with the right topology. Amenta, Choi, and Kolluri [5] provided an algorithm that also avoids manifold extraction and its output is always guaranteed to be the boundary of a three-dimensional solid.

Distance functions methods have been used by Edelsbrunner [38], Chaine [20], and Giesen and John [43] to reconstruct Σ from P . Although all three of these algorithms work well in practice, they provide no guarantees for the geometric quality of their output in the ε -sampling framework. The algorithm presented in this section is a modified version of that Giesen and John's algorithm [43]. In Chapter 6 we introduce a modification of similar spirit to Edelsbrunner WRAP algorithm [38] for which topological and geometric guarantees can be provided.

Giesen and John [42] were first to observe that when P is a dense enough sample of a surface Σ , the bounded volume S enclosed by Σ and its complement S^* can each be approximated as a unions of stable manifolds of a subset of critical points of the distance function h_P . They provided a heuristic method for picking a suitable set of critical points of h_P the stable manifolds of which combined would resemble S . Their heuristic considers distance of the critical points to P (critical values) and is based on a pairing of critical points based on incidence graph of the stable flow complex. Being a heuristic, this method of Giesen and John supplies no theoretical guarantees although it works reasonably well in practice.

Theorem 4.20 is a formalization and slight generalization for weighted points of Lieutier and Chazal's technique for establishing homotopy equivalence through invocation of Proposition 1.8 as applied in [56, 22, 23]. The development of some of the the machinery of Chapter 2 was also inspired in part by their work.

Chapter 5

A Core for Approximating the Medial Axis

We turn now to the problem of medial axis approximation. Specifically, we wish to approximate the medial axis of a bounded open shape S whose boundary Σ is a smooth surface, using an ε -sample P of Σ . We introduce a piece-wise linear structure we call the inner medial axis core of the sample P . This core is then shown to have the same homotopy type as the medial axis of the inner shape S of the target surface Σ .

Once the core is computed, it can be extended to better capture the geometry of the medial axis while maintaining its homotopy type. To extend the core one picks, using any algorithm at hand, a set of points that approximates a subset of the medial axis of the target shape and adds this set to the computed core along with its *flow closure*. We show that as long as the chosen approximating set of points does not get too close to the surface of the shape, adding the closure keeps the homotopy type of the core. Because of this property, computing the core can be used to augment most medial axis approximation algorithm into a topologically accurate one. Perhaps the most relevant example of such an algorithm is the one devised by Dey and Zhao [37]. Their algorithm filters out some Voronoi faces that closely approximate a large part of the medial axis of a sampled surface. This filtering is done entirely based on local conditions. Although their algorithm performs well geometrically it cannot guarantee the topological type of its output. The core extension method covered in this chapter is specially suitable in conjunction to this algorithm because computing flow closure of Voronoi facets is easy.

5.1 Definition of Core

Given a sample P of the surface, we compute a piece-wise linear cell complex, which we call the *inner medial axis core* of h_P . The core is extracted from — and in fact is a subcomplex of — the unstable flow complex induced by the sample which itself can be regarded as a refinement of the Voronoi $(n - 1)$ -skeleton of the sample. As before, the relevant analysis is done in the relative ε -sampling framework.

Thus we assume we are given an ε -sample P from a smooth hypersurface Σ in \mathbb{R}^n where the value of ε is determined later but a priori is assumed to be less than $1/\sqrt{3}$ to make sure it is possible to distinguish between surface and medial axis critical points of h_P using Corollary




Figure 5.1: An example of the core in 2D. Inner medial axis critical points are shown in red. Hollow points represent 1-saddles and solid ones are used for maxima. The inner core is shown in blue.

3.4.

With these assumptions, the inner medial axis core or simply the *core* of h_P is defined as the union of the unstable manifolds of the inner medial axis critical points of h_P . In other words, if we denote by C_M^- the set of all inner medial axis critical points of h_P , the core of h_P denoted \mathcal{C}_P is defined as

$$\mathcal{C}_P = \bigcup_{c \in C_M^-} \text{Um}(c).$$

Figure 5.1 depicts a 2-dimensional example of the core. Our primary goal in this section is to show that for a sufficiently dense samples of Σ , the medial axis core of h_P is homotopy equivalent to the medial axis $M(S)$ of S .

Remark. For technical reasons we refrain from considering the outer medial axis core of h_P which can be defined in a way similar to the inner version only using the outer medial axis critical points. The main reason for this is the difficulty brought about by the unboundedness of S^* and the unstable manifolds of some of outer medial axis critical points. The homotopy equivalence method of this chapter is, as expected, based on flow methods and in particular Proposition 1.8 which fail for unbounded sets.

In practice however, one workaround for this is to enclose the surface Σ with a huge $(n - 1)$ -sphere Σ_0 and treat the sphere as an extra component of the surface Σ . The sphere Σ_0 is sampled in addition to the original surface. This will introduce a component of the medial axis corresponding to the volume enclosed between Σ and Σ_0 . If Σ_0 is chosen large enough,

the medial axis of this component agrees with the original outer medial axis of Σ , i.e. $M(S^*)$, near Σ and the larger the outer sphere is chosen the farther from Σ this agreement is extended. This will introduce a set of medial axis critical points trapped between the surface and the the sphere that can be treated as outer medial axis critical points for which a core similar to the inner medial axis core can be computed. This core can be shown to be homotopy equivalent to the component of the medial axis between Σ and the outer sphere in exactly the same manner as we are about to see for the inner medial axis core.

5.2 Homotopy Type of the Core

In this section we first show that the inner medial axis core \mathcal{C}_P is homotopy equivalent to the inner medial axis $M(S)$ which is itself homotopy equivalent to S by Theorem 1.11.

After that we prove that for any set of points $W \subset S_\delta$ for large enough δ , the set $\mathcal{C}_P \cup \phi_P(W)$ is also homotopy equivalent to the core. A consequence of this is that any algorithm for geometric approximation of $M(S)$ can be turned into one that always captures the topology of the medial axis.

5.2.1 Homotopy Equivalence of Core and Medial Axis

The homotopy equivalence proof presented in this section is also based on Theorem 4.20. As before, we refer to the sets X and Y in this theorem the source set and the sink set, respectively. To establish the homotopy equivalence between the source and sink sets using a flow map ϕ_P , Theorem 4.20 requires both of these sets to be flow-tight for ϕ_P . In addition, it is required that the speed of the flow at points of source set outside the sink set be bounded from below. Here, our sink set is the core \mathcal{C}_P which is flow-tight by definition. For source set, we choose the closure of S_δ for a value of δ that turns $\text{cl } S_\delta$ flow-tight for ϕ_P . By Lemma 3.17, it is enough to choose $\varepsilon^2 \leq \delta \leq 10\varepsilon^2$ and take $\varepsilon \leq 0.14$. Of course we need to make sure that the sink set is a subset of the source.

Lemma 5.1 *If $\varepsilon \leq 0.14$ and $\varepsilon^2 \leq \delta \leq 10\varepsilon^2$ then $\mathcal{C}_P \subset S_\delta$.*

Proof. By Corollary 3.4, the inner medial axis critical points of h_P are all contained in $S_{1-2\varepsilon^2} \subset S_\delta$. In fact, $1 - 2\varepsilon^2 > 10\varepsilon^2$ implies that for every inner medial axis critical point c , some neighborhood of c is contained in S_δ in addition to c . Every point of $x \in \mathcal{C}_P$ is on the flow orbit $\phi_P(y)$ of a point y infinitesimally close to some inner medial axis critical point c and is therefore in S_δ . On the other hand, by Lemma 3.17, S_δ is flow-tight for our choices ε and δ . Therefore x being on the orbit of a point $y \in S_\delta$ has to be in S_δ as well. ■

Thus all left to do is to bound the speed of the flow outside the core.

Lemma 5.2 For $\varepsilon \leq 0.14$, if P is an ε -sample of Σ , then there is a constant $\nu > 0$ such that $\|v_P(x)\| \geq \nu$ for all $x \in \text{cl } S_\delta$ with $2\varepsilon^2 \leq \delta \leq 10\varepsilon^2$.

Proof. Let $\zeta > 0$ be the minimum distance between any pair of non-intersecting dual faces $D \in \text{Del } P$ and $V \in \text{Vor } P$. Consider the reduced shape $\text{cl } S_\delta$ for $\delta = 2\varepsilon^2$. By Lemma 3.17, $\text{cl } S_\delta$ is closed under the flow ϕ_P . Therefore, every flow line of ϕ_P in $\text{cl } S_\delta$ ends in some inner medial axis critical point of h_P in the limit. Consider now any point $x \in \text{cl } S_\delta \setminus \mathcal{C}_P$ and let $V_P(x)$ be the lowest dimensional face of $\text{Vor } P$ that contains x and let $D_P(x)$ be its dual face in $\text{Del } P$. There are two cases to consider depending on whether $V_P(x)$ and $D_P(x)$ intersect.

Case 1. $V_P(x) \cap D_P(x) = \emptyset$. In this case, the distance $x - d_P(x) \geq \zeta$ since $d_P(x) \in D_P(x)$ and $\text{dist}(V_P(x), D_P(x)) \geq \zeta$. Therefore we get

$$\|v_P(x)\| = 2\|x - d_P(x)\| \geq 2 \text{dist}(V_P(x), D_P(x)) \geq 2\zeta.$$

Case 2. $V(x) \cap D(x) = \{c\}$, where c is a critical point of d_P . Notice that in this case $V_P(x) = V_P(c)$. Recalls that for a critical point c , $V_P(c)$ is contained in the unstable manifold $\text{Um}(c)$ of c .

First we note that c cannot be a medial axis critical point since otherwise $x \in V_P(x) = V_P(c) \subseteq \text{Um}(c) \subseteq \mathcal{C}_P$ and this contradicts our choice of x .

Therefore, c must be a surface critical point and we have $c \in \Sigma_{\varepsilon^2}$ by Corollary 3.4. However, $x \in \text{cl } S_\delta$ where $\delta \geq 2\varepsilon^2$. With a similar argument to the one used in Lemma 3.14, if we use $\delta = 2\varepsilon^2$ and $\lambda = \varepsilon^2$, the open ball with center c and radius

$$\frac{(\delta - \lambda)(1 - \lambda)}{1 - \lambda + 2\delta} f(\hat{c})$$

is entirely contained in Σ_δ and therefore $\|x - c\| \geq \xi$ where

$$\xi = \frac{(1 - \varepsilon^2)\varepsilon^2}{1 + 3\varepsilon^2} \cdot f_0,$$

where f_0 is the reach of Σ , i.e. the minimum of the local feature size on Σ which we assumed is strictly positive. Thus we have

$$\|v_P(x)\| = 2\|x - c\| \geq 2\xi.$$

Thus for every point $x \in \text{cl } S_\delta \setminus \mathcal{C}_P$, $\|v_P(x)\| \geq \nu$ where

$$\nu = 2 \min\{\xi, \zeta\}. \quad \blacksquare$$

We are now ready to prove that the core captures the topology of the medial axis.

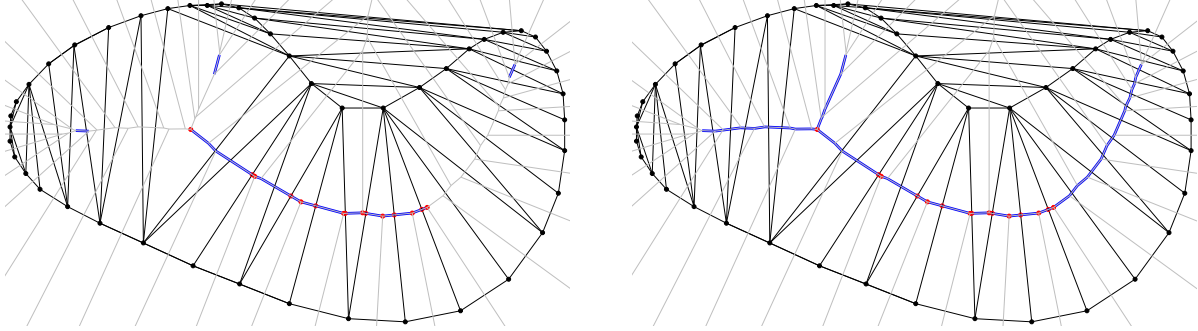


Figure 5.2: Left: a set W of Voronoi faces (in blue) that approximate parts of the inner medial axis are added to the core (also in blue). Right: flow closures $\phi_P(W)$ of W is added.

Theorem 5.3 *Let $\varepsilon \leq 0.14$ and let P be an ε -sample of a smooth surface Σ in \mathbb{R}^n that encloses a bounded shape S . Then, the shape S , its medial axis $M(S)$, and the inner medial axis core \mathcal{C}_P induced by P consisting of the the union of unstable manifolds of inner medial axis critical points of h_P , are all homotopy equivalent.*

Proof. The homotopy equivalence of S and $M(S)$ holds by Theorem 1.11 for every bounded open subsets of \mathbb{R}^n . By Lemma 3.18, $S \simeq \text{cl } S_\delta$ for any $0 \leq \delta < 1$ and and by Lemmas 5.1 and 5.2 all the requirements of Theorem 4.20 are met which entails that $\mathcal{C}_P \simeq \text{cl } S_\delta$ for $2\varepsilon^2 \leq \delta \leq 10\varepsilon^2$. Therefore we have

$$M(S) \simeq S \simeq \text{cl } S_\delta \simeq \mathcal{C}_P. \quad \blacksquare$$

5.2.2 Extending the Core

As we mentioned in the beginning of this chapter, one of the most appealing properties of the core is its flexibility for being used with other medial axis approximation algorithms. The following theorem makes this precise.

Theorem 5.4 *Under the the assumptions of Theorem 5.3, let $W \subset S_{2\varepsilon^2}$ be any set of points. Then $\mathcal{C}_P \cup \phi_P(W)$ is homotopy equivalent to S , where $\phi_P(W)$ is the flow-closure of W under ϕ_P .*

Proof. By Lemma 5.2 all points in $\text{cl } S_{2\varepsilon^2}$ including those in W flow into \mathcal{C}_P in finite time. Also, by Lemma 3.17, $\phi_P(W) \subseteq \text{cl } S_{2\varepsilon^2}$. On the other hand, by definition $\mathcal{C}_P \cup \phi_P(W)$ is flow-tight for ϕ_P . It can be easily verified that with these observations, $\mathcal{C}_P \cup \phi_P(W)$ can replace \mathcal{C}_P in the proof of Theorem 5.3 with no change in the argument. \blacksquare

Computing the flow closure of all the points in W can be computationally difficult depending on the nature of W . However, If W is a sub-complex of the Voronoi complex $\text{Vor } P$ (this is for example the case in the algorithm of Dey and Zhao [37]), computing the flow closure of all the points in W can be done in bulk by computing the flow closures of a whole face at a time.

5.3 Geometric Approximation

Although Theorem 3.3 ensures that the medial axis critical points lie very close to the medial axis, it provides no guarantee for the closeness of medial axis and the paths connecting medial axis critical points on their unstable manifolds. The same concern exists when the core is extended with a set of points close the medial axis; to guarantee the topology one must include the flow closures of added points but it is not clear that this closure stays close to the medial axis as well.

Notation. For the rest of this chapter it is convenient to use a handier notation than $\sqrt{h_P(x)}$ for the distance of a point x to the set P . We thus define

$$\rho(x) = \sqrt{h_P(x)} = \text{dist}(x, P).$$

For a point x in the core \mathcal{C}_P or any other approximation of the medial axis, the *relative approximation error* at x can be considered to be the ratio between $\text{dist}(x, M(S))$ and $\rho(x)$. Ideally, we would like that $\text{dist}(x, M(S))$ grew linearly (or slower) as a function of $\rho(x)$, when x traced a flow orbit. This would correspond to maintaining a constant relative error throughout the flow path. However, proving this, if it is true at all, appears elusive. Nevertheless, in this section we show that if we start from a point x close to the medial axis and follow the flow line $\phi_P(x)$, the distance to $M(S)$ along this path grows as a function of ρ at a rate that is slightly super-linear at worst. More precisely we will show that if we scale the space so that $\rho(x)$ has unit length, then if x has a medial axis point within distance $O(\sqrt{\varepsilon})$, every point y in the flow line starting at x will have a medial axis point within distance $O(\sqrt{\varepsilon})\rho(y)^{1+O(\sqrt{\varepsilon})}$.

Proposition 5.5 *For every point $x \in S$ and for every $p \in A_P(x)$,*

$$\|x - p\|^2 \leq s(x)^2 + \varepsilon^2 f(\hat{x})^2 + \varepsilon^2 s(x) f(\hat{x}). \quad (5.1)$$

Proof. Let $\lambda = \|x - \hat{x}\|/f(\hat{x})$. By Lemma 3.2 we have

$$\|x - p\|^2 = (\rho(x))^2 \leq (\lambda^2 + \varepsilon^2(1 + \lambda))f(\hat{x})^2.$$

replacing $\|x - \hat{x}\|$ with $s(x)$ gives us Equation (5.1). ■

The *driving angle* of a point x with respect to the sample P is defined as

$$\theta_P(x) = \angle(-v_P(x), p - x)$$

for any $p \in A_P(x)$. A critical point has driving angle equal to $\pi/2$ while points in the interior of a full-dimensional Voronoi cell V_p of $\text{Vor } P$ have driving angles equal to 0.

The following lemma generalizes Theorem 3.3 although for simplicity it aims for a weaker bound.

Lemma 5.6 *Let x be a point with driving angle $0 < \theta \leq \pi/2$.*

$$\min\{\|x - \hat{x}\|, \|x - \check{x}\|\} \leq \frac{2\varepsilon}{1 - \cos \theta} \cdot \mu(x).$$

Proof. Let $\lambda = \|x - \check{x}\|/\mu(x)$. The ball $B = B(\check{x}, \mu(x))$ contains no sample points. On the other hand, there is a sample point within distance $\varepsilon\mu(x)$ from \hat{x} . By triangle inequality, this implies that there is a sample point within distance $(1 - \lambda + \varepsilon)\mu(x)$ from x . Consider the ball $B' = B(x, (1 - \lambda + \varepsilon)\mu(x))$. All sample points in $A_P(x)$ are contained in $B' \setminus B$. This implies that the angle $\angle(-v(x), y - x)$ is at least θ for any $y \in \partial B' \cap \partial B$. Thus using the cosine rule on the triangle $\check{x}xy$ for any such y , we get the following inequality.

$$\mu(x)^2 \leq \lambda^2 \mu(x)^2 + (1 - \lambda + \varepsilon)^2 \mu(x)^2 + 2\lambda\mu(x)(1 - \lambda + \varepsilon)\mu(x) \cos \theta,$$

or equivalently

$$2(1 - \cos \theta)\lambda^2 - 2(1 + \varepsilon)(1 - \cos \theta)\lambda + \varepsilon(2 + \varepsilon) \geq 0.$$

Solving this inequality for λ we conclude that either,

$$\begin{aligned} \lambda &\leq \frac{1 + \varepsilon}{2} \left(1 - \sqrt{1 - \frac{2\varepsilon(2 + \varepsilon)}{(1 - \cos \theta)(1 + \varepsilon)^2}} \right) \\ &\leq \frac{1 + \varepsilon}{2} \cdot \frac{2\varepsilon(2 + \varepsilon)}{(1 - \cos \theta)(1 + \varepsilon)^2} \\ &\leq \frac{2\varepsilon}{1 - \cos \theta}, \end{aligned}$$

or

$$\begin{aligned}
\lambda &\geq \frac{1+\varepsilon}{2} \left(1 + \sqrt{1 - \frac{2\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)^2}} \right) \\
&\geq \frac{1+\varepsilon}{2} \left(2 - \frac{2\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)^2} \right) \\
&\geq 1 + \varepsilon - \frac{\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)} \\
&\geq 1 - \frac{2\varepsilon}{1-\cos\theta},
\end{aligned}$$

■

Lemma 5.7 *Let $B = B(x, R)$ be a ball empty of sample points with at least one sample point on its boundary and containing at least one medial axis point. Then the ball $B(x, (1 - 4\varepsilon^2)R)$ does not intersect Σ .*

Proof. Take any point $y \in B \cap \Sigma$. Since B intersects the medial axis, $f(y) \leq 2R$. Thus there is a sample point within distance $\varepsilon f(y) \leq 2\varepsilon R$ from y . Since B contains no sample points, y must be within distance $2\varepsilon R$ from ∂B . We grow a ball B' centered at x until its boundary touches Σ . Let R' be the radius of B' . By the above argument $R' \geq (1 - 2\varepsilon)R$. Let y be a point of tangency of B' and Σ . As indicated above, $f(y) \leq 2R$. Let B_o be the tangent ball of radius $f(y)$ at the opposite side of Σ with respect to x . With an argument similar to that of Proposition 5.5 we get for $s(x) = \|x - y\|$,

$$R^2 \leq s(x)^2 + \varepsilon^2 f(y)^2 + \varepsilon^2 s(x) f(y).$$

Using $f(y) \leq 2R$ we get

$$R^2 \leq s(x)^2 + 4\varepsilon^2 R^2 + 2\varepsilon^2 R s(x),$$

which by rearranging gives the following quadratic inequality for $s(x)$:

$$s(x)^2 + 2\varepsilon^2 R s(x) - (1 - 4\varepsilon^2)R^2 \geq 0.$$

Since $s(x) \geq (1 - 2\varepsilon)R$, the only valid range for $s(x)$ in the above inequality is

$$s(x) \geq \varepsilon^2 R^2 + \sqrt{\varepsilon^4 R^2 + (1 - 4\varepsilon^2)R^2} \geq (1 - 4\varepsilon^2)R.$$

■

Recall that the driver $d_P(x)$ is the same for all points x of the same Voronoi face in $\text{Vor } P$. In fact the flow orbit $\phi_P(x)$ turns exactly whenever it reaches the relative interior of a new Voronoi face. Consider a point $x \in S$ that lies on the $(n - 1)$ -skeleton of $\text{Vor } P$, i.e., $|A_P(x)| \geq 2$, and consider a line segment L in the flow orbit $\phi_P(x)$. The distance $\rho(x)$ monotonically increases along a flow orbit. We can therefore parameterize this line segment using the distance to the sample set P . Let $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be a non-decreasing real valued differentiable function and consider for each point $y \in L$, the ball $B_g(y) = B(y, g(\rho(y)))$.

We define the set $D(y)$ as those points on the boundary of $B_g(y)$ that are left outside $B_g(y)$ when y moves infinitesimally in the direction of $v_P(y)$. In other words if we take $\eta > 0$ such that $d_P(\phi_P(\eta, y)) = d_P(y)$ (meaning that $\phi_P(\eta, y)$ is on the same line segment of $\phi_P(x)$ as y), then

$$D(y) = \bigcap_{0 < \alpha < \eta} (\partial B_g(y) \setminus \text{cl } B_g(\phi_P(\alpha, y))).$$

Lemma 5.8 $D(x)$ consists of those points $y \in \partial B_g(x)$ that satisfy $\angle(y - x, d_P(x) - x) \leq \psi_0$ where $\cos \psi_0 = dg/d\rho$ at x .

Proof. Let $x' = \phi_P(t, x)$ be a point on L , satisfying $\|x - x'\| = \tau$ where τ is infinitesimally small. By definition, $A_P(x') = A_P(x)$. Let y be any point at distance $g(\rho(x))$ from x making an angle of ψ with $d_P(x) - x$. We have for the distance of y to x' :

$$\begin{aligned} \|y - x'\|^2 &= \|y - x\|^2 + \tau^2 + 2\tau\|x - y\| \cos \psi \\ &= g(\rho(x))^2 + \tau^2 + 2\tau g(\rho(x)) \cos \psi. \end{aligned}$$

For y not to be contained in $\text{cl } B_g(x')$ it must hold that $\|y - x'\| > g(\rho(x'))$, or equivalently:

$$g(\rho(x))^2 + \tau^2 + 2\tau g(\rho(x)) \cos \psi > g(\rho(x'))^2.$$

By rearranging we get

$$\tau + 2g(\rho(x)) \cos \psi > \frac{g(\rho(x) + \tau)^2 - g(\rho(x))^2}{\tau}.$$

In the limit when $\tau \rightarrow 0$ we get

$$\lim_{\tau \rightarrow 0} \tau + 2g(\rho(x)) \cos \psi \geq \lim_{\tau \rightarrow 0} \frac{g(\rho(x) + \tau)^2 - g(\rho(x))^2}{\tau},$$

which gives us

$$\begin{aligned} 2g(\rho(x)) \cos \psi &> \frac{d}{d\rho} (g(\rho(x)))^2 \\ &= 2g(\rho(x)) \frac{d}{d\rho} g(\rho(x)). \end{aligned}$$

Thus $\cos \psi > \cos \psi_0$ or $\psi < \psi_0$. ■

The main result of this section is stated in Theorem 5.10. Before we prove this theorem we state some known results about a very related flow map ϕ that is induced by the distance function $s(\cdot)$ which assigns to each point in space its distance (and not the squared distance) to Σ . The gradient ∇s of s is defined only on $S \setminus M(S)$. Even though Σ is a not finite, it is possible to define a vector field v that generalizes ∇s to all of S . In fact this can be done in a manner rather similar to the way ∇h_P for finite sets P is generalized into vector field v_P . It is however, important to note that $s(\cdot) = \text{dist}(\cdot, \Sigma)$ is the actual distance to Σ and not the

square of it. In particular, the gradient $\nabla s(x)$ is a unit vector at every point $x \in S \setminus M(S)$ where the gradient is defined.

Recall that for any $x \in \Sigma$, $A(x)$ represents the set of closest points to x in Σ . For a point $x \in S$ define $d(x)$ as the center of the smallest enclosing ball of $A(x)$. The point $d(x)$ is the *driver* of point x . The steepest ascent vector field $v : S \rightarrow \mathbb{R}^n$ (which generalizes ∇s) is defined as

$$v(x) = \frac{x - d(x)}{s(x)}.$$

Observe that for points $x \in S \setminus M(S)$ where x has a unique closest point \hat{x} in Σ and therefore $A(x) = \{\hat{x}\}$, the smallest enclosing ball of $A(x)$ has radius zero and therefore $d(x) = \hat{x}$. Thus we get

$$\|v(x)\| = \frac{\|x - \hat{x}\|}{s(x)} = 1,$$

as expected.

It was proven by Lieutier [56] that the vector field $v(x)$ can be integrate on S to give a flow $\phi : \mathbb{R}^+ \times S \rightarrow S$. In fact the map $t \mapsto v(\phi(t, x))$ is shown to be the right derivative of $x \mapsto \phi(t, x)$ and therefore at any point $x \in S$, $v(x)$ is tangent to the curve $t \mapsto \phi(t, x)$.

In particular Lieutier proved the following lemma which we shall use in the proof of Theorem 5.10.

Lemma 5.9 [56] *If x is a point in $M(S)$, then $\phi(t, x) \in M(S)$ for all $t \geq 0$. In other words $M(S)$ is flow-tight for the flow ϕ .*

Theorem 5.10 *Let $x_0 \in S$ be a point with $|A_P(x)| \geq 2$ and $\rho(x_0) = \rho_0$, and let $x_1 = \phi_P(t_1, x_0)$ be such that for all $0 \leq t \leq t_1$, $\cos(\theta_P(\phi_P(t, x_0))) \geq c$. If there is a medial axis point within distance g_0 from x_0 , then there is a medial axis point within distance $g(x_1)$ where*

$$g(x) = g_0 \left(\frac{\rho(x)}{\rho_0} \right)^\xi$$

from x_1 , provided that

$$\xi \geq \frac{1}{c} \left(1 + 4\varepsilon^2 \frac{\rho_0^2}{g_0^2} \right).$$

Proof. We prove the theorem by showing that as x moves along the flow line, there always remains a medial axis point within distance $g_0(\rho(x)/\rho(x_0))^\xi$ from x . We do this by showing that this proposition is maintained when x moves infinitesimally along the flow line. To this end, we first recall that that the ball $B_g(x)$ is an open ball by definition. If $B_g(x)$ contains a medial axis point z , then for any direction vector v , there is a small enough real number $\tau > 0$, such that the translated ball $B_g(x) + \tau v = \{y + \tau v \mid y \in B_g(x)\}$ contains z as well. Since g is increasing, this implies that $z \in B_g(x + \tau v)$. In particular this implies by choosing $v = v_P(x)$ that if $B_g(x)$ contains a medial axis point, so does $B_g(x + \tau v_P(x))$ for τ sufficiently

small.

We thus only need to consider the case where $B_g(x)$ contains no medial axis point while its boundary does. Thus, let $z \in \partial B_g(x)$ be a medial axis point. By Lemma 5.8, if $\cos(\angle(z - x, -v_P(x))) \leq \frac{d}{d\rho}g(\rho(x))$, then z is contained in $B_g(x + \tau v_P(x))$ for a sufficiently small $\tau > 0$ and therefore we have nothing to prove. Thus we only need to consider the case when

$$\begin{aligned} \cos(\angle(z - x, -v_P(x))) &> \frac{d}{d\rho}(g(\rho(x))) \\ &= \frac{\xi g_0}{\rho(x_0)} \left(\frac{\rho(x)}{\rho(x_0)} \right)^{\xi-1}. \end{aligned}$$

We denote this maximum angle by ψ . We show that in this case, the flow vector $v(z)$ points to the interior of the ball $B_g(x)$. This implies that the flow $\phi(z)$ enters $B_g(x)$ and therefore by Lemma 5.9 $B_g(x)$ must contain a medial axis point, contradicting our choice of z .

The ball $B(x, \rho(x))$ contains no sample points but includes z , a medial axis point, and therefore by Lemma 5.7, the ball $B_0 = B(x, (1 - 4\varepsilon^2)\rho(x))$ does not intersect Σ . Consider the plane Π tangent to $\partial B_g(x)$ at z . This plane, intersects the ball B_0 in a disk of radius

$$R_0 \geq \sqrt{((1 - 4\varepsilon^2)\rho(x))^2 - g(\rho(x))^2}.$$

We will show that there are surface points at distance less than R_0 from z , i.e. $s(z) < R_0$. It can be easily observed that any such surface point must lie on the side of Π opposite to the side containing $B_g(x)$ and therefore $\text{conv}(A(x))$ resides on the side of Π opposite to $B_g(x)$. Since $d(x) \in \text{conv}(A(x))$, this implies that $v(x)$ points to the interior of $B_g(x)$, as desired. So, all left to show is that $s(z) < R_0$. To prove this, we show that for at least one of the points $y \in A_P(x)$, $\|z - y\| < R_0$. Since $|A_P(x)| \geq 2$, every plane containing the points x and $d_P(x)$ that does not intersect $A_P(x)$ must contain at least one point of $A_P(x)$ on each side. The maximum distance between z and $A_P(x)$ can therefore occur when $|A_P(x)| = 2$ and z is on the bisector plane of the segment connecting the two points in $A_P(x)$. To find this maximum distance we use a change of coordinates. We perform the calculations in 3D. However, they can be easily extended to higher dimensions. Denoting the three coordinate directions by u_1, u_2 , and u_3 , we put the origin at x , the driver $d_P(x)$ on the u_1 -axis, the two points p_1, p_2 in $A_P(x)$ on the u_1u_2 -plane, and z on the u_1u_3 -plane, see Figure 5.3. We can calculate the coordinates of p_1, p_2 , and z as follows:

$$\begin{aligned} p_1 &= (\rho \cos \theta, \rho \sin \theta, 0), \\ p_2 &= (\rho \cos \theta, -\rho \sin \theta, 0), \\ z &= (g(\rho) \cos \psi, 0, g(\rho) \sin \psi), \end{aligned}$$

where θ denotes the the driving angle $\theta_P(x)$ and ρ is the shorthand for $\rho(x)$. Thus we get




Figure 5.3: Proof of Theorem 5.10. The gray cap represents $D(x)$.

for the distance between z and p_1 (same distance between z and p_2):

$$\begin{aligned}\|z - p_1\|^2 &= (\rho \cos \theta - g(\rho) \cos \psi)^2 + (\rho \sin \theta)^2 + (g(\rho) \sin \psi)^2 \\ &= \rho^2 + g(\rho)^2 - 2\rho g(\rho) \cos \theta \cos \psi\end{aligned}$$

Denoting $\rho(x_0)$ by ρ_0 and $\rho(x)$ by ρ and using the lower bounds for $\cos \psi$ and $\cos \theta$ we get:

$$\begin{aligned}\|z - p_1\|^2 &< \rho^2 + g_0^2 \left(\frac{\rho}{\rho_0}\right)^{2\xi} - 2c\rho g_0 \left(\frac{\rho}{\rho_0}\right)^\xi \frac{\xi g_0}{\rho_0} \left(\frac{\rho}{\rho_0}\right)^{\xi-1} \\ &= \rho^2 + (1 - 2c\xi)g_0^2 \left(\frac{\rho}{\rho_0}\right)^{2\xi}\end{aligned}$$

Thus in order for $\|z - p_1\| < R_0$, it suffices to have

$$\rho^2 + (1 - 2c\xi)g_0^2 \left(\frac{\rho}{\rho_0}\right)^{2\xi} \leq (1 - 4\varepsilon^2)\rho^2 - g_0^2 \left(\frac{\rho}{\rho_0}\right)^{2\xi}.$$

Since $(1 - 4\varepsilon^2)^2 > 1 - 8\varepsilon^2$, the above inequality is satisfied when the following one is:

$$4\varepsilon^2\rho^2 \leq (c\xi - 1)g_0^2 \left(\frac{\rho}{\rho_0}\right)^{2\xi},$$

or identically,

$$\frac{4\varepsilon^2}{c\xi - 1} \leq \frac{g_0^2}{\rho_0^2} \left(\frac{\rho}{\rho_0}\right)^{2\xi-2}.$$

Since $\rho/\rho_0 \geq 1$, the above inequality holds if

$$g_0 \geq \left(\frac{2\varepsilon}{\sqrt{c\xi - 1}}\right)\rho_0.$$

This is guaranteed by the bound on ξ prescribed in the statement of the Theorem. ■

Remark It may appear at first that when g_0 goes to zero the above theorem must guarantee a tighter bound. However, the reader must notice that g_0 also appears in the denominator of the exponent of the given bound. As a result when the flow line starting at a given x_0 is followed the best bound is not necessarily obtained by using the medial axis point nearest to x_0 . In other words, a larger g_0 may lead to a better bound on g .

Corollary 5.11 *Let x be a point in the $(n - 1)$ -skeleton of $\text{Vor } P$ such that*

$$\|x - \check{x}\| \leq \frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}\rho(x),$$

then for every point y on the flow path $\phi_P(x)$, there is a medial axis point within distance

$$\frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}\rho(x) \left(\frac{\rho(y)}{\rho(x)} \right)^\xi$$

from y , where $\xi = 1 + O(\sqrt{\varepsilon})$.

Proof. We first consider the case where y has a driving angle $\theta = \theta_P(y)$ with $\cos \theta \leq 1 - \sqrt{\varepsilon}$. By Lemma 5.6, $\|y - \check{y}\| \leq 2\sqrt{\varepsilon}\mu(y)$, or equivalently, $h_S(y) \geq (1 - 2\sqrt{\varepsilon})\mu(y)$. Since $\rho(y) \geq h_S(y)$, we get

$$\|y - \check{y}\| \leq 2\sqrt{\varepsilon}\mu(y) \leq \frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}h_S(y) \leq \frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}\rho(y).$$

When the cosine of the medial angle grows above $1 - \sqrt{\varepsilon}$ in a point y along the flow line $\phi_P(x)$, by Theorem 5.10 there always is a point within distance

$$\frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}\rho(x) \left(\frac{\rho(y)}{\rho(x)} \right)^\xi,$$

from y , where

$$\xi = \frac{1}{1 - \sqrt{\varepsilon}} \left(1 + 4\varepsilon^2 \frac{(1 - 2\sqrt{\varepsilon})^2}{(2\sqrt{\varepsilon})^2} \right) \leq 1 + O(\sqrt{\varepsilon}).$$

■

An immediate consequence of this corollary is that the core and the flow closures converge to being contained in the medial axis as $\varepsilon \rightarrow 0$. As a result, when the core is used and extended using the filtering conditions of [37], the computed approximate medial axis converges to the true medial axis in the limit.




Figure 5.4: Left: Core computed for the 3-holes model. The red lines are either unstable manifolds of index-2 saddle points or the one dimensional parts (hairs) of index-1 saddle points. Middle: Filtered Voronoi facets based on a condition similar to the angle condition of Dey and Zhao. Right: Extended core, i.e., the core, plus the flow closures of the facets from the middle picture.

5.4 Experiments

Figure 5.4 demonstrates a result obtained with an implementation of the unstable flow complex data structure, the core, and the extension of the core using Voronoi facets filtered using conditions similar to those in [37]. As can be observed, the flow closure has filled the holes in the junctions of the geometric approximation of the medial axis with filtered Voronoi facets.

Bibliography and Remarks.

As a geometric object the medial axis is unstable since very small changes in the boundary of the shape can cause comparatively enormous changes in its medial axis. This instability of the medial axis bears two consequences. First, it makes the medial axis hard to compute exactly due to numerical instabilities and consequently, exact computation of medial axis has only been attempted for a small number of limited classes of shapes (See for example [28]). Second, the complete medial axis may be less interesting than an approximation of it that is more stable under small geometric changes of the shape and still carries the same topological type. Chazal and Lieutier [22] defined the λ -medial axis, which is a subset of the medial axis made by trimming out the points for which all closest surface points are contained in a ball of radius λ . The λ -medial axis has the desired properties of ignoring small features of the surface that can cause major changes in the medial axis, and is therefore stable to small perturbations of the surface, but is guaranteed to have the same homotopy type as the medial axis only for suitably small values of λ . The largest topologically safe λ is the smallest

critical value of the distance function s induced by Σ . This critical value can of course be which can be very small, leaving very little room for filtering. It is also difficult to compute algorithmically [7].

Like the surface reconstruction problem, approximations of the medial axis of a shape is often desired at the presence of a sample of the boundary of the shape. Chazal and Lieutier [22] gave a simple algorithm that approximates the λ -medial axis of a shape as a sub-complex of the Voronoi complex of a given (uniform) noisy ε -sample of the boundary of the shape.

In the relative ε -sampling framework, Amenta and Bern [2], and Boissonnat and Cazals [13] established that a subset of the Voronoi vertices in the Voronoi diagram of the sample points lie close to the medial axis. Later Amenta, Choi and Kolluri [5] designed an algorithm that computes an approximation of the medial axis as a cell complex, which is homotopy equivalent to the medial axis of the shape provided it has a smooth boundary and the sample is dense enough. The output of this algorithm tends to be noisy and has to be cleaned up by heuristics. Furthermore the algorithm needs the computation of the Voronoi diagram of a set that is more than twice as large as the set of sample points. Dey and Zhao [37] addressed these shortcomings by designing an algorithm that approximates a subset of the medial axis of a sampled surface by a sub-complex of the Voronoi diagram of the sample. This often geometrically pleasing output is guaranteed to converge in relative distance to the medial axis of the sampled shape when the sample grows infinitely dense (i.e. when ε approaches zero) but suffers the lack of any topological guarantees. In essence, this algorithm uses two local conditions, known as the *ratio* condition and the *angle condition*, for filtering out the “interesting” Voronoi facets. The authors prove that as the sampling density ε goes to zero, the selected Voronoi facets converge to the true medial axis of the target surface. In fact, this happens to hold for the ratio condition alone. However, using the combination of ratio and angle conditions, and using properly chosen parameters, the algorithm succeeds to produce an output that is highly satisfactory from a practical point of view. In fact, with a well chosen set of parameters, this algorithm may sometimes succeeds to capture the topology of the medial axis of the target surface.

In [7], Attali, Boissonnat, and Edelsbrunner excellently cover various aspects of the medial axis approximation problem and describes the existing practical and theoretical challenges involved in this problem and survey the state of the art in this area.

Chapter 6

The Wrap Reconstruction Algorithm

In [38], Herbert Edelsbrunner described an algorithm for solving the surface reconstruction problem by a method he described casually as “wrapping finite sets in space”. This algorithm, to which we generically refer as “WRAP”, is successfully implemented and commercialized (as Raindrop Geomagic WRAP¹). The WRAP algorithm does not make assumptions on the input data and therefore can only make very general statements about the type of output. Specifically, for the main variant of his algorithm, the output is shown to always be a *pinched sphere*. The significance of WRAP as argued by Edelsbrunner, is that it is based on discrete methods that are inspired by concepts in continuous mathematics, such as Morse functions and gradient fields; in fact, WRAP was arguably the first algorithm that aimed at finding a solid mathematical basis for addressing the shape reconstruction problem in a systematic way through flow maps and related Morse theoretic concepts.

Edelsbrunner’s approach employs the distance function h_Q induced by a set Q of weighted points. The set Q is defined as the set of Voronoi vertices in $\text{Vor } P$ with a positive weight w_q assigned to each point in Q . For a point $q \in Q$, w_q is the square of the distance from q to its closest point in P , i.e. the circumradius of the Delaunay tetrahedron dual to q^2 . Observe that this is exactly the weighted point-set Q we constructed through polarity (Section 1.4.4) whose Voronoi complex is “similar” to the Delaunay complex of P . In addition, the algorithm adds to Q a symbolic critical point (a minimum) at infinity. This point is assumed to have infinite weight. This also happens to be consistent with the idea of polarity as described in Section 1.4.4.

The WRAP reconstruction algorithm essentially works by “approximating” the unstable manifolds of critical points of h_Q by subcomplexes of the Delaunay complex $\text{Del } P$ of the input sample P . In fact, the original WRAP algorithm is only concerned about approximating the unstable manifold of the minimum at infinity by a collection of Delaunay simplices (added to the complement of $\text{conv } P$). Then by removing this collection it leaves an output which is a Delaunay subcomplex.

It turns out that the separation of critical points (Theorem 3.3) is identically valid for the critical points of h_Q . This allows us to modify WRAP into an algorithm for approximation

¹<http://www.geomagic.com/en/products/studio/modules/wrap/>

²Edelsbrunner actually considers $-h_Q$. For technical reasons, we prefer to work with h_Q ; in particular, this makes our approach consistent with the previous Chapters in this thesis.

of the volume S bounded by Σ for which geometric and topological guarantees can be established. More precisely, the output of this modified WRAP turns out to differ from S only within a thin tubular neighborhood of Σ and can be shown to have the same homotopy type as S . Proving these guarantees relies heavily on our understanding of the geometric behavior of the flow map ϕ_Q .

6.1 Description of WRAP

Our presentation of the algorithm differs slightly from that of Edelsbrunner and this is due to our description's reliance on integral lines of the flow map ϕ_Q induced by Q which correspond to *limit curves* studied by Edelsbrunner only with opposite orientation.

As usual, a target surface Σ in \mathbb{R}^n is known through a finite ε -sample $P \subset \Sigma$. We define a set Q of weighted points consisting of the Voronoi vertices in $\text{Vor } P$ with every vertex $q \in Q$ given the weight $w_q = h_P(q)$, i.e. the square of the distance between q and its closest points in P . Moreover, we denote by Q^+ the set Q with addition of a symbolic point of infinite weight at infinity.

As we saw in Section 1.4.4, with this definition of Q^+ , we can write

$$\text{Vor } P = \text{Del } Q^+ \quad \text{and} \quad \text{Del } P = \text{Vor } Q^+.$$

Notice the above equalities still hold when Q^+ is replaced with Q if we restricted ourselves to the convex hulls of the sets whose Delaunay complex is being considered. In other words, if we clip Voronoi faces of $\text{Vor } P$ by intersecting them with $\text{conv } Q$, $\text{Vor } P = \text{Del } Q$ and likewise when restricted to $\text{conv } P$, $\text{Del } P = \text{Vor } Q$ and in this sense the point-sets P (unweighted) and Q (weighted), as well as their Voronoi and Delaunay complexes are *dual* to each other.

An immediate consequence of this duality is that the critical points of h_Q are exactly the same as those of h_P . We call those simplices of $\text{Del } P$ that contain a critical point, i.e. those that intersect their dual faces in $\text{Vor } P$, *critical simplices* (or *centered* in Edelsbrunner's terminology).

Recall that the flow orbit $\phi_Q(x)$ of any point $x \in \mathbb{R}^n$ is a piece-wise linear curve that does not turn in the relative interior of a Voronoi face of $\text{Vor } Q$. Thus, the intersection of a flow orbit of ϕ_Q and a full-dimensional simplex of $\text{Del } P = \text{Vor } Q$ is a line segment (if not empty). With lower dimensional simplices there is a second possibility, namely, a flow line can *cross* the simplex and thus intersect it in a single point. In such a case, the simplex is called *transversal* (or *equivocal* according to Edelsbrunner) relative to the flow map ϕ_Q . Of course, the flow line can just as well intersect a non-full-dimensional simplex in a line segment in which case we say that the flow is *tangential* on the simplex in question or simply call the simplex *tangential* with respect to ϕ_Q (Edelsbrunner calls such simplices *confident*). Notice

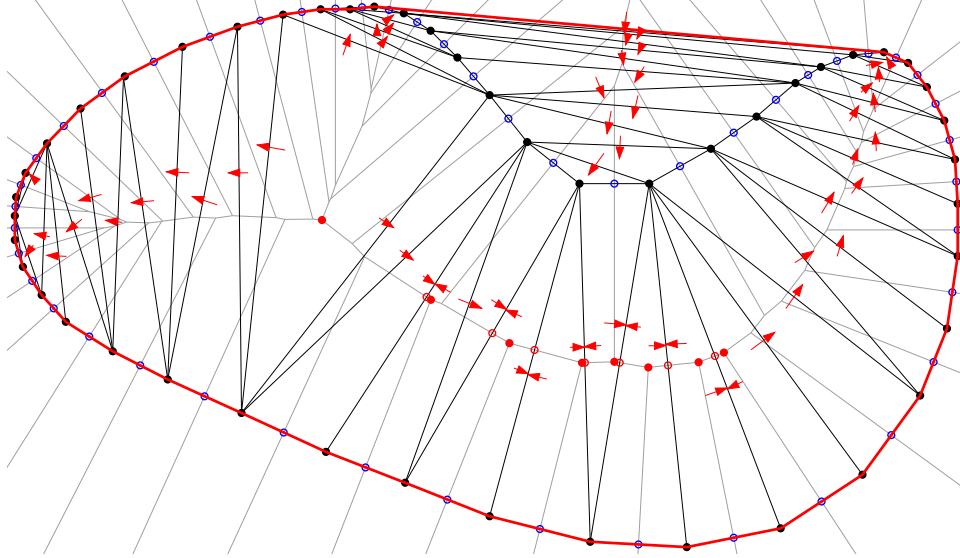


Figure 6.1: Precedence relation between Delaunay simplices is shown by arrows. An arrow crossing an edge signifies that the flow on that edge is transversal while arrows ending on an edge show that the flow on the edge is tangential.

Algorithm (original)WRAP(sample point-set P)

- 1 Let $\Delta \subseteq D$ be the set of *critical* simplices.
 - 2 Let $O = \{\tau \in D : \omega \trianglelefteq \tau \text{ and } \forall \sigma \in \Delta : \sigma \not\trianglelefteq \tau\}$.
 - 3 Return $I = D \setminus O$.
-

Figure 6.2: The original WRAP algorithm.

that critical simplices are in fact tangential simplices that contain a critical point.

We say a simplex τ precedes a simplex σ and denote it by $\tau \prec \sigma$ if τ and σ are incident simplices, i.e. τ is either a face or a coface of σ , and some flow line of ϕ_Q enters the relative interiors of σ immediately after leaving the relative interior of τ . More formally, when τ is a coface of σ , $\tau \prec \sigma$ if there exists a point x and a time $t_0 > 0$ and a real number $0 < \alpha < t_0$ such that $\phi_Q(t_0, x) \in \text{rel int } \tau$ and $\phi_Q(t, x) \in \text{rel int } \sigma$ for $\alpha < t < t_0$. Similarly, when τ is a proper face of σ , $\tau \prec \sigma$ if for some point x there exist time $t_0 > 0$ and real $\alpha > t_0$ such that $\phi_Q(t_0, x) \in \text{rel int } \tau$ and $\phi_Q(t, x) \in \text{rel int } \sigma$ for $t_0 < t < \alpha$. We define the relation “ \preceq ” as the reflexive transitive closure of “ \prec ”, namely, $\tau \preceq \sigma$ if there is a sequence $\tau = \tau_0 \prec \cdots \prec \tau_k = \sigma$ with $k \geq 0$. Figure 6.1 shows the precedence relation between Delaunay simplices.




Figure 6.3: An two-dimensional example of the execution of WRAP. The removed Delaunay simplices are shown in red.

Remark Edelsbrunner’s definition of the precedence relation, which we denote by “ \triangleleft ” is slightly different from ours, in that $\tau \triangleleft \sigma$ if $\tau \prec \sigma$, and in addition, the flow on one of τ or σ is transversal. This definition thus invalidates $\tau \triangleleft \sigma$ in the case where $\tau \prec \sigma$ but the flow is tangential in τ and reaches the face σ of τ and continues tangentially on σ . Note that no other case is possible; flow cannot cross two incident simplices transversally and cannot move tangentially from a face to a coface. The reflexive transitive closure of “ \triangleleft ” is denoted by “ \trianglelefteq ”.

Some subtlety is associated with having a point (a minimum) at infinity in Q^+ . In practice this point has no role in the computation of the Delaunay triangulation of Q . However, because of it, $\mathbb{R}^n \setminus \text{conv } Q$ become the Voronoi cell in $\text{Vor } Q^+$ associated to the critical point (minimum) at infinity. Consequently, the WRAP algorithm must treat $\mathbb{R}^n \setminus \text{conv } P$ as a special *abstract* critical simplex ω that contains this critical point. Since this critical point is infinitely far away, every simplex τ of $\text{Del } P$ that is contained in the boundary of $\text{conv } P$ is considered preceded by ω .

With these preliminaries covered, Edelsbrunner’s WRAP algorithm can now be stated as shown in Figure 6.2.

As stated, the output of WRAP is not guaranteed to agree topologically with the sampled surface Σ . In fact, Edelsbrunner proves that the produced output I is the boundary of a contractible volume. However, Edelsbrunner also suggests methods for extending the algorithm in order to allow production of non-contractible output. For example, he suggests to consider, in addition to the simplices that are preceded by ω , those that are preceded by other “significant critical simplices”. The results in rest of this chapter are witness to the accuracy of this insightful intuition; the other “significant critical simplices” turn out to be critical simplices associated to outer medial axis critical points of h_Q .




Figure 6.4: Extension of the sample points for simulation of a critical point at infinity. Left: Delaunay triangulation of a set of points. Right: the point-set on the left is enclosed in a large enough ball and the boundary of the ball is sampled. The Delaunay triangulation of the original point-set is a subcomplex of the Delaunay triangulation of the extended one.

In Figure 6.5, we present a modified version of WRAP which can capture the topology of Σ or rather the bounded volume S enclosed by it. As mentioned above, the modification rests primarily on separation of critical points which allows us in particular to filter out the so-called *surface critical points* which are in essence (at least from a topological standpoint) the *artifacts of discretization* of the surface. Essentially, our algorithm amends WRAP by adding to ω all other outer medial axis critical simplices.

The rest of this chapter is dedicated to proving that this modified version of WRAP produces an output that is geometrically close to S and has the same homotopy type as S provided that the input P to the algorithm is an ε -sample of Σ for a sufficiently small value of ε . In the rest of this chapter WRAP refers to this modified version.

6.2 Simulating the Minimum at Infinity

Before we proceed we address a technical difficulty related to the assumption of a minimum at infinity. Our proofs heavily rely on the continuity of the flow map ϕ_Q induced by the set Q of weighted Voronoi vertices in $\text{Vor } P$. Per se, the argument given for continuity of ϕ_Q in Section 2.3.1 has no way of handling “a symbolic point of infinite weight at infinity”. However, we cannot ignore this symbolic point either. Since ϕ_Q is a steepest ascent flow, and because h_Q without the extra critical point at infinity, goes to infinity when the distance to $\text{conv } P$ grows infinitely large, the flow ϕ_Q escapes $\text{conv } P$ outside of it. This is in contrast to the behavior h_{Q^+} suggests in which the steepest ascent direction outside the convex hull of P is toward $\text{conv } P$.

Several of our proofs in what follows are interested in the behavior of ϕ_{Q^+} on the outer boundary of Σ_δ . These proofs are based on the assumption that every point on this boundary

is contained in some Delaunay n -simplex. This assumption trivially fails when the point in question is outside $\text{conv } P$. We handle these problem by simulating the behavior of the critical point at infinity through adding a component Σ_0 to the surface Σ in such a way that the local feature size of every point in Σ with respect to $M(\Sigma \cup \Sigma_0)$ is the same as its local feature size with respect to $M(\Sigma)$ alone. This will then imply that if P_0 is an ε -sample of Σ_0 , $P_1 = P \cup P_0$ is an ε -sample of $\Sigma_1 = \Sigma \cup \Sigma_0$. Moreover, the component of the tubular neighborhood Σ_δ which is the thickening of Σ in the two-component surface $\Sigma \cup \Sigma_0$ will be identical to the one for the surface Σ alone.

Thus consider a ball $B = B(c, R)$ enclosing Σ and therefore the sample $P \subset \Sigma$. Let R_D be the circumradius of the largest Delaunay ball in $\text{Del } P$. Consider the ball $B_0 = B(c, R_0)$ where

$$R_0 = 2R + 2 \max\{R_D, F\}$$

in which $F = \sup_{x \in \Sigma} f(x)$. let $\Sigma_0 = \partial B_0$ be the new component of the surface and let P_0 be an ε -sample of it. Since $R_0 \geq R + 2R_D$, every Delaunay ball of $\text{Del } P$ is entirely contained in the interior of B_0 and therefore remains empty of the points in P_0 . Consequently, $\text{Del } P$ is a subcomplex of $\text{Del } P_1$ where $P_1 = P \cup P_0$ is the extended point-set which samples the two-component surface $\Sigma_1 = \Sigma \cup \Sigma_0$ (see Figure 6.4).

Let z be a point of $M_1 = M(\Sigma_1)$, i.e. the medial axis of the extended surface. If $z \in B(c, R_0/2)$, then every closest point of Σ_1 to z has to be in Σ and therefore $z \in M = M(\Sigma)$. Thus, for a point $x \in \Sigma$, the distance to M_1 is at least as large as the distance to M since the points outside $B(c, R_0/2)$ are at least $R + F$ distance away from x while x has a point within distance F in M . Thus $P_1 = P \cup P_0$ is a valid ε -sample of Σ_1 . This in particular ensures that the separation of critical point given by Theorem 3.3 remains valid for separation of critical points of h_{P_1} . We intend to use the set Q_1 of all Voronoi vertices in $\text{Vor } P_1$, weighted as above, instead of the Q solely for our analysis of the WRAP algorithm. Since the critical points of h_{P_1} and h_{Q_1} are the same, Theorem 3.3 implies that the critical points of h_{Q_1} are contained either in the ε^2 tubular neighborhood of Σ_1 , or near its medial axis.

Proposition 6.1 *The surface critical points of h_{Q_1} contained in Σ_{ε^2} are exactly the same as the (surface) critical points of h_Q contained in Σ_{ε^2} .*

Proof. Let $c \in \Sigma_{\varepsilon^2}$ be a critical point of h_{Q_1} . Since $c \in \Sigma_{\varepsilon^2}$, $\|c - \hat{c}\| \leq \varepsilon^2 f(\hat{c})$. By the ε -sampling condition, \hat{c} has a point $p \in P$ within distance $\varepsilon f(\hat{c})$. Thus

$$\text{dist}(c, P_1) \leq \text{dist}(c, P) \leq \varepsilon^2 f(\hat{c}) + \varepsilon f(\hat{c}) \leq 2\varepsilon F < 2\varepsilon R_0.$$

Thus $A_{P_1}(c) \subset P$. Since c is a critical point, $c \in \text{conv } A_{P_1}(c) \subset \text{conv } P$. Thus c is a critical point of h_P . ■

Algorithm (modified)WRAP(sample point-set P)

- 1 $C \leftarrow$ set of the critical points of h_P (or h_Q).
 - 2 $(C_M^+, C_M^-, C_\Sigma) \leftarrow \text{SEPARATE}(C)$.
 - 3 Let $\Delta \subset D$ be the set of critical simplices.
 - 4 Let Δ^* be the set of critical simplices corresponding to C_M^+ plus ω .
 - 5 Let $O = \{\tau \in D : \exists \sigma \in \Delta^*, \sigma \preceq \tau \text{ and } \forall \sigma \in \Delta \setminus \Delta^* : \sigma \not\preceq \tau\}$.
 - 6 Return $I = D \setminus O$.
-

Figure 6.5: The modified WRAP algorithm.

6.3 Flow on Tubular Neighborhoods

As mentioned before, our topological proofs hinge on using the continuous map ϕ_{Q_1} as a homotopy in Theorem 4.20. The main result of this section is that the tubular neighborhood Σ_δ of Σ is tight for the flow ϕ_{Q_1} when δ is in a suitable range.

Recall from Section 1.3 that the power of a point x with respect to a ball $B = B(c, R)$ is denoted by $\pi_B(x)$ and is defined as

$$\pi_B(x) = \|x - c\|^2 - R^2.$$

In other words, the power of a point x with respect to a ball of radius R centered at c is equivalent to the square of the distance between x and the point c with weight R^2 . Thus $\pi_B(x)$ is positive outside B , zero on ∂B and negative inside B .

The following proposition is a well-known result on the structure of the Delaunay complex. The proof is included for completeness.

Proposition 6.2 *Let P be a set of points in \mathbb{R}^n and let x be a point in a Delaunay n -simplex $\tau \in \text{Del } P$. Let B be the ball circumscribing τ and let B' be an arbitrary empty ball. Then $\pi_B(x) \leq \pi_{B'}(x)$. In other words, of all the empty balls, B is the one with respect to which the power of x is the smallest.*

Proof. Consider a mapping from $\mathbb{R}^n \rightarrow \mathbb{R}^{n+1}$ given by $x \mapsto x^*$ where

$$x^* = \left(x_1, \dots, x_n, \sum_{i=1}^n x_i^2 \right),$$

when $x = (x_1, \dots, x_n)$. In other words, every point is lifted to the unit paraboloid Π_0 with equation $x_{n+1} = x_1^2 + \dots + x_n^2$. It can be verified that lifted image of a sphere ρ in \mathbb{R}^n lies in a hyperplane H_ρ in \mathbb{R}^{n+1} and a sphere ρ in \mathbb{R}^n is empty of P if and only if H_ρ is below the




Figure 6.6: Illustration of the proof of Lemma 6.3.

image p^* of every $p \in P$. Furthermore, given a sphere ρ and a point x , $\pi_\rho(x)$ is the signed vertical distance between x^* and H_ρ . So, for a given point x , the empty ball B containing x that minimizes $\pi_B(x)$ must have $H_{\partial B}$ below every p^* where $p \in P$, but must vertically be as far away as possible above x^* . This makes $H_{\partial B}$ a supporting hyperplane of the lower hull of $P^* = \{p^* : p \in P\}$ and thus B must be a Delaunay ball. \blacksquare

Lemma 6.3 Consider an ε -sample of the surface Σ_1 with $\varepsilon \leq 1/10$. For a $0 < \delta < 1$, let x be a point on the boundary of Σ_δ , i.e. $\|x - \hat{x}\|/f(\hat{x}) = \delta$ with $\hat{x} \in \Sigma$. Let τ be a Delaunay n -simplex in $\text{Del } P_1$ that contains x and let c be the circumcenter of τ . If c is in the same side of Σ as x , then for the angle $\alpha = \angle(c - x, x - \hat{x})$, $\cos \alpha \geq 1 - \delta - 6\varepsilon$.

Proof. Let u and v be the two points on the line normal to Σ at \hat{x} and at distance $f(\hat{x})$ from it, with u being on the same side of Σ as x . The balls $B_u = B(u, f(\hat{x}))$ and $B_v = B(v, f(\hat{x}))$ are disjoint from Σ and their boundaries are tangent to it (see Figure 6.6).

Consider the line segment \overline{cv} . Since we assumed c is in the same side of Σ as x (and u), \overline{cv} intersects Σ in at least one point. Let z be an arbitrary point in $\overline{cv} \cap \Sigma$. Let us denote $\|c - v\|$ by d and $\|z - v\|$ by b . Notice that $d \geq b \geq f(\hat{x})$.

Since the local feature size function is 1-Lipschitz, using the triangle inequality we get

$$f(z) \leq \|z - \hat{x}\| + f(\hat{x}) \leq \|z - v\| + \|v - \hat{x}\| + f(\hat{x}) = b + 2f(\hat{x}).$$

By the ε -sampling condition, there must be a sample point in P within $\varepsilon f(z) \leq \varepsilon(b + 2f(\hat{x}))$ from z . The ball $B_c = B(c, R)$ circumscribing τ does not intersect P and therefore, $\|z - c\| + \varepsilon f(z) \geq R$, or else z does not meet the sampling condition. Using $\|z - c\| = d - b$ and

the above upper bound for $f(z)$, we get

$$\begin{aligned}
R &\leq d - b + \varepsilon(b + 2f(\hat{x})) \\
&= d - (1 - \varepsilon)b + 2\varepsilon f(\hat{x}) \\
&\leq d - (1 - \varepsilon)f(\hat{x}) + 2\varepsilon f(\hat{x}) \\
&= d - f(\hat{x}) + 3\varepsilon f(\hat{x}).
\end{aligned} \tag{6.1}$$

Consequently if B'_v denote the ball centered at v with radius $(1 - 3\varepsilon)f(\hat{x})$, then $B_c \cap B'_v = \emptyset$.

On the other, using Proposition 6.2 for Delaunay ball B and empty ball B_u we can write

$$\|x - c\|^2 - R^2 \leq (1 - \delta)^2 f(\hat{x})^2 - f(\hat{x})^2. \tag{6.2}$$

Using the cosine rule on the triangle cxv for angle $\angle(x - c, x - v) = \pi - \alpha$ we obtain

$$\cos(\pi - \alpha) = \frac{\|x - v\|^2 + \|x - c\|^2 - \|c - v\|^2}{2\|x - v\|\|x - c\|} = \frac{(1 + \delta)^2 f(\hat{x})^2 + \|x - c\|^2 - \|c - v\|^2}{2(1 + \delta)f(\hat{x})\|x - c\|}.$$

Combining this with inequalities (6.1) and (6.2), and defining $r = R/f(\hat{x})$ results

$$\cos \alpha \geq \frac{2r(1 - 3\varepsilon) - 6\varepsilon - 2\delta^2 + 9\varepsilon^2}{2(1 + \delta)\sqrt{r^2 + (1 - \delta)^2 - 1}}.$$

The right hand side of the of the above inequality is a function of r (taking ε and δ as constants) defined for $r^2 > 1 - (1 - \delta)^2$ (notice that $r = R/f(\hat{x})$ is always positive). It can be verified that its derivative has a unique root corresponding to a global minimum. Calculating the value of the function at this minimum gives us

$$\cos \alpha \geq \frac{-2\delta^2 - 6\varepsilon + 9\varepsilon^2 + \frac{4(1-3\varepsilon)(2-9\varepsilon+9\varepsilon^2)}{2+3\varepsilon}}{2(1 + \delta)\sqrt{-1 + (1 - \delta)^2 + \frac{4(2-9\varepsilon+9\varepsilon^2)^2}{(2+3\varepsilon)^2}}}$$

Elementary algebraic simplifications assuming that $\varepsilon \leq 1/10$ entails the statement of the Lemma. ■

Lemma 6.4 *Let P_1 be an ε -sample of the surface Σ_1 with $\varepsilon < 1/3$. Let x be a point on the boundary of Σ_δ for $\delta > 9\varepsilon^2$. Finally, let τ be a Delaunay n -simplex in $\text{Del } P_1$ containing x and let $B = B(c, R)$ be the ball enclosed by the circumsphere of τ . Then c and x are on the same side of Σ .*

Proof. Suppose to the contrary that x and c are in opposite sides of Σ . Let u be the point on normal to Σ at \hat{x} at the same side of Σ as x , satisfying $\|\hat{x} - u\| = f(\hat{x})$. Since c is on the opposite side of Σ from x (and u), the segment \overline{cu} must intersect Σ at a some point z . Let $d = \|c - u\|$ and $b = \|z - u\| \geq f(\hat{x})$. Using the 1-Lipschitzness of local feature size, we can

bound $f(z)$ as

$$f(z) \leq \|z - \hat{x}\| + f(\hat{x}) \leq \|z - u\| + \|u - \hat{x}\| + f(\hat{x}) = b + 2f(\hat{x}).$$

Thus z must have a sample point within $\varepsilon f(z) \leq \varepsilon(b + 2f(\hat{x}))$ in P . This puts an upper bound on the radius R of the empty ball B :

$$R \leq \|c - z\| + \varepsilon f(z) \leq d - b + \varepsilon(b + 2f(\hat{x})).$$

Since $b \geq f(\hat{x})$ we get:

$$R \leq (d - 1 + 3\varepsilon)f(\hat{x}). \quad (6.3)$$

On the other hand, since x is in τ and $B_u = B(u, f(\hat{x}))$ is empty, $\pi_{B_c}(x) \leq \pi_{B_u}(x)$, i.e.

$$\|x - c\|^2 - R^2 \leq \|x - u\|^2 - f(\hat{x})^2, \quad (6.4)$$

Using $\|x - u\| = (1 - \delta)f(\hat{x})$ and the cosine rule on the triangle cux we can write $\|x - c\|^2$ as

$$\|c - x\|^2 = (1 - \delta)^2 f(\hat{x})^2 + d^2 - 2d(1 - \delta)f(\hat{x}) \cos \alpha,$$

where $\alpha = \angle(u - c, u - x)$. Combining with Equation (6.4) and using $\cos \alpha \leq 1$ gives us

$$R^2 \geq d^2 - 2d(1 - \delta)f(\hat{x}) + f(\hat{x})^2. \quad (6.5)$$

Combining Equations (6.3) and (6.5) plus using the fact that $d \geq f(\hat{x})$ implies that $\delta \leq 9\varepsilon^2$, which is a contradiction. \blacksquare

Lemma 6.5 *Let x be a point on the boundary of Σ_δ for $0 < \delta < 1$ and let v be a vector satisfying*

$$\tan \alpha \leq \frac{1 - \delta}{2\delta},$$

where $\alpha = \angle(x - \hat{x}, v)$. Then there is a real number $t_0 > 0$, such that $x + vt \in \Sigma_\delta$ for every $0 \leq t \leq t_0$.

Proof. For the sake of simplicity we take $f(\hat{x})$ as unit length. Let c be a point between \hat{x} and \tilde{x} satisfying $\|c - \hat{x}\| = f(\hat{x}) = 1$. Consider a point $y = x + tv$, close enough to x so that $y \in B_c = B(c, 1)$. Let θ represent the angle $\angle(c - x, c - y)$. It is easy to see that the parameter t in the statement of the lemma can be replaced with the angle θ corresponding to t . Indeed, we prove that there is a θ_0 such that all points on the segment $\overline{xy_0}$ where y_0 represents the point corresponding to angle θ_0 , are all in Σ_δ .

Since \hat{y} can be no farther away from y than \hat{x} , \hat{y} must be inside the ball $B_y = B(y, \|y - \hat{x}\|)$. Considering the fact that B_c is disjoint from Σ , $\|\hat{x} - \hat{y}\|$ cannot be larger than the diameter of the spherical cap $\partial B_y \setminus B_c$. Thus $\|\hat{x} - \hat{y}\| \leq 2 \sin \theta$ and therefore,

$$f(\hat{y}) \geq f(\hat{x}) - \|\hat{x} - \hat{y}\| \geq 1 - 2 \sin \theta. \quad (6.6)$$




Figure 6.7: Illustration of the proof of Lemma 6.5.

Let us denote the distance between c and y by ℓ . Also, assume without loss of generality that v is a unit vector and therefore $\|x - y\| = t$. From the sine law on the triangle $cx y$, we have

$$\ell = (1 - \delta) \sin \alpha / \sin(\alpha - \theta).$$

Since $\|y - \hat{y}\| \leq \|y - \hat{x}\|$, using the cosine law on triangle $cy\hat{x}$ we have

$$\begin{aligned} \|y - \hat{y}\|^2 &\leq \|y - \hat{x}\|^2 \\ &= 1 + \ell^2 - 2\ell \cos \theta \\ &= 1 + \frac{(1 - \delta)^2 \sin^2 \alpha}{\sin^2(\alpha - \theta)} - \frac{2(1 - \delta) \sin \alpha \cos \theta}{\sin(\alpha - \theta)}. \end{aligned} \quad (6.7)$$

Let us denote the right hand side of Equation (6.7) as $g(\theta)$ (taking α and δ as constants). Thus from (6.6) and (6.7) we have

$$\frac{\|y - \hat{y}\|^2}{f(\hat{y})^2} \leq \frac{g(\theta)}{(1 - 2 \sin \theta)^2}.$$

The statement of the Lemma follows if the function $h(\theta) = g(\theta)/(1 - 2 \sin \theta)^2 \leq \delta^2$ when θ belongs to some interval $[0, \theta_0)$. However, since $h(0) = \delta^2$, this amounts to asserting that $(dh/d\theta)(0) \leq 0$ and $(d^2h/d\theta^2)(0) < 0$ when $(dh/d\theta)(0) = 0$. These claims can be verified algebraically when the specified bound on α is applied. \blacksquare

Corollary 6.6 For $\varepsilon \leq 1/10$ and $9\varepsilon^2 < \delta \leq 3/10 - 2\varepsilon$, no flow line of ϕ_{Q_1} leaves the δ -tubular neighborhood Σ_δ . In other words, Σ_δ is flow-tight for ϕ_Q .

Proof. By Lemmas 6.5 and 6.4, it suffices to choose δ in such a way that $\beta = \angle(\hat{x} - x, v_Q(x))$ is smaller than the angle α required in Lemma 6.3. In other words, it suffices to have

$$\tan \beta \leq \frac{1 - \delta}{2\delta},$$

or equivalently

$$\frac{1}{\cos^2 \beta} \leq 1 + \frac{(1 - \delta)^2}{4\delta^2}.$$

By Lemma 6.3, $\cos \beta \geq 1 - \delta - 6\varepsilon$. Thus it suffices to choose δ in such a way that

$$\frac{1}{(1 - \delta - 6\varepsilon)^2} \leq 1 + \frac{(1 - \delta)^2}{4\delta^2}.$$

It can be verified that for $\varepsilon \leq 1/10$, the inequality is enforced when $9\varepsilon^2 < \delta \leq 3/10 - 2\varepsilon$. ■

6.4 Geometric Quality

The purpose of this section is to prove that the set O of simplices removed from $\text{Del } P$ by the WRAP algorithm advances close to the actual surface Σ . In particular, this entails that the symmetric difference between the output I of WRAP and the original shape S is contained in the tubular neighborhood $\Sigma_{9\varepsilon^2}$. One use of this result is showing a gap between the points in I and exterior medial axis critical points which is important in our homotopy proof. The supplied analysis uses geometric results that are only known to hold in 3D. As a result, our guarantees for the WRAP algorithm are only valid for surfaces in \mathbb{R}^3 .

Lemma 6.7 *Let P_1 be an ε -sample of the surface Σ_1 for $\varepsilon \leq 1/10$. For a simplex τ in $\text{Del } P_1$, if the face of $\text{Vor } P_1$ dual to τ intersects Σ_λ for $\lambda \leq 1/2$, then the simplex τ is contained entirely in Σ_δ for $\delta \geq \frac{5}{2}\varepsilon^2$.*

Proof. We use the setting of the proof of Lemma 3.2 (refer to Figure 3.1). Let $r = r(\varepsilon, \lambda)$ and $\ell = \ell(\varepsilon, \lambda)$. Consider the hyperplane H^+ passing through z and parallel to the tangent plane to Σ at \hat{x} . Let H^- be the hyperplane parallel to H^+ and symmetric to it with respect to \hat{x} . The hyperplanes H^+ and H^- intersect the balls B^+ and B^- respectively in two similar disks C^+ and C^- both centered on the segment $\overline{c^+c^-}$. The region $L(x)$ is contained in the convex hull of $C^+ \cup C^-$ which is a cylinder L_2 whose central axis lies on c^+c^- . Since L_2 and τ are both convex, $\tau \subset L_2$. The important observation is that since the disks C^+ and C^- are respectively contained in the balls B^+ and B^- in opposite sides of Σ , every segment parallel to c^+c^- connecting a point from C^+ to its corresponding point in C^- must intersect Σ . This is in particular true for those of such segments that intersect τ . Thus the length of these segments is an upper bound on the distance between any point in τ and Σ . This length, i.e. the distance between C^+ and C^- is exactly $2(1 - \cos \beta)$. Thus for every point

$u \in \tau$:

$$\|u - \hat{u}\| \leq 2(1 - \cos \beta) = \varepsilon^2 \cdot \frac{1 + \lambda}{1 - \lambda}.$$

On the other hand by the triangle inequality

$$\|\hat{u} - \hat{x}\| \leq \|\hat{u} - u\| + \|u - x\| \leq \|\hat{u} - u\| + r \leq \varepsilon^2 \cdot \frac{1 + \lambda}{1 - \lambda} + \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda}}.$$

Since the local feature size is a 1-Lipschitz function

$$f(\hat{u}) \geq f(\hat{x}) - \|\hat{x} - \hat{u}\| \geq 1 - \varepsilon^2 \cdot \frac{1 + \lambda}{1 - \lambda} + \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda}}.$$

Therefore

$$\frac{\|u - \hat{u}\|}{f(\hat{u})} \leq \frac{\varepsilon^2 \cdot \frac{1 + \lambda}{1 - \lambda}}{1 - \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda}} - \varepsilon^2 \cdot \frac{1 + \lambda}{1 - \lambda}} \leq \frac{5}{2} \varepsilon^2,$$

for $\lambda < 1/2$ and $\varepsilon \leq 1/10$. ■

Lemma 6.8 *Let $\lambda < 1/2$ and let \mathcal{D} be the subcomplex of $\text{Del } P_1$ consisting of all n -simplices whose circumcenters are contained in S_λ^* , along with all proper faces of such simplices. Then $|\mathcal{D}|$ covers $\text{conv } P_1 \cap S_\delta^*$ for $\delta > 9\varepsilon^2$.*

Proof. Every point $x \in \text{conv } P_1$ is in at least one Delaunay n -simplex. To prove the lemma, we show that for every $x \in \text{conv } P_1 \cap S_\delta^*$, the circumcenter of every n -simplex $\tau \in \text{Del } P_1$ containing x is in S_λ^* which implies that $\tau \in \mathcal{D}$.

Thus assume $x \in \text{conv } P_1 \cap S_\delta^*$ and let τ be a Delaunay n -simplex that contains x . Since $x \in S_\delta^*$, $\|x - \hat{x}\| > 9\varepsilon^2 f(\hat{x})$. Thus by Lemma 6.4, the circumcenter c of τ is at the same side of Σ as x and therefore $c \in S^*$. Now, by Lemma 6.7, $\|c - \hat{c}\| \geq \frac{1}{2} f(\hat{c})$ since otherwise $\tau \subset \Sigma_{2.5\varepsilon^2}$ while $x \in \tau$ is in $S_{9\varepsilon^2}^*$, a contradiction. Thus, $c \in S_\lambda^*$ and $\tau \in \mathcal{D}$ and therefore the point x is covered by $|\mathcal{D}|$. Since this argument holds for every $x \in \text{conv } P_1 \cap S_\delta^*$, the proof is complete. ■

Orienting Triangle Normals. In general, for a triangle τ let n_τ denote the direction of the line normal to the plane of τ . It is possible to orient n_τ according to vector n_p^+ for a vertex p of τ . Thus of the two vectors parallel to the direction n_τ , we take the one with angle smaller than $\pi/2$ from n_p^+ as n_τ^+ and the other one as n_τ^- . It is of course generally possible to get a conflicting orientation for n_τ when we repeat this using a different vertex of τ . We call a Delaunay triangle τ *flat* (to the surface) if the orientations of n_τ with respect to surface normals at all three of its vertices are consistent. For flat triangles n_τ^+ and n_τ^- are well-defined without reference to any particular vertex.

Next we show that all Delaunay simplices whose dual intersect S_λ^* for a large enough λ , are

put in O by WRAP. This means that WRAP progresses in removing simplices from $\text{Del } P$ (and putting them in O) and reaches a close neighborhood of the surface. The proof of the main result of this section, i.e. Theorem 6.11 depends on the following lemma the proof of which makes use of Lemma 4.9 from [35] which states that a Delaunay triangles with small circumradius is almost tangent to surface at its vertices. This lemma is only proven for \mathbb{R}^3 and is believed not to generalize to higher dimensions. Because of this, the guarantees provided in this chapter only hold in three dimensions. However, this is the only weak link in the provided chain of arguments and a proof of the following lemma for arbitrary dimensions generalizes all the other guarantees in this paper.

Lemma 6.9 *Let $\varepsilon \leq 0.03$ and $\delta = \varepsilon$. Let \mathcal{D} be as defined in Lemma 6.8 for $\lambda = \delta = \varepsilon$. Let τ be a Delaunay triangle on the boundary of $|\mathcal{D}|$ and let v_i be the vertex of the Voronoi edge dual to τ that is contained in S_δ^* . Furthermore, let x be any intersection point of e and the outer boundary of Σ_δ . Then (1) τ is flat, (2) for every point $y \in \tau$, the angle between the vectors $x - y$ and n_τ^+ (n_τ^-), is at most*

$$\frac{r}{1-3r} + \beta \left(\frac{\ell}{1-r} \right) + \arcsin \left(\frac{r}{\ell} \right),$$

where $r = r(\varepsilon, \delta)$ and $\ell = \ell(\varepsilon, \delta)$ are defined in Lemma 3.2. Finally (3) v_i is on the same side of the plane of τ as x and is farther away from this plane than x .

Proof. For simplicity, we assume that $f(\hat{x}) = 1$. By Lemma 3.2 $h_P(x) \leq \ell = \ell(\varepsilon, \delta)$. Since x is at equal distance from the vertices of τ , $h_P(x)$ is an upper bound for the circumradius of τ . Moreover, by Lemma 3.2, $\|\hat{x} - p\| \leq r = r(\varepsilon, \delta)$ which results, using Lipschitz property of the local feature size, that $f(p) \geq 1 - r$. Therefore, by Lemma 4.9

$$\angle(n_\tau, n_p) \leq \beta \left(\frac{\ell}{1-r} \right). \quad (6.8)$$

Let n_τ^+ be the orientation of n_τ that makes an acute angle with n_p^+ .

For any vertex q of τ , $\|q - \hat{x}\| \leq r$ and therefore using the Lemma 1.12 we obtain

$$\angle(n_q^+, n_{\hat{x}}^+) \leq \frac{r}{1-3r}. \quad (6.9)$$

Since (6.9) holds for every vertex of τ , it also holds for p and therefore

$$\angle(n_p^+, n_q^+) \leq \angle(n_p^+, n_{\hat{x}}^+) + \angle(n_q^+, n_{\hat{x}}^+) \leq \frac{2r}{1-3r}.$$

Thus we get

$$\angle(n_\tau^+, n_q^+) \leq \beta \left(\frac{\ell}{1-r} \right) + \frac{2r}{1-3r} \leq 13^\circ,$$

for our choices of ε and δ . Thus n_τ^+ and n_τ^- are well-defined and τ is flat.

For a point $y \in \tau$, let $\alpha(y)$ be the angle between n_τ^+ and the vector $x - y$. Let x_0 be the intersection point of the affine hulls of e and τ . Since e and τ are orthogonal, $\cot \alpha(y)$ is exactly $\|x - x_0\|/\|y - x_0\|$. Thus $\alpha(y)$ depends only on the distance between y and x_0 and grows monotonically when $\|y - x_0\|$ grows. Therefore, over τ , $\alpha(y)$ achieves its maximum at every one of the three vertices of τ .

Therefore, it suffices to prove the statement of the lemma only for $y = p$ where p is the vertex of τ with the largest face angle in τ . Equations (6.8) and (6.9) put a bound on the angle between n_τ^+ and $n_{\hat{x}}$:

$$\angle(n_\tau^+, n_{\hat{x}}^+) \leq \frac{r}{1-3r} + \beta \left(\frac{\ell}{1-r} \right) < 11^\circ,$$

for $\delta = \varepsilon \leq 0.03$. By Lemma 3.2 the angle between $x - p$ and $n_{\hat{x}}$ is at most

$$\angle(n_{\hat{x}}^+, x - p) \leq \arcsin \left(\frac{r}{\ell} \right).$$

Combining, we get

$$\angle(n_\tau^+, x - p) \leq \angle(n_\tau^+, n_{\hat{x}}^+) + \angle(n_{\hat{x}}^+, x - p) \leq \frac{r}{1-3r} + \beta \left(\frac{\ell}{1-r} \right) + \arcsin \left(\frac{r}{\ell} \right).$$

Next we show that x and v_i are on the same side of the plane Π of τ as x . First observe that the above argument remains valid for every intersection point x of e and the outer boundary of Σ_δ . Since for any such x , the angle between $x - p$ and n_τ^+ is less than 90° , all these points are on the same side of Π . Among all such intersection points, we take x to be the one closest to v_i and prove that v_i is on the same side of Π as x . First notice that since $v_i \in S_\delta^*$ and x is the closest point to v_i on the outer boundary of Σ^δ , the segment $\overline{xv_i}$ is entirely contained in S_δ^* . On the other hand, as shown above, the angle between the segment $\overline{xv_i}$ (parallel to n_τ) and $n_{\hat{x}}$ is at most

$$\frac{r}{1-3r} + \beta \left(\frac{\ell}{1-r} \right).$$

For our choices of ε and δ , this is a smaller angle than $\theta = \arctan \left(\frac{1-\delta}{2\delta} \right)$ and therefore the segment $\overline{xv_i}$ falls inside the double-cone with apex x , axis $n_{\hat{x}}$ and half angle $\arctan \theta$. On the other hand, by Lemma 6.5, xv_i cannot be inside the cone opening toward $n_{\hat{x}}^-$ since every segment \overline{xz} in this cone must have an initial segment $\overline{xz'}$ in Σ_δ while xv_i is entirely in S_δ^* . Thus v_i has to be in the cone opening toward $n_{\hat{x}}^+$. Since we showed above that $\angle(n_{\hat{x}}^+, n_\tau^+)$ differ by at most 11° , v_i has to be on the same side of Π as x and farther away from Π than x . ■

Lemma 6.10 *Let P_1 be an ε -sample of a surface Σ_1 in \mathbb{R}^3 for $\varepsilon \leq 0.03$. Let \mathcal{D} be as defined in Lemma 6.8. Then no flow line of ϕ_{Q_1} starting from a simplex in $\text{Del } P_1 \setminus \mathcal{D}$ enters the interior of $|\mathcal{D}|$.*

Proof. We prove the theorem by analyzing the flow direction on the boundary of $|\mathcal{D}|$. Specifically, for every point x on a Delaunay triangle or edge on $\partial|\mathcal{D}|$, we show that $v_Q(x)$ is either tangent to $\partial|\mathcal{D}|$ or points toward its exterior. Notice that Delaunay vertices, i.e. points in P , are maxima (and thus sinks) of ϕ_Q and thus need not to be analyzed.

Let τ be a Delaunay triangle on $\partial|\mathcal{D}|$. By definition the Voronoi edge e dual to τ has one vertex v_o in Σ_δ and another v_i in S_δ^* (thus intersecting the outer boundary of Σ_δ at some point x). Pick x arbitrarily if there are multiple such intersections. By Lemma 6.9, τ is flat. Take the plane Π of τ as the horizontal plane and the direction of e as vertical. Of the two Delaunay tetrahedra incident to τ , let σ_i be the one included in \mathcal{D} and let σ_o be the other one. Clearly, σ_i is dual to the Voronoi vertex v_i and σ_o is dual to v_o . If v_i and v_o are on opposite sides of Π , $d_Q(\tau)$, the closest point to σ on e will be the intersection of the affine hulls of τ and e and the flow on τ will be tangential in which case there is nothing to prove. Thus assume that v_i and v_o are on the same side of Π .

The three points v_i , x and v_o are on a line orthogonal to Π with x between v_i and v_o . The vertical order of σ_i and σ_o (in the direction of e) agrees with that of v_i and v_o . Thus the flow on τ is toward σ_o if and only if v_i is farther from Π than x . Lemma 6.9 guarantees that this is indeed the case.

Next, consider a Delaunay edge e on the boundary of $|\mathcal{D}|$ and let τ and τ' be two Delaunay triangles incident to e and on the boundary of $|\mathcal{D}|$. Let H and H' be half-planes, respectively supporting τ and τ' , and sharing the line through e for boundary. Let W be the set of all tetrahedra incident to e but not in \mathcal{D} and between τ and τ' . $|W|$ is a polytope that has e , τ , and τ' on its boundary. Let ϕ be the dihedral angle between τ and τ' measured from inside W . We will also refer to the wedge made by τ and τ' and contained locally inside W simply as W since our analysis only concerned about this wedge of W .

Let p be an endpoint of e and thus a common vertex of τ and τ' . It is shown in the proof of Lemma 6.9 that each of the two angle $\angle(n_\tau^+, n_p^+)$ and $\angle(n_{\tau'}^+, n_p^+)$ are less than 13° when $\varepsilon \leq 0.03$. Thus the angle between n_τ^+ and $n_{\tau'}^+$ is at most 26° . Both n_τ^+ and $n_{\tau'}^+$ point toward the interior of W and therefore consistently orient τ and τ' . Thus ϕ , the dihedral angle between τ and τ' , can only be between $180 - 26 = 154^\circ$ and $180 + 26 = 206^\circ$ (note that an inconsistent orientation of τ and τ' puts ϕ in one of the ranges $0-26^\circ$ or $334-360^\circ$).

In the rest of this proof, we analyze this setting restricted only to the bisector plane Π of e . Note that the Voronoi facet dual to e is a planar polygon P contained in Π and so is the driver $d_Q(e)$, which is the closest point to e (or equivalently the mid-point m of e) on P . Let t and t' be intersections of τ and τ' with Π respectively. Thus t and t' are line segments in Π , each incident at one end to m . Voronoi edges s and s' respectively dual to τ and τ' are edges of the polygon P and are contained in Π . Also, let ℓ , ℓ' , λ , and λ' be the supporting lines of t , t' , s , and s' in Π respectively.

To simplify the argument, in drawing this arrangement on Π , we place m at the origin and draw the lines ℓ and ℓ' with a small (less than 26°) angle from the horizontal axis of the




Figure 6.8: Illustration of the proof of Lemma 6.10 for the case where $\phi \geq 180^\circ$

plane. The lines λ and λ' make similar angles with the vertical axis. It is thus meaningful to talk about the *top* or *bottom* planar wedge made by ℓ and ℓ' or by λ and λ' . In the sequel, the top wedge made by ℓ and ℓ' is denoted by $\ell \uparrow \ell'$ and the bottom wedge by $\ell \downarrow \ell'$. The wedge W corresponds to the planar wedge made by ℓ and ℓ' as determined by t and t' .

Observation 1 A first observation is that regardless of the position of t and t' along ℓ and ℓ' with respect to m , both s and s' , the Voronoi edges dual to τ and τ' respectively, intersect $\ell \uparrow \ell'$. This is a consequence of Lemma 6.9 which places an upper-bound on the angle between $x - m$ and n_τ^+ , where x is an intersection point of s and the outer boundary of Σ_δ .

$$\angle(x - m, n_\tau^+) \leq \frac{r}{1 - 3r} + \beta \left(\frac{\ell}{1 - r} \right) + \arcsin \left(\frac{r}{\ell} \right) < 57^\circ,$$

for our choices of ε and δ . Thus then angle between $x - m$ and ℓ is at least $90 - 57 = 33^\circ$. Recalling that n_τ^+ is the vector we get by orienting λ upward, shows that x is in $\ell \uparrow \ell'$. Similarly, any intersection point x' between s' and the outer boundary of Σ_δ lies in the same wedge.

The second key observation is based on the definitions of Delaunay and Voronoi complexes.

Observation 2 The order of λ and P along the direction ℓ agrees with that of t and m . The same holds with λ replaced with λ' and t with t' .

We first look at the case where $\phi \geq 180^\circ$. Refer to Figure 6.8. In this figure the obtuse angle made by t and t' corresponds to W . We want to show that the driver $d_Q(m)$ of m , i.e. the closest point to m on P is *not* in $\ell \downarrow \ell'$. This would imply that $v_Q(m)$ points toward the interior or W , which is what we wish to prove. By Observation 2, Since t is to the left of m , P must be on the right of λ . Similarly, P must be on the left of λ' . Thus P must be in $\lambda \downarrow \lambda'$ (grayed in the figure). Since by Observation 1, P intersects $\ell \uparrow \ell'$, this wedge




Figure 6.9: Illustration of the proof of Lemma 6.10 for the case where $\phi < 180^\circ$

and $\lambda \uparrow \lambda'$ must intersect. Thus, the left case of Figure 6.8 cannot happen. Thus suppose $\lambda \uparrow \lambda'$ intersects $\ell \uparrow \ell'$. We consider two cases depending on whether m is in $\lambda \uparrow \lambda'$ or not. If not (Figure 6.8 middle), suppose to the contrary that the driver $d = d_Q(m)$ is in $\ell \downarrow \ell'$. Consider the line among λ and λ' that separates m and d , say λ' as in the figure, and consider the angle between $m - d$ and $x' - d$, where x' is a point of intersection of s' and the outer boundary of Σ_δ . It is easy to observe that this angle is acute. However, P being a convex polygon, the segment $\overline{dx'}$ is contained in P and this segment making an acute angle with mx is in contradiction with d being the closest point of P to m . In the case where m is contained in the wedge $\lambda \downarrow \lambda'$, suppose again to the contrary that the driver d of m is the wedge $\ell \downarrow \ell'$. It is easy to observe that in this case, at least one of the angles $\angle(x - d, m - d)$ and $\angle(x' - d, m - d)$ is acute. This results a contradiction an in the previous case.

The case $\phi < 180^\circ$ is similarly handled. Refer to Figure 6.9. Here the wedge W determined by t and t' is exactly $\ell \downarrow \ell'$ and we want to prove that the driver $d = d_Q(m)$ is in the wedge $\ell \uparrow \ell'$ which entails that the flow at m enters W . Notice that in this case, Observation 2 implies that $P \subset \lambda \uparrow \lambda'$. Thus if $\lambda \uparrow \lambda'$ is contained in $\ell \uparrow \ell'$ (left figure), there is nothing to prove. If $\lambda \uparrow \lambda'$ is not contained in $\ell \uparrow \ell'$ but it doesn't include m (middle figure), then let λ' be the line that separates P and m and let x' be an intersection point of s' and the outer boundary of Σ_δ . If $d \in W$, then $\angle(m - d, x' - d)$ is acute leading to a contradiction as in the previous case. If m is contained in $\lambda \uparrow \lambda'$ and $d \in W$, then similar to the previous case one can observe that at least one of the angles $\angle(m - d, x' - d)$ or $\angle(m - d, x - d)$ is acute, resulting a contradiction with d being the driver of m . ■

Using the previous lemma, the geometric guarantee of the WRAP algorithm is given by the following theorem.

Theorem 6.11 *For $\varepsilon \leq 0.03$, the output I of the WRAP algorithm is contained in $S \cup \Sigma_{9\varepsilon^2}$ and includes $S_{9\varepsilon^2} = S \setminus \Sigma_{9\varepsilon^2}$.*

Proof. Lemma 6.10 shows that no simplices of the subcomplex \mathcal{D} of $D = \text{Del } P_1$ is preceded

by a surface or an interior medial axis critical simplex. This means that all simplices in $\mathcal{D} \cap \text{Del } P$ are picked by WRAP and put in O . Since by Lemma 6.8, all simplices in $D \setminus \mathcal{D}$ are contained in $S \cup \Sigma_{9\varepsilon^2}$, the same is true for $I = \text{Del } P \setminus O$, i.e. $|I| \subset S \cup \Sigma_{9\varepsilon^2}$.

To show that $I \supset S \setminus \Sigma_{9\varepsilon^2}$, assume to the contrary that a simplex $\tau \in O$ intersects $S \setminus \Sigma_{9\varepsilon^2}$. Since by Corollary 6.6 $\Sigma_{9\varepsilon^2}$ is flow-tight, any point $x \in \tau \cap S_{9\varepsilon^2}$ can only be on the unstable manifold of inner medial axis critical points and therefore τ is preceded by some inner medial axis critical simplex. This contradicts the choice of τ . ■

6.5 Topological Correctness

As mentioned before and similar to our approach in earlier chapters, we intend to use the flow map ϕ_{Q_1} for the role of H in Proposition 1.8. In particular, we will use the machinery of Theorem 4.20 for establishing the homotopy equivalence of S and the output of WRAP.

As a first step, we show that if we enlarge S by adding the tubular neighborhood Σ_δ to it, the resulting set preserves the homotopy type of S .

Lemma 6.12 *The two sets $S' = S \cup \Sigma_\delta$ and $\text{cl } S = S \cup \Sigma$ are homotopy equivalent, for any $0 < \delta < 1$. In fact, the latter is a strong deformation retract of the former.*

Proof. Consider the retraction map $r : S' \rightarrow \text{cl } S$ given by

$$r(x) = \begin{cases} \hat{x} & x \in S' \setminus \text{cl } S \\ x & x \in \text{cl } S \end{cases}$$

The map r is continuous on Σ_δ because $\Sigma_\delta \cap M(\Sigma) = \emptyset$ and $M(\Sigma)$ consists of the only points in space where the map $x \mapsto \hat{x}$ is not continuous (in fact undefined). Since the map r is identity on $\text{cl } S$, r is continuous on all of its domain.

If we now define the map $R : [0, 1] \times S' \rightarrow S'$ as

$$R(t, x) = \begin{cases} (1-t)x + t\hat{x} & x \in S' \setminus \text{cl } S \\ x & x \in \text{cl } S, \end{cases}$$

the map R is a straight-line homotopy from the identity on S' to the retraction r . ■

Lemma 6.13 *Let E be the union of unstable manifolds of all surface and inner medial axis critical points of h_Q . Then $S \cup \Sigma_\delta$ is homotopy equivalent to E for $9\varepsilon^2 < \delta < 3/10 - 2\varepsilon$.*

Proof. By Corollary 6.6, $\Sigma_\delta \cup S$ is flow-tight for ϕ_{Q_1} . On the other hand E is also flow-tight for ϕ_{Q_1} since by definition $\phi_{Q_1}(E) = E$. Therefore Theorem 4.20 implies the desired homotopy

equivalence if we only show that $\|v_{Q_1}(x)\|$ is bounded from below for every $x \in (S \cup \Sigma_\delta) \setminus E$. But $\|v_{Q_1}(x)\| = 2\|x - d_{Q_1}(x)\|$. When $D_{Q_1}(x) \cap V_{Q_1}(x) = \emptyset$, the distance between x and $d_{Q_1}(x)$ is at least the distance between $V_{Q_1}(x)$ and $D_{Q_1}(x)$. Since both $\text{Vor } Q_1$ and $\text{Del } Q_1$ are finite complexes, there is a lower bound on the distance between pairs of dual faces from the two complexes that do not intersect.

On the other hand, if $V_{Q_1}(x) \cap D_{Q_1}(x)$ do intersect, then their intersection is a critical point c coinciding with $d_{Q_1}(x)$. Thus x is by definition on the unstable manifold of c . If c is a surface or inner medial axis critical point, then $\text{Um}(c) \subset E$ and thus $x \in E$ contradicting the choice of c . Thus c can only be an outer medial axis critical point. But then, by Theorem 3.3, c is not contained in $\Sigma_{1-2\varepsilon^2}$. The distance between $\Sigma_{1-2\varepsilon^2}$ and Σ_δ being strictly positive puts a lower bound on the distance between x and c . ■

Lemma 6.14 *Let E be as in Lemma 6.13. If $\varepsilon \leq 0.03$, then the output I of the algorithm WRAP is homotopy equivalent to E .*

Proof. We first observe that I is flow-tight for ϕ_{Q_1} . By definition, the set O in the WRAP algorithm consists of those simplices in $\text{Del } P$ that are only preceded by exterior medial axis critical simplices (including ω). If a flow line leaves I it enters O . But by definition, simplices in I are reachable from critical simplices other than the outer medial axis critical ones. Therefore if a flow line enters O from I , it provides a path from a critical simplex of surface or inner medial axis to some simplex in O contradicting the definition of O .

On the other hand, E is also flow-tight for ϕ_Q as was shown in the proof of Lemma 6.13. We observe here that $E \subseteq I$. This is because E is by definition the locus of all points reachable by some flow path starting at a critical simplex of inner medial axis or the surface. Trivially, these critical simplices are themselves in I . For every other point $x \in E$, there is a flow path $\phi_Q(y)$ starting at a point y in the relative interior of some critical simplex τ in the mentioned group reaching x . Let $\tau = \tau_0, \dots, \tau_k$ be the set of simplices intersecting $\phi_Q(y)$ until it reaches x with τ_k being the simplex containing x in its relative interior. The existence of this path implies that $\tau_0 \prec \dots \prec \tau_k$. Therefore τ_k cannot be placed in O by the WRAP algorithm and consequently $\tau_k \in I$. So, we have shown that $E \subseteq I$.

Thus in order to complete the proof we only need to show that at every $x \in I \setminus E$, $\|v_Q(x)\| > c$ for some $c > 0$ and then the homotopy equivalence between E and I follows from Theorem 4.20. Similar to the proof of Lemma 6.13, we only need to show that the points in $I \setminus E$ are at a positive distance from every outer medial axis critical point. This follows by a similar argument to the one in the proof of Lemma 6.13 using the fact that $I \subset S \cup \Sigma_\delta$ for $\delta = 9\varepsilon^2$ as is shown by Theorem 6.11. ■

The topological guarantee of the wrap algorithm follows from Lemmas 6.12, 6.13, 6.14.

Theorem 6.15 *Let P be an ε -sample of a surface Σ embedded in \mathbb{R}^3 with $\varepsilon < 0.03$. Then the output I of the WRAP algorithm is homotopy equivalent to the bounded shape S enclosed by Σ .*

Bibliography and Remarks.

Although the WRAP algorithm was originally implemented in 1996 at Raindrop Geomagic, the details of the algorithm were only published much later [38] after issuance of a patent in 2002. One of the main beauties of this algorithm is that its output is always a unique collection of Delaunay simplices that does not depend on the sampling density of the input, i.e. the algorithm doesn't use ε as a parameter. Although the mathematical foundations behind the algorithm are involved, the description and implementation of the algorithm is rather straight forward — certainly more so than any other Delaunay-based reconstruction algorithm. Our presentation of the algorithm focuses on the descriptive characterization of the output and trivializes the involved algorithmic computation. This is also to some extent the case in the original paper of Edelsbrunner. For a simplified procedural description of the WRAP algorithm refer to Joe O'Rourke's introduction to the subject [59].

References

- [1] M. ALEXA, J. BEHR, D. COHEN-OR, S. FLEISHMAN, D. LEVIN, AND C. T. SILVA, *Point set surfaces*, in IEEE Visualization, 2001.
- [2] N. AMENTA AND M. W. BERN, *Surface reconstruction by Voronoi filtering*, Discrete & Computational Geometry, 22 (1999), pp. 481–504.
- [3] N. AMENTA, M. W. BERN, AND D. EPPSTEIN, *The crust and the β -skeleton: Combinatorial curve reconstruction*, Graphical Models and Image Processing, 60 (1998), pp. 125–135.
- [4] N. AMENTA, S. CHOI, T. K. DEY, AND N. LEEKHA, *A simple algorithm for homeomorphic surface reconstruction*, International Journal of Computational Geometry and Applications, 12 (2002), pp. 125–141.
- [5] N. AMENTA, S. CHOI, AND R. K. KOLLURI, *The power crust, unions of balls, and the medial axis transform*, Computational Geometry, 19 (2001), pp. 127–153.
- [6] N. AMENTA AND Y. J. KIL, *Defining point-set surfaces*, ACM Transactions in Graphics, 23 (2004), pp. 264–270.
- [7] D. ATTALI, J.-D. BOISSONNAT, AND H. EDELSBRUNNER, *Stability and computation of the medial axis — a state-of-the-art report*, Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration, (2004).
- [8] F. AURENHAMMER, *Voronoi diagrams — a study of a fundamental geometric data structure*, ACM Computing Surveys, 23 (1991), pp. 345–405.
- [9] A. BARVINOK, *A Course in Convexity*, American Mathematical Society, 2002.
- [10] R. BLANDING, C. BROOKING, MARK GANTER, AND D. STORTI, *A skeletal-based solid editor*, in ACM Symposium on solid and Physical Modeling, 1999.
- [11] H. BLUM, *A transformatoin for extracting new descriptors of shapes*, in Models for the Perception of Speech and Visual Form, W. Wathen-Dunn, ed., MIT Press, 1967, pp. 362–380.
- [12] J.-D. BOISSONNAT, *Geometric structures for three-dimensional shape representation*, ACM Transactions in Graphics, 3 (1984), pp. 266–286.

- [13] J.-D. BOISSONNAT AND F. CAZALS, *Smooth surface reconstruction via natural neighbour interpolation of distance functions*, Computational Geometry, 22 (2002), pp. 185–203.
- [14] J.-D. BOISSONNAT AND S. OUDOT, *Provably good sampling and meshing of Lipschitz surfaces*, in Symposium on Computational Geometry (SCG), 2006, pp. 337–346.
- [15] W. M. BOOTHBY, *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Academic Press, 2003.
- [16] J. BRANDT AND V. R. ALGAZI, *Continuous skeleton computation by Voronoi diagrams*, Computer Vision, Graphics, and Image Processing, (1992), pp. 329–338.
- [17] P.-T. BREMER AND J. C. HART, *A sampling theorem for MLS surfaces*, in Point Based Graphics, 2005.
- [18] K. BUCHIN AND J. GIESEN, *Flow complex: General structure and algorithms*, in Canadian Conference on Computational Geometry (CCCG), 2005, pp. 270–273.
- [19] J. C. CARR, R. K. BEATSON, J. B. CHERRIE, T. J. MITCHELL, W. R. FRIGHT, B. C. MCCALLUM, AND T. R. EVANS, *Reconstruction and representation of 3D objects with radial basis functions*, in International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH), 2001, pp. 67–76.
- [20] R. CHAINE, *A geometric-based convection approach of 3D reconstruction*, in Eurographics Symposium on Geometry Processing, 2003, pp. 218–229.
- [21] F. CHAZAL, D. COHEN-STEINER, AND A. LIEUTIER, *A sampling theory for compacts in Euclidean space*, in Symposium on Computational Geometry (SCG), 2006, pp. 319–326.
- [22] F. CHAZAL AND A. LIEUTIER, *The λ -medial axis*, Graphical Models and Image Processing, 67 (2005), pp. 304–331.
- [23] —, *Topology guaranteeing manifold reconstruction using distance function to noisy data*, in Symposium on Computational Geometry (SCG), 2006, pp. 112–118.
- [24] F. CHAZAL AND R. SOUFFLET, *Stability and finiteness properties of medial axis and skeleton*, Journal of Dynamical and Control Systems archive Journal of Dynamical and Control Systems, 10 (2004), pp. 149–170.
- [25] H. I. CHOI, S. W. CHOI, AND H. P. MOON, *Mathematical theory of medial axis transform*, Pacific Journal of Mathematics, 181 (1997), pp. 57–88.
- [26] F. H. CLARKE, *Generalized gradients and applications*, Transactions of the American Mathematical Society, 205 (1975), pp. 247–262.
- [27] —, *Optimization and Non-Smooth Analysis*, Wiley-Interscience, 1983.

- [28] T. CULVER, J. KEYSER, AND D. MANOCHA, *Exact computation of the medial axis of a polyhedron*, Computer Aided Geometric Design, 21 (2004), pp. 65–98.
- [29] B. CURLESS AND M. LEVOY, *A volumetric method for building complex models from range images*, in Annual Conference on Computer Graphics (SIGGRAPH), 1996, pp. 303–312.
- [30] C. W. CURTIS, *Linear Algebra, An Introductory Approach*, Springer-Verlag, 1984.
- [31] B. DELAUNAY, *Sur la sphère vide*, Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk, 7 (1934), pp. 793–800.
- [32] T. K. DEY, J. GIESEN, E. A. RAMOS, AND B. SADRI, *Critical points of the distance to an epsilon-sampling of a surface and flow-complex-based surface reconstruction*, in Symposium on Computational Geometry (SCG), 2005, pp. 218–227.
- [33] —, *Separation of critical points and flow-complex based surface reconstruction*, International Journal of Computational Geometry and Applications, (to appear).
- [34] T. K. DEY, S. GOSWAMI, AND J. SUN, *Extremal surface based projections converge and reconstruct with isotopy*, tech. rep., The Ohio State University, 2005.
- [35] T. K. DEY AND P. KUMAR, *A simple provable algorithm for curve reconstruction*, in Symposium on Discrete Algorithms (SODA), 1999, pp. 893–894.
- [36] T. K. DEY AND J. SUN, *An adaptive MLS surface for reconstruction with guarantees*, in Symposium on Geometry Processing, 2005, pp. 43–52.
- [37] T. K. DEY AND W. ZHAO, *Approximating the medial axis from the Voronoi diagram with a convergence guarantee*, Algorithmica, 38 (2003), pp. 179–200.
- [38] H. EDELSBRUNNER, *Surface reconstruction by wrapping finite point sets in space*, Discrete & Computational Geometry, 32 (2004), pp. 231–244.
- [39] M. FOSKEY, M. GARBER, M. LIN, AND D. MANOCHA, *A Voronoi-based hybrid motion planner*, in IEEE/RSJ International Conference on Intelligent Robots and Systems, 2001, pp. 55–60.
- [40] T. FRANKEL, *The Geometry of Physics*, Cambridge University Press, 2004.
- [41] J. H. G. FU, *Tubular neighborhoods in Euclidean spaces*, Duke Mathematical Journal, 52 (1985), pp. 1025–1046.
- [42] J. GIESEN AND M. JOHN, *Surface reconstruction based on a dynamical system*, Computer Graphics Forume, 21 (2002), pp. 363–371.
- [43] —, *The flow complex: A data structure for geometric modeling*, in Symposium on Discrete Algorithms (SODA), 2003, pp. 285–294.

- [44] J. GIESEN, E. A. RAMOS, AND B. SADRI, *Medial axis approximation and unstable flow complex*, in Symposium on Computational Geometry, 2006, pp. 327 – 336.
- [45] ———, *Medial axis approximation and unstable flow complex*, International Journal of Computational Geometry and Applications, (to appear).
- [46] K. GROVE, *Critical point theory for distance functions*, Symposia in Pure Mathematics, 54 (1993), pp. 357–385.
- [47] K. GROVE AND K. SHIOHAMA, *A generalized sphere theorem*, Annals of Mathematics, 106 (1977), pp. 201–211.
- [48] A. HATCHER, *Algebraic Topology*, Cambridge University Press, 2001.
- [49] M. W. HIRSCH AND S. SMALE, *Differential Equations, Dynamical Systems, and Linear Algebra*, Academic Press, 1974.
- [50] H. HOPPE, T. DEROSE, T. DUCHAMP, J. A. McDONALD, AND W. STUETZLE, *Surface reconstruction from unorganized points*, in Annual Conference on Computer Graphics (SIGGRAPH), 1992, pp. 71–78.
- [51] M. JOHN, *Flow Complexes: Structure, Algorithms, and Applications*, PhD thesis, Swiss Federal Institute of Technology, ETH Zürich, 2003.
- [52] R. KAYE AND R. WILSON, *Linear Algebra*, Oxford University Press, 1998.
- [53] R. K. KOLLURI, *Provably good moving least squares*, in Symposium on Discrete Algorithms (SODA), 2005, pp. 1008–1017.
- [54] J. M. LEE, *Introduction to Topological Manifolds*, Springer-Verlag, 2000.
- [55] D. LEVIN, *The approximation power of moving least-squares*, in Mathematics of Computation, vol. 67, 1998, pp. 1517–1531.
- [56] A. LIEUTIER, *Any bounded open subset of \mathbb{R}^n has the same homotopy type as its medial axis*, Computer-Aided Design, 36 (2004), pp. 1029–1046.
- [57] J. R. MUNKRES, *Elements of Algebraic Topology*, Addison Wesley, 1993.
- [58] Y.-L. O, A. TOET, D. FOSTER, AND H. J. A. M. HEIJMANS, *Shape in Picture: Mathematical Description of Shape in Grey-Level Images (Nato a S I Series Series III, Computer and Systems Sciences)*, Springer-Verlag, 1994.
- [59] J. O’ROURKE, *Computational geometry column 45*, ArXiv Computer Science e-prints, (2004).
- [60] S. M. PIZER, K. SIDDIQI, G. SZÉKELY, J. N. DAMON, AND S. W. ZUCKER, *Multiscale medial loci and their properties*, International Journal of Computer Vision, 55 (2003), pp. 155–179.

- [61] E. A. RAMOS AND B. SADRI, *Geometric and topological guarantees for the wrap reconstruction algorithm*, in Symposium on Discrete Algorithms (SODA), 2007, p. to appear.
- [62] G. P. ROBINSON, A. C. F. COLCHESTER, L. D. GRIFFIN, AND D. J. HAWKES, *Integrated skeleton and boundary shape representation for medical image interpretation*, in European Conference on Computer Vision, 1992, pp. 725–729.
- [63] B. SCHÖLKOPF, J. GIESEN, AND S. SPALINGER, *Kernel methods for implicit surface modeling*, in Annual Conference on Neural Information Processing Systems (NIPS), 2004, pp. 1193–1200.
- [64] A. SHEFFER, M. ETZION, A. RAPPOPORT, AND M. BERCOVIER, *Hexahedral mesh generation using the embedded Voronoi graph*, in 7th International Meshing Roundtable, 1998, pp. 347–364.
- [65] K. SURESH, *Automating the CAD/CAE dimension reduction process*, in ACM Symposium on solid and Physical Modeling, 2003, pp. 76–85.
- [66] R. TAM AND W. HEIDRICH, *Shape simplification based on the medial axis transform*, in Proceedings of the IEEE Visualization (VIS), 2003, p. 63.
- [67] G. VORONOI, *Nouvelles application des paramètres continus à la théorie des formes quadratique*, Journal für die reine und angewandte Mathematik, 134 (1908), pp. 198–287.
- [68] J. H. C. WHITEHEAD, *Simplicial spaces, nuclei, and m -groups*, Proceedings of the London Mathematical Society, 45 (1939), pp. 243–327.
- [69] F.-E. WOLTER, *Cut locus and medial axis in global shape interrogation and representation*, Tech. Rep. 92-2, MIT Department of Ocean Engineering and MIT Design Laboratory, 1992.

Author's Biography

Bardia Sadri was born in Tehran, Iran, on August 1, 1977. He graduated from Sharif University of Technology at Tehran in June 1999 with a B.S. degree in computer engineering. He joined the computer science program at the University of Illinois at Urbana-Champaign in August 1999 and earned his M.S. and Ph.D. degrees in May 2004 and December 2006 respectively. Following the completion of his Ph.D., Bardia will begin working as a post-doctoral fellow at the department of computer science at Duke University in Durham, North Carolina.