

Geometric and Topological Inference

Jean-Daniel Boissonnat Frédéric Chazal
Mariette Yvinec

August 17, 2017

Contents

I	Topological preliminaries	17
1	Topological spaces	19
1.1	Topological spaces	20
1.2	Comparing topological spaces	22
1.2.1	Homeomorphism, isotopy and homotopy equivalence .	22
1.2.2	Hausdorff distance	25
1.3	Exercises	26
1.4	Bibliographical notes	27
2	Simplicial complexes	29
2.1	Geometric simplicial complexes	30
2.2	Abstract simplicial complexes	31
2.3	Nerve	33
2.4	Filtrations of simplicial complexes	34
2.5	Vietoris-Rips and Čech filtrations	35
2.6	Combinatorial manifolds triangulations	37
2.7	Representation of simplicial complexes	38
2.8	Exercises	40

2.9	Bibliographical notes	41
II	Delaunay complexes	43
3	Convex polytopes	45
3.1	Definitions	46
3.1.1	Convex polytopes	46
3.1.2	Facial structure of polytopes	46
3.1.3	Convex polyhedra	48
3.1.4	Simplicial polytopes and simple convex polyhedra	48
3.1.5	The boundary complex	49
3.2	Duality	50
3.2.1	Point-hyperplane duality	50
3.2.2	Lower hulls and upper envelope	51
3.3	Combinatorial bounds	53
3.4	Convex hull algorithms	55
3.4.1	An incremental algorithm	55
3.4.2	A randomized algorithm	60
3.5	Exercises	66
3.6	Bibliographical notes	69
4	Delaunay complexes	71
4.1	Lower envelopes and minimization diagrams	72
4.2	Voronoi diagrams	73
4.3	Delaunay complexes	76

<i>CONTENTS</i>	5
-----------------	---

4.4	Weighted Delaunay complexes	81
4.4.1	Weighted points and weighted distance	81
4.4.2	Weighted Voronoi diagrams	81
4.4.3	Weighted Delaunay complexes	83
4.4.4	Complexity of weighted Delaunay complexes	86
4.5	Examples of weighted Voronoi diagrams	86
4.5.1	k -order Voronoi diagrams	86
4.5.2	Bregman diagrams	88
4.6	Exercises	95
4.7	Bibliographical notes	99

5	Good triangulations	101
----------	----------------------------	------------

5.1	Nets	102
5.1.1	Nets in Euclidean space	102
5.1.2	Delaunay complex of a net	104
5.1.3	Nets in discrete metric spaces	108
5.2	Thick simplices	111
5.2.1	Thickness and singular value	112
5.2.2	Whitney's angle bound	114
5.3	Thick triangulations via weighting	116
5.3.1	Weighting schemes	117
5.3.2	Θ_0 -thickness and flakes	119
5.3.3	The weight range of a flake with small radius	120
5.3.4	Lovász local lemma	123
5.3.5	Applying the Lovász local lemma to remove flakes	124

5.3.6	Algorithm	126
5.4	Protection	128
5.5	Exercises	131
5.6	Bibliographical notes	132
6	Delaunay filtrations	135
6.1	Alpha complexes	137
6.1.1	Definitions	137
6.1.2	Computing alpha complexes and filtrations	139
6.1.3	Weighted alpha complex	140
6.1.4	Application to union of balls	141
6.2	Witness complexes	142
6.2.1	Identity of witness and Delaunay complexes	144
6.2.2	Computing witness complexes	147
6.2.3	Relaxed witness complexes	148
6.3	Exercises	150
6.4	Bibliographical notes	152
III	Reconstruction of smooth submanifolds	153
7	Triangulation of submanifolds	157
7.1	Reach and ε -nets on submanifolds	158
7.1.1	Submanifolds	158
7.1.2	Projection map, medial axis and reach	159
7.1.3	ε -nets on a submanifold	163

7.2	Projection maps	164
7.3	Triangulation of submanifolds	171
7.3.1	Proof of statement 1	172
7.3.2	Whitney's lemma	173
7.3.3	Satisfying the conditions of Whitney's lemma	174
7.3.4	Local homeomorphism	177
7.3.5	The restriction of Π to $\hat{\mathcal{M}}$ is injective	178
7.3.6	The restriction of Π to $\hat{\mathcal{M}}$ is surjective	179
7.3.7	End of proof of statement 1	179
7.3.8	Proof of statements 2 and 3, and of Theorem 7.16	179
7.4	Exercises	180
7.5	Bibliographical notes	182
8	Reconstruction of submanifolds	183
8.1	Homotopy reconstruction using alpha-shapes	185
8.2	Tangential Delaunay complex	189
8.2.1	Definition	190
8.2.2	Inconsistent simplices	192
8.2.3	Geometric properties of inconsistency triggers	194
8.3	Submanifold reconstruction	198
8.3.1	Weight assignment	199
8.3.2	Reconstruction algorithm	202
8.3.3	Guarantees on the reconstruction	204
8.4	Exercises	206
8.5	Bibliographical notes	208

IV	Distance-based inference	211
9	Stability of distance functions	213
9.1	Distance function and Hausdorff distance	215
9.2	Critical points of distance functions	215
9.3	Topology of the offsets	219
9.4	Stability of critical points	221
9.5	The critical function of a compact set	225
9.6	Sampling conditions and μ -reach	227
9.7	Offset reconstruction	229
9.8	Exercises	231
9.9	Bibliographical notes	232
10	Distance to probability measures	235
10.1	Extending the sampling theory for compact sets	236
10.2	Measures and the Wasserstein distance W_2	240
10.2.1	Replacing compact sets by measures	240
10.2.2	The Wasserstein distance W_2	241
10.3	Distance function to a probability measure	243
10.3.1	Definition	243
10.3.2	Distance function to empirical measures	244
10.3.3	Equivalent formulation	245
10.3.4	Stability of the distance function to a measure	247
10.3.5	The distance to a measure is a distance-like function.	248
10.4	Applications to geometric inference	250

10.4.1	Distance to a measure vs. distance to its support . . .	250
10.4.2	Shape reconstruction from noisy data	251
10.5	Exercises	252
10.6	Bibliographical notes	254
11	Homology inference	255
11.1	Simplicial homology and Betti numbers	256
11.1.1	The space of k -chains	256
11.1.2	The boundary operator and homology groups	257
11.2	An algorithm to compute Betti numbers	260
11.3	Singular homology and topological invariance	263
11.4	Betti numbers inference	264
11.5	Persistent homology	267
11.5.1	A simple example	268
11.5.2	Persistent homology of a filtration	270
11.5.3	Persistence diagrams and bottleneck distance	275
11.5.4	Persistence modules, interleaving and stability	276
11.5.5	Persistence stability for functions	279
11.5.6	Persistence stability for compact sets and complexes built on top of point clouds	280
11.6	Exercises	283
11.7	Bibliographical notes	284

Introduction

Motivation and goals. In many practical situations, geometric objects are only known through a finite set of possibly noisy sample points. A natural question is then to recover the geometry and the topology of the unknown object from this information. The most classical example is probably surface reconstruction, where the points are measured on the surface of a real world object. A perhaps more surprising example is the study of the large scale structure formed by the galaxies, which cosmologists believe to be an interconnected network of walls and filaments. In other applications, the shape of interest may be a low dimensional object embedded in a higher-dimensional space, which is the basic assumption in *manifold learning* [105]. This is for example the case in time series analysis, when the shape of study is the attractor of a dynamical system sampled by a sequence of observations. When these structures are highly nonlinear and have a non trivial topology as it is often the case, simple dimensionality reduction techniques do not suffice and must be complemented with more geometric and topological techniques.

A lot of research was done in this direction, originating from several sources. A bunch of contributions came from the field of *computational geometry*, where much effort was done to elaborate provably correct surface reconstruction algorithms, under a suitable sampling condition. We refer to [65] for a thorough review of this approach. However most of this research focused on the case of sampled smooth surfaces in $\mathbb{R}^3$, which is by now fairly well covered. Extending these results to higher dimensional submanifolds and to non smooth objects is one of the objectives of this book. Such an extension requires new data structures to walk around the curse of dimensionality. Handling more general geometric shapes also requires concepts from topology and has provoked an interest in the subject of *computational topology*. Computational topology has recently gained a lot of momentum and has been very successful at providing qualitative invariants and efficient algorithms to compute them. Its application to data analysis led to the rapidly evolving field of *topological data analysis* that provides a general framework to analyze the shape of data and has been applied to various types of data across many fields.

This book. This book intends to cover various aspects of geometric and topological inference, from data representation and combinatorial questions to persistent homology, an adaptation of homology to point cloud data. The aim of this book is not to provide a comprehensive treatment of topological data analysis but to describe the mathematical and algorithmic foundations

of the subject.

Two main concepts will play a central role in this book : simplicial complexes and distance functions. *Simplicial complexes* generalize the notion of triangulation of a surface and are constructed by gluing together simplices : points, line segments, triangles and their higher dimensional counterparts. Simplicial complexes can be considered, at the same time, as continuous objects carrying topological and geometric information and as combinatorial data structures that can be efficiently implemented. Simplicial complexes can be used to produce fine meshes leading to faithful approximations well suited to scientific computing purposes, or much coarser approximations, still useful to infer important features of shapes such as their homology or some local geometric properties.

Simplicial complexes have been known and studied for a long time in mathematics but only used in low dimensions due to their high complexity. In this book, we will address the complexity issues by focusing on the inherent, usually unknown, structure in the data which we assume to be of relative low intrinsic dimension. We will put emphasis on output-sensitive algorithms, introduce new simplicial complexes with low complexity, and describe approximation algorithms that scale well with the dimension.

Another central concept in this book is the notion of *distance function*. All the simplicial complexes used in this book encode proximity relationships between the data points. A prominent role is taken by Voronoi diagrams, their dual Delaunay complexes and variants of those, but other simplicial complexes based on distances like the Čech, the Vietoris-Rips or the witness complexes will also be considered.

The book is subdivided into four parts.

Part I contains two chapters that present background material on topological spaces and simplicial complexes.

Part II introduces Delaunay complexes and their variants. Since Delaunay complexes are closely related to polytopes, the main combinatorial and algorithmic properties of polytopes are presented first in Chapter 3.

Delaunay complexes, to be introduced in Chapter 4 are defined from Voronoi diagrams which are natural space partitions induced by the distance function to a sample. Delaunay complexes appear as the underlying basic data structure for manifold reconstruction. The extensions of Voronoi diagrams

and Delaunay complexes to weighted distances are also presented together with their relevant applications to k th-nearest neighbor search and Bregman divergences which are used in information theory, image processing and statistical analysis.

Although Delaunay triangulations have many beautiful properties, their size depends exponentially on the dimension of the space in the worst-case. It is thus important to exhibit realistic assumptions under which the complexity of the Delaunay triangulation does not undergo such a bad behaviour. This will be done through the notion of nets. Another issue comes from the fact that, in dimensions greater than 2, Delaunay simplices may have an arbitrarily small volume, even if their vertices are well distributed. Avoiding such bad simplices is a major issue and the importance of thick triangulations has been recognized since the early days of differential topology. They play a central role in numerical simulations to ensure the convergence of numerical methods solving partial differential equations. They also play a central role in the triangulation of manifolds and, in particular, the reconstruction of submanifolds of high dimensional spaces as shown in Chapter 8. Chapter 5 defines thick triangulations and introduces a random perturbation technique to construct thick Delaunay triangulations in Euclidean space.

Chapter 6 introduces two filtrations of simplicial complexes. Filtrations are nested sequences of subcomplexes that allow to compute persistent homology as described in Chapter 11. We first introduce alpha-complexes and show that they provide natural filtrations of Delaunay and weighted Delaunay complexes. We then introduce witness complexes and their filtrations. The witness complex is a weak version of the Delaunay complex that can be constructed in general metric spaces using only pairwise distances between the points, without a need for coordinates. We will also introduce a filtration of the witness complex.

Part III is devoted to the problem of reconstructing a submanifold $\mathcal{M}$ of $\mathbb{R}^d$ from a finite point sample $P \in \mathcal{M}$. The ultimate goal is to compute a triangulation of $\mathcal{M}$, i.e. a simplicial complex that is homeomorphic to $\mathcal{M}$. This is a demanding quest and, in this part, we will restrict our attention to the case where $\mathcal{M}$ is a smooth submanifold of $\mathbb{R}^d$.

In Chapter 7, we introduce the basic concepts and results, and state a theorem that provides conditions for a simplicial complex $\hat{\mathcal{M}}$ with vertex set $P \subset \mathcal{M}$ to be both a triangulation and a good geometric approximation of $\mathcal{M}$.

Chapter 8 is devoted to the problem of reconstructing submanifolds from point samples. This problem is of primary importance when $\mathcal{M}$ is a surface of $\mathbb{R}^3$ (it is then known as the surface reconstruction problem). It also finds applications in higher dimensions in the context of data analysis where data are considered as points in some Euclidean space, of possibly high dimension. In this chapter, we first exhibit conditions under which the alpha-complex of $P \subset \mathcal{M}$ has the same homotopy type as $\mathcal{M}$, a weaker property than being homeomorphic to $\mathcal{M}$. We then consider the problem of reconstructing a smooth submanifold $\mathcal{M}$ embedded in a space of possibly high dimension d . We then cannot afford to triangulate the ambient space as is being routinely done when working in low dimensions. A way to walk around this difficulty is to assume, as is common practice in data analysis and machine learning, that the intrinsic dimension k of $\mathcal{M}$ is small, even if the dimension of the ambient space may be very large. Chapter 8 takes advantage of this assumption and presents a reconstruction algorithm whose complexity is linear in d and exponential only in k .

The assumptions made in Part III are very demanding: the geometric structures of the data should be smooth submanifolds, the amount of noise in the data should be small and the sampling density should be high. These assumptions may not be satisfied in practical situations. Part IV aims at weakening the assumptions. Chapter 9 studies the stability properties of the sublevel sets of distance functions and provide sampling conditions to infer the underlying geometry and topology of data.

Approximations in Chapter 9 are with respect to the Hausdorff distance. This is a too strong limitation when the data contain outliers that are far away from the underlying structure we want to infer. To overcome this problem, Chapter 10 introduces a new framework where data are no longer considered as points but as distributions of mass or, more precisely probability measures. It is shown that the distance function approach can be extended to this more general framework.

Although Chapters 9 and 10 provide strong results on the topology of the sublevel sets of distance functions, computing and manipulating such sublevel sets is limited in practice to low dimensions. To go beyond these limitations, we restrict our quest to the inference of some topological invariants of the level sets, namely their homology and the associated Betti numbers. Chapter 11 introduces persistent homology and provides tools to robustly infer the homology of sampled shapes.

Efficient implementations of most of the algorithms described in this book can be found in the CGAL library (<http://www.cgal.org/>) or in the GUDHI library (<http://gudhi.gforge.inria.fr/>).

Acknowledgments. The material in this book originates from a course that the authors gave at MPRI (*Master Parisien de Recherche en Informatique*). It includes results obtained through joint work with several researchers the authors want to warmly thank: David Cohen-Steiner, Ramsay Dyer, Vin de Silva, Kunal Dutta, Arijit Ghosh, Marc Glisse, Leonidas Guibas, André Lieutier, Quentin Mérigot, Frank Nielsen, Steve Oudot, Primoz Skraba, Mathijs Wintraecken. The authors are also grateful to Claire Brécheteau for kindly reading Part III and suggesting improvements.

The research leading to this book has been partially supported by the Agence Nationale de la Recherche under the project TopData and by the European Research Council under the Advanced Grant GUDHI (Geometric Understanding in Higher Dimensions).

Part I

Topological preliminaries

Chapter 1

Topological spaces

Basic mathematical notions useful in the sequel of this book are given in this chapter. For conciseness, the definitions and results are not always given in their full generality. They are restricted to the simplest version necessary to follow and understand the results and proofs of this book.

1.1 Topological spaces

This section quickly recalls a few basic notions and definitions from general topology. Most of the topological objects encountered in this book are metric spaces whose definition is also recalled.

A *topology* on a set X is a family $\mathcal{O}$ of subsets of X that satisfies the three following conditions: encountered in this book are metric spaces whose definition is also recalled.

Definition 1.1 (Topological space) *A topology on a set X is a family $\mathcal{O}$ of subsets of X that satisfies the three following conditions:*

1. *the empty set $\emptyset$ and X are elements of $\mathcal{O}$,*
2. *any union of elements of $\mathcal{O}$ is an element of $\mathcal{O}$,*
3. *any finite intersection of elements of $\mathcal{O}$ is an element of $\mathcal{O}$.*

The set X together with the family $\mathcal{O}$, whose elements are called open sets, is a topological space. A subset C of X is closed if its complement is an open set. If $Y \subset X$ is a subset of X , then the family $\mathcal{O}_Y = \{O \cap Y : O \in \mathcal{O}\}$ is a topology on Y , called the induced topology.

Definition 1.2 (Closure, interior and boundary) *Let S be a subset of a topological space X . The closure $\bar{S}$ of S , is the smallest closed set containing S . The interior $\overset{\circ}{S}$ of S is the largest open set contained in S . The boundary ∂S of S is the set difference $\partial S = \bar{S} \setminus \overset{\circ}{S}$.*

Definition 1.3 (Metric space) *A metric (or distance) on a set X is a map $d : X \times X \rightarrow [0, +\infty)$ such that:*

1. *for any $x, y \in X$, $d(x, y) = d(y, x)$,*
2. *for any $x, y \in X$, $d(x, y) = 0$ if and only if $x = y$,*

3. for any $x, y, z \in X$, $d(x, z) \leq d(x, y) + d(y, z)$.

The set X together with d is a metric space. The smallest topology containing all the open balls $B(x, r) = \{y \in X : d(x, y) < r\}$ is called the metric topology on X induced by d .

Definition 1.4 (Continuous map) A map $f : X \rightarrow X'$ between two topological spaces X and X' is continuous if and only if the pre-image $f^{-1}(O') = \{x \in X : f(x) \in O'\}$ of any open set $O' \subset X'$ is an open set of X . Equivalently, f is continuous if and only if the pre-image of any closed set in X' is a closed set in X .

Definition 1.5 (Compact space) A topological space X is a compact space if any open cover of X admits a finite subcover, i.e. for any family $\{U_i\}_{i \in I}$ of open sets such that $X = \cup_{i \in I} U_i$ there exists a finite subset $J \subseteq I$ of the index set I such that $X = \cup_{j \in J} U_j$.

For metric spaces, compactness is characterized using sequences: a metric space X is compact if and only if any sequence in X has a convergent subsequence.

Definition 1.6 (Connected spaces) A topological space X is connected if it is not the union of two disjoint open sets: if O_1, O_2 are two open sets such that $X = O_1 \cup O_2$ then $O_1 = \emptyset$ or $O_2 = \emptyset$.

A topological space X is path-connected if for any $x, y \in X$ there exists a continuous map $\gamma : [0, 1] \rightarrow X$ such that $\gamma(0) = x$ and $\gamma(1) = y$.

A path-connected space is always connected, but the reverse is not true in general. See Exercise 1.1.

Euclidean spaces. The space $\mathbb{R}^d$, $d \geq 1$ and its subsets are examples of particular interest. All along the book, for $x = (x_1, \dots, x_d) \in \mathbb{R}^d$

$$\|x\| = \sum_{i=1}^d x_i^2$$

denotes the *Euclidean norm* on $\mathbb{R}^d$. It induces the *Euclidean metric* on $\mathbb{R}^d$: $d(x, y) = \|x - y\|$. The standard topology on $\mathbb{R}^d$ is the one induced by the Euclidean metric.

A subset $K \subset \mathbb{R}^d$ (endowed with the topology induced from the Euclidean one) is compact if and only if it is closed and bounded (Heine-Borel theorem).

1.2 Comparing topological spaces

There are many ways of measuring how close two objects are. We distinguish between topological and geometric criteria.

1.2.1 Homeomorphism, isotopy and homotopy equivalence

In topology, two topological spaces are considered to be the same when they are *homeomorphic*.

Definition 1.7 (Homeomorphism) *Two topological spaces X and Y are homeomorphic if there exists a continuous bijective map $h : X \rightarrow Y$ such that its inverse h^{-1} is also continuous. The map h is called a homeomorphism.*

As an example, a circle and a simple closed polygonal curve are homeomorphic. On the other hand, a circle and a segment are not homeomorphic. See Exercise 1.6.

The continuity of the inverse map in the definition is automatic in some cases. If U is an open subset of $\mathbb{R}^d$ and $f : U \rightarrow \mathbb{R}^d$ is an injective continuous map, then $V = f(U)$ is open and f is a homeomorphism between U and V by Brouwer's invariance of domain.¹ The domain invariance theorem may be generalized to manifolds: If M and N are topological k -manifolds without boundary and $f : U \rightarrow N$ is an injective continuous map from an open subset of M to N , then f is open and is a homeomorphism between U and $f(U)$.

If X is homeomorphic to the standard unit ball of $\mathbb{R}^d$, X is called a *topological ball*.

¹See T. Tao's blog <https://terrytao.wordpress.com/2011/06/13/brouwers-fixed-point-and-invariance-of-domain-theorems-and-hilberts-fifth-problem/>

The notions of compactity and connexity are preserved by homeomorphism. See Exercise 1.4.

Let h be a map between two topological spaces X and Y . If h is a homeomorphism onto its image, it is called an *embedding* of X in Y .

When the considered spaces X and Y are subspaces of $\mathbb{R}^d$, the notion of *isotopy* is stronger than the notion of homeomorphism to distinguish between spaces.

Definition 1.8 (Ambient isotopy) *An ambient isotopy between $X \subset \mathbb{R}^d$ and $Y \subset \mathbb{R}^d$ is a map $F : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ such that $F(., 0)$ is the identity map on $\mathbb{R}^d$, $F(X, 1) = Y$ and for any $t \in [0, 1]$, F is a homeomorphism of $\mathbb{R}^d$.*

Intuitively, the previous definition means that X can be continuously deformed into Y without creating any self-intersection or topological changes. The notion of isotopy is stronger than the notion of homeomorphism in the sense that if X and Y are isotopic, then they are obviously homeomorphic. Conversely two subspaces of $\mathbb{R}^d$ that are homeomorphic may not be isotopic. This is the case for a knotted and an unknotted torus embedded in $\mathbb{R}^3$ as the ones in Figure 1.1. Note that, although intuitively obvious, proving that these two surfaces are not isotopic is a non obvious exercise that requires some background in algebraic topology.




Figure 1.1: Two surfaces embedded in $\mathbb{R}^3$ homeomorphic to a torus that are not isotopic.

In general, deciding whether two spaces are homeomorphic is a very difficult task. It is sometimes more convenient to work with a weaker notion of

equivalence between spaces called *homotopy equivalence*.

Given two topological spaces X and Y , two maps $f_0, f_1 : X \rightarrow Y$ are *homotopic* if there exists a continuous map $H : [0, 1] \times X \rightarrow Y$ such that for all $x \in X$, $H(0, x) = f_0(x)$ and $H(1, x) = f_1(x)$. Homotopy equivalence is defined in the following way.

Definition 1.9 (Homotopy equivalence) *Two topological spaces X and Y have the same homotopy type (or are homotopy equivalent) if there exist two continuous maps $f : X \rightarrow Y$ and $g : Y \rightarrow X$ such that $g \circ f$ is homotopic to the identity map in X and $f \circ g$ is homotopic to the identity map in Y .*

As an example, the unit ball in an Euclidean space and a point are homotopy equivalent but not homeomorphic. A circle and an annulus are also homotopy equivalent - see Figure 1.2 and Exercises 1.8.




Figure 1.2: An example of two maps that are homotopic (left) and examples of spaces that are homotopy equivalent, but not homeomorphic (right).

Definition 1.10 (Contractible space) *A contractible space is a space that has the same homotopy type as a single point.*

For example, a segment, or more generally any ball in an Euclidean space $\mathbb{R}^d$ is contractible - see Exercise 1.7.

It is often difficult to prove homotopy equivalence directly from the definition. When Y is a subset of X , the following criterion reveals useful to prove homotopy equivalence between X and Y .

Proposition 1.11 *If $Y \subset X$ and if there exists a continuous map $H : [0, 1] \times X \rightarrow X$ such that:*

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

then X and Y are homotopy equivalent.

Definition 1.12 (Deformation retract) *If, in Proposition 1.11, the last property of H is replaced by the following stronger one*

$$\forall y \in Y, \forall t \in [0, 1], H(t, y) = y,$$

then H is called a deformation retract of X to Y .

A classical way to characterize and quantify topological properties and features of spaces is to consider their *topological invariants*. These are mathematical objects (numbers, groups, polynomials,...) associated to each topological space that have the property of being the same for homeomorphic spaces. The homotopy type is clearly a topological invariant: two homeomorphic spaces are homotopy equivalent. The converse is false: for example, a point and a segment are homotopy equivalent but are not homeomorphic. See Exercise 1.7. Moreover, most of the topological invariants considered in the sequel are indeed homotopy invariants, i.e. they are the same for spaces that are homotopy equivalent.

1.2.2 Hausdorff distance

The set of compact subsets of a metric space can be endowed with a metric, called the Hausdorff distance, that allows to measure how two compact subsets are far away from each other. We give here the definition for compact subspaces of $\mathbb{R}^d$ but it immediately adapts to the compact subsets of any metric space.

Definition 1.13 (Offset) *Given a compact set X of $\mathbb{R}^d$, the tubular neighborhood or offset X^ε of X of radius ε , i.e. the set of all points at distance at most ε from X :*

$$X^\varepsilon = \{y \in \mathbb{R}^d : \inf_{x \in X} \|x - y\| \leq \varepsilon\} = \bigcup_{x \in X} \bar{B}(x, \varepsilon)$$

where $\bar{B}(x, \varepsilon)$ denotes the closed ball $\{y \in \mathbb{R}^d : \|x - y\| \leq \varepsilon\}$.

Definition 1.14 The Hausdorff distance $d_H(X, Y)$ between two closed subsets X and Y of $\mathbb{R}^d$ is the infimum of the $\varepsilon \geq 0$ such that $X \subset Y^\varepsilon$ and $Y \subset X^\varepsilon$. Equivalently,

$$d_H(X, Y) = \max \left(\sup_{y \in Y} \left(\inf_{x \in X} \|x - y\| \right), \sup_{x \in X} \left(\inf_{y \in Y} \|x - y\| \right) \right).$$

The Hausdorff distance defines a distance on the space of compact subsets of $\mathbb{R}^d$. See Exercise 1.10.

1.3 Exercises

Exercise 1.1 Let X be a path connected space. Show that X is connected. Let $X \subset \mathbb{R}^2$ be the union of the vertical closed segment $\{0\} \times [-1, 1]$ and the curve $\{(t, \sin(\frac{1}{t})) \in \mathbb{R}^2 : t \in (0, 1]\}$. Show that X is compact and connected but not path-connected.

Exercise 1.2 Let S be a subset of a metric space X . Show that:

1. $x \in X \in \bar{S}$ if and only if for any $r > 0$, $B(x, r) \cap S \neq \emptyset$.
2. $x \in X \in \overset{\circ}{S}$ if and only if there exists $r > 0$ such that $B(x, r) \subset S$.

Exercise 1.3 Let X be a metric space. Given $x \in X$ and $r > 0$, show that the set $\bar{B}(x, r) = \{y \in X : d(x, y) \leq r\}$ is a closed set which is indeed the closure of the open ball $B(x, r) = \{y \in X : d(x, y) < r\}$.

Exercise 1.4 Let X, Y two homeomorphic topological spaces. Prove the following equivalences:

1. X is compact if and only if Y is compact.
2. X is connected (resp. path-connected) if and only if Y is connected (resp. path-connected).

Exercise 1.5 Show that the Euclidean space is not compact (without using the Heine-Borel theorem).

Exercise 1.6 A continuous polygonal curve $P \subset \mathbb{R}^2$ with consecutive edges $e_1 = [p_1, p_2], e_2 = [p_2, p_3], \dots, e_n = [p_n, p_{n+1}]$ is simple and closed if and only if $e_i \cap e_j = \emptyset$ whenever $2 \leq |i-j| \pmod{n}$, $e_i \cap e_{i+1} = p_{i+1}$ for $i = 1, \dots, n-1$ and $e_n \cap e_1 = p_1$. Show that P is homeomorphic to a circle. Show that a circle and a segment are not homeomorphic.

Exercise 1.7 Let X be a segment (i.e. a space homeomorphic to $[0, 1]$) and let Y be a point. Prove that X and Y are homotopy equivalent but not homeomorphic. More generally prove that any ball in $\mathbb{R}^d$ is contractible.

Exercise 1.8 Let X be the unit circle in $\mathbb{R}^2$ and let $Y \subset \mathbb{R}^2$ be the annulus of inner radius 1 and outer radius 2. Prove that X and Y are homotopy equivalent.

Exercise 1.9 Let X and Y be two topological spaces that are homotopy equivalent. Show that if X is path-connected, then Y is also path-connected.

Exercise 1.10 Show that the Hausdorff distance is a distance on the space of compact subsets of $\mathbb{R}^d$. Show that this is no longer true if we extend the definition to non compact sets (give an example of two different sets that are at distance 0 from each other).

1.4 Bibliographical notes

All the notions introduced in this chapter are classical, but fundamental, and presented with many details in the classical mathematical literature. For more details about basic topology, the reader may refer to any standard book on general topology such as, e.g. [114]. The geometry of metric spaces is a wide subject in mathematics. The reader interested in the topics may have a look at [30]. More details and results about the notions of homotopy and homotopy equivalence can be found in [87, pp. 171-172] or [113, p. 108].

Chapter 2

Simplicial complexes

Geometric shapes like curves, surfaces or their generalization in higher dimensions are “continuous” mathematical objects that cannot be directly encoded as a finite discrete structure usable by computers or computing devices. It is thus necessary to find representations of these shapes that are rich enough to capture their geometric structure and to comply with the constraints inherent to the discrete and finite nature of implementable data structures. On another side, when the only available data are point clouds sampled around unknown shapes, it is necessary to be able to build some continuous space on top of the data that faithfully encode the topology and the geometry of the underlying shape. Simplicial complexes offer a classical and flexible solution to overcome these difficulties.

2.1 Geometric simplicial complexes

The points of a finite set $P = \{p_0, p_1, \dots, p_k\}$ in $\mathbb{R}^d$ are said to be *affinely independent* if they are not contained in any affine subspace of dimension less than k .

Definition 2.1 (Simplex) *Given a set $P = \{p_0, \dots, p_k\} \subset \mathbb{R}^d$ of $k + 1$ affinely independent points, the k -dimensional simplex σ , or k -simplex for short, spanned by P is the set of convex combinations*

$$\sum_{i=0}^k \lambda_i p_i, \quad \text{with} \quad \sum_{i=0}^k \lambda_i = 1 \quad \text{and} \quad \lambda_i \geq 0.$$

The points $p_0, \dots, p_k$ are called the vertices of σ .

Notice that σ is the convex hull of the points P , i.e. the smallest convex subset of $\mathbb{R}^d$ containing $p_0, p_1, \dots, p_k$. A 0-simplex is a point, a 1-simplex is a line segment, a 2-simplex is a triangle, a 3-simplex is a tetrahedron.

The *faces* of the simplex σ whose vertex set is P are the simplices spanned by the subsets of P . Any face different from σ is called a *proper face* of σ .

For example, the faces of a triangle spanned by 3 independent points $\{p_0, p_1, p_2\} \in \mathbb{R}^d$ are the simplices $\emptyset, [p_0], [p_1], [p_0, p_1], [p_1, p_2], [p_2, p_0]$ and $[p_0, p_1, p_2]$. Observe that, by convention, $\emptyset$ is usually added to the faces as the simplex spanned by the empty subset of the vertices.

Definition 2.2 (Simplicial complex) A (finite) simplicial complex K in $\mathbb{R}^d$ is a (finite) collection of simplices such that:

1. any face of a simplex of K is a simplex of K ,
2. the intersection of any two simplices of K is either empty or a common face of both.

All the simplicial complexes considered in this book are finite. The simplices of K are called the *faces* of K . The *dimension* of K is the highest dimension of its simplices. A complex of dimension k is also called a k -complex. A subset of the simplices of K which is itself a simplicial complex is called a *subcomplex* of K . The j -skeleton $\text{Sk}_j(K)$ of K is the subcomplex of K consisting of the simplices of dimension at most j .




Figure 2.1: Left: an example of a simplicial complex. Right: a union of simplices which is not a simplicial complex

For a simplicial complex K in $\mathbb{R}^d$, its *underlying space* $|K| \subset \mathbb{R}^d$ is the union of the simplices of K . The *topology of K* is the topology induced on $|K|$ by the standard topology in $\mathbb{R}^d$. When there is no risk of confusion, we do not clearly make the distinction between a complex in $\mathbb{R}^d$ and its underlying space.

2.2 Abstract simplicial complexes

Notice that when its vertex set is known, a simplicial complex in $\mathbb{R}^d$ is fully and combinatorially characterized by the list of its simplices. This leads to the following notion of abstract simplicial complex.

Definition 2.3 (Abstract simplicial complex) Let $V = \{v_1, \dots, v_n\}$ be a finite set. An abstract simplicial complex $\tilde{K}$ with vertex set V is a set of finite subsets of V satisfying the two conditions :

1. The elements of V belong to $\tilde{K}$.
2. If $\tau \in \tilde{K}$ and $\sigma \subseteq \tau$, then $\sigma \in \tilde{K}$.

The elements of $\tilde{K}$ are called the simplices or the faces of $\tilde{K}$. If $\sigma \in \tilde{K}$ has precisely $k + 1$ elements, the dimension of σ is k and we say that σ is a k -simplex. The dimension of $\tilde{K}$ is the maximal dimension of its simplices.

Any simplicial complex K in $\mathbb{R}^d$ naturally determines an abstract simplicial complex $\tilde{K}$, called the *vertex scheme* of K : K and $\tilde{K}$ have the same set of vertices and the simplices of $\tilde{K}$ are the sets of vertices of the simplices of K . Conversely, if an abstract complex $\tilde{K}$ is the vertex scheme of a complex K in $\mathbb{R}^d$, then K is called a *geometric realization* of $\tilde{K}$. Notice that any finite abstract simplicial complex $\tilde{K}$ has a geometric realization in an Euclidean space in the following way. Let $\{v_1, v_2, \dots, v_n\}$ be the vertex set of $\tilde{K}$ where n is the number of vertices of $\tilde{K}$, and let $\sigma \subset \mathbb{R}^n$ be the simplex spanned by $\{e_1, e_2, \dots, e_n\}$ where, for any $i = 1, \dots, n$, e_i is the vector whose coordinates are all 0 except the i^{th} one which is equal to 1. Then K is the subcomplex of σ defined by $[e_{i_0}, \dots, e_{i_k}]$ is a k -simplex of K if and only if $[v_{i_0}, \dots, v_{i_k}]$ is a simplex of K . It can also be proven that any finite abstract simplicial complex of dimension d can be realized as a simplicial complex in $\mathbb{R}^{2d+1}$ (Exercice 2.3).

Definition 2.4 (Isomorphism of abstract simplicial complexes) Two abstract simplicial complexes $\tilde{K}, \tilde{K}'$ with vertex sets V and V' are isomorphic if there exists a bijection $\phi : V \rightarrow V'$ such that $\{v_0, \dots, v_k\} \in \tilde{K}$ if and only if $\{\phi(v_0), \dots, \phi(v_k)\} \in \tilde{K}'$.

The relation of isomorphism between two abstract simplicial complexes induces homeomorphism between their geometric realizations.

Proposition 2.5 If two simplicial complexes K, K' are the geometric realizations of two isomorphic abstract simplicial complexes $\tilde{K}, \tilde{K}'$, then $|K|$ and $|K'|$ are homeomorphic topological spaces. In particular, the underlying spaces of any two geometric realizations of an abstract simplicial complex are homeomorphic.

In this book, we will often encounter abstract simplicial complexes whose vertices are points in $\mathbb{R}^d$. Let $\tilde{K}$ be an abstract complex with vertex set $V \subset \mathbb{R}^d$. If the convex hull of each k -simplex $\sigma = \{v_0, \dots, v_k\} \in \tilde{K}$ is a geometric k -simplex in $\mathbb{R}^d$ and if the collection of these simplices defines a simplicial complex K , then we say that $\tilde{K}$ *naturally embeds* in $\mathbb{R}^d$ and that K is the *natural embedding* of $\tilde{K}$. When there is no ambiguity, the same notation is used for $\tilde{K}$ and K .

An important remark about terminology: as the underlying spaces of all geometric realizations of an abstract simplicial complex are homeomorphic to each other, it is usual to relate the topological properties of these underlying spaces to the complex itself. For example, when one claims that an abstract simplicial complex K is homeomorphic or homotopy equivalent to a topological space X , it is meant that the underlying space of any geometric realization of K is homeomorphic or homotopy equivalent to X .

2.3 Nerve

As noticed in previous the section, simplicial complexes can be seen at the same time as topological spaces and as purely combinatorial objects.

Definition 2.6 (Covers) *An open cover of a topological space X is a collection $\mathcal{U} = (U_i)_{i \in I}$ of open subsets $U_i \subseteq X$, $i \in I$ where I is a set, such that $X = \cup_{i \in I} U_i$. Similarly, a closed cover of X is a collection of closed sets whose union is X .*

Definition 2.7 (Nerve) *Given a cover of a topological space X , $\mathcal{U} = (U_i)_{i \in I}$ we associate an abstract simplicial complex $C(\mathcal{U})$ whose vertex set is $\mathcal{U}$ and such that*

$$\sigma = [U_{i_0}, U_{i_1}, \dots, U_{i_k}] \in C(\mathcal{U}) \text{ if and only if } \cap_{j=0}^k U_{i_j} \neq \emptyset.$$

Such a simplicial complex is called the nerve of the cover $\mathcal{U}$.

When all the sets U_i are open and all their finite intersections are contractible, i.e. are homotopy equivalent to a point, the Nerve Theorem relates the topology of X and $C(\mathcal{U})$.

Theorem 2.8 (Nerve Theorem) *Let $\mathcal{U} = (U_i)_{i \in I}$ be a finite open cover of a subset X of $\mathbb{R}^d$ such that any intersection of the U_i 's is either empty or contractible. Then X and $C(\mathcal{U})$ are homotopy equivalent.*

The nerve theorem also holds for closed covers under a slightly more restrictive assumption on X . The following version is general enough for our purpose.

Theorem 2.9 (Nerve Theorem for convex covers) *Let $X \subset \mathbb{R}^d$ be a finite union of closed convex sets $\mathcal{F} = (F_i)_{i \in I}$ in $\mathbb{R}^d$. Then X and $C(\mathcal{F})$ are homotopy equivalent.*

A cover satisfying the assumptions of the Nerve Theorem is sometimes called a *good cover*. The Nerve Theorem is of fundamental importance in computational topology and geometric inference: it provides a way to encode the homotopy type of continuous topological space X by a simplicial complex describing the intersection pattern of a good cover. In particular, when X is a (finite) union of (closed or open) balls in $\mathbb{R}^d$, it is homotopy equivalent to the nerve of this union of balls.

2.4 Filtrations of simplicial complexes

Simplicial complexes often come with a specific ordering of their simplices that plays a fundamental role in geometry inference.

Definition 2.10 *A filtration of a finite simplicial complex K is a nested sequence of sub-complexes $\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$ such that*

$$K^{i+1} = K^i \cup \sigma^{i+1} \quad \text{where } \sigma^{i+1} \text{ is a simplex of } K.$$

Equivalently, a filtration of K can be seen as an ordering of the simplices such that for any $i \geq 0$, the collection of the first i simplices is a simplicial complex. To ensure this later condition, it is sufficient to know that every simplex σ^i appears in the filtration after all its faces.

As a filtration of K is just an ordering of the simplices, in some cases, it might be more natural to index the simplices by an increasing sequence

$(\alpha_i)_{i=1}^m$ of real numbers: $\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \dots \subset K^{\alpha_m} = K$. In this case, it is often convenient to extend the filtration to the whole set of real numbers by defining $K^\alpha = K^{\alpha_i}$ for $\alpha \in [\alpha_i, \alpha_{i+1})$, $K^\alpha = \emptyset$ for $\alpha < \alpha_0$ and $K^\alpha = K$ for $\alpha \geq K^{\alpha_m}$.

For example, when a function is defined on the vertices of K , one can define a sublevel set filtration in the following way.

Filtration associated to a function defined on the vertices of a complex. Let K be a simplicial complex and let f be a real valued function defined on the vertices of K . For any simplex $\sigma = \{v_0, \dots, v_k\}$ one defines $f(\sigma)$ by

$$f(\sigma) = \max_{i=0 \dots k} f(v_i)$$

Ordering the simplices of K according to the values of each simplex defines a filtration of K . Note that different simplices can have the same value. In this case, they are ordered according to increasing dimension and simplices of the same dimension with same value can be ordered arbitrarily. The filtration induced by f is the filtration by the sublevel sets $f^{-1}(]-\infty; t])$ of f .

2.5 Vietoris-Rips and Čech filtrations

Filtrations are often built on top of finite sets of points to reveal the underlying topological structure of data (see Chapter 11). Let $P \subset \mathbb{R}^d$ be a (finite) set of points.

Definition 2.11 *Given $\alpha > 0$, the Čech complex with vertex set P and parameter α is the nerve $\check{\text{Cech}}(P, \alpha)$ of the unions of balls centered on P with radius α . The simplices of $\check{\text{Cech}}(P, \alpha)$ are characterized by the following condition:*

$$\{x_0, x_1, \dots, x_k\} \in \check{\text{Cech}}(P, \alpha) \quad \Leftrightarrow \quad \bigcap_{i=0}^k B(x_i, \alpha) \neq \emptyset.$$

As α goes from 0 to $+\infty$, the nested sequence of complexes $\check{\text{Cech}}(P, \alpha)$ defines the Čech complex filtration.

Given a k -dimensional face σ of the simplex of dimension $|P|-1$, the smallest α such that $\sigma \in \check{\text{Cech}}(P, \alpha)$ is the radius of the smallest ball enclosing the vertices of σ . As a consequence, the k -dimensional skeleton of the Čech filtration can be computed by computing the $O(|P|^k)$ minimum enclosing balls of all the subsets of at most k points of P . Although the computation of the minimum ball enclosing a set of k points can be done in time $O(k)$ (see the bibliographic notes), the computation of the whole Čech filtration quickly becomes intractable in practice. Given $\alpha > 0$, the computation of the k -skeleton of $\check{\text{Cech}}(P, \alpha)$ can be done by first computing all the cliques of at most $(k+1)$ vertices of the 1-skeleton of $\check{\text{Cech}}(P, \alpha)$ which is a graph, and second by selecting the cliques whose minimum enclosing ball has its radius upper bounded by α .




Figure 2.2: The Čech (left) and Vietoris-Rips (right) complexes built on top of a finite set of points in $\mathbb{R}^2$. Note that they both contain a 3-simplex and are thus not embedded in $\mathbb{R}^2$.

A simplicial complex which is closely related to the Čech filtration is the *Vietoris-Rips filtration*, $\text{Rips}(P)$.

Definition 2.12 Given $\alpha > 0$, the *Vietoris-Rips complex* $\text{Rips}(P, \alpha)$ with vertex set P and parameter α is defined by the following condition

$$\{x_0, x_1, \dots, x_k\} \in \text{Rips}(P, \alpha) \Leftrightarrow \|x_i - x_j\| \leq \alpha \quad \text{for all } i, j \in \{0, \dots, k\}.$$

As α goes from 0 to $+\infty$, the nested sequence of complexes $\text{Rips}(P, \alpha)$ defines the *Vietoris-Rips filtration*.

The Vietoris-Rips complex is much simpler to compute than the Čech filtration as it just involves distance comparisons. The Vietoris-Rips complex is the largest simplicial complex that has the same 1-skeleton as the Čech complex. It is thus completely characterized by its 1-skeleton. The whole

k -dimensional skeleton of the Vietoris-Rips filtration can be computed by computing the diameter of all the subsets of at most k points of P .

The Čech and the Vietoris-Rips filtrations are related by the following interleaving property that plays a fundamental role in Chapter 11.

Lemma 2.13 *Let P be a finite set of points in $\mathbb{R}^d$. for any $\alpha \geq 0$,*

$$\text{Rips}(P, \alpha) \subseteq \check{\text{Cech}}(P, \alpha) \subseteq \text{Rips}(P, 2\alpha)$$

Proof If $\sigma = \{x_0, x_1, \dots, x_k\} \in \text{Rips}(P, \alpha)$ then $x_0 \in \bigcap_{i=0}^k B(x_i, \alpha)$. So, $\sigma \in \check{\text{Cech}}(P, \alpha)$. This proves the first inclusion.

Now, if $\sigma = \{x_0, x_1, \dots, x_k\} \in \check{\text{Cech}}(P, \alpha)$, there exists $y \in \mathbb{R}^d$ such that $y \in \bigcap_{i=0}^k B(x_i, \alpha)$, i.e. $\|x_i - y\| \leq \alpha$ for any $i = 0, \dots, k$. As a consequence, for all $i, j \in \{0, \dots, k\}$, $\|x_i - x_j\| \leq 2\alpha$ and $\sigma \in \text{Rips}(P, 2\alpha)$. $\square$

Remark that the Čech and Vietoris-Rips filtrations can be defined for a set of points in any metric space and that the above interleaving property still holds. When the points P are in $\mathbb{R}^d$, the interleaving of Lemma 2.13 is not tight and can be slightly improved (see Exercise 11.3).

2.6 Combinatorial manifolds triangulations

Definition 2.14 (Star and link) *Let K be a simplicial complex with vertex set P . The star of $p \in P$ is the set of simplices of K that have p as a vertex. We denote it $\text{star}(p, K)$. The link of p is the set of simplices $\tau \subset \sigma$ such that $\sigma \in \text{star}(p, K)$ but $\tau \notin \text{star}(p, K)$. We denote it by $\text{link}(p, K)$.*

Observe that the star of a simplex is *not* a complex while the link is. We will use the name *closed star* of p in K to denote the subcomplex of K that consists of the simplices of $\text{star}(p, K)$ and their subfaces.

Definition 2.15 (Pure complex) *A simplicial k -complex K is pure if every simplex in K is the face of a k -simplex.*

Definition 2.16 (Boundary complex) *Let K be a pure simplicial k -complex. The boundary of K , denoted ∂K is the $(k-1)$ -subcomplex of K whose*

$(k-1)$ -simplices are the $(k-1)$ -simplices of K that are incident to only one face of dimension k .

Definition 2.17 (Combinatorial manifold) *A simplicial complex K is a combinatorial k -manifold if*

1. K is pure k complex,
2. the link of any vertex of $K \setminus \partial K$ is a triangulated $(k-1)$ -sphere
3. the link of any vertex of ∂K is a triangulated $(k-1)$ -ball.

We define the *adjacency graph* of a combinatorial k -manifold K as the graph whose nodes are the k -simplices of K and two nodes are joined by an edge in the graph if the two simplices associated to the two nodes have a $(k-1)$ -simplex in common.

An example of a combinatorial manifold is the boundary complex of a polytope (see Section 3.1).

Definition 2.18 (Triangulation of a point set) *A triangulation of a finite point set $P \in \mathbb{R}^d$ is a geometric simplicial complex K whose vertex set is P and whose underlying space is the convex hull of P .*

The triangulation of a finite point set $P \in \mathbb{R}^d$ is a combinatorial manifold whose boundary is the boundary complex of the convex hull of P (see Exercise 3.9).

Definition 2.19 (Triangulation of a topological space) *A triangulation of a topological space X is a simplicial complex K and a homeomorphism $h : K \rightarrow |X|$.*

2.7 Representation of simplicial complexes

To represent a simplicial complex K , we need a data structure that represents the simplices of the complex and is able to provide efficient implementations of elementary operations such as face and coface retrieval, and

maintenance of the data structure upon elementary modifications of the complex, insertion or removal of a new simplex. In addition, we may want to attach a filtration to each simplex and iterate over the simplices of a filtered complex by increasing values of filtration. We present below several data structures to represent simplicial complexes.

We say that two simplices σ and τ of a simplicial complex are *incident* if one is included in the other. We say that σ and τ are *adjacent* if they are incident and if their dimensions differ by exactly 1. The *Hasse diagram* of a simplicial complex K is a graph whose nodes represent the simplices (of all dimensions) and two nodes are joined by an arc if the associated simplices are adjacent.

The Hasse diagram provides an explicit representation of all the simplices of K . If the dimension of the complex is considered as a constant, it is easy to see that all elementary operations on K can be performed efficiently. However the size of the Hasse diagram may be problematic in practice when considering big complexes. A more compact data structure that still provides an explicit representation of all the simplices of K is the so-called *simplex tree* which is a minimal spanning tree of the Hasse diagram. The simplex tree is constructed as a prefix tree of the simplices considered as words on the alphabet $v_1, \dots, v_n$ where $v_1, \dots, v_n$ are the labels of the vertices of K (see Figure 2.3 and Exercice 2.1).

Both the Hasse diagram and the simplex tree are convenient to store and retrieve information attached to a simplex, e.g. a filtration value. Nevertheless, in some applications, we may content ourselves with less expressive but more compact representations. Instead of representing all simplices, one can represent only the maximal ones, i.e. the simplices that have no coface in the complex. This may lead to a dramatic saving in memory size as can be seen from the following simple example : the complex consisting of a unique simplex has one maximal simplex but 2^d simplices in total.

Representing only the maximal simplices is especially well suited for combinatorial manifold. A combinatorial manifold K can be represented by its maximal simplices, together with the adjacency graph of those maximal simplices. We will simply call this graph the *adjacency graph of the combinatorial manifold*. When K is of dimension d , the adjacency graph of the maximal simplices is a connected graph of degree $d + 1$ that be efficiently traversed. Such a graph will be used in Section 3.4 when constructing convex hulls.




Figure 2.3: A Hasse diagram and a Simplex Tree with their associated simplicial complexes.

2.8 Exercises

Exercise 2.1 A simplicial complex C is said to be (path-)connected if for any pair of points $(x, y) \in C$ there exists a continuous path $\gamma : [0, 1] \rightarrow C$ such that $\gamma(0) = x$ and $\gamma(1) = y$. Prove that a simplicial complex C is connected if and only if its 1-skeleton is connected.

Exercise 2.2 Give examples of simplicial complexes in $\mathbb{R}^3$ that are homeomorphic to a ball, a sphere, and a torus.

Exercise 2.3 Prove that any abstract simplicial complex K of dimension d can be realized as a geometric simplicial complex in $\mathbb{R}^{2d+1}$. (Hint : map the vertices of K to points on the moment curve $C = \{(x, x^2, \dots, x^{2d+1}) \in \mathbb{R}^{2d+1}, x \in \mathbb{R}\}$. Show that any subset of $2d + 2$ points on C are affinely independent and that the image of K is a realization of K in $\mathbb{R}^{2d+1}$. See also Exercise 3.12.)

Exercise 2.1 (Simplex Tree) Let K be a d -simplicial complex whose vertices are labelled $1, \dots, n$. Each simplex is represented by a word which is

the sorted list of its labels. We then store all simplices of K as a prefix tree called the *simplex tree* of K . The simplex tree is characterized by the following properties:

1. The nodes of the simplex tree are in bijection with the simplices (of all dimensions) of the complex. The root is associated to the empty face.
2. Each node of the tree, except the root, stores the label of a vertex. Specifically, a node associated to a simplex σ stores the last label of σ .
3. The vertices whose labels are encountered along a path from the root to a node associated to a simplex σ are the vertices of σ . Along such a path, the labels are sorted by increasing order and each label appears no more than once.

Show that the simplex tree of K is a spanning tree of the Hasse diagram of K . Let σ be a i -simplex. Show that we can decide whether σ is in K or insert σ in K (assuming all its subfaces are in K) in time $O(i \log n)$. Show also how to locate the faces and the cofaces of σ in the simplex tree and how to remove σ from K .

Exercise 2.2 Show how to compute the Vietoris-Rips filtration of a set of points P in some metric space using the simplex tree.

2.9 Bibliographical notes

Our presentation of simplicial complexes follows the one in Munkres [112]. The nerve theorem and its variants are classical results in algebraic topology. A proof is given in Hatcher [97, Section 4G].

The computation of the Čech filtration of a finite set of points relies on the computation of minimum enclosing balls. Welzl has proposed an elegant randomized algorithm of linear complexity to compute the minimal ball enclosing a set of n points. The algorithm can be adapted to compute the minimal enclosing ellipsoid [133]. The Vietoris-Rips filtration can be obtained by computing the cliques in the 1-skeleton of the Vietoris-Rips complex. This is an NP-complete problem but efficient solutions exist for sparse graphs [79]. The Vietoris-Rips filtration can be constructed and stored using a simplex tree, a data structure proposed by Boissonnat and Maria [23].

The simplex tree is implemented in the GUDHI library [106]. More compact data structures can be found in [5, 15].

The problem of deciding if two abstract complexes are homeomorphic is undecidable Markov 1958. Deciding if a d -simplicial complex is a homeomorphic to a d -sphere is undecidable for $d \geq 5$ [Novikov]. It is trivial for $d = 2$ and decidable for $d = 3$. The question remains open $d = 4$.

Part II

Delaunay complexes

Chapter 3

Convex polytopes

Convex polytopes and convex polyhedra are fundamental structures that play an essential role in computational geometry, linear programming, optimisation and many other fields. In this book, they also play a central role because of their close relationship with Voronoi diagrams and Delaunay complexes, introduced in Chapter 4 and used throughout the book. This chapter presents here the combinatorial and algorithmic aspects of convex polyhedra to be used in the following chapters. For a more complete treatment of the rich theory of convex polytopes, the reader is referred to classical textbooks (see the bibliographic notes).

3.1 Definitions

3.1.1 Convex polytopes

Definition 3.1 *A convex polytope in $\mathbb{R}^d$ is the convex hull of a finite set of points. In this book, we only consider convex polytopes and the word polytope is used as a synonym for convex polytope.*

Hence, a convex polytope is a closed bounded subset of $\mathbb{R}^d$. The *dimension* of a convex polytope is the dimension of the affine subspace spanned by the polytope. Simplices (see Section 2) are particular cases of polytopes.

3.1.2 Facial structure of polytopes

A hyperplane h of $\mathbb{R}^d$ is a subset of $\mathbb{R}^d$ defined by a linear equation :

$$h = \{x \in \mathbb{R}^d : h(x) = a \cdot x + b = 0\},$$

where $a \in \mathbb{R}^d, b \in \mathbb{R}$. A hyperplane h divides $\mathbb{R}^d$ in two half-spaces:

$$\begin{aligned} h^+ &= \{x \in \mathbb{R}^d : h(x) = a \cdot x + b \geq 0\} \\ h^- &= \{x \in \mathbb{R}^d : h(x) = a \cdot x + b \leq 0\} \end{aligned}$$

Note the half-spaces h^+ and h^- are defined as closed subsets, so that h , h^+ and h^- do not form a partition of $\mathbb{R}^d$ but h , $h^+ \setminus h$ and $h^- \setminus h$ do.

Definition 3.2 *A hyperplane h of $\mathbb{R}^d$ is a supporting hyperplane of the polytope $\mathcal{P}$ iff the intersection $\mathcal{P} \cap h$ is non empty and $\mathcal{P}$ is included in one of the two half-spaces defined by h .*

The intersection of $\mathcal{P}$ with a supporting hyperplane h is called a face of $\mathcal{P}$.

Lemma 3.3 *A face of a polytope is a polytope.*

Proof Let $P = \{p_1, \dots, p_n\}$ be a finite set of points in $\mathbb{R}^d$ and let $\mathcal{P} = \text{conv}(P)$ be the polytope which is the convex hull of P . Any face f of $\mathcal{P}$ is the intersection of $\mathcal{P}$ with a supporting hyperplane h . We prove that the face $f = \mathcal{P} \cap h$ is the convex hull $\text{conv}(P \cap h)$ of the subset of points of P included in h . The inclusion $\text{conv}(P \cap h) \subset f$ comes from the convexity of f . Indeed, f is convex since it is the intersection of two convex sets. Furthermore f contains $P \cap h$, so it contains the convex hull $\text{conv}(P \cap h)$. To prove the reverse inclusion, we assume without loss of generality that $\mathcal{P} \subset h^+$. Then we have $h(p_i) \geq 0$ for any $p_i \in P$. Any point p included in $\mathcal{P}$ is a convex combination $p = \sum_{i=1}^n \lambda_i p_i$ of points in P . If furthermore p belongs to f , $h(p) = \sum_{i=1}^n \lambda_i h(p_i)$ vanishes and this can happen only if $\lambda_i = 0$ for all points p_i that are not in h . Thus, if p belongs to f , it belongs to $\text{conv}(P \cap h)$. $\square$

The proof of the previous lemma, shows that the faces of the polytope $\mathcal{P}$ are the convex hulls of subsets of P . A polytope has therefore a finite number of faces. The boundary of the polytope $\mathcal{P}$ is the union of its faces. The faces of dimension 0 are called *vertices*. The faces of dimension 1 are called *edges*. If $\mathcal{P}$ has dimension d , the faces of dimension $d-1$ and $d-2$ are called respectively *facets* and *ridges*.

The vertices of the polytope $\mathcal{P} = \text{conv}(P)$ are points of P . The following lemma, whose proof is left as an exercise (Exercise 3.1) is a well known result of the theory of polytopes.

Lemma 3.4 *Any polytope is the convex hull of its vertices.*

The facial structure of a simplex can be easily described. Indeed, if σ is a simplex that is the convex hull of the set $S = \{p_1, \dots, p_k\}$ of $k+1$ independent points, any subset of S is a set of independent points whose convex hull is a simplex and, except in the case of S itself, this simplex is a face of σ .

Therefore a k -simplex has $k+1$ -vertices, and $\binom{k+1}{j+1}$ faces of dimension j , for $j = 0$ to $k-1$.

Another fundamental result of the theory of polytopes is the following lemma whose proof is also left as an exercise (see Exercises 3.4 and 3.5).

Lemma 3.5 *Any polytope is the intersection of a finite number of half-spaces. Reciprocally, any intersection of a finite number of half-spaces that is bounded, is a polytope.*

Therefore any polytope can be described either as the convex hull of a finite set of points or as the intersection of a finite number of half-spaces. More precisely, the proof of Lemma 3.5 shows that the minimal set of half-spaces whose intersection is identical to the polytope $\mathcal{P}$ is the set of half-spaces bounded by the hyperplanes supporting $\mathcal{P}$ along its facets and containing $\mathcal{P}$.

3.1.3 Convex polyhedra

In the sequel, we extend the notion of polytopes to describe unbounded intersections of half-spaces as well as polytopes.

Definition 3.6 *A convex polyhedron is the intersection of a finite number of half-spaces.*

A convex polyhedron may be bounded or not. From Lemma 3.5, a convex polytope is just a special case of convex polyhedron. The notion of supporting hyperplanes and faces introduced above for convex polytopes extend naturally to convex polyhedra. The faces of a convex polyhedron are themselves convex polyhedra and may be unbounded if the convex polyhedron is unbounded.

3.1.4 Simplicial polytopes and simple convex polyhedra

A point set P in $\mathbb{R}^d$ is said to be in *general position* when any subset of P with size at most $d + 1$ is a set of affinely independent points. When the points of P are in general position, any hyperplane h includes at most d points of P and the points in $P \cap h$ are affinely independent. Therefore all the faces of the polytope $\mathcal{P} = \text{conv}(P)$ are simplices and the polytope $\mathcal{P}$ is called a *simplicial polytope*.

A set of n hyperplanes in $\mathbb{R}^d$ is said to be *in general position* if the intersection of any subset of k of them is an affine subspace of dimension $d - k$. A

convex polyhedron defined as the intersection of n half-spaces bounded by hyperplanes in general position is called *simple*.

3.1.5 The boundary complex

Let $\mathcal{P}$ be a convex polyhedron. Any face of a face of $\mathcal{P}$ is a face of $\mathcal{P}$ and the intersection of two faces of $\mathcal{P}$ is either empty or a common face of both faces.

Therefore, if $\mathcal{P}$ is simplicial polytope, its faces form a geometric simplicial complex (see Chapter 2). This complex is called the *boundary complex* of $\mathcal{P}$. It can be shown (see Exercise 3.9) that the boundary complex of a simplicial polytope is a combinatorial manifold (see Definition 2.17).

If $\mathcal{P}$ is not a simplicial polytope, its faces may not be simplices. Still, they are convex polyhedra and they form a cell complex as defined now.

Definition 3.7 (Cell complex) *A cell complex is a set C of convex polyhedra, called the faces of C , that satisfies the two following properties*

- *Any face of a face of C is a face of C .*
- *The intersection of any two faces of C is either empty or a common face of both faces.*

The cell complex formed by the faces of a convex polyhedron $\mathcal{P}$ is still called the *boundary complex* of $\mathcal{P}$. We adapt to cell complexes and therefore to convex polyhedra the notions of incidence and adjacency defined in Section 2.7:

- Two faces of a convex polyhedron are said to be *incident* if one is included in the other.
- Two facets of a convex polyhedron are said to be *adjacent* if they share a $(d - 2)$ -subface.

3.2 Duality

3.2.1 Point-hyperplane duality

We introduce a duality between points and hyperplanes that makes use of the unit paraboloid of $\mathbb{R}^d$ and gives a special role to the last coordinate axis, called the *vertical axis*. We denote by $x(i), i = 1, \dots, d$, the coordinates of a point $x \in \mathbb{R}^d$.

The unit paraboloid $\mathcal{Q}$ is defined as :

$$\mathcal{Q} = \left\{ x \in \mathbb{R}^d : x(d) = \sum_{i=1}^{d-1} x(i)^2 \right\}.$$

In the rest of this section, we write for short $\sum'$ for $\sum_{i=1}^{d-1}$.

Let p be a point of the paraboloid $\mathcal{Q}$. The duality associates to point p the hyperplane p^* tangent to $\mathcal{Q}$ at p :

$$\begin{aligned} p^* &= \{x \in \mathbb{R}^d : x(d) = 2\sum' p(i)(x(i) - p(i)) + p(d)\} \\ &= \{x \in \mathbb{R}^d : x(d) - 2\sum' p(i)x(i) + p(d) = 0\}. \end{aligned}$$

More generally, duality associates to any point p of $\mathbb{R}^d$ the non vertical hyperplane p^* defined by:

$$p^* = \{x \in \mathbb{R}^d : x(d) - 2\sum' p(i)x(i) + p(d) = 0\}.$$

Conversely, let h be a non vertical hyperplane of $\mathbb{R}^d$. The equation of h can be written in normal form:

$$h = \{x \in \mathbb{R}^d : x(d) + \sum' h(i)x(i) + h(d) = 0\},$$

and duality associates to h the point h^* with coordinates

$$\left(\frac{-h(1)}{2}, \dots, \frac{-h(d-1)}{2}, h(d) \right).$$

Since $p^{**} = p$, duality is an involutive bijection between points of $\mathbb{R}^d$ and non vertical hyperplanes of $\mathbb{R}^d$. Duality preserves incidences of points and

hyperplanes: if p and q are two points in $\mathbb{R}^d$ with dual hyperplanes p^* and q^* respectively, we have

$$q \in p^* \iff q(d) - 2\sum' p(d)q(d) + p(d) = 0 \iff p \in q^*.$$

Let h be a non vertical hyperplane whose equation in normal form is $h(x) = 0$. We say that point p is above h or that h is below p if $h(p) > 0$. We say that point p is below h or that h is above p if $h(p) < 0$. For a non vertical hyperplane h , we call *upper half-space* and denote by h^+ the half-space bounded by h that is above h . We call *lower half-space* and denote by h^- the half-space below h :

$$\begin{aligned} h^+ &= \{x \in \mathbb{R}^d : h(x) > 0\} \\ h^- &= \{x \in \mathbb{R}^d : h(x) < 0\}. \end{aligned}$$

Duality reverses the above-below relations between points and hyperplanes : if p and q are two points in $\mathbb{R}^d$ with dual hyperplanes p^* and q^* respectively, we have:

$$\begin{aligned} q \in p^{*+} &\iff q(d) - 2\sum' p(d)q(d) + p(d) > 0 \iff p \in q^{*+} \\ q \in p^{*-} &\iff q(d) - 2\sum' p(d)q(d) + p(d) < 0 \iff p \in q^{*-}. \end{aligned}$$

3.2.2 Lower hulls and upper envelope

Let $P = \{p_1, \dots, p_n\}$ be a set of n points in $\mathbb{R}^d$. The lower hull of P , denoted by $\text{lowerhull}(P)$, is a subcomplex of the boundary complex of the polytope $\mathcal{P} = \text{conv}(P)$. Let $H(P)$ be the set of hyperplanes supporting the polytope $\mathcal{P}$. We distinguish the subset $H_+(P)$ of lower supporting hyperplanes where a hyperplane h of $H(P)$ is a lower supporting hyperplane if h is a non vertical hyperplane and $\mathcal{P}$ is included in the upper half-space h^+ . The *lower hull* of P is then defined as the subcomplex of the convex hull boundary, formed by faces of $\mathcal{P}$ included in lower supporting hyperplanes:

$$\text{lowerhull}(P) = \{\mathcal{P} \cap h : h \in H_+(P)\}$$

Consider now a set $H = \{h_1, \dots, h_n\}$ of non vertical hyperplanes in $\mathbb{R}^d$, and let $\mathcal{H}^+$ be the convex polyhedron which is the intersection of the upper half-spaces bounded by the hyperplanes of H :

$$\mathcal{H}^+ = h_1^+ \cap h_2^+ \dots \cap h_n^+.$$

The boundary complex of the convex polyhedron $\mathcal{H}^+$ is called the *upper envelope* of the set of hyperplanes H .

Two complexes K and K^* (simplicial or cellular complexes) of dimension d are said to be *dual complexes* if there is a bijective correspondance between the faces of K and the faces of K^* such that

- the k -faces of K correspond to the $(d - k)$ -faces of K^*
- the correspondance preserves incidences and reverses inclusion relations, meaning that if $f \subset g$ are incident faces of K , the corresponding faces f^* and g^* of K^* are incident and such that $g^* \subset f^*$.

Observe that the lower envelope of a set of points in $\mathbb{R}^d$, and the upper envelope of a set of hyperplanes are both $(d - 1)$ -complexes. The duality between points and hyperplanes of $\mathbb{R}^d$ introduced in Section 3.2.1 yields a duality between the lower hull of a set of points and the upper envelope of the set of dual hyperplanes.

Lemma 3.8 (Lower hull - upper envelope duality) *Let P be a set of points in $\mathbb{R}^d$ and let P^* be the set of dual hyperplanes. The lower hull of P and the upper envelope of P^* are dual $(d - 1)$ -complexes.*

Proof We assume here for simplicity that the points of P are in general position and leave the proof for the general case as an exercise (Exercise 3.10).

Let $\mathcal{P} = \text{conv}(P)$ be the convex hull of P and by $\mathcal{P}^{*+}$ be the convex polyhedron which is the intersection of the upper half-spaces bounded by the dual hyperplanes:

$$\mathcal{P}^{*+} = p_1^{*+} \cap p_2^{*+} \dots \cap p_n^{*+}.$$

Let f be a k -face of $\text{lowerhull}(P)$. Face f is a k -face of the polytope $\mathcal{P}$ and, since general position is assumed, f includes $(k + 1)$ -vertices of $\mathcal{P}$ which are points of P . Let us assume, without loss of generality, that $\{p_1, \dots, p_{k+1}\}$ are included in f . We consider a lower supporting hyperplane h that intersects

$\mathcal{P}$ along f . Hyperplane h includes $\{p_1, \dots, p_{k+1}\}$ and, since it is a lower supporting hyperplane, the other points of P are included in the upper half-space h^+ . Let h^* be the point dual to h . Then, from the properties of the point hyperplane duality introduced above, we get:

$$p_i \in h^+ \iff h^* \in p_i^{*+}, \quad \forall i = 1, \dots, n \quad (3.1)$$

$$p_i \in h \iff h^* \in p_i^*, \quad \forall i = 1, \dots, k+1 \quad (3.2)$$

$$p_i \notin h \iff h^* \notin p_i^*, \quad \forall i = k+2, \dots, n \quad (3.3)$$

Equations 3.1, 3.2 and 3.3 show that the point h^* belongs to the $(d-1-k)$ -face f^* of $\mathcal{P}^{*+}$ which is the intersection of $\mathcal{P}^{*+}$ with the $k+1$ hyperplanes $\{p_i^*, i = 1, \dots, k+1\}$. Therefore, duality maps the k -face $f = \text{conv}(\{p_1, \dots, p_{k+1}\})$ of $\text{lowerhull}(P)$ to the $(d-1-k)$ -face $f^* = p_1^* \cap p_2^* \cap \dots \cap p_{k+1}^* \cap \mathcal{P}^{*+}$ of $\mathcal{P}^{*+}$.

In the other way, any $(d-1-k)$ -face of $\mathcal{P}^{*+}$ is the intersection of $\mathcal{P}^{*+}$ with $k+1$ hyperplanes in P^* (see Exercise 3.6). Let us consider a $(d-1-k)$ -face f^* of $\mathcal{P}^{*+}$ and say that f^* is the intersection of $\mathcal{P}^{*+}$ with the k hyperplanes $\{p_i^*, i = 1, \dots, k+1\}$. Equations 3.1, 3.2 and 3.3 show that any point h^* in f^* is the dual of a hyperplane h which supports $\mathcal{P}$ along the face $f = \text{conv}(\{p_1, \dots, p_{k+1}\})$. Moreover, h is a lower supporting hyperplane. Therefore, duality maps back f^* to f and the correspondance is bijective. $\square$

3.3 Combinatorial bounds

The following theorem bounds the total number of faces of a convex polyhedron in $\mathbb{R}^d$ which is known to have either n facets or n vertices. The theorem is known as the *upper bound theorem*, though the bound is tight in the worst case.

Theorem 3.9 (Upper Bound Theorem) *The total number of faces of a convex polyhedron in $\mathbb{R}^d$, defined as the intersection of n half-spaces or as the convex hull of n points, is $\Theta\left(n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

Proof We prove here the upper bound. The lower bound is the topic of Exercise 3.12.

Let $\mathcal{P}$ be a convex polyhedron defined as the intersection of n half-spaces of $\mathbb{R}^d$. To prove the upper bound on the number of faces of $\mathcal{P}$, we may assume that the hyperplanes bounding the half-spaces defining $\mathcal{P}$ are in general position. Indeed, otherwise, we can slightly perturb those hyperplanes to bring them in general position. During this process, the number of faces of $\mathcal{P}$ can only increase, see Exercise 3.11. Hence, any upper bound that is valid for hyperplanes in general position holds for any set of hyperplanes. We also assume that the polyhedron has at least one vertex. Indeed, otherwise the number of bounding hyperplanes is at most $d - 1$ and the total number of faces is at most 2^{d-1} . In addition, we may assume, without loss of generality, that two vertices of $\mathcal{P}$ do not have the same x_d coordinate.

We first bound the number of vertices of $\mathcal{P}$ and then extend this bound to faces of any dimension. Let p be one of the vertices of $\mathcal{P}$. Because general position is assumed, p , as any other vertex of $\mathcal{P}$, is included in exactly d of the bounding hyperplanes. Therefore p is incident to exactly d facets and d edges of $\mathcal{P}$. Thus at least $\lceil \frac{d}{2} \rceil$ edges incident to p are included in either in the half-space $h^+ : x_d \geq x_d(p)$ or in the half-space $h^- : x_d \leq x_d(p)$. Since the bounding hyperplanes are in general position, $\mathcal{P}$ is simple and the affine hull of any subset of $k < d$ edges incident to a vertex of $\mathcal{P}$ contains a k -face of $\mathcal{P}$ (Exercise 3.7). Therefore, each vertex p of $\mathcal{P}$ is a vertex with extremal x_d -coordinate for at least one face of dimension $\lceil \frac{d}{2} \rceil$. Since any face has at most one vertex of maximal x_d coordinate and one vertex of minimal x_d -coordinate, the number of vertices of $\mathcal{P}$ is at most twice the number of $\lceil \frac{d}{2} \rceil$ -faces of $\mathcal{P}$.

From the general position assumption, a k -face of $\mathcal{P}$ is the intersection of $d - k$ of the bounding hyperplanes (see Exercise 3.6). We deduce that the number of k -faces is at most $\binom{n}{d - k} = O(n^{d-k})$, which is $O(n^{\lfloor \frac{d}{2} \rfloor})$ for $k = \lceil \frac{d}{2} \rceil$. From the above discussion, we conclude that the number of vertices of $\mathcal{P}$ is $O(n^{\lfloor \frac{d}{2} \rfloor})$.

Let us bound now the number of k -faces for $k > 0$. The number of k -faces incident to any vertex of $\mathcal{P}$ is $\binom{d}{d - k}$, which is a constant for fixed d .

Hence the upper bound $O(n^{\lfloor \frac{d}{2} \rfloor})$ holds also for the number of faces of any dimension.

The duality introduced in Section 3.2 immediately implies that the same upper bound holds for the lower hull of a set of n points and finally (by

symmetry) to the number of faces of a polytope defined as the convex hull of n points. $\square$

3.4 Convex hull algorithms

We present here two algorithms to build the convex hull of a set of points P . Both algorithms are incremental, meaning that they insert the points of P one by one, while maintaining the convex hull of the currently subset of inserted points. Convex hulls are represented by the Hasse diagram of their boundary complexes (see Section 2.7). Incremental algorithms make also use of the adjacency graph of convex hull facets which, as noticed in Section 2.7), is encoded in the Hasse diagram. Both algorithms also work while maintaining only the adjacency graph of the facets of the convex hull. To make the description of the algorithms simple and to focus on the main ideas, we assume here that the input set of points P is in general position. However this is not a limitation of the presented convex hull algorithms.

3.4.1 An incremental algorithm

Assuming that points in P are added in the order $p_1, p_2, \dots, p_n$, we denote by P_i the subset of the first i -points. Before presenting the whole algorithm and its analysis, we explain how the convex hull is updated when inserting point p_{i+1} .

From $\text{conv}(P_i)$ to $\text{conv}(P_{i+1})$

When point p_{i+1} is considered, the faces of $\text{conv}(P_i)$ may be classified in the following way (see Figure 3.1):

- A facet f of $\text{conv}(P_i)$ is *red* if the hyperplane h_f supporting $\text{conv}(P_i)$ along f separates p_{i+1} from $\text{conv}(P_i)$ meaning that p_i belongs to the open half-space h_f^+ that does not intersect $\text{conv}(P_i)$. Otherwise, as general position is assumed, p_i is included in the half-space h_f^- whose closure contains $\text{conv}(P_i)$, and the facet f is said to be *blue*.




Figure 3.1: Incremental convex hull algorithm. Top: $d = 2$. Bottom: $d = 3$. The old edges are colored in red and blue and purple. The new edges are in black.

- A k -face with $k < d$ is said to be *red* if all its incident facets are red, *blue* if all its incident facets are blue, and *purple* if it is incident to both red and blue facets.

The incremental algorithm and its analysis relies on the three following facts, all related to the transformation of $\text{conv}(P_i)$ into $\text{conv}(P_{i+1})$ when adding point p_{i+1} .

1. The set of red facets induces a connected subgraph of the adjacency graph of $\text{conv}(P_i)$.
2. The set of faces of $\text{conv}(P_{i+1})$ includes the blue and purple faces of $\text{conv}(P_i)$ together with additional *new faces* which are all incident to p_{i+1} . Each new face is formed by the convex hull $\text{conv}(g \cup p_{i+1})$ where g is a purple face of $\text{conv}(P_i)$. See Figure 3.1.




Figure 3.2: When adding point p_{i+1} , the set of purple facets of $\text{conv}(P_i)$ is isomorphic to a $(d - 1)$ -polytope with at most i vertices.

3. The set of purple faces of $\text{conv}(P_i)$ with their incidence relation is isomorphic to the set of new faces of $\text{conv}(P_{i+1})$. Both sets are isomorphic to the set of faces of a $(d - 1)$ -polytope with at most i vertices.

Fact 1 and Fact 2 are illustrated on Figure 3.1. Fact 3 is illustrated on Figure 3.2 showing that the purple faces of $\text{conv}(P_i)$ and the set of faces of $\text{conv}(P_{i+1})$ incident to p_{i+1} are both isomorphic to the set of faces of a $(d - 1)$ -polytope obtained as the intersection of $\text{conv}(P_{i+1})$ with any hyperplane h that separates $\text{conv}(P_i)$ from p_{i+1} .

In summary, updating $\text{conv}(P_i)$ to $\text{conv}(P_{i+1})$ requires to identify the red and purple faces of $\text{conv}(P_i)$. The Hasse diagram of $\text{conv}(P_{i+1})$ is then obtained from the Hasse diagram of $\text{conv}(P_i)$ by removing red faces of $\text{conv}(P_i)$ and creating a new face $\text{conv}(g \cup p_{i+1})$ for each purple face g of $\text{conv}(P_i)$. Incidence relations among new faces are deduced from incidence relations among purple faces in $\text{conv}(P_i)$.

The algorithm

The incremental convex hull algorithm is summarized in the pseudo-code of Algorithm 1.

The incremental algorithm sorts the points in P by lexicographic order of

Algorithm 1 Incremental convex hull algorithm

Input: A set P of n points in $\mathbb{R}^d$
Sort the points of P by lexicographic order of coordinates
Let $p_1, p_2, \dots, p_n$ be the points of P in lexicographic order
Let P_i be the subset of first i points
Build the Hasse diagram of $\text{conv}(P_{d+1})$
for $i = d+1$ **to** $i=n-1$ **do**
 Find a red facet of $\text{conv}(P_i)$ {with respect to p_{i+1} }
 Find all red and purple faces of $\text{conv}(P_i)$
 Update the Hasse of $\text{conv}(P_i)$ into the Hasse diagram of $\text{conv}(P_{i+1})$.
Output: The Hasse diagram of the convex hull of P

their coordinates and insert them in that order. Let $p_1, p_2, \dots, p_n$ be the sorted sequence of points in P and let P_i be the subset of first i points.

The lexicographic ordering ensures that when point p_{i+1} is inserted, at least one of the facet of $\text{conv}(P_i)$ incident to p_i is red. Indeed point p_{i+1} is separated from $\text{conv}(P_i)$ by a hyperplane h through p_i . Let h^+ be the half-space bounded by h that does not contain $\text{conv}(P_i)$. Let, in addition, H denote the set of hyperplanes which are the affine hulls of the facets of $\text{conv}(P_i)$ incident to p_i . Write H^+ for the union of the half-spaces not containing $\text{conv}(P_i)$ that are bounded by the hyperplanes in H . Since h^+ is included in H^+ , p_{i+1} is separated from $\text{conv}(P_i)$ by at least one of the hyperplanes of H . Thus to find a first red facet, the algorithm walks in the adjacency graph of the facets of $\text{conv}(P_i)$, visiting only facets of $\text{conv}(P_i)$ that are incident to p_i , until a red facet is encountered.

Then to find all red and purple faces, the algorithm traverses the adjacency graph visiting all red facets of $\text{conv}(P_i)$ and, from each read facet, it traverses the Hasse diagram to identify all red and purple faces of $\text{conv}(P_i)$. The Hasse diagram is then updated as explained in the previous paragraph.

Complexity Analysis

Theorem 3.10 *The incremental convex hull algorithm computes the convex hull of n points in $\mathbb{R}^d$ in $\Theta\left(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor}\right)$ time.*

Proof First notice that the only numerical predicates involved in Algo-

gorithm 1 are the *orientation predicates* called to decide if a facet of $\text{conv}(P_i)$ is red or blue when inserting point p_{i+1} . The orientation predicate amounts to evaluate the sign of a $(d+1) \times (d+1)$ determinant (see Exercise 3.16).

To find the first red facet, the algorithm visits only facets of $\text{conv}(P_i)$ incident to the last inserted point p_i . Therefore each facet will be visited only once for this purpose during the whole algorithm.

Identifying red and purple faces is also clearly proportionnal to the number of these faces. As a red face disappears from the convex hull, each face is visited only once as a red face, and the cost of visiting a purple face g can be charged on the new face $\text{conv}(g \cup p_{i+1})$.

Updating the Hasse diagram has a complexity proportional to the number of new and removed faces which are respectively the number of purple and red faces of $\text{conv}(P_i)$.

In summary, the cost of step $i+1$ where point p_{i+1} is inserted is proportional to the number of new and removed faces. Since a face is created only once and removed at most once, the total cost of the incremental algorithm, except for the initial sorting of the points, is proportionnal to the total number of created faces.

From Fact 3 above and the upper bound theorem (Theorem 3.9), the number of faces created when inserting p_{i+1} is $O\left(i^{\lfloor \frac{d-1}{2} \rfloor}\right)$ and the total cost of updating the Hasse diagram is:

$$\sum_{i=d+1}^{n-1} O\left(i^{\lfloor \frac{d-1}{2} \rfloor}\right) = O\left(n^{\lfloor \frac{d+1}{2} \rfloor}\right).$$

Taking into account the initial sorting of the points according to the lexicographic order, the complexity of the incremental algorithm is:

$$\begin{aligned} O\left(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor}\right) &= O(n \log n) \text{ if } d = 2 \\ &= O\left(n^{\lfloor \frac{d+1}{2} \rfloor}\right) \text{ if } d > 2 \end{aligned}$$

Furthermore, the complexity bound $O\left(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor}\right)$ is tight for any incremental convex hull algorithm. See Exercise 3.14 for examples of points where the incremental convex hull requires $\Omega(n^{\lfloor \frac{d+1}{2} \rfloor})$ times. $\square$

The upper bound theorem gives a lower bound of $\Omega\left(n^{\lfloor \frac{d}{2} \rfloor}\right)$ for computing

the convex hull of n points in $\mathbb{R}^d$.

Moreover, $\Omega(n \log n)$ is also a lower bound of complexity for computing the convex hull of n points since it is known that sorting n number reduces in linear time to the computation of the convex hull of n points (See Exercise 3.15).

We deduce that computing the convex hull of n points in $\mathbb{R}^d$ has a complexity which is at least $\Omega\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right)$. The incremental algorithm is therefore worst-case optimal in even dimensions. However, it is not optimal in odd dimensions.

3.4.2 A randomized algorithm

The algorithm

The randomized incremental algorithm is quite similar to the above incremental algorithm : points are inserted one by one in the convex hull. At each insertion the set of red and purple faces of the current hull are identified and the convex hull is updated accordingly. The main difference with respect to the incremental algorithm of Section 3.4.1 is that the points are no longer introduced in lexicographic order but in random order. We will show that the expected complexity of the random algorithm matches the lower complexity bound for convex hull computation. Expectation here concerns only the random insertion order and assumes that, for an input set of n points, all the $n!$ possible insertion sequences occur with the same probability.

In the following, points in P are assumed to be indexed according to their insertion order. We denote by $P_i = \{p_1, \dots, p_i\}$ the subset formed by the i first inserted points. As in the deterministic algorithm presented before, the randomized algorithm has to find a first red facet. However, we cannot rely now on the lexicographic order of the input points and restrict our attention to the faces incident to the lastly inserted point. To walk around this issue, we will use an additional data structure called the *conflict graph*.

A facet f of the current convex hull $\text{conv}(P_i)$ is said to be in *conflict* with the not yet inserted point p_j with $j > i$ iff the hyperplane h_f supporting $\text{conv}(P_i)$ along f separates p_j from $\text{conv}(P_i)$. The conflict graph maintained by the algorithm is a bipartite graph including for each not yet inserted point p_j an edge between this point and a facet of the current hull in conflict with

p_j . Note that when a new point p_{i+1} is inserted, the current hull is $\text{conv}(P_i)$ and the facets of $\text{conv}(P_i)$ in conflict with p_{i+1} are precisely the red facets of $\text{conv}(P_i)$. Therefore, when inserting point p_{i+1} , the conflict edge incident to p_{i+1} gives in constant time access to a red facet of $\text{conv}(P_i)$. Then the algorithm finds all the red and purple faces of $\text{conv}(P_i)$ and updates the Hasse diagram of the convex hull exactly as in the deterministic algorithm. The only difference is that now, after each insertion, the algorithm needs, in addition, to update the conflict graph. Before explaining how this is done, we summarize the randomized algorithm as Algorithm 2.

Algorithm 2 Randomized convex hull algorithm

Input: A set P of n points in $\mathbb{R}^d$
Choose randomly a subset P_{d+1} of $d+1$ points in P
Build the Hasse diagram of $\text{conv}(P_{d+1})$
Initialize the conflict graph
for $i = d+1$ **to** $i=n-1$ **do**
 Choose randomly $p_{i+1} \in P \setminus P_i$
 $P_{i+1} = P_i \cup \{p_{i+1}\}$
 Follow the conflict edge of p_{i+1} to find a first red facet
 Find all red and purple faces of $\text{conv}(P_i)$
 Update the Hasse and compute the Hasse diagram of $\text{conv}(P_{i+1})$.
 Update the conflict graph
Output: The Hasse diagram of the convex hull of P

Updating the conflict graph

The algorithm needs to restore conflict edges between the facets of $\text{conv}(P_{i+1})$ and the points of $P \setminus P_{i+1}$. Nothing needs to be done for the points of $P \setminus P_{i+1}$ that were previously in conflict with a blue facet of $\text{conv}(P_i)$ since such a facet is still a facet of $\text{conv}(P_{i+1})$. Let p_j be a point of $P \setminus P_{i+1}$ that was previously in conflict with a red facet f_j of $\text{conv}(P_i)$. We need to find a facet of $\text{conv}(P_{i+1})$ in conflict with p_j . Let R denote the set of red facets of $\text{conv}(P_i)$ and let R_j denote the set of facets of $\text{conv}(P_i)$ in conflict with p_j . As noted previously, the set R induces a connected subgraph of the adjacency graph of $\text{conv}(P_i)$. The same is true for R_j and for the subset $F_j = R \cap R_j$ of red facets in conflict with p_j . Note that the boundary of F_j includes red ridges which are on the boundary of R_j but not on the boundary of R and purple ridges which are on the boundary of R . See Figure 3.3 for an illustration in

$\mathbb{R}^2$.

To find a facet of $\text{conv}(P_{i+1})$ in conflict with p_j , the algorithm starts at f_j , the facet that was in conflict with p_j before the insertion of p_{i+1} , and walk on the adjacency graph of $\text{conv}(P_i)$, visiting the facets of F_j . At each purple ridge h encountered on the boundary of F_j , the algorithm checks if one of the two following cases occurs:

Case 1 : the blue facet g of $\text{conv}(P_i)$ incident on h is in conflict with p_{i+1} ,

Case 2 : the new facet $f' = \text{conv}(h, p_{i+1})$ of $\text{conv}(P_{i+1})$, is in conflict with p_{i+1} .

In both cases, a facet of $\text{conv}(P_{i+1})$ in conflict with p_{i+1} has been found and the walk stops. If the walk traverses all facets of F_j without encountering one of the above two cases, p_j lies inside $\text{conv}(P_{i+1})$ and can be discarded. See Figure 3.3.

Analysis of the randomized algorithm

We insist that the randomized algorithm presented here always computes the actual convex hull of the set of input points. The random choices performed by the algorithm concern only the order in which the points are inserted and the performances of the algorithm. The analysis below will bound the expected complexity of the algorithm assuming that all insertion sequences occur with the same probability.

Theorem 3.11 *The randomized incremental algorithm computes the convex hull of n points in $\mathbb{R}^d$ in expected time $O\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

Before we give the proof of the theorem, we recall an important result of Clarkson and Shor [59], known as the random sampling theorem.

The random sampling theorem. We call *configuration* a subset of d independent points in $\mathbb{R}^d$. A configuration σ is said to be defined on a set of points P if the points in σ belong to P . Let us choose as the origin a point o in $\text{conv}(P)$. A configuration is said to be in conflict with a point p if the hyperplane which is the affine hull of σ separates o from p . A configuration




Figure 3.3: Randomized construction of the convex hull: updating the conflicting graph. The up left picture show the set of red facets when inserting p_{i+1} , the three other pictures show, in three different cases, the set of facets of $\text{conv}(P_i)$ in conflict with p_{i+1} and p_j . In upper right picture, p_j is internal to $\text{conv}(P_{i+1})$, in lower left picture, case 1 and case 2 both occur, in lower right picture only case 2 occurs.

is said to have j conflicts on P if it is in conflict with j points of P . We denote by $C(P)$ the set of configurations defined on P and by $C_j(P)$ (resp. $C_{\leq k}(P)$), the set of configurations defined on P and with j conflicts (resp. at most k conflicts) on P .

In the following, we consider random samples S of P and we denote by $\Gamma_0(s, P)$ the expected number of configurations defined on S and without conflict on S , where S is a random sample of size s of P .

Theorem 3.12 (The random sampling theorem) *Let P be a set of n points. The number $|C_{\leq k}(P)|$ of configurations defined on P and with at most k -conflicts on P is bounded as follows :*

$$|C_{\leq k}(P)| \leq 4(d+1)^d k^d \Gamma_0\left(\left\lfloor \frac{n}{k} \right\rfloor, P\right). \quad (3.4)$$

In our context, a configuration that is defined on P and without conflict on P , is just the vertex set of a facet of $\text{conv}(P)$ so that there is a bijection between $C_0(P)$ and the facets of $\text{conv}(P)$. We then deduce from the upper bound theorem (Theorem 3.9) that the number of configurations defined and without conflict on a sample of size s of P is at most $O\left(s^{\lfloor \frac{d}{2} \rfloor}\right)$. Therefore:

$$\Gamma_0(s, P) = O\left(s^{\lfloor \frac{d}{2} \rfloor}\right). \quad (3.5)$$

Plugging Equation 3.5 into Equation 3.4, we get that for any set P of n points the number of configurations defined and with at most k conflicts over P is at most :

$$|C_{\leq k}(P)| = O\left(k^{\lceil \frac{d}{2} \rceil} n^{\lfloor \frac{d}{2} \rfloor}\right). \quad (3.6)$$

Proof of Theorem 3.11. We can now analyze the cost of updating the Hasse diagram of the convex hull and the cost of maintaining the conflict graph.

Cost of updating the Hasse diagram. As in the case of the incremental algorithm, the cost of updating the Hasse diagram is proportional to the total number of convex hull facets that are created by the algorithm.

Let us bound the expected number $n(i+1)$ of facets that are created at Step $i+1$ when point p_{i+1} is inserted. Since the algorithm inserts the points of

P in random order, P_{i+1} is a random subset of P of size $i+1$. Notice that a facet created at step $i+1$ corresponds to a configuration of $C_0(P_{i+1})$. Given P_{i+1} , a configuration σ of $C_0(P_{i+1})$ corresponds to a facet created at step $i+1$ iff one of the points in this configuration is the point P_{i+1} inserted at step i , which happens with probability $\frac{d}{i+1}$. Thus, we have

$$\begin{aligned} n(i+1) &= \sum_{\sigma \in C(P)} \text{proba}(\sigma \in C_0(P_{i+1})) \times \frac{d}{i+1} \\ &= \frac{d}{i+1} \Gamma_0(i+1, P) \\ &= O\left((i+1)^{\lfloor \frac{d}{2} \rfloor - 1}\right). \end{aligned}$$

By summing over all steps i and using linearity of expectation, we bound the expected total number of facets created by the algorithm and therefore the expected cost of updating the Hasse diagram by :

$$\sum_i^n n(i) = O(n^{\lfloor \frac{d}{2} \rfloor}). \quad (3.7)$$

Cost of updating the conflict graph. We now bound the cost of updating the conflict graph. As explained above, when inserting p_{i+1} at step $i+1$, the algorithm has to restore a conflict edge for each point p_j with $j > i+1$. This is done by traversing the incidence graph of $\text{conv}(P_i)$, visiting the facets in conflict with both p_{i+1} and p_j . The cost of this procedure is proportional to the number of visited facets which we analyse now.

For any $p_j \in P \setminus P_{i+1}$, the subset $S = P_{i+1} \cup \{p_j\}$ is a random sample of P of size $i+2$. The facets visited to restore the conflict for p_j at step $i+1$ correspond to configurations in $C_2(S)$. Assume that a subset S of P of size $i+2$ is given and that $S = P_{i+1} \cup \{p_j\}$. Then, any configuration σ in $C_2(S)$ is a facet of $\text{conv}(P_i)$ in conflict with p_{i+1} and p_j iff p_{i+1} and p_j are the two elements of S in conflict with σ , which happens with probability $\frac{2}{(i+1)(i+2)}$. Given S , the expected number $n(i+1, j, S)$ of facets visited to restore the conflict for p_j at step $i+1$ is

$$\begin{aligned} n(i+1, j, S) &= \sum_{\sigma \in C(P)} \text{proba}(\sigma \in C_2(S)) \times \frac{2}{(i+1)(i+2)} \\ &\leq \frac{2}{(i+1)(i+2)} |C_{\leq 2}(S)|. \end{aligned}$$

Then, using Equation 3.6,

$$n(i+1, j, S) \leq O\left((i+1)^{\lfloor \frac{d}{2} \rfloor - 2}\right).$$

Since this is true for any subset S of P of size $i+2$, we get that the expected number $n(i+1, j)$ of facets visited to restore a conflict edge for p_j at step $i+1$ is also $O\left((i+1)^{\lfloor \frac{d}{2} \rfloor - 2}\right)$. Then, the expected total cost for updating the conflict graph is

$$\sum_{i=1}^n \sum_{j=i+1}^n n(i, j) = \sum_{i=1}^n (n-i) O\left(i^{\lfloor \frac{d}{2} \rfloor - 2}\right) = O\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right). \quad (3.8)$$

At each insertion, finding a first red facet takes constant time using the conflict graph. Finding all the red and purple faces then takes a time proportional to the number of red and purple faces and the total cost of this operation is thus also given by Equation 3.7. In summary, the expected combinatorial complexity of the randomized incremental construction of the convex hull is given by Equations 3.7 and 3.8, which achieves the proof of Theorem 3.11. $\square$

The randomized version of the incremental construction of a convex hull has therefore an expected complexity which is better than the complexity of the deterministic incremental construction. Since this expected complexity matches the complexity of the convex hull, the randomized incremental construction of a convex hull is optimal.

3.5 Exercises

Exercise 3.1 Show that a convex polytope $\mathcal{P}$ is the convex hull of its vertices.

Hint: one of the inclusion is trivial. To prove the other one, consider the minimal subset $P' \subseteq P$ such that $\mathcal{P} = \text{conv}(P) = \text{conv}(P')$ and prove that each point in P' is a vertex of $\mathcal{P}$.

Exercise 3.2 Show that the intersection of any finite set of faces of a polytope is also a face of the polytope.

Exercise 3.3 Show that any face of a convex polytope $\mathcal{P}$ is the intersection of facets of $\mathcal{P}$.

Show that a $(d-2)$ -face of a polytope $\mathcal{P}$ is the intersection of two facets of $\mathcal{P}$.

Exercise 3.4 Let $\mathcal{P}$ be a convex polytope and let H be the set of hyperplanes that support $\mathcal{P}$ along its facets. To each hyperplane $h \in H$, we associate the half-space h^+ bounded by h that contains $\mathcal{P}$. Show that $\mathcal{P} = \bigcap_{h \in H} h^+$.

Exercise 3.5 (Polytopes and intersection of half-spaces) Show that, if it is bounded, the intersection of a finite set of half-spaces is a polytope.

Exercise 3.6 (Faces of a convex polyhedron) Let H be a set of n hyperplanes $h_1, \dots, h_n$ and $\mathcal{P}$ be the polyhedron defined as the intersection of the n half-spaces $h_1^+, \dots, h_n^+$ where h_i^+ is the half-space bounded by h_i that contains a given point o . Let I be any subset of the indices $1, \dots, n$ and $F_I = \bigcap_{i \in I} h_i$. Show that, if it is non empty, the intersection $\mathcal{P} \cap F_I$ is a face of $\mathcal{P}$ and that all the faces of $\mathcal{P}$ can be obtained this way, i.e. as the intersection with $\mathcal{P}$ of the hyperplanes of a subset of H . Show, in addition, that, if H is in general position, $\mathcal{P} \cap F_I$ is a face of dimension $d-k$ if $|I| = k$.

Exercise 3.7 (Faces of a simple convex polyhedron) Prove that if $\mathcal{P}$ is a simple convex polyhedron and p is a vertex of $\mathcal{P}$, the affine hull of any subset of $k < d$ edges incident to p contains a face of $\mathcal{P}$ of dimension $d-k$. Hint: use Exercise 3.2 and duality.

Exercise 3.8 (General position and duality) Show that n hyperplanes are in general position iff their dual points are in general position.

Exercise 3.9 Show that the boundary complex of a polytope is a combinatorial manifold.

Hint: Let p be a vertex of the d -polytope $\mathcal{P}$. We show that the link of p in the boundary complex of $\mathcal{P}$ is a $(d-2)$ -topological sphere. Indeed, any hyperplane h that separates p from the other vertices of $\mathcal{P}$ (see Figure 3.2) intersects $\mathcal{P}$ along a polytope of dimension $(d-1)$ whose boundary complex is isomorphic to the link of p in the boundary complex of $\mathcal{P}$.

Exercise 3.10 Prove Lemma 3.8 in the general case.

Hint : Proves that if k' points in $\mathbb{R}^d$ span an affine hull of dimension $k < k'$, their dual hyperplanes intersect along a $(k - d)$ -dimensional affine space.

Exercise 3.11 Let $\mathcal{P}$ be a convex polyhedron. Prove that perturbing the hyperplanes bounding the half-spaces that define $\mathcal{P}$ can only increase the number of faces of $\mathcal{P}$.

Exercise 3.12 (Cyclic polytopes) A cyclic polytope is a polytope in $\mathbb{R}^d$ that is the convex hull of points lying on the moment curve $\mathcal{M}_d$ defined by the parametric representation

$$\mathcal{M}_d = \{x(t, t^2, \dots, t^d), t \in \mathbb{R}\}$$

Show that a cyclic polytope of $\mathbb{R}^d$ with n vertices has $\binom{n}{k} (k - 1)$ -faces for any k such that $0 \leq k \leq d/2$.

Exercise 3.13 (Polarity) Polarity has been introduced with the paraboloid $\mathcal{Q}$. A sphere could have been used instead as shown in this exercise. Let o be a point of $\mathbb{R}^d$. Polarity associates to any point p of $\mathbb{R}^d$ distinct of o the hyperplane $p^* = \{x : (x - o) \cdot (p - o) = 1\}$. Show that polarity is a bijection between points of $\mathbb{R}^d$ distinct de o and hyperplanes not passing through o . Let P be a set of point of $\mathbb{R}^d$ whose convex hull includes o . Show that polarity induces a duality between the boundary complex of the polytope $(P) = \text{conv}(P)$ and the boundary complex of the polytope $(P^{*+}) = \bigcap_{p \in P} p^{*+}$ where p^{*+} is the half-space bounded by p^* not containing o .

Exercise 3.14 (Lower bound for incremental convex hull computation)

Show that any incremental algorithm that constructs the convex hull of n points of $\mathbb{R}^d$ takes $\Omega(n^{\lfloor \frac{d+1}{2} \rfloor})$ time in the worst-case.

Hint : take the points on the moment curve $\mathcal{M}_d$ (see Exercise 3.12) and insert them by increasing values of their first coordinate.

Exercise 3.15 (Lower bound for convex hull computation) Show that sorting n number reduces in linear time to the computation of the convex hull of n points in $\mathbb{R}^2$.

Hint: Take n real numbers $x_1, \dots, x_n$. We associate to each x_i the point $p_i = (x_i, x_i^2)$ on the unit parabola. If we know the convex hull of the p_i , we can deduce in linear time the list of the x_i sorted by increasing values.

Exercise 3.16 (Orientation predicate) Show that the numerical operation required to decide if a facet $f = \{p_{i_1}, \dots, p_{i_d}\}$ of $\text{conv}(P_i)$ is red or blue when inserting point p_{i+1} amounts to evaluate the sign of a $(d+1) \times (d+1)$ -determinant

$$\begin{vmatrix} 1 & \cdots & 1 & 1 \\ p_{i_1} & \cdots & p_{i_d} & p_{i+1} \end{vmatrix}.$$

Exercise 3.17 (Vertical projection of a lower hull) Let P be a set of point in general position in $\mathbb{R}^d$. We call vertical projection the projection Π onto the hyperplane of equation $x(d) = 0$. Show that the restriction of the vertical projection to the lower hull of P is 1 to 1. Show that the vertical projection naturally embeds the lower hull of P as a triangulation of $\Pi(P')$, where P' is the vertex set of the lower hull of P .

Exercise 3.18 (Minkowski formula) Show that if $\mathcal{P}$ is a convex polytope and f_j , $j \in J$, are its facets, we have $\sum_{j \in J} \text{vol}(f_j) n_j = 0$, where n_j is the unit normal vector to f_j oriented towards the outside of $\mathcal{P}$.

Hint : let $x \in \mathcal{P}$. Compute the volume of $\mathcal{P}$ by summing the volumes of the pyramids (x, f_j) . The results follows by observing that the volume of $\mathcal{P}$ does not depend on x .

3.6 Bibliographical notes

A modern introduction to the theory of polytopes can be found in Ziegler's book [137]. The original proof of the upper bound theorem has been established by McMullen in 1970. The simple asymptotic version given in Theorem 3.9 is due to Seidel [124].

The incremental convex hull algorithm is due to Seidel [125]. A complete description of this algorithm handling degenerated cases is given in Edelsbrunner's book [73]. The random sampling theorem is due to Clarkson and Shor [59]. The same paper proposes the first randomized algorithm to build convex hulls. This algorithm solves in fact the dual problem of computing

the intersection of half-spaces. Chazelle [52] has proposed a *deterministic* algorithm to compute the convex hull of a finite point set that is worst case optimal in any dimension. Obtained through derandomisation of the randomized algorithm this algorithm is however mostly of theoretical interest and no implementation is known.

The theory of randomized algorithms is well-developed and finds applications in many areas of computer science. See the book by Motwani and Raghavan for a broad perspective [110], and the books of Boissonnat and Yvinec [27] and of Mulmuley [111] for a geometric viewpoint.

Chapter 4

Delaunay complexes

Delaunay complexes are fundamental data structures that have been extensively studied in computational geometry and used in many application areas.

The Delaunay complex of a finite set of points $P \in \mathbb{R}^d$ is defined as the nerve of the Voronoi diagram of P which we define first. We prove Delaunay's theorem that states that, when the points of P are in general position, the Delaunay complex of P has a natural embedding in $\mathbb{R}^d$ called the Delaunay triangulation of P .

The proof relies on the so-called lifting map that associates to each point of $\mathbb{R}^d$ a point in $\mathbb{R}^{d+1}$. We show that the Delaunay complex of P is the projection onto $\mathbb{R}^d$ of the lower hull of the lifted points. Using then the results of Chapter 3, we will bound the combinatorial complexity of Delaunay complexes and Voronoi diagrams, and obtain optimal algorithms for their construction.

Voronoi diagrams and Delaunay complexes are a special case of more general structures called *weighted* Voronoi diagrams and Delaunay complexes. The class of weighted Voronoi diagrams include the class of Euclidean Voronoi diagrams and, as we will see, most of the properties of Voronoi diagrams still hold for weighted diagrams. Weighted Voronoi diagrams can be found under various disguises in various applications and the dual weighted Delaunay complexes will play an important role in Chapters 5 and 8.

4.1 Lower envelopes and minimization diagrams

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be a set of d -variate continuous functions defined over $\mathbb{R}^d$. The *lower envelope* of $\mathcal{F}$ is defined as

$$\mathcal{F}^- = \min_{1 \leq i \leq n} f_i.$$

From $\mathcal{F}$ and $\mathcal{F}^-$, we define a partition of $\mathbb{R}^d$ called the *minimization diagram* of $\mathcal{F}$. For a point $x \in \mathbb{R}^d$, we define the *index set* $I(x)$ of x as the set of all indices i such that $\mathcal{F}^-(x) = f_i(x)$. An equivalence relation noted $\equiv$ can then be defined on the points of $\mathbb{R}^d$: two points are equivalent if they have the same index set:

$$x \equiv y \iff I(x) = I(y).$$

The equivalence classes are relatively open sets that cover $\mathbb{R}^d$. Their closures are called the *faces* of the minimization diagram of $\mathcal{F}$ and the collection of




Figure 4.1: The lower envelope of a set of univariate functions. The minimization diagram is drawn on the horizontal line with the corresponding indices. The face of index $\{1\}$ consists of two components.

all those faces constitutes the minimization diagram of $\mathcal{F}$ (see Figure 4.1). The index set of a face is defined as the largest subset of indices common to all the points of the face. Conversely, the face of index set I is the set of all points x such that $I \subset I(x)$.

Upper envelopes and *maximization diagrams* can be defined analogously. Upper envelopes have been defined in a geometric way in Section 3.2.2 for sets of non vertical hyperplanes. The two notions are in fact closely related. Let $H = \{h_1, \dots, h_n\}$ be a set of non vertical hyperplanes in $\mathbb{R}^{d+1}$. Each hyperplane h_i in H has an equation which in normal form reads $x(d+1) = f_i(x)$ where $x(d+1)$ is the last coordinate of a point in $\mathbb{R}^{d+1}$ and $f_i(x)$ is an affine function of the d first coordinates. Therefore h_i can be regarded as the graph of the affine function $f_i(x)$. Let $\mathcal{F}$ be the set of affine functions $\{f_1, \dots, f_n\}$. The upper envelope defined in Section 3.2.2 for the set H is the graph of the upper envelope of the set $\mathcal{F}$, $\mathcal{F}^+ = \max_{1 \leq i \leq n} f_i$.

4.2 Voronoi diagrams

Let $P = \{p_1, \dots, p_n\}$ be a set of points of $\mathbb{R}^d$. To each p_i , we associate its *Voronoi cell* $V(p_i, P)$, or simply $V(p_i)$ when there is no ambiguity on the set P :

$$V(p_i, P) = \{x \in \mathbb{R}^d : \|x - p_i\| \leq \|x - p_j\|, \forall p_j \in P\}.$$




Figure 4.2: The Voronoi diagram of a set of 9 points.


The cell $V(p_i)$ is the intersection of the $n - 1$ half-spaces bounded by the bisecting hyperplanes of p_i and each of the other points of P . $V(p_i)$ is therefore a convex polyhedron, possibly unbounded, that contains p_i . The Voronoi cells of P have disjoint interiors and, since any point of $\mathbb{R}^d$ belongs to at least one Voronoi cell, they cover the entire space $\mathbb{R}^d$. The collection of the Voronoi cells and their faces constitute a *cell complex* called the Voronoi diagram of P and denoted by $\text{Vor}(P)$.

The Voronoi diagram of P is the minimization diagram of the set of distance functions $\{\delta_i, \dots, \delta_n\}$, where

$$\delta_i(x) = \|x - p_i\|.$$

Since minimizing $\|x - p_i\|$ over i is the same as minimizing $(x - p_i)^2$, the Voronoi diagram of P can alternatively be defined as the minimization diagram of the smooth functions $(x - p_i)^2$. The graphs of those functions are translated copies of the vertical paraboloid of revolution of $\mathbb{R}^{d+1}$ of equation $x(d+1) = x^2$.

Observing further that, for any x , $\arg \min_i (x - p_i)^2 = \arg \max_i (2p_i \cdot x - p_i^2)$, we obtain that the Voronoi diagram of P is the maximization diagram of a

Figure 4.3: $\mathcal{V}(P)$, with one of its faces projected onto $\mathbb{R}^d$.

set of affine functions, namely the functions

$$f_i(x) = 2p_i \cdot x - p_i^2.$$

We can rewrite this construction in more geometric terms and establish a link with convex polyhedra. For $i = 1, \dots, n$, let h_i be the hyperplane of $\mathbb{R}^{d+1}$ that is the graph of the affine function $f_i(x)$:

$$h_i = \{(x, x(d+1)) \in \mathbb{R}^d \times \mathbb{R}, x(d+1) = 2p_i \cdot x - p_i^2\}.$$

Let $H = \{h_1, \dots, h_n\}$ and denote by $\mathcal{V}(P)$ the upper envelope of hyperplanes in H . By definition (see Section 3.2.2), the upper envelope $\mathcal{V}(P)$ is the boundary complex of the convex polyhedra $h_1^+ \cap h_2^+ \dots \cap h_n^+$ which is the intersection of the upper half-spaces bounded by the hyperplanes in H . Since the Voronoi diagram $\text{Vor}(P)$ is the maximization diagram of the set of functions $\{f_1, \dots, f_n\}$, it is the projection onto the space $\mathbb{R}^d = \{x, x(d+1) = 0\}$ of the upper envelope $\mathcal{V}(P)$, see Figure 4.3.

We deduce that the combinatorial complexity of the Voronoi diagram of n points of $\mathbb{R}^d$ is at most the combinatorial complexity of a polyhedron defined as the intersection of n half-spaces of $\mathbb{R}^{d+1}$, which is $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ as shown in Section 3.3. This bound is tight (Exercise 4.3). Moreover, we can construct $\text{Vor}(P)$ by constructing $\mathcal{V}(P) \subset \mathbb{R}^{d+1}$ and then projecting its faces onto $\mathbb{R}^d$. Thanks to Theorem 3.11, we conclude that $\text{Vor}(P)$ can be constructed in time $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.




Figure 4.4: The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines).

4.3 Delaunay complexes

Let P be a finite set of points in $\mathbb{R}^d$ and write as before $\text{Vor}(P)$ for its Voronoi diagram. The collection of all Voronoi cells of $\text{Vor}(P)$ form a cover of $\mathbb{R}^d$ by closed convex sets, and the nerve of this cover is an abstract simplicial complex called the *Delaunay complex* of P . Specifically, let f be a face of dimension k of $\text{Vor}(P)$ and let $\text{int} f$ be the set of points of f that do not belong to any proper subspace of f . All the points of $\text{int} f$ have the same subset σ of closest points in P and f is the intersection of the Voronoi cells of the points in σ . Accordingly, σ is a simplex in the Delaunay complex. See Figure 4.4. We denote by $\text{Del}(P)$ the Delaunay complex of the set P .

This definition can be rephrased in terms of empty balls. A ball $B \in \mathbb{R}^d$ is said to be *empty* of points of P if the interior of B includes no points of P . We say that a d -ball circumscribes a finite subset of points if the sphere bounding B passes through all the points of the subset. The following lemma is just another view of the definition of the Delaunay complex.

Lemma 4.1 (The empty ball property) *Any subset $\sigma \subset P$ is a simplex of the Delaunay complex of P iff it has a circumscribing (open) ball empty*

of points of P . Such a ball is called a Delaunay ball.

The Delaunay complex cannot always be embedded in $\mathbb{R}^d$. Consider, for example, the case of a set P consisting of $m > d + 1$ points lying on a same hypersphere. The center of the hypersphere belongs to the Voronoi cells of all the m points, which implies that the Delaunay complex contains the $(m - 1)$ -simplex whose vertex set is P . This simplex cannot be embedded in $\mathbb{R}^d$ since its dimension $m - 1$ is greater than d .

However, as shown below, the Delaunay complex can be embedded in $\mathbb{R}^d$ when the points of P are in *general position wrt spheres*.

Definition 4.2 (General position wrt spheres) *We say that a finite set of points P is in general position wrt spheres if no subset of $d + 2$ points of P lie on a same hypersphere.*

We can now state Delaunay's fundamental result.

Theorem 4.3 (Delaunay triangulation) *If a finite set of points $P \in \mathbb{R}^d$ is in general position wrt spheres, then the Delaunay complex $\text{Del}(P)$ has a natural embedding¹ in $\mathbb{R}^d$. This embedding is a triangulation of P called the Delaunay triangulation of P .*

Proof We identify $\mathbb{R}^d$ with the hyperplane of $\mathbb{R}^{d+1}$ of equation $x(d+1) = 0$ and introduce a lifting map ϕ that maps points and balls of $\mathbb{R}^d$ to points of $\mathbb{R}^{d+1}$. We then consider the lower hull of the set of lifted points $\phi(P)$ and show that the projection onto $\mathbb{R}^d$ embeds this lower hull as a geometric simplicial complex of $\mathbb{R}^d$. We will see that this complex is identical to the natural embedding of the Delaunay complex $\text{Del}(P)$. Further, we will prove that it is a triangulation of P .

Let $b = b(c, r)$ be a d -ball of $\mathbb{R}^d$ with center c and radius r . The lifting map ϕ associates to b the point $\phi(b) = (c, c^2 - r^2)$ of $\mathbb{R}^{d+1}$. A point x of $\mathbb{R}^d$ can be considered as a ball with null radius and is mapped to the point $\phi(x) = (x, x^2)$ of $\mathbb{R}^{d+1}$. Observe that the lift of a point lies on the paraboloid of revolution $\mathcal{Q} = \{x \in \mathbb{R}^{d+1} : x(d+1) = x^2\}$. We then use in $\mathbb{R}^{d+1}$ the

¹Defined in Section 2.2.




Figure 4.5: The empty ball property.

point hyperplane duality defined in Section 3.2.1 and denote by $\phi(b)^*$ the hyperplane of $\mathbb{R}^{d+1}$ that is the dual of $\phi(b)$:

$$\phi(b)^* = \{x \in \mathbb{R}^{d+1} : x(d+1) - 2cx + c^2 - r^2 = 0\}.$$

The sphere ∂b bounding $b(c, r)$ is the projection onto $\mathbb{R}^d$ of the intersection of $\mathcal{Q}$ with the hyperplane $\phi(b)^*$ (see Figure 4.6). Furthermore, the point x belongs to the (open) ball b iff $\phi(x)$ belongs to the halfspace $\phi(b)^{*-}$ below $\phi(b)^*$. Therefore, we have:

$$x \in \partial b \iff \phi(x) \in \phi(b)^* \quad (4.1)$$

$$x \in b \iff \phi(x) \in \phi(b)^{*-}. \quad (4.2)$$

Consider now a set P of n points in general position wrt spheres and let σ be a subset of $d+1$ points of P . We write b for the ball that circumscribes σ . According to Lemma 4.1, σ is a Delaunay simplex iff the ball b is empty of points of P . Then, from Equations 4.1 and 4.2, the hyperplane $\phi(b)^*$ contains $\phi(\sigma)$ and the halfspace $\phi(b)^{*-}$ does not contain any point of $\phi(P)$. We conclude that σ is a Delaunay simplex iff the convex hull of $\phi(\sigma)$ is a facet of the lower hull $\mathcal{D}(P)$ of the lifted points $\phi(P)$.

Since P is in general position wrt spheres, $\phi(P)$ is in general position wrt hyperplanes (i.e., in the usual sense). It follows that $\mathcal{D}(P)$ is a simplicial complex embedded in $\mathbb{R}^{d+1}$. Consider now the projection onto $\mathbb{R}^d$, called the vertical projection. The restriction of the vertical projection to $\mathcal{D}(P)$ is 1-1 (see Exercise 3.17), and therefore the vertical projection of $\mathcal{D}(P)$ is




Figure 4.6: The polar hyperplane of a ball.

a simplicial complex embedded in $\mathbb{R}^d$. From the previous paragraph, the facets of $\mathcal{D}(P)$ are in bijective correspondence with the Delaunay simplices of $\text{Del}(P)$. Furthermore, for each simplex σ of $\text{Del}(P)$, the corresponding facet of $\mathcal{D}(P)$ is the convex hull of the lift $\phi(\sigma)$ of σ so that its vertical projection is the convex hull of σ . It follows that the projection of the lower hull $\mathcal{D}(P)$ is the natural embedding of the Delaunay complex $\text{Del}(P)$.

Let us further show that the natural embedding of $\text{Del}(P)$ is a triangulation of P . By definition, the set of vertices of $\text{Del}(P)$ is P . Furthermore, since $\mathcal{D}(P)$ is the lower hull of $\phi(P)$, its vertical projection coincides with the vertical projection of the convex hull of $\phi(P)$ which is just the convex hull of P . It follows that the natural embedding of $\text{Del}(P)$ is a triangulation of P . This concludes the proof of Delaunay triangulation theorem. $\square$

Since the Delaunay complex $\text{Del}(P)$ is defined as the nerve of the cell complex $\text{Vor}(P)$, there is a dual correspondence between the faces of these two complexes, i.e. a bijective correspondence between their faces that preserves incidences and reverses inclusions. The duality between lower hulls and upper envelopes introduced in Section 3.2 yields another dual correspondence between $\text{Vor}(P)$ and $\text{Del}(P)$. Indeed, notice that the set of hyperplanes $\{h_1, h_2, \dots, h_n\}$ defining the upper envelope $\mathcal{V}(P)$ are the duals $\{\phi(p_1)^*, \phi(p_2)^*, \dots, \phi(p_n)^*\}$ of the lifted points $\phi(P)$. Therefore, by the results of Section 3.2, the upper envelope $\mathcal{V}(P)$ and the lower hull $\mathcal{D}(P)$ are

dual complexes. Since the vertical projection induces a bijection between the cellular complexes $\mathcal{V}(P)$ and $\text{Vor}(P)$, and since the lifting map induces another bijection between the complexes $\mathcal{D}(P)$ and $\text{Del}(P)$, we get a dual correspondance between $\text{Vor}(P)$ and $\text{Del}(P)$ through the duality between $\mathcal{V}(P)$ and $\mathcal{D}(P)$. In fact the proof of the Delaunay triangulation theorem shows that the two correspondences between $\text{Vor}(P)$ and $\text{Del}(P)$ coincide, so that the following diagram commutes:

$$\begin{array}{ccc}
 \mathcal{V}(P) & \xrightarrow{\text{duality}} & \mathcal{D}(P) \\
 \uparrow & & \downarrow \\
 \text{Voronoi Diagram } \text{Vor}(P) & \xrightarrow{\text{nerve}} & \text{Delaunay Complex } \text{Del}(P)
 \end{array}$$

It follows that the combinatorial complexity of the Delaunay complex of n points is the same as the combinatorial complexity of its dual Voronoi diagram. Moreover, the Delaunay complex of n points of $\mathbb{R}^d$ can be deduced from the dual Voronoi diagram and vice versa in time proportional to their size. We also deduce from what precedes that computing the Delaunay complex of n points of $\mathbb{R}^d$ reduces to constructing the convex hull of n points of $\mathbb{R}^{d+1}$. The following theorem is then a direct consequence of Theorems 3.9 and 3.11.

Theorem 4.4 *The combinatorial complexity of the Voronoi diagram of a set of n points of $\mathbb{R}^d$ and of its Delaunay complex is $\Theta\left(n^{\lceil \frac{d}{2} \rceil}\right)$. Both structures can be computed in optimal time $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.*

The bounds in this theorem are tight in the worst case. In particular, the Voronoi diagram of n points of $\mathbb{R}^3$ may be quadratic (see Exercise 4.3). However, under some assumptions on the point distribution, better bounds can be obtained (see Section 5.1).

If the points of P are not in general position wrt spheres, the Delaunay complex may be of dimension greater than the dimension d of the embedding space and therefore cannot be embedded in $\mathbb{R}^d$. Accordingly the hyperplane set $\phi(P)^*$ and the point set $\phi(P)$ are not in general position. Hence $\mathcal{V}(P)$ is not simple and $\mathcal{D}(P)$ is not simplicial. The diagram above won't commute.

Nevertheless, the projection of $\mathcal{D}(P)$ will still be a cell complex. Let f be a j -face of $\mathcal{D}(P)$ that is not a simplex. Its projection is no longer a simplex

but is still a convex j -polytope we denote by f' . Moreover, f' is still the convex hull of the vertices of a k -simplex σ of the nerve of $\text{Vor}(P)$ with $k > j$. We say that f' is the *shadow* of σ . An embedded triangulation can then be obtained by triangulating the shadows that are not simplices (see Exercise 4.6). Any such triangulation is called a *Delaunay triangulation*. Since there are several ways of triangulating the faces of a polytope, P admits several Delaunay triangulations.

4.4 Weighted Delaunay complexes

4.4.1 Weighted points and weighted distance

A weighted point $\hat{p} = (p, w)$ is an element of $\mathbb{R}^d \times \mathbb{R}$ where the point $p \in \mathbb{R}^d$ is called the *center* of $\hat{p}$ and $w \in \mathbb{R}$ is called its *weight*. A point is identified with a weighted point of weight 0. When w is non negative, $\hat{p}$ can be considered as the ball $b(p, \sqrt{w})$ centered at p and of squared radius w . Since considering negative weights will cause no problem in our developments, we favour the name of weighted point.

We define the weighted distance between two weighted points $\hat{p}_1 = (p_1, w_1)$ and $\hat{p}_2 = (p_2, w_2)$ as

$$D(\hat{p}_1, \hat{p}_2) = (p_1 - p_2)^2 - w_1 - w_2.$$

Note that the weighted distance is not a distance in the usual sense since it is not positive and does not satisfy the triangular inequality. Two weighted points are said to be *orthogonal* if their distance is zero. The term orthogonal comes from the observation that, when the weights are non negative, two weighted points are orthogonal iff the spheres bounding the associated balls are orthogonal spheres (see Figure 4.7). The weighted distance $D(x, \hat{p})$ from an unweighted point x to a weighted point $\hat{p} = (p, w)$ whose weight w is non negative is also called the *power* of point x with respect to the sphere bounding the ball $b(p, \sqrt{w})$.

4.4.2 Weighted Voronoi diagrams

Let $\hat{P} = \{\hat{p}_1, \dots, \hat{p}_n\}$ be a set of weighted points of $\mathbb{R}^d \times \mathbb{R}$. To each $\hat{p}_i = (p_i, w_i)$, we associate the cell $V(\hat{p}_i) \subset \mathbb{R}^d$ consisting of the points x of $\mathbb{R}^d$




Figure 4.7: The weighted distance: two weighted points with respectively a negative (left), null (middle) and positive (right) weighted distance.

whose weighted distance to $\hat{p}_i$ is not larger than its weighted distance to any other weighted points of $\hat{P}$:

$$V(\hat{p}_i) = \{x \in \mathbb{R}^d : D(x, \hat{p}_i) \leq D(x, \hat{p}_j), \forall \hat{p}_j \in \hat{P}\}.$$

The set of points of $\mathbb{R}^d$ that are at equal distance from two weighted points $\hat{p}_i$ and $\hat{p}_j$ is a hyperplane called the bisecting hyperplane of $\hat{p}_i$ and $\hat{p}_j$. This hyperplane is orthogonal to the line joining the centers of $\hat{p}_i$ and $\hat{p}_j$. The cell $V(\hat{p}_i)$ is the intersection of the $n - 1$ half-spaces bounded by the bisecting hyperplanes of $\hat{p}_i$ and each of the other weighted points of $\hat{P}$. If this intersection is not empty, it is a convex polyhedron, possibly unbounded. We call *weighted Voronoi cells* the non empty cells $V(\hat{p}_i)$, $i \in [1 : n]$.

We define the *weighted Voronoi diagram* of $\hat{P}$, noted $\text{Vor}(\hat{P})$, as the cell complex whose faces are the weighted Voronoi cells and their faces. Note that the set of weighted Voronoi diagrams includes Voronoi diagrams. Indeed, when all weighted points have the same weight, their weighted Voronoi diagram is identical to the Voronoi diagram of their centers.

Equivalently, the weighted Voronoi diagram of $\hat{P}$ can be defined as the minimization diagram of the set of functions $\{D(x, \hat{p}_1), \dots, D(x, \hat{p}_n)\}$. Observing that, for any x , $\arg \min_i D(x, \hat{p}_i) = \arg \max_i (2p_i \cdot x - p_i^2 + w_i)$, we obtain that the weighted Voronoi diagram of $\hat{P}$ is the maximization diagram of the set of affine functions $\mathcal{F} = \{f_1, \dots, f_n\}$, where

$$f_i(x) = 2p_i \cdot x - p_i^2 + w_i.$$

The graph of $f_i(x)$ is a non vertical hyperplane of $\mathbb{R}^{d+1}$ that we denote by h_i :

$$h_i = \{(x, x(d+1)) \in \mathbb{R}^d \times \mathbb{R}, x(d+1) = 2p_i \cdot x - p_i^2\}.$$

Let $H = \{h_1, \dots, h_n\}$ and let us denote by $\mathcal{V}(P)$ the upper envelope of the hyperplanes of H . The maximization diagram of $\mathcal{F}$ is obtained by projecting vertically the faces of the upper envelope $\mathcal{V}(\hat{P})$.




Figure 4.8: A weighted Voronoi diagram.

Hence, the faces of the weighted Voronoi diagram $\text{Vor}(\hat{P})$ are the vertical projections of the faces of $\mathcal{V}(\hat{P})$.

Weighted Voronoi diagrams are very similar to Voronoi diagrams: the main difference is that some weighted point $\hat{p}_i$ in $\hat{P}$ may have an empty Voronoi cell in $\text{Vor}(\hat{P})$ (see the small circle in the upper left corner of Figure 4.8). Equivalently, the corresponding hyperplane h_i does not contribute to a face of the upper envelope $\mathcal{V}(\hat{P})$. Notice however that if the weights are non negative and small enough so that the weighted points in $\hat{P}$ correspond to disjoint balls, then all the weighted points have a non empty Voronoi cell, (Exercise 4.13).

4.4.3 Weighted Delaunay complexes

Let $\hat{P}$ be a finite set of weighted points of $\mathbb{R}^d$ and write as before $\text{Vor}(\hat{P})$ for its Voronoi diagram. The nerve of the collection of all cells of $\text{Vor}(\hat{P})$ is an abstract simplicial complex denoted by $\text{Del}(\hat{P})$ and called the *weighted Delaunay complex* of $\hat{P}$. Specifically, let f be a face of dimension k of $\text{Vor}(\hat{P})$ and let $\text{int} f$ be the set of points of f that do not belong to any proper subface of f . All the points of $\text{int} f$ have the same subset $\hat{\sigma}$ of closest weighted points in $\hat{P}$ (for the weighted distance), and f is the intersection of the weighted Voronoi cells of the weighted points in $\hat{\sigma}$. Accordingly, $\hat{\sigma}$ is a simplex in the weighted Delaunay complex.

Recall that two weighted points are said to be *orthogonal* if their weighted distance is zero. A weighted point $\hat{x}$ is said to be orthogonal to a finite set




Figure 4.9: The weighted delaunay triangulation associated to the weighted Voronoi diagram of Figure 4.8.

$\hat{\sigma}$ of weighted points if $\hat{x}$ is orthogonal to all the weighted points of $\hat{\sigma}$.

We say that the set $\hat{P}$ of weighted points is in *general position* when no weighted point of $\mathbb{R}^d \times \mathbb{R}$ is at equal weighted distance from more than $d+1$ weighted points of $\hat{P}$. In the following, we assume that the set $\hat{P}$ of weighted points is in general position.

Then, if $\hat{\sigma} = \{\hat{p}_0, \dots, \hat{p}_d\}$ is a subset of $d+1$ weighted points in $\hat{P}$, there exists a unique weighted point $\hat{c}$ in $\mathbb{R}^d \times \mathbb{R}$ orthogonal to $\hat{\sigma}$. The center $c(\sigma)$ of $\hat{b}$ is called the weighted center of $\hat{\sigma}$, and the square root of the weight of $\hat{c}$, denoted by $r(\hat{\sigma})$, is called the weighted radius of $\hat{\sigma}$. Observe that the weighted radius may be imaginary which causes no problem.

A weighted point $\hat{x}$ is said to be *free* of weighted points of $\hat{P}$ if no point of $\hat{P}$ has a negative weighted distance to $\hat{x}$. Observe that the notion of free weighted points generalizes the notion of empty balls (see Lemma 4.1). Indeed, when the weights of points in $\hat{P}$ are 0 and the weight of $\hat{x}$ is positive, the ball corresponding to $\hat{x}$ is an empty ball for P .

Let v be a vertex of $\text{Vor}(\hat{P})$. The weighted Delaunay simplex $\hat{\sigma}$ associated to v is a d -simplex. The vertex v is at equal weighted distance w from all the weighted points of $\hat{\sigma}$, and this distance is smaller than the weighted distance from v to all other weighted points in $\hat{P} \setminus \hat{\sigma}$. Therefore the weighted point (v, w) is orthogonal to $\hat{\sigma}$ and free of points of $\hat{P}$. This property generalizes to faces of $\text{Vor}(\hat{P})$ of any dimension and the following lemma is just another view of the definition of the weighted Delaunay complex.

Lemma 4.5 (The free weighted point property) *Any subset $\hat{\sigma} \subset \hat{P}$ is*

a simplex of the weighted Delaunay complex $\text{Del}(\hat{P})$ iff there is a weighted point orthogonal to $\hat{\sigma}$ and free of weighted points of $\hat{P}$.

We proceed now as in Section 4.3 and prove the weighted version of Delaunay's triangulation theorem. With a slight abuse of language, we call natural embedding of $\text{Del}(\hat{P})$ in $\mathbb{R}^d$, a geometric realization of $\text{Del}(\hat{P})$ in which each simplex $\hat{\sigma}$ of $\text{Del}(\hat{P})$ is embedded as the convex hull of the set σ of the centers of the weighted points in $\hat{\sigma}$.

Theorem 4.6 (Weighted Delaunay triangulation) *If a finite set of weighted points $\hat{P} \in \mathbb{R}^d \times \mathbb{R}$ is in general position, then the weighted Delaunay complex $\text{Del}(\hat{P})$ has a natural embedding in $\mathbb{R}^d$ which is a triangulation of a subset P' of P . P' is the set of centers of the weighted points in $\hat{P}$ that have a non empty cell in $\text{Vor}(\hat{P})$.*

Proof We first define the lifting map for weighted points as follows: to a weighted point $\hat{p} = (p, w)$ of $\mathbb{R}^d$, we associate the point $\phi(\hat{p}) = (p, p^2 - w)$ of $\mathbb{R}^{d+1}$. Then we consider the lower hull of points of $\phi(\hat{P})$, denoted by $\mathcal{D}(\hat{P})$. Arguing as in Section 4.3, we get that the lower hull $\mathcal{D}(\hat{P})$ is a geometric simplicial complex embedded in $\mathbb{R}^{d+1}$ whose vertical projection onto $\mathbb{R}^d$ is the natural embedding of $\text{Del}(\hat{P})$. Let $\hat{P}'$ be the subset of weighted points in $\hat{P}$ that have non empty cells in $\text{Vor}(\hat{P})$. The vertex set of $\mathcal{D}(\hat{P})$ is the set $\phi(\hat{P}') \subseteq \phi(\hat{P})$ and the vertical projection of $\mathcal{D}(\hat{P})$ is a triangulation of P' , the set of centers of the weighted points of $\hat{P}'$. $\square$

Observe that the hyperplanes $h_i, i = 1, \dots, n$ defined in Subsection 4.4.2 are the hyperplanes dual to the points in $\phi(\hat{P})$. Therefore, by the results of Section 3.2, the upper envelope $\mathcal{V}(P)$ and the lower hull $\mathcal{D}(P)$ are dual complexes. Arguing once more as in Section 4.3, we get the following diagram which commutes when the weighted points are in general position :

$$\begin{array}{ccc} \mathcal{V}(\hat{P}) & \xrightarrow{\text{duality}} & \mathcal{D}(\hat{P}) \\ \uparrow & & \downarrow \\ \text{Vor}(\hat{P}) & \xrightarrow{\text{nerve}} & \text{Del}(\hat{P}) \end{array}$$

If the weighted points are not in general position, we can triangulate the non simplicial faces $\mathcal{D}(\hat{P})$ as described in Section 4.3. The vertical projection of $\mathcal{D}(\hat{P})$ will then be a triangulation of P' .

4.4.4 Complexity of weighted Delaunay complexes

The following theorem states that computing the weighted Voronoi diagram of n weighted points of $\mathbb{R}^d \times \mathbb{R}$ (or equivalently its dual weighted Delaunay triangulation) has the same asymptotic complexity as computing the Euclidean Voronoi diagram or the Delaunay triangulation of n points of $\mathbb{R}^d$. The theorem is a direct consequence of Section 4.4.3 and of results on convex hulls (Theorems 3.9 and 3.11).

Theorem 4.7 *The combinatorial complexity of the weighted Voronoi diagram of n weighted points of $\mathbb{R}^d \times \mathbb{R}$ and of its dual weighted Delaunay triangulation are $\Theta\left(n^{\lceil \frac{d}{2} \rceil}\right)$. Both structures can be computed in optimal time $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.*

4.5 Examples of weighted Voronoi diagrams

We have seen that the weighted Voronoi diagram of n weighted points is the maximization diagram of n affine functions. The converse is also true (see Exercise 4.14). We give in this section two examples of weighted Voronoi diagrams that are of interest in the context of data analysis.

4.5.1 k -order Voronoi diagrams

Let P be a set of n points of $\mathbb{R}^d$ and let P_k be the set of all subsets of k points of P for some fixed $k \in [1 : n - 1]$. We define the Voronoi cell of a subset $K \in P_k$ as the set of points of $\mathbb{R}^d$ that are closer to all the sites in K than to any other site in $P \setminus K$:

$$V_k(K) = \{x \in \mathbb{R}^d : \forall p_i \in K, \forall p_j \in P \setminus K, \|x - p_i\| \leq \|x - p_j\|\}.$$

Let us consider the subsets of P_k whose Voronoi cells are not empty. These cells are convex polyhedra and form a cell complex whose domain is $\mathbb{R}^d$ called the k -order diagram of P (see Figure 4.10). For $k = 1$, we obtain the usual Voronoi diagram.




Figure 4.10: The 2-order Voronoi diagram of a set of points (in bold line) and the corresponding 1-order Voronoi diagram (in thin line).

Theorem 4.8 *The k -order diagram of P is the weighted Voronoi diagram of a set of $\binom{n}{k}$ weighted points of $\mathbb{R}^d$.*

Proof Let $K_1, \dots, K_s$ be the $s = \binom{n}{k}$ subsets of k points of P . For any point $x \in \mathbb{R}^d$, we have

$$\begin{aligned}
 x \in V_k(K_i) &\iff \frac{1}{k} \sum_{p \in K_i} (x - p)^2 \leq \frac{1}{k} \sum_{q \in K_j} (x - q)^2 \quad \forall j, 1 \leq j \leq s \\
 &\iff x^2 - 2 \left(\frac{1}{k} \sum_{p \in K_i} p \right) \cdot x + \frac{1}{k} \sum_{p \in K_i} p^2 \\
 &\quad \leq x^2 - 2 \left(\frac{1}{k} \sum_{q \in K_j} q \right) \cdot x + \frac{1}{k} \sum_{q \in K_j} q^2 \\
 &\iff D(x, \hat{p}_i) \leq D(x, \hat{p}_j)
 \end{aligned}$$

where $\hat{p}_i = (c_i, w_i)$ is the weighted point centered at the center of mass $c_i = \frac{1}{k} \sum_{p \in K_i} p$ of K_i of weight $w_i = c_i^2 - \rho_i^2$ with $\rho_i^2 = \frac{1}{k} \sum_{p \in K_i} p^2$. Hence,

$x \in V_k(K_i)$ iff x lies in the weighted Voronoi cell of $\hat{p}_i$. $\square$

It follows that the k -order Voronoi diagram of a finite set of points P of $\mathbb{R}^d$ is a cell complex of $\mathbb{R}^d$ whose nerve is a triangulation of $\mathbb{R}^d$ under a general position assumption.

The k -order Voronoi diagram of P can be used as a data structure to compute in sublinear time the k points of P that are closest to a query point x . The query can be answered by identifying the cell of the k -order Voronoi diagram that contains x and reading the associated k -nearest points of P .

4.5.2 Bregman diagrams

Let Ω be a convex domain of $\mathbb{R}^d$ and F a strictly convex and differentiable function F (called the *generator function* of the divergence) defined over Ω . For any two points $p = (p(1), \dots, p(d))$ and $q = (q(1), \dots, q(d))$ of Ω , the Bregman divergence $D_F(p||q) : \Omega \times \Omega \mapsto \mathbb{R}$ associated to F is defined as

$$D_F(p||q) = F(p) - F(q) - \nabla F(q) \cdot (p - q) \quad (4.3)$$

where $\nabla F = [\frac{\partial F}{\partial x_1} \dots \frac{\partial F}{\partial x_d}]^T$ denotes the gradient operator.

Informally speaking, Bregman divergence D_F is the *tail* of the Taylor expansion of F . Geometrically, the Bregman divergence $D_F(p||q)$ is measured as the vertical distance (i.e. along the $(d+1)$ -axis) between $\hat{p} = (p, F(p))$ and the hyperplane H_q tangent to the graph $\mathcal{F}$ of F at point $\hat{q}$: $D_F(p||q) = F(p) - H_q(p)$. See Figure 4.11.

We now give some basic properties of Bregman divergences. First, observe that, for most functions F , the associated Bregman divergence is *not* symmetric, i.e. $D_F(p||q) \neq D_F(q||p)$ (the symbol $||$ is put to emphasize this point). Hence, it is not a distance. Nevertheless, the strict convexity of the generator function F implies that, for any p and q in Ω , $D_F(p||q) \geq 0$, with $D_F(p||q) = 0$ if and only if $p = q$.

Examples of Bregman divergences

Examples of Bregman divergences are the squared Euclidean distance (obtained with $F(x) = x^2$) and the quadratic distance function $F(x) = x^T Q x$




Figure 4.11: The Bregman divergence.

where Q is a symmetric positive definite matrix.

The notion of Bregman divergence encapsulates various information measures based on entropic functions such as the Kullback-Leibler divergence based on the Shannon entropy which is widely used in information theory, image processing and various fields. Let p be a discrete probability distribution so that $\sum_{i=1}^d p(i) = 1$. The Shannon entropy is defined as $F(p) = \sum_i p(i) \log_2 p(i)$. F is a convex function and the associated Bregman divergence between two probability distributions p and q is easily shown to be

$$\begin{aligned}
 D_F(p||q) &= \sum_{i=1}^d p(i) \log_2 p(i) - \sum_{i=1}^d q(i) \log_2 q(i) - (p - q) \cdot \nabla F(q) \\
 &= \sum_{i=1}^d p(i) \log_2 \left(\frac{p(i)}{q(i)} \right) \quad \left(\text{since } \sum_{i=1}^d p(i) = \sum_{i=1}^d q(i) = 1 \right) \\
 &\stackrel{\text{def}}{=} KL(p||q).
 \end{aligned}$$

$KL(p||q)$ is called the *Kullback-Leibler divergence* or the relative entropy of the two probability distributions p and q .

Bregman diagrams

Let $P = \{p_1, \dots, p_n\}$ be a finite point set of $\Omega \subset \mathbb{R}^d$. We associate to each site p_i the distance function, $D_i(x) = D_F(x||p_i)$. The *minimization diagram* of the D_i , $i = 1, \dots, n$, is called the Bregman Voronoi diagram of P , which we denote by $\text{Vor}_F(P)$. The d -dimensional cells of this diagram are in *1-1 correspondence* with the sites p_i and the d -dimensional cell of p_i is defined as

$$V_F(p_i) \stackrel{\text{def}}{=} \{x \in \Omega \mid D_F(x||p_i) \leq D_F(x||p_j) \ \forall p_j \in P\}.$$

It is easy to see that the minimization diagram of the n functions $D_F(x||p_i)$ $i = 1, \dots, n$, is the maximization of the n affine functions

$$h_i(x) = (x - p_i) \cdot \nabla F(p_i) + F(p_i), \quad i = 1, \dots, n.$$

Hence, Bregman diagrams are weighted Voronoi diagrams. More precisely, we have

Theorem 4.9 *The Bregman Voronoi diagram of n sites is identical to the restriction to Ω of the weighted Voronoi diagram of the n weighted points (p'_i, w_i) , $i = 1, \dots, n$, where $p'_i = \nabla F(p_i)$ and $w_i = p_i'^2 + 2(F(p_i) - p_i \cdot p'_i)$.*

Proof $D_F(x||p_i) \leq D_F(x||p_j)$ iff

$$-F(p_i) - (x - p_i) \cdot p'_i \leq -F(p_j) - (x - p_j) \cdot p'_j.$$

Multiplying by 2 the two sides of the inequality and adding x^2 to both sides yields

$$\begin{aligned} x^2 - 2x \cdot p'_i - 2F(p_i) + 2p_i \cdot p'_i &\leq x^2 - 2x \cdot p'_j - 2F(p_j) + 2p_j \cdot p'_j \\ \iff (x - p'_i) \cdot (x - p'_i) - w_i &\leq (x - p'_j) \cdot (x - p'_j) - w_j, \end{aligned}$$

where $w_i = p_i'^2 + 2(F(p_i) - p_i \cdot p'_i)$ and $w_j = p_j'^2 + 2(F(p_j) - p_j \cdot p'_j)$. The last inequality means that the weighted distance of x to the weighted point (p'_i, w_i) is no more than its weighted distance to the weighted point (p'_j, w_j) . $\square$

It is to be observed that not all weighted Voronoi diagrams are Bregman Voronoi diagrams. Indeed, in weighted Voronoi diagrams, some weighted points may have empty cells while each site has necessarily a non empty cell in a Bregman Voronoi diagram.

Bregman complexes

We define the Bregman complex of P as the nerve of the Bregman diagram of P , $\text{Vor}_F(P)$. We denote it by $\text{Del}_F(P)$. An analogue of Delaunay's triangulation theorem also exists in this context.

Let $\hat{P}$ be the lifted image of P on the graph $\mathcal{F}$ of F , i.e. $\hat{P} = \{(p, F(p)), p \in P\} \in \mathbb{R}^{d+1}$. Write $\mathcal{T}$ for the lower hull of $\hat{P}$. We assume in this section that P is in *general position*, meaning that there is no point $x \in \Omega$ whose divergences to $d+2$ points of P are equal. Equivalently, P is in general position if $\hat{P}$ contains no subset of $d+2$ points on a same hyperplane.

Assume that P is in general position. Then, for the same reasons as for Delaunay triangulations (see Section 4.3), $\mathcal{T}$ is a simplicial complex and the vertical projection of $\mathcal{T}$ onto $\mathbb{R}^d$ is a triangulation that is equal to $\text{Del}_F(P)$ of P by the lifting map argument given in the beginning of Section 4.5.2 (see Figurefig:bregman-lift). In other words, $\text{Del}_F(P)$ naturally embeds in $\Omega \subseteq \mathbb{R}^d$. We call $\text{Del}_F(P)$ the *Bregman triangulation* of P . When $F(x) = x^2$ and $\Omega = \mathbb{R}^d$, $\text{Del}_F(P)$ is the Delaunay triangulation of P .

We now show that the empty ball property of Delaunay triangulations (Lemma 4.1) naturally extends to Bregman triangulations. We define the *Bregman ball* centered at c and of radius r as

$$b_F(c, r) = \{x \in \Omega \mid D_F(x||c) < r\}.$$

It is easy to see that any Bregman ball b_F is obtained as the vertical projection of the intersection of $\mathcal{F}$ with a halfspace below a non-vertical hyperplane (see Figure 4.6 for the case where $F = x^2$).

A Bregman ball is said to be *empty* if it does not contain any point of P . Let $\sigma = [p_0, \dots, p_d]$ be a (geometric) d -simplex of $\text{Breg}_F(P)$. The affine hull of the lifted points $\hat{p}_0, \dots, \hat{p}_d$ is a hyperplane h_σ of $\mathbb{R}^{d+1}$ whose intersection with $\mathcal{F}$ projects vertically onto the boundary of the (unique) Bregman ball b_σ that circumscribes σ . Since, by construction, $\text{conv}(\hat{p}_0, \dots, \hat{p}_d) \in h_\sigma$ is a facet of the lower hull of $\hat{P}$, b_σ must be empty.

It also follows from the discussion that, under the general position assumption, Bregman complexes naturally embeds in $\mathbb{R}^d$. This is another extension of Delaunay's triangulation theorem.

Bregman diagrams and complexes of the second type

We have defined Bregman diagrams as the minimization diagram of the functions $D_F(x||p_i)$. A symmetric definition can be given when exchanging the variable x and the site p_i and considering $D_F(p_i||x)$. The bisectors are no longer hyperplanes and the diagram is no longer a weighted Voronoi diagram. Still, Legendre duality, an essential notion in convex analysis, allows to transform a diagram of the second type into a diagram of the first type as explained next (Lemma 4.11).

Legendre duality

Let F be a strictly convex and differentiable real-valued function on Ω . The gradient of F , ∇F , is well defined as well as its inverse $\nabla^{-1}F$, and $\nabla F \circ \nabla^{-1}F = \nabla^{-1}F \circ \nabla F$ are identity maps. We write x' for $\nabla F(x)$ and Ω' for the *gradient space* $\{\nabla F(x)|x \in \Omega\}$.

The Legendre transformation associates to F a *convex conjugate* function $F^* : \Omega' \mapsto \mathbb{R}$ given by [122]:

$$F^*(x') = x \cdot x' - F(x). \quad (4.4)$$

Taking the derivative of Equation 4.4, we get

$$\nabla F^*(x') \cdot dx' = x \cdot dx' + x' \cdot dx - \nabla F(x) \cdot dx = x \cdot dx' = \nabla^{-1}F(x') \cdot dx',$$

from which we deduce that $\nabla F^* = \nabla^{-1}F$.

Figure 4.12 gives a geometric interpretation of the Legendre transformation. Consider the hyperplane h_x tangent to the graph of $\mathcal{F}$ at $\hat{x}$. This hyperplane intersects the $(d+1)$ -axis at the point $(0, -F^*(x'))$. Indeed, the equation of h_x is $y(d+1) = x' \cdot (y - x) + F(x) = x' \cdot y - F^*(x')$. Hence, the $(d+1)$ -axis intercept of h_x is equal to $-F^*(x')$. Any hyperplane passing through another point of $\mathcal{F}$ and parallel to h_x necessarily intersects the z -axis above $-F^*(x')$.

To ensure that D_{F^*} is a Bregman divergence, we need Ω' to be convex. This is trivial if Ω has no boundary and, in particular, when $\Omega = \mathbb{R}^d$. Otherwise, we will further require that F is a function of Legendre type, i.e. that the norm of the gradient of F goes to infinity whenever we approach the




Figure 4.12: The z -intercept $(0, -F^*(x'))$ of the tangent hyperplane h_x of F at $\hat{x}$ defines the value of the Legendre transform F^* for the dual coordinate x' .

boundary of Ω (see [122]). Under the additional assumption that F is of Legendre type, we can associate to the conjugate function F^* a Bregman divergence D_{F^*} . We have the following remarkable result:

Lemma 4.10 $D_F(p||q) = F(p) + F^*(q') - p \cdot q' = D_{F^*}(q'||p')$

Proof By Equation 4.3, $D_F(p||q) = F(p) - F(q) - (p - q) \cdot q'$, and, according to Equation 4.4, we have $F(p) = p' \cdot p - F^*(p')$ and $F(q) = q' \cdot q - F^*(q')$. Hence,

$$\begin{aligned} D_F(p||q) &= F(p) - F(q) - (p - q) \cdot q' \\ &= p' \cdot p - F^*(p') - p \cdot q' + F^*(q') \\ &= F^*(q') - F^*(p') - p \cdot (q' - p') \\ &= D_{F^*}(q'||p') \end{aligned}$$

where the last equality holds since $p = \nabla F^{-1} \nabla F(p) = \nabla F^*(p')$. $\square$

Observe that, when D_F is symmetric, D_{F^*} is also symmetric.

The Legendre transform of the quadratic form $F(x) = \frac{1}{2}x^T Q x$, where Q is a symmetric invertible matrix, is $F^*(x') = \frac{1}{2}x'^T Q^{-1}x'$. We say that F is self dual. Observe that the corresponding divergences D_F and D_{F^*} are both generalized quadratic distances.

To compute F^* , we can use the fact that $\nabla F^* = \nabla F^{-1}$ (see Exercise 4.16). However, integrating functions symbolically may be difficult or even not possible, and, in some cases, it will be required to approximate numerically the inverse gradient $\nabla^{-1}F(x)$.

We can now define Bregman diagrams of the second type where the cell of p_i is defined as

$$\tilde{V}_F(p_i) \stackrel{\text{def}}{=} \{x \in \Omega \mid D_F(p_i||x) \leq D_F(p_j||x) \ \forall p_j \in P\}.$$

In contrast with the diagram of the first-type $\text{Vor}_F(P)$, the diagram of the second type $\widetilde{\text{Vor}}(P)$ has, in general, curved faces. From the Legendre duality between divergences, we deduce correspondences between the diagrams of the first and the second types. As usual, F^* denotes the convex conjugate of F .

Lemma 4.11 *We have*

$$\widetilde{\text{Vor}}_F(P) = \nabla^{-1}F(\text{Vor}_{F^*}(P')) \quad \text{and} \quad \text{Vor}_F(P) = \nabla^{-1}F(\widetilde{\text{Vor}}_{F^*}(P')).$$

Proof By Lemma 4.10, we have $D_F(x||y) = D_{F^*}(y'||x')$, which gives $V_F(p_i) = \{x \in \Omega \mid D_{F^*}(p'_i||x') \leq D_{F^*}(p'_j||x') \ \forall p'_j \in P'\} = \nabla^{-1}F(\tilde{V}_{F^*}(p'_i))$. This proves the second part of the lemma. The proof of the first part follows the same path. $\square$

Hence, constructing $\widetilde{\text{Vor}}_F(P)$, the (curved) Bregman diagram of the second-type, reduces to constructing an affine diagram in the gradient space Ω' (and mapping the cells by $\nabla^{-1}F$). The nerve of the Bregman diagram of the second-type is a simplicial complex called the Bregman triangulation of the second type. It follows from Lemma 4.11 that the Bregman triangulation of the second type of P can be realized as the (curved) image by $\nabla^{-1}F$ of the Bregman Delaunay triangulation of the first type of P' .

4.6 Exercises

Exercise 4.1 (Space of spheres) Let b be a d -ball of $\mathbb{R}^d$ of center c and radius r and let $s = c^2 - r^2$. We associate to b the point of $\mathbb{R}^{d+1}$ $\phi(\sigma) = (c, s)$. Show that the image by ϕ of a point, considered as a ball of radius 0, is a point of the paraboloid $\mathcal{Q}$ of equation $x_{d+1} = x^2$.

Show that the image by ϕ of the balls whose bounding spheres pass through a given point p of $\mathbb{R}^d$ is the *hyperplane* h_p of $\mathbb{R}^{d+1}$ of equation $x_{d+1} = 2p \cdot x - p^2$, which is tangent to $\mathcal{Q}$ at $\phi(p) = (p, p^2)$.

What are the preimages by ϕ of the points of $\mathbb{R}^{d+1}$ that lie above $\mathcal{Q}$, on the boundary of $\mathcal{V}(P)$, in the interior of $\mathcal{V}(P)$, on a line ?

Exercise 4.2 (Farthest point diagram) Consider the diagram obtained by projecting the faces of $h_{p_1}^- \cap \dots \cap h_{p_n}^-$ vertically, where the h_{p_i} are defined as in Exercise 4.1. Characterize the points that belong to a face of this diagram.

Dually, project vertically the faces of the *upper* convex hull of the $\phi(p_i) = (p_i, p_i^2)$, $i = 1, \dots, n$. Show that we obtain a triangulation of the vertices of $\text{conv}(P)$ such that each ball circumscribing a simplex contains all the points of P .

Exercise 4.3 (Upper bound) Consider first the case of a set of points P lying on two non coplanar lines of $\mathbb{R}^3$. Let $n_1 + 1$ and $n_2 + 1$ be the numbers of points on each of the lines. Show that the Delaunay complex of P has $n_1 n_2$ tetrahedra (or, equivalently, that their Voronoi diagram has $n_1 n_2$ vertices).

Show that the moment curve $\mathcal{M}_d \subset \mathbb{R}^d$ (defined in Exercise 3.12) can be drawn on the paraboloid $\mathcal{Q}$ of equation $x_{d+1} = x^2$. Deduce then from Exercise 3.12 that the bound in Theorem 4.4 is tight.

Exercise 4.4 (Triangulation of linear size) Prove that any set of n points of $\mathbb{R}^d$ in general position admits a triangulation of size $O(n)$.

Exercise 4.5 (in_ball predicate) Let B be a ball of $\mathbb{R}^d$ whose bounding sphere S passes through $d + 1$ points $p_0, \dots, p_d$. Show that a point p_{d+1} of $\mathbb{R}^d$ lies on S , in the interior of B or outside B , depending whether the determinant of the $(d + 2) \times (d + 2)$ matrix

$$\text{in_sphere}(p_0, \dots, p_{d+1}) = \begin{vmatrix} 1 & \cdots & 1 \\ p_0 & \cdots & p_{d+1} \\ p_0^2 & \cdots & p_{d+1}^2 \end{vmatrix}$$

is 0, negative or positive. Show that this predicate is the only numerical operation that is required to check if a triangulation is a Delaunay triangulation.

Show that the only numerical operation that is required to check if a triangulation is the weighted Delaunay triangulation of a set of weighted points $\hat{p}_1, \dots, \hat{p}_n$ is the evaluation of signs of determinants of $(d+2) \times (d+2)$ matrices of the form

$$\text{power_test}(\hat{p}_{i_1}, \dots, \hat{p}_{i_{d+2}}) = \begin{vmatrix} 1 & \cdots & 1 \\ p_{i_1} & \cdots & p_{i_{d+2}} \\ p_{i_0}^2 - w_{i_0} & \cdots & p_{i_{d+1}}^2 - w_{i_{d+1}} \end{vmatrix}$$

where p_i and w_i are respectively the center and the weight of $\hat{p}_i$. (Hint : use the lifting map and Exercise 3.16).

Exercise 4.6 (Triangulation of a convex polytope) Describe an algorithm to triangulate a convex polytope (Hint: proceed by faces of increasing dimensions).

Exercise 4.7 (Minimal spanning tree) Let P be a finite set of points of $\mathbb{R}^d$. A spanning tree of P is a tree whose vertices are the points of P . A spanning tree is called a minimum spanning tree (MST) if the sum of the lengths of its edges is minimal among all spanning trees. Show that $\text{MST}(P) \subset \text{Del}(P)$.

Exercise 4.8 (Delaunay complexes contain a triangulation) Let P be any finite set of points of $\mathbb{R}^d$. Show that there exists a subcomplex $K \subseteq \text{Del}(P)$ that is a triangulation of P .

Exercise 4.9 (Natural coordinates) Let $P = \{p_1, \dots, p_n\}$ be a finite set of points of $\mathbb{R}^d$. As usual, $or(P)$ denotes the Voronoi diagram of P and $V(p_i)$ the cell of p_i in $or(P)$. Given a point $x \in \text{conv}(P)$, we write $P^+ = P \cup \{x\}$, $V^+(x)$ for the Voronoi cell of x in $or(P^+)$, $V^+(x, p_i) = V^+(x) \cap V(p_i)$ and $W(x, p_i) = V^+(x) \cap V(p_i) = V(p_i) \setminus V^+(x)$. Now we define $v_i(x) = \text{vol}(V^+(x, p_i))$, $\bar{v}_i(x) = v_i(x)/\|x - p_i\|$ and $\bar{v}(x) = \sum_{i=1}^n \bar{v}_i(x)$. In addition, we define $w_i(x) = \text{vol}(W(x, p_i))$ and $w(x) = \sum_{i=1}^n w_i(x)$.

We call Laplace coordinates the n functions $\lambda_1, \dots, \lambda_n$ defined by $\lambda_i(x) = \bar{v}_i(x)/\bar{v}(x)$ for $x \notin P$, and $\lambda_i(p_j) = \delta_{ij}$ otherwise, where δ_{ij} is the Kronecker delta. We call Sibson's coordinates the n functions $arsigma_i(x) = w_i(x)/w(x)$, $i = 1, \dots, n$. Show that the set of λ_i is a partition of unity, i.e. $x = \sum_i \lambda_i(x) p_i$. Same question for the set of $arsigma_i$.

(Hint : for Laplace coordinates, apply Exercise 3.18 to $V^+(x)$. For Sibson's coordinates, apply Exercise 3.18 to the polytope of $\mathbb{R}^{d+1}$ whose boundary is $\mathcal{V}(P^+) \setminus \mathcal{V}(P)$, where $\mathcal{V}(P)$ is defined in Section 4.3).

Exercise 4.10 (Conservation law) Prove the following conservation law for flows entering a weighted Voronoi cell $L(\sigma_i)$ normally to the facets of the cell: if f_{ij} , $j \in J$, are the facets of $L(\sigma_i)$, we have $\sum_{j \in J} ol(f_{ij}) \frac{c_j - c_i}{\|c_j - c_i\|} = 0$. This property makes Voronoi and weighted Voronoi diagrams useful when applying finite volume methods in fluid dynamics. (Hint : use Exercise 3.18).

Exercise 4.11 (Section of a Voronoi diagram) Let H be a k -dimensional affine space of $\mathbb{R}^d$. Show that the intersection L_H of H with the weighted Voronoi diagram of n weighted points $\mathbb{R}^d$ is the weighted Voronoi diagram of n weighted points of H .

Exercise 4.12 (Union of balls) Show that the combinatorial complexity of the union of n balls of $\mathbb{R}^d$ is the same as the combinatorial complexity of their weighted Voronoi diagram. Design a worst-case optimal algorithm to compute such a union of balls. (Hint : bound first the number of vertices.)

Exercise 4.13 (Centered triangulation)) Let $\hat{P}$ be a finite set of weighted points with non negative weights. Let P be the set of their centers and let $l = \min_{p,q \in P} \|p - q\|$ be the minimal distance between two centers. Prove that, if the weights are less than $\frac{l}{4}$, all weighted points of $\hat{P}$ have a non empty weighted Voronoi cell and each center belongs to the cell of its own weighted point.

Exercise 4.14 (Maximization diagrams) Show that the maximization diagram of n affine functions defined over $\mathbb{R}^d$ is the weighted Voronoi diagram of n weighted points of $\mathbb{R}^d$.

Exercise 4.15 (Complexity of k -order Voronoi diagrams) Use Theorems 3.12 and 3.9 to show that the combinatorial complexity of all $\leq k$ -order Voronoi diagrams is $\Theta(k^{\lceil \frac{d+1}{2} \rceil} n^{\lfloor \frac{d+1}{2} \rfloor})$. Propose an efficient algorithm to compute all these diagrams.

Exercise 4.16 Show that the Hellinger distance between two distributions p and q , which is defined as $\sqrt{\int (\sqrt{f(x)} - \sqrt{g(x)})^2 dx}$, is a Bregman divergence when setting $F(x) = -\sqrt{1-x^2}$. Show that the inverse gradient is $\frac{x}{\sqrt{1+x^2}}$ and that the conjugate function is $\int \frac{x dx}{\sqrt{1+x^2}} = \sqrt{1+x^2}$.

Show that, for the exponential function $F(x) = \exp x$, we have $F^*(y) = y \log y - y$ (the unnormalized Shannon entropy) and, for the dual bit entropy $F(x) = \log(1+\exp x)$, we have $F^*(y) = y \log \frac{y}{1-y} + \log(1-y)$, the bit entropy.

Show that, for the Burg entropy $F(x) = -\log x$, the gradient and inverse gradient are identical, so that it is *self-dual*, i.e. $F = F^*$.

Exercise 4.17 We define the Bregman ball of the second type centered at c and of radius r as

$$\tilde{b}_F = \{x \in \Omega \mid D_F(c \parallel x) < r\}.$$

Show that if, F is of Legendre type, $\tilde{b}_F$ is contractible.

Exercise 4.18 Define the α -Bregman complex as the restriction of the Bregman complex to the union of the Bregman balls of radius α . Consider the two types associated to the two types of balls. Show that, in both cases, the α -Bregman complex has the same homotopy type as the union of the Bregman balls.

4.7 Bibliographical notes

Voronoi diagrams are very natural constructions that have been discovered several times and appear in the literature under various names like Dirichlet tessellations or Thiessen diagrams. The observation that Voronoi diagrams of $\mathbb{R}^d$ are projections of convex polyhedra of $\mathbb{R}^{d+1}$ goes back to Voronoi himself. Delaunay has defined the triangulations that bear his name and proved Theorem 4.3 in his seminal paper [64].

Voronoi diagrams and Delaunay triangulations are fundamental geometric structures that have received a lot of attention. Main results can be found in most textbooks on discrete and computational geometry [61, 27, 88] and more comprehensive treatments can be found in the books by Okabe et al. [116] and Aurenhammer et al. [7].

The space of spheres we referred to in Exercise 4.1 is fully developed in the books by Pedoe [118] and Berger [9].

Weighted Voronoi diagrams appear also in the literature under the names of Laguerre diagrams or power diagrams. Weighted Delaunay triangulations are also named regular triangulations. They were first studied in a systematic way by Aurenhammer [6] who proved Theorem 4.8. The solution to Exercise 4.15 is due to Clarkson and Shor [59]. Bregman Voronoi diagrams were introduced by Boissonnat, Nielsen and Nock [24]. The fact that the α -Bregman complex has the same homotopy type as the union of the Bregman balls (Exercise 4.18) has been first observed by Edelsbrunner and Wagner [78]. A recent survey on affine and curved Voronoi diagrams can be found in [26].

Natural coordinates (Exercise 4.9) have been introduced by Sibson [128, 127].

The CGAL library [17] (www.cgal.org) offers fully reliable and efficient implementations of algorithms to construct Delaunay and weighted Delaunay

triangulations in arbitrary dimensions.

Chapter 5

Good triangulations

In this chapter, we are interested in constructing simplicial complexes with a guaranteed quality. On one hand, we would like to control the density of their vertices and, on the other hand, we would like to control the shape of the simplices and avoid simplices that are too flat.

In Section 5.1, we introduce ε -nets in order to capture a notion of *good* sample of a bounded subset $\Omega \subset \mathbb{R}^d$. We will see that, for fixed d , the complexity of Voronoi diagrams and Delaunay complexes of an ε -net of Ω of size n is linear in n . This is to be compared with the bound $\Theta(n^{\lceil \frac{d}{2} \rceil})$ given in Chapter 4 which is tight in the worst case.

Although Delaunay triangulations have many beautiful properties, their simplices, in dimension greater than 2, may have an arbitrarily small volume even if their vertices are well distributed. Avoiding such bad simplices is a major issue and the importance of thick triangulations, to be introduced in Section 5.2, has been recognized since the early days of differential topology. Thick triangulations play a central role in many works on the triangulation of manifolds (see Part III) and appear to be crucial in scientific computing to ensure the convergence of numerical methods.

In order to improve the quality of the simplices of a Delaunay complex, one can perturb the position of the vertices or the metric of the space. We introduce in Section 5.3 a perturbation scheme that associates to each point a weight and replaces the Delaunay complex by its weighted version. We show that the weight assigned to each point can be computed so that the resulting weighted Delaunay complex has some guaranteed thickness. The method is an algorithmic application of the Lovász local lemma which is recalled in Section 5.3.4.

5.1 Nets

5.1.1 Nets in Euclidean space

We consider a bounded subset Ω of $\mathbb{R}^d$ and denote by P a finite set of points in Ω . The Hausdorff distance $d_H(P, \Omega)$ is called the *sampling radius* of P and denoted by ε . We also say that P is an ε -dense sample of Ω . We further call $\eta = \min_{p, q \in P} \|p - q\|$ the *separation* of P and $\bar{\eta} = \eta/\varepsilon$ the *separation ratio* of P .

Any finite point set P of Ω whose sampling radius is ε and whose separation ratio is $\bar{\eta}$ is called an $(\varepsilon, \bar{\eta})$ -net of Ω . Note that for *any* finite set of distinct points $P \subset \Omega$, there is some positive ε and $\bar{\eta}$ such that P is an $(\varepsilon, \bar{\eta})$ -net for Ω . Thus ε and $\bar{\eta}$ are simply parameters that describe properties of P in Ω .

In the sequel, we will often consider the subset of $(\varepsilon, \bar{\eta})$ -samples of Ω where $\bar{\eta}$ is lower bounded by a positive constant. Such point sets will be called ε -nets for short.

Lemma 5.1 *Let P be an $(\varepsilon, \bar{\eta})$ -net of Ω . If the radius of the smallest ball enclosing any connected component of Ω is greater than ε , then for any point $p \in P$, the distance $L(p)$ from p to its nearest neighbor in $P \setminus \{p\}$ is at most 2ε . Therefore we must have $\bar{\eta} \leq 2$.*

Proof Since P is ε -dense in Ω , the union of the balls $B(p, \varepsilon)$, $p \in P$, covers Ω . Now let $r > \varepsilon$ and assume for a contradiction that there exists a point $p \in P$ such that $B(p, 2r)$ does not contain any point of P other than p . Then Ω intersects the spherical shell $B(p, r) \setminus B(p, \varepsilon)$ and this intersection is not covered by any ball $B(q, \varepsilon)$, $q \in P$, violating the hypothesis that P is ε -dense in Ω . $\square$

The next lemma shows that nets exist.

Lemma 5.2 (Existence of ε -nets) *Let Ω be a bounded subset of $\mathbb{R}^d$ and ε be any positive real. Then Ω admits an $(\varepsilon, 1)$ -net.*

Proof We apply the following procedure : while there exists a point $p \in \Omega$ at distance at least ε from P , insert p in P . Since the domain is compact and the algorithm inserts no point at distance less than ε from a previously inserted point, the algorithm terminates. Upon termination, any point of Ω is at distance less than ε from P . $\square$

Lemma 5.3 (Size of an ε -net) *Let Ω be a bounded domain of $\mathbb{R}^d$. The number of points of an $(\varepsilon, \bar{\eta})$ -net satisfies*

$$\frac{\text{vol}_d(\Omega)}{\text{vol}_d(B(\varepsilon))} \leq n(\varepsilon, \bar{\eta}) \leq \frac{\text{vol}_d(\Omega^{\frac{\bar{\eta}}{2}})}{\text{vol}_d(B(\frac{\bar{\eta}}{2}))}$$

where $B(r)$ denotes a d -ball of radius r , $\eta = \bar{\eta}\varepsilon$, and

$$\Omega^{\frac{\eta}{2}} = \left\{ x \in \mathbb{R}^d, d(x, \Omega) \leq \frac{\eta}{2} \right\}.$$

If $\text{vol}_d(\Omega) > 0$, then $n(\varepsilon, \bar{\eta}) = \Theta\left(\frac{1}{\varepsilon^d}\right)$ where the constant in the Θ depends on the geometry of Ω and on $\bar{\eta}^d$.

Proof The ball $B(p, \varepsilon)$ of radius ε that are centered at the points $p \in P$ cover Ω . This yields the left inequality. The balls $B(p, \frac{\eta}{2})$ of radius $\frac{\eta}{2}$ that are centered at the points $p \in P$ are disjoint and they are all contained in $\Omega^{\frac{\eta}{2}}$. This leads to the right inequality. $\square$

5.1.2 Delaunay complex of a net

Let Ω be a subset of $\mathbb{R}^d$ of positive d -volume and let P a finite set of points in Ω . The nerve of the covering of Ω by the Voronoi cells of P is called the *restriction* of $\text{Del}(P)$ to Ω and denoted by $\text{Del}_{|\Omega}(P)$. Equivalently, $\text{Del}_{|\Omega}(P)$ is the subcomplex of $\text{Del}(P)$ whose dual Voronoi faces intersect Ω .

We consider now the case where P is an $(\varepsilon, \bar{\eta})$ -net of Ω . We first observe that every simplex σ of $\text{Del}_{|\Omega}(P)$ must have a circumradius not greater than ε . Indeed, since $\sigma \in \text{Del}_{|\Omega}(P)$, $\text{Vor}(\sigma) \cap \Omega \neq \emptyset$, and since P is ε -dense, we have for any $c \in \text{Vor}(\sigma) \cap \Omega$ and any $q \in \sigma$

$$\|p - q\| \leq \|p - c\| + \|c - q\| \leq 2\varepsilon.$$

Observe also that all Delaunay simplices with a vertex at distance greater than 2ε from the boundary of Ω belong to $\text{Del}_{|\Omega}(P)$ (see Exercise 5.1).

Lemma 5.4 (Delaunay complex of a net) *Let Ω be a bounded subset of $\mathbb{R}^d$, P an $(\varepsilon, \bar{\eta})$ -net of Ω , and assume that d and $\bar{\eta}$ are positive constants. The restriction of the Delaunay complex of P to Ω has size $2^{O(d^2)}|P|$ (i.e. the size is linear for fixed d).*

Proof Let p be a point of P . We first bound the number of neighbors of p , i.e. the vertices of the link of p in $\text{Del}_{|\Omega}(P)$. Let σ be a simplex of $\text{Del}_{|\Omega}(P)$ in the star of p . As observed above, the diameter of σ is at most

2ε . Moreover, all the open balls $B(q, \frac{\bar{\eta}\varepsilon}{2})$, $q \in P$, are disjoint by definition of the separation. Hence, the number of neighbors of p is at most

$$n_p = \frac{\text{vol}_d(B(2\varepsilon + \frac{\bar{\eta}\varepsilon}{2}))}{\text{vol}_d(B(\frac{\bar{\eta}\varepsilon}{2}))} = \left(1 + \frac{4}{\bar{\eta}}\right)^d$$

where $B(r)$ denotes a ball of radius r .

Assuming without real loss of generality that $n_p \geq d + 1$, we deduce that the number of simplices in the star of p is at most

$$\sum_{i=1}^{d+1} \binom{n_p}{i} = 2^{O(d^2)}.$$

□

We will show now that the randomized incremental construction (see Theorems 4.4 and 3.11) constructs the Delaunay complex of any net P in time $O(n \log n)$ where $n = |P|$. Remarkably, we don't need to modify the algorithm whose behaviour will automatically adapt to the fact that the input point set is a net. To analyze the complexity of the algorithm in this context, it is sufficient to show that the expected complexity of the Delaunay complex of a random sample S of P has linear size $O(|S|)$. Indeed, the expected number of simplices that appear in the Delaunay complex of a point set $P \in \mathbb{R}^d$ during the randomized incremental construction, is

$$O\left(\sum_{i=1}^n \frac{d}{i} \Gamma_0(i, P)\right),$$

where $\Gamma_0(i, P)$ is the expected size of the Delaunay complex of a random sample of size i , drawn from the point set P . It will then immediately follow from the analysis of the complexity of the randomized algorithm that the complexity of the algorithm is $O(n \log n)$ when P is a net.

Lemma 5.5 *A random subsample S of an ε -net P has a Delaunay complex of expected size $2^{O(d^2)}|S|$ (i.e. the expected size is linear for fixed d).*

Proof We first give the proof in the case where S is a Bernoulli sample, i.e. it is obtained from P by picking every point $p \in P$ with probability ϖ . For convenience, we will write $\varpi = (\varepsilon/\delta)^d$.

Let us fix a point $p \in P$; we shall upper bound the size of $\text{star}(p, \text{Del}(S))$, i.e. the set of simplices of $\text{Del}(S)$ with vertex p . Let $I_k = [2^{k-1}\delta, 2^k\delta)$, $k \in \mathbb{N}$. Consider a d -tuple τ of points of P such that the d -simplex formed by the points in $\sigma = \tau \cup \{p\}$, has circumcentre c_σ and circumradius r_σ . Then, given that $p \in S$, the event $E^{(\tau)} = (\sigma \in \text{Del}(S))$ could occur only if the following two events occur:

- (i) $E_1^{(\sigma)} = (\forall p' \in \sigma, p' \in S)$
- (ii) $E_2^{(\sigma)} = (B(c_\sigma, r_\sigma) \cap S = \emptyset)$, where $B(x, r)$ is the open ball centered at x with radius r .

Given that $p \in S$, the probability that $\sigma \in \text{Del}(S)$ can be therefore upper-bounded as follows. We write $n_\sigma = P \cap B(p, r_\sigma) \geq (r_\sigma/\varepsilon)^d$ (by a packing argument).

$$\begin{aligned}
\text{proba}[E^{(\tau)}] &= \text{proba}[E_1^{(\sigma)} \wedge E_2^{(\sigma)}] \\
&= \varpi^d \text{proba}[E_2^{(\sigma)}] \\
&\leq \varpi^d (1 - \varpi)^{n_\sigma} \\
&\leq \varpi^d e^{-\varpi n_\sigma} \\
&\leq \left(\frac{\varepsilon}{\delta}\right)^{d^2} e^{-(\varepsilon/\delta)^d (r_\sigma/\varepsilon)^d} \\
&\leq \left(\frac{\varepsilon}{\delta}\right)^{d^2} e^{-(r_\sigma/\delta)^d}
\end{aligned} \tag{5.1}$$

If σ has circumradius $r_\sigma \in I_k$, we get

$$\text{proba}[E^{(\tau)}] \leq (\varepsilon/\delta)^{d^2} e^{-2^{(k-1)d}}$$

By the triangle inequality, if $\sigma \in \text{Del}(S)$ has a circumradius r_σ , then all the points in σ must lie in the ball $B(p, 2r_\sigma)$. Therefore, the number of potential d -tuples which can contribute to $\text{star}(p, \text{Del}(S))$ is at most $((4r_\sigma/\varepsilon)^d)^d = (4r_\sigma/\varepsilon)^{d^2} \leq (2^{k+2}\delta/\varepsilon)^{d^2}$. Let

$$Z_p(k) := \begin{cases} |\{\sigma \in \text{Del}(S) : p \in \sigma, r_\sigma \in [2^{k-1}\delta, 2^k\delta)\}|, & p \in S, \\ 0, & \text{otherwise,} \end{cases}$$

denote the number of Delaunay simplices which contain $p \in S$ and have

circumradius $r_\sigma \in I_k$. Then we get

$$\begin{aligned}\mathbb{E}[Z_p(k)] &\leq \sum_{\tau \in (P \cap B(p, 2r))^d} \text{proba}[E^{(\tau)}] \\ &\leq (2^{k+2}\delta/\varepsilon)^{d^2} (\varepsilon/\delta)^{d^2} e^{-2^{(k-1)d}} \\ &= 4^{d^2} 2^{kd^2} e^{-2^{(k-1)d}}\end{aligned}$$

Define $Y_p := \sum_{k=1}^{\infty} Z_p(k)$ to be the number of Delaunay simplices in $\text{star}(p, \text{Del}(S))$. Summing k over all allowed ranges of r , Summing over all allowed ranges of r , we get

$$\begin{aligned}\mathbb{E}[Y_p] &\leq \sum_{k=1}^{\infty} 4^{d^2} e^{-2^{(k-1)d}} \cdot 2^{kd^2} \\ &= 4^{d^2} \sum_{k'=1}^{\infty} e^{-k'} \cdot k'^d \cdot 2^{d^2} \quad (\text{with } k' = 2^{(k-1)d}) \\ &= 8^{d^2} \sum_{k'=1}^{\infty} e^{-k'} \cdot k'^d \\ &\leq 8^{d^2} \int_{x=0}^{\infty} e^{-x} \cdot x^d dx = 8^{d^2} \Gamma(d+1) \\ &= 2^{O(d^2)}.\end{aligned}$$

where $\Gamma(t) := \int_0^{\infty} e^{-x} x^{t-1} dx$ denotes the *Gamma* function. Therefore, the expected size of $\text{Del}(S)$ is given by

$$\mathbb{E}[|\text{Del}(S)|] \leq \sum_{p \in P} \text{proba}[p \in S] \mathbb{E}[Y_p] \leq n\varpi \cdot 2^{O(d^2)} = \mathbb{E}[|S|] \cdot 2^{O(d^2)}.$$

The previous result does not apply directly to the randomized construction of the Delaunay complex of P . When the points are inserted in a random order, the i -th point to be inserted is a random point in a uniform sample of size i of P . We thus consider now the case of a random subset $S \subseteq P$ of size s . Given that $p \in S$, choosing the rest of the random subsample S is equivalent to choosing a random sample of $s-1$ elements from $P \setminus \{p\}$. Now from Equation 5.1 and Exercise 5.2, we get:

$$\begin{aligned}\text{proba}[E^{(\tau)}] &\leq \left(\frac{s-1}{n-1}\right)^d \left(1 - \frac{s-1}{n-1}\right)^{(r_\sigma/\varepsilon)^d} \\ &\approx \left(\frac{s}{n}\right)^d \left(1 - \frac{s}{n}\right)^{(r_\sigma/\varepsilon)^d} \\ &= \varpi^d (1 - \varpi)^{(r_\sigma/\varepsilon)^d},\end{aligned}$$

where $\varpi = s/n$ is the probability for a point of P to belong to S . The rest of the computations follows as previously leading to the same asymptotic bound on $\mathbb{E}[|\text{Del}(S)|]$. $\square$

We conclude from the discussion above and Lemma 5.5:

Theorem 5.6 *Let P be an ε -net in $\mathbb{R}^d$. The randomized incremental construction of a Delaunay complex of P requires on expectation $2^{O(d^2)}|P|$ memory space and $2^{O(d^2)}|P| \log |P|$ time.*

5.1.3 Nets in discrete metric spaces

So far, we have considered nets in subsets of Euclidean spaces. However nets can be defined in more general metric spaces. In this section, we consider the case where Ω is a finite set of points that we rename W to emphasize the distinction. We do not assume that W is embedded in Euclidean space and the points are not given a location. Instead, we only assume to know the distance matrix of W , i.e. the $|W| \times |W|$ matrix M whose element $m_{i,j}$ is the distance between the points w_i and w_j . W together with M defines a *discrete metric space*. Nets can be defined in discrete metric spaces in very much the same way they have been defined in Euclidean space.

Extracting from W a coarser sample that is a net of W will allow to benefit from the nice properties of nets. It will also allow to represent data at various resolutions, to cluster data and to construct witness complexes, a weak variant of Delaunay complexes that can be defined and constructed in any discrete metric space (Section 6.2).

We first prove the existence of nets in the context of discrete metric spaces (see Lemma 5.2 for its analog in the Euclidean case).

Lemma 5.7 (Existence of nets) *Let W be a finite set of points such that the distance of any point $q \in W$ to $W \setminus \{q\}$ is at most ε and let $\lambda \geq \varepsilon$. One can extract from W a subsample L that is a $(\lambda, 1)$ -net of W .*

In preparation for Section 6.2, we will often refer to the points of L as landmarks and to the points of W as witnesses.

Proof We construct the sample L by inserting points of W one by one. Initially $L := \emptyset$ and $W' := W$. At each step, we pick a point of W' , say p ,

insert it in L , and remove p from W' as well as all the points of W' whose distance to p is less than λ . We stop when W' is empty. The algorithm necessarily terminates since W is finite. Upon termination, all the points of W are at distance at most λ from a point of L (since otherwise it would be inserted in L) and two points of L are at distance at least λ (since we never insert a point at distance less than λ from the current set L). It follows that L is a $(\lambda, 1)$ -net of W . $\square$

We can improve the algorithm described in the proof of Lemma 5.7 by inserting, at each step, a point $p \in W$ that is farthest from the current set of landmarks L (see Algorithm 3 (Case 1)).

We maintain for each point w of W' its closest point $L(w)$ in L and, at each step, we select the point $w \in W'$ which is most distant from its closest landmark $L(w)$. The algorithm requires $O(|W|)$ storage and its time complexity is easily seen to be $O(|L| \times |W|)$.

Algorithm 3 Farthest point insertion

Input: the distance matrix of a finite point set W and either a positive constant λ (Case 1) or an integer k (Case 2)

$L := \emptyset$

$W' := W$

$L(w) := p_\infty$ for all $w \in W'$ $\{p_\infty$ is a fake point at infinite distance from $W\}$

$\lambda^* := \max_{w \in W'} \|w - L(w)\|$

$w^* :=$ a point $p \in W'$ such that $\|p - L(p)\| = \lambda^*$

while either $\lambda^* > \lambda$ (Case 1) or $|L| < k$ (Case 2) **do**

 add w^* to L and remove w^* from W'

for each point w of W' such that $\|w - w^*\| < \|w - L(w)\|$ **do**

$L(w) := w^*$

 update λ^* and w^*

Output : $L \subseteq W$, a $(\lambda, 1)$ -net of W (Case 1), an approximate solution to the k -centers problem (Case 2)

The next lemma shows that this strategy implies that the separation ratio remains constant over the refinement.

For any $i > 0$, let L_i and λ_i denote respectively L and $\lambda^* = \max_{w \in W'} \|w - L(w)\|$ at the end of the i -th iteration of the main loop of the algorithm. If we label the points by their insertion order, we have $L_i = \{p_1, \dots, p_i\}$ and

$\lambda_i = d(p_i, L_{i-1})$. Since L_i grows with i , λ_i is a decreasing function of i . Moreover:

Lemma 5.8 *At each iteration $i > 0$, L_i is a $(\lambda_i, 1)$ -net of W .*

Proof The fact that L_i is λ_i -dense in W follows directly from the definition of λ_i . Let us show that L_i is λ_i separated. Consider a closest pair of points $p_a p_b$ of L_i and assume that p_b has been inserted after p_a . Then $\|p_a - p_b\| = \lambda_b \geq \lambda_i$. This implies that L_i is λ_i -separated. $\square$

In some applications, we are interested in selecting a given number k of points. We are then interested in the following *k-centers problem* : select from W a subset L of k points that are as far apart from each other as possible. More precisely, we want to choose a subset L of k points of W in such a way as to maximize the separation of L , i.e. the minimum pairwise distance between the points of L . We can use Algorithm 3 (Case 2) for that purpose. Interestingly, this simple greedy algorithm provides a 2-approximation to the problem.

Lemma 5.9 (Approximation of the k -centers problem) *Algorithm 3 (Case 2) provides an approximation ratio of 2 for the k -centers problem.*

Proof Let $L_k = l_1, \dots, l_k$ be the set of the k first selected points by the algorithm, labelled according to the order they have been selected. Let λ_k denote the distance of the k -th point from all previously selected points l_i , $i = 1, \dots, k-1$. Consider the set of balls $B(l_i, \lambda_k)$, $i = 1, \dots, k-1$. Their union contains all the points of W and therefore all the points of any optimal solution L_{opt} to the k -centers problem. Since $|L_{\text{opt}}| = k$ and there are $k-1$ balls, there must be two points of L_{opt} that fall in the same ball $B(l_i, \lambda_k)$ for some $i \leq k-1$. It follows that there exists a pair of points in L_{opt} that is at distance at most $2\lambda_k$ by the triangular inequality.

On the other hand, the distance between any two points of L_k is at least λ_k by Lemma 5.8. It follows that Algorithm 3 (Case 2) provides an approximation ratio of 2 for the k -centers problem. $\square$




Figure 5.1: The four vertices of a squared face f of a uniform grid are cocircular and can be circumscribed by a sphere centered at the center of the face. This sphere does not enclose any other vertex of the grid. We can then slightly perturb the vertices of f so that the convex hull of the vertices of f is now a tetrahedron of positive volume whose circumscribing sphere does not include any other vertex of the grid. Hence, it is a tetrahedron in the Delaunay triangulation of the (perturbed) vertices of the grid.

5.2 Thick simplices

For a given set of points $P \in \mathbb{R}^2$, $\text{Del}(P)$ maximizes, over all possible triangulations of P , the smallest angle of the triangles (Exercise 5.6) and it can be easily shown that, if P is a net, then the angles of the triangles are lower bounded by some positive constant (Exercise 5.5). However, the property does not hold for higher dimensional Delaunay triangulations and one cannot bound the dihedral angles of higher dimensional simplices as shown in Figure 5.1.

For any vertex p of a simplex σ , the *face opposite* p , denoted by σ_p , is the convex hull of the other vertices of σ . The *altitude* of p in σ is the distance $D(p, \sigma) = d(p, \text{aff}(\sigma_p))$ from p to the affine space $\text{aff}(\sigma_p)$ spanned by σ_p . The altitude $D(\sigma)$ of σ is the minimum over all vertices p of σ of $D(p, \sigma)$. A poorly-shaped simplex can be characterized by the existence of a relatively small altitude. The *thickness* of a j -simplex σ is the dimensionless quantity

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \frac{D(\sigma)}{j\Delta(\sigma)} & \text{otherwise,} \end{cases}$$

where $\Delta(\sigma)$ denotes the *diameter* of σ , i.e. the length of its longest edge.

5.2.1 Thickness and singular value

We will show in the next lemma that the thickness of a simplex is related to the singular values of a matrix. Before stating the lemma, we recall some well known results on matrices and their singular values and refer the reader to the book of Trefethen and Bau [131] for an excellent introduction to singular values. If A is a $d \times j$ matrix with $j \leq d$, we assume that the singular values of A are ordered according to decreasing absolute values and we denote the i^{th} singular value by $s_i(A)$. We have $s_1(A) = \|A\| = \sup_{\|x\|=1} \|Ax\|$ and $s_j(A) = \inf_{\|x\|=1} \|Ax\|$. We will employ the following standard observation:

Lemma 5.10 *If $\mu > 0$ is an upper bound on the norms of the columns of A , then $s_1(A) = \|A\| \leq \sqrt{j}\mu$.*

From the given definitions, one can verify that if A is an invertible $d \times d$ matrix, then $s_1(A^{-1}) = s_d(A)^{-1}$, but it is convenient to also accommodate non-square matrices, corresponding to simplices that are not full dimensional. If A is a $d \times j$ matrix of rank $j \leq d$, then the *pseudo-inverse* $A^\dagger = (A^\top A)^{-1} A^\top$ is the unique left inverse of A whose kernel is the orthogonal complement of the column space of A . We have the following general observation [131]:

Lemma 5.11 *Let A be a $d \times j$ matrix of rank $j \leq d$ and let $A^\dagger$ be its pseudo inverse $= (A^\top A)^{-1} A^\top$. We have*

$$s_i(A^\dagger) = s_{j-i+1}(A)^{-1}, \text{ for } i = 1, \dots, j.$$

In particular, $s_j(A) = s_1(A^\dagger)^{-1}$.

The columns of A form a basis for the column space of A . The pseudo-inverse can also be described in terms of the *dual basis*. If we denote the columns of A by $\{a_i\}$, then the i^{th} dual vector, w_i , is the unique vector in




Figure 5.2: Choosing p_0 as the origin, the edges emanating from p_0 in $\sigma = [p_0, \dots, p_j]$ form a basis for $\text{aff}(\sigma)$. The proof of Lemma 5.12 demonstrates that the dual basis $\{w_i\}$ consists of vectors that are orthogonal to the facets, and with magnitude equal to the inverse of the corresponding altitude.

the column space of A such that $w_i^\top a_i = 1$ and $w_i^\top a_j = 0$ if $i \neq j$. Then $A^\dagger$ is the $j \times d$ matrix whose i^{th} row is $w_i^\top$.

By exploiting a close connection between the altitudes of a simplex and the vectors dual to a basis defined by the simplex, we obtain the following key lemma that relates the thickness of a simplex to the smallest singular value of an associated matrix:

Lemma 5.12 (Thickness and singular value) *Let $\sigma = \text{conv}(p_0, \dots, p_j)$ be a non-degenerate j -simplex in $\mathbb{R}^d$, with $j > 0$, and let P be the $d \times j$ matrix whose i^{th} column is $p_i - p_0$. Then*

$$s_j(P) \geq D(\sigma)/\sqrt{j} = \sqrt{j} \Theta(\sigma) \Delta(\sigma).$$

Proof We first show that the i^{th} row of $P^\dagger$ is given by $w_i^\top$, where w_i is orthogonal to $\text{aff}(\sigma_{p_i})$, and

$$\|w_i\| = D(p_i, \sigma)^{-1}.$$

Indeed, by the definition of $P^\dagger$, it follows that w_i belongs to the column space of P , and it is orthogonal to all $(p_{i'} - p_0)$ for $i' \neq i$. Let $u_i = w_i/\|w_i\|$. By the definition of w_i , we have $w_i^\top (p_i - p_0) = 1 = \|w_i\| u_i^\top (p_i - p_0)$. By the

definition of the altitude of a vertex, we have $u_i^\top(p_i - p_0) = D(p_i, \sigma)$. Thus $\|w_i\| = D(p_i, \sigma)^{-1}$. Since

$$\max_{1 \leq i \leq j} D(p_i, \sigma)^{-1} = \left(\min_{1 \leq i \leq j} D(p_i, \sigma) \right)^{-1} = (j\Theta(\sigma)\Delta(\sigma))^{-1},$$

Lemma 5.10, yields

$$s_1(\mathbf{P}^\dagger) \leq (\sqrt{j}\Theta(\sigma)\Delta(\sigma))^{-1}.$$

The stated bound on $s_j(\mathbf{P})$ follows from Lemma 5.11. $\square$

The proof of Lemma 5.12 shows that the pseudoinverse of $\mathbf{P}$ has a natural geometric interpretation in terms of the altitudes of σ , and thus the altitudes provide a convenient lower bound on $s_j(\mathbf{P})$. By Lemma 5.10, $s_1(\mathbf{P}) \leq \sqrt{j}\Delta(\sigma)$, and thus

$\Theta(\sigma) \leq \frac{s_j(\mathbf{P})}{s_1(\mathbf{P})}$. In other words, $\Theta(\sigma)^{-1}$ provides a convenient upper bound on the *condition number* of $\mathbf{P}$. Roughly speaking, thickness imparts a kind of stability on the geometric properties of a simplex. This is exactly what is required when we want to show that a small change in a simplex will not yield a large change in some geometric quantity of interest.

5.2.2 Whitney's angle bound

The following lemma is due to Whitney. It shows that, if the vertices of a simplex σ are at small relative distance from an affine space H , and if the thickness of the simplex is bounded away from 0, then the angle between the affine hull of σ and H is small. Before stating the lemma, we define angles between vector spaces.

Definition 5.13 (Angles between subspaces) *If U and V are vector subspaces of $\mathbb{R}^d$, with $\dim U \leq \dim V$, the angle between them is defined by*

$$\sin \angle(U, V) = \sup_{\substack{u \in U \\ \|u\|=1}} \|u - \pi_V u\|, \quad (5.2)$$

where π_V is the orthogonal projection onto V .

Alternatively, the angle between vector subspaces U and V can be defined as:

$$\angle(U, V) = \max_{u \in U} \{ \min_{v \in V} \angle(u, v) \} \quad (5.3)$$

The angle between affine subspaces K and H is defined as the angle between the corresponding parallel vector subspaces.

Lemma 5.14 (Whitney's angle bound) Suppose σ is a j -simplex of $\mathbb{R}^d$, $j < d$, whose vertices all lie within a distance δ from a k -dimensional affine space $h \subset \mathbb{R}^d$ with $k \geq j$. Then

$$\sin \angle(\text{aff}(\sigma), h) \leq \frac{2j\delta}{D(\sigma)} = \frac{2\delta}{\Theta(\sigma)\Delta(\sigma)}.$$

Proof Suppose $\sigma = \text{conv}(p_0, \dots, p_j)$. Choose p_0 as the origin of $\mathbb{R}^d$ and let $\Pi_h : \mathbb{R}^d \rightarrow h$ be the orthogonal projection onto h . Let u be any unit vector in $\text{aff}(\sigma)$. Since the vectors $v_i = (p_i - p_0)$, $i \in \{1, \dots, j\}$ form a basis for $\text{aff}(\sigma)$, we may write $u = Pa$, where P is the $d \times j$ matrix whose i^{th} column is v_i , and $a \in \mathbb{R}^j$ is the vector of coefficients. Then, defining $X = P - \Pi_h P$, we get

$$\|u - \Pi_h u\| = \|Xa\| \leq \|X\| \|a\|.$$

Since $d(p_i, h) \leq \delta$ for all $0 \leq i \leq j$, $\|v_i - \Pi_h v_i\| \leq \delta$. It follows then from Lemma 5.10 that

$$\|X\| \leq 2\sqrt{j}\delta.$$

Observing that $1 = \|u\| = \|Pa\| \geq \|a\| \inf_{\|x\|=1} \|Px\| = \|a\| s_j(P)$, we find

$$\|a\| \leq \frac{1}{s_j(P)},$$

and the result follows from Lemma 5.12. $\square$

Whitney's angle bound will be especially useful in Chapter 8. There, h will be the tangent space T_p at a point p of a smooth manifold $\mathcal{M}$, and σ will be a thick simplex whose vertices are close (relatively to the diameter of the simplex) to T_p . Whitney's lemma asserts that the affine hull of σ makes a small angle with T_p . Thickness plays a crucial role here as the following example shows. The Schwarz lantern is a polyhedral surface inscribed in a cylinder as shown in Figure 5.3. By increasing the number of vertices of the lantern, we can make the Hausdorff distance between the lantern and the cylinder arbitrarily small but increasing the sampling density does not guarantee that the planes of the facets of the lantern provide a good approximation of the tangent planes of the cylinder. In fact, the angle between the normal to a facet and the normal to the cylinder at any of the




Figure 5.3: The Schwarz lantern.

vertices of the facet can be made arbitrarily close to $\pi/2$. Such a situation cannot happen if the facets have a bounded thickness as stated by Whitney's lemma.

5.3 Thick triangulations via weighting

The notion of thickness introduced in Section 5.2 is an important measure of the shape of a simplex. A simplicial complex is thick when all its simplices are thick. We have seen (Figure 5.1) that Delaunay triangulations are not necessarily thick even if the vertices form an $(\varepsilon, \bar{\eta})$ -net. The goal of this section is to show that a thick simplicial complex can nevertheless be obtained from an $(\varepsilon, \bar{\eta})$ -net by assigning (relatively small) weights to the points of the net and considering the weighted Delaunay triangulation of the resulting set of weighted points.

To keep the exposition simple, we depart from the rest of this book and will work in this section in the flat torus $\mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$ instead of $\mathbb{R}^d$. As a consequence, if P is a finite set of points in $\mathbb{T}^d$, the Delaunay complex and weighted Delaunay complexes of P have no boundary. Boundary issues obscure the central properties we want to develop. It is not difficult to extend the results to the case of a bounded domain of $\mathbb{R}^d$, provided that we only look sufficiently far away from the boundary of the domain.

5.3.1 Weighting schemes

We consider a set of points P that is an $(\varepsilon, \bar{\eta})$ -net in $\mathbb{T}^d$. A *weighting scheme* on P is a function w from P to $\mathbb{R}$ which assigns to each point $p \in P$ a weight $w(p) \in \mathbb{R}$. We denote by $\hat{P}$ the resulting set of weighted points, i. e. $\hat{P} = \{(p, w(p)) : p \in P\}$, and by $\text{Del}(\hat{P})$ the corresponding weighted Delaunay triangulation.

We restrict the weighting scheme to non negative weights. The relative amplitude $\tilde{w}$ of the weighting scheme w is defined as

$$\tilde{w} = \max_{p \in P} \frac{w(p)}{L^2(p)},$$

where $L(p)$ is the distance from p to its nearest neighbor in $P \setminus p$. We also wish that each point of P appears as a vertex of $\text{Del}(\hat{P})$. As shown in Section 4.4.3, this condition is ensured if, for each $p \in P$, $w(p) < \frac{L^2(p)}{4}$. From now on, we will only consider weighting schemes that have a relative amplitude smaller than $\frac{1}{4}$, that is $\tilde{w} \leq \tilde{w}_0$ where $\tilde{w}_0$ is a constant smaller than $\frac{1}{4}$.

The main result of this section is that, for any $(\varepsilon, \bar{\eta})$ -net P of $\mathbb{T}^d$, given a small enough constant Θ_0 , there are weighting schemes with relative amplitude smaller than $\tilde{w}_0 < \frac{1}{4}$ and such that the weighted Delaunay triangulation $\text{Del}(\hat{P})$ has no j -simplex with thickness less than Θ_0^j , for $j = 1, \dots, d$. Moreover, as we will see in Section 5.3, such weighting schemes can be computed by a simple randomized algorithm.

In the rest of this section, we use the same notation σ for a geometric simplex with vertices in P and for its abstract counterpart which is just the subset of P formed by the vertices of σ . Given a weighting scheme w defined on P , each simplex $\sigma \subset P$ corresponds to a subset $\hat{\sigma}$ of $\hat{P}$: $\hat{\sigma} = \{(p, w(p)) : p \in \sigma\}$. Two weighted points are said to be *orthogonal* if their weighted distance is zero.

A weighted point orthogonal to all the weighted points of $\hat{\sigma}$ is said to be orthogonal to $\hat{\sigma}$. If σ is a simplex of dimension j , the weighted points orthogonal to $\hat{\sigma}$ are centered on an affine subspace of dimension $d - j$ orthogonal to the affine subspace $\text{aff}(\sigma)$. The intersection point of these two subspaces is denoted by $c(\hat{\sigma})$. Let $(c(\hat{\sigma}), R^2(\hat{\sigma}))$ be the weighted point centered on $c(\hat{\sigma})$ and orthogonal to $\hat{\sigma}$. The weighted point $(c(\hat{\sigma}), R^2(\hat{\sigma}))$ is the weighted point with minimal weight among all the weighted points orthogonal to $\hat{\sigma}$. The

point $c(\hat{\sigma})$ is called the *weighted center* of the simplex σ and $R(\hat{\sigma})$ is called the *weighted radius* of σ .

Note that weighted center $c(\hat{\sigma})$ and the weighted radius $R(\hat{\sigma})$ depend on the weights assigned to the vertices of σ and are different from the center $c(\sigma)$ and circumradius $R(\sigma)$ of σ . Still, the two following lemmas provides upper and lower bounds for the weighted radii. Note the the first one applies to any simplex with vertices in P when the relative amplitude of the weighting scheme is small while the second one applies only to simplices in the weighted Delaunay triangulation but assumes only that the weighting scheme is positive.

We will also make use of the following lemma that bounds the weighted radii and the diameters of the simplexes in $\text{Del}(\hat{P})$.

Lemma 5.15 (Weighted radii and diameters of simplices in $\text{Del}(\hat{P})$)

Let P be an $(\varepsilon, \bar{\eta})$ -net of $\mathbb{T}^d$ and assume that a non negative weighting scheme with relative amplitude smaller than $\tilde{w}_0 < \frac{1}{4}$ has been defined on P . Then, any simplex σ of $\text{Del}(\hat{P})$ has a weighted radius $R(\hat{\sigma})$ that is at most ε and a diameter $\Delta(\tau)$ that is at most $2\sqrt{2}\varepsilon$.

Proof Let σ be a simplex of $\text{Del}(\hat{P})$. The simplex σ being included in some d -simplex τ of $\text{Del}(\hat{P})$, the weighted radius $R(\hat{\sigma})$ of σ is at most the weighted radius $R(\hat{\tau})$ of τ . We now prove that $R(\hat{\tau})$ is at most ε . Indeed otherwise, since P is ε -dense, the ball $B(c(\hat{\tau}), R(\hat{\tau}))$ associated to the weighted center $(c(\hat{\tau}), R^2(\hat{\tau}))$ of τ would include a point q of P . Then the weighted distance from q to $(c(\hat{\tau}), R^2(\hat{\tau}))$ is negative, and this contradicts the fact that σ belongs to $\text{Del}(\hat{P})$.

Let us consider now the Euclidean distance $d(c(\hat{\sigma}), p)$ from the weighted circumcenter $c(\hat{\sigma})$ to any vertex p of σ . We have:

$$\begin{aligned} d(c(\hat{\sigma}), p)^2 &= R(\hat{\sigma})^2 + w(p) \\ &\leq \varepsilon^2 + \tilde{w}_0 L(p)^2. \end{aligned}$$

where $L(p)$ is the distance from p to its nearest neighbor in $P \setminus \{p\}$. From Lemma 5.1, we have $L(p) \leq 2\varepsilon$ and thus $d(c(\hat{\sigma}), p) \leq \varepsilon\sqrt{1 + 4\tilde{w}_0} \leq \sqrt{2}\varepsilon$. Since this bound holds for any vertex p of σ , we conclude that the diameter $\Delta(\sigma)$ of σ is at most $2\sqrt{2}\varepsilon$. $\square$

5.3.2 Θ_0 -thickness and flakes

A simplex that is not thick has a relatively small altitude. We focus here on a special class of non-thick simplices, called *flakes* in which *all* the altitudes are relatively small. Let Θ_0 be a constant smaller than 1.

Definition 5.16 (Θ_0 -flakes) *A j -simplex σ is Θ_0 -thick if $\Theta(\sigma) \geq \Theta_0^j$. A Θ_0 -flake is a simplex that is not Θ_0 -thick but whose proper faces are all Θ_0 -thick.*

Observe that a flake must have dimension at least 2, since $\Theta(\sigma) = 1$ for any simplex σ with dimension $j < 2$.

A simplicial complex whose simplices are all Θ_0 -thick is said to be Θ_0 -thick. Observe that a simplicial complex is Θ_0 -thick iff it includes no Θ_0 -flake. Indeed, if σ is not Θ_0 -thick, then either it is a Θ_0 -flake or it has a proper j -face $\sigma_j \subset \sigma$ that is not Θ_0 -thick. By considering such a face with minimal dimension, we arrive at the following observation:

Lemma 5.17 *A simplex is not Θ_0 -thick if and only if it has a face that is a Θ_0 -flake.*

We show in Lemma 5.19 below an upper bound on the altitudes of a Θ_0 -flake. First, we provide a general relationship between the altitudes of a simplex:

Lemma 5.18 *Let σ be a j -simplex with $j \geq 2$. If p and q are two vertices of σ , we note σ_p the subface $\sigma \setminus \{p\}$ of σ and σ_q the subface $\sigma \setminus \{q\}$. The altitudes $D(p, \sigma)$ and $D(p, \sigma_q)$ of p within respectively σ and σ_q and the altitudes $D(q, \sigma)$ and $D(q, \sigma_p)$ of q within respectively σ and σ_p satisfy the following relation*

$$\frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}.$$

Proof The proof follows from a volume computation. Let $\text{vol}_j(\sigma)$, $\text{vol}_{j-1}(\sigma_p)$ and $\text{vol}_{j-2}(\sigma_{pq})$ be the volumes of the simplices σ , σ_p and $\sigma_{pq} = \sigma \setminus \{p, q\}$ respectively. We have

$$\text{vol}_j(\sigma) = \frac{1}{j} D(p, \sigma) \text{vol}_{j-1}(\sigma_p) = \frac{1}{j(j-1)} D(p, \sigma) D(q, \sigma_p) \text{vol}_{j-2}(\sigma_{pq}).$$

The similar relation, obtained replacing p by q , obviously holds and both relations together prove the lemma. $\square$

We arrive at the following important observation about flake simplices:

Lemma 5.19 (Flakes have small altitudes) *If a k -simplex σ is a Θ_0 -flake, then for every vertex $p \in \sigma$, the altitude $D(p, \sigma)$ satisfies the bound*

$$D(p, \sigma) < \frac{k}{k-1} \frac{\Delta^2(\sigma)\Theta_0}{L(\sigma)} < \frac{\Delta^2(\sigma)\Theta_0}{L(\sigma)},$$

where $\Delta(\sigma)$ and $L(\sigma)$ are the lengths of the longest and shortest edges of σ .

Proof Recalling Lemma 5.18, we have

$$D(p, \sigma) = \frac{D(q, \sigma)D(p, \sigma_q)}{D(q, \sigma_p)}, \quad (5.4)$$

and taking q to be a vertex with minimal altitude in σ , we have

$$D(q, \sigma) = k\Theta(\sigma)\Delta(\sigma) < k\Theta_0^k\Delta(\sigma). \quad (5.5)$$

Moreover, since σ_p is Θ_0 -thick, we have:

$$D(q, \sigma_p) \geq (k-1)\Theta(\sigma_p)\Delta(\sigma_p) \geq (k-1)\Theta_0^{k-1}L(\sigma). \quad (5.6)$$

Furthermore:

$$D(p, \sigma_q) \leq \Delta(\sigma_q) \leq \Delta(\sigma). \quad (5.7)$$

Plugging Equations 5.5, 5.6 and 5.7 into Equation 5.4 yields the claimed bound. $\square$

5.3.3 The weight range of a flake with small radius

Let σ be a Θ_0 flake of dimension j . If we assign weights to the vertices of σ , the weighted radius $R(\hat{\sigma})$ of σ depends on the weights of its vertices. We show here that to keep the weighted radius $R(\hat{\sigma})$ smaller than a given ε we have to choose the weight of each vertex of σ within a small interval whose measure is linear in Θ_0 .

Lemma 5.20 *Let σ be a Θ_0 -flake and assume that we are given a weighting scheme with non negative weights on the vertices of σ . If the weighted radius $R(\hat{\sigma})$ of σ is smaller than ε , the weight $w(p)$ of any vertex p of σ belongs to an interval $I(\sigma, p)$ whose measure $|I(\sigma, p)|$ satisfies:*

$$|I(\sigma, p)| \leq 4 \frac{\Delta(\sigma)^2}{L(\sigma)} \Theta_0 \varepsilon.$$

If furthermore the vertices of σ belongs to an $(\varepsilon, \bar{\eta})$ -net, and the weighting scheme has a relative amplitude smaller than $\tilde{w}_0 < \frac{1}{4}$, the measure of $I(\sigma, p)$ satisfies:

$$|I(\sigma, p)| \leq |I| \stackrel{\text{def}}{=} 32 \Theta_0 \frac{\varepsilon^2}{\bar{\eta}}.$$

Proof Let $c(\hat{\sigma})$ and $R(\hat{\sigma})$ be respectively the weighted center and weighted radius of σ . Likewise, we use $c(\hat{\sigma}_p)$ and $R(\hat{\sigma}_p)$ for respectively the weighted center and the weighted radius of σ_p , where σ_p is the subface $\sigma \setminus p$ of σ . Referring to Figure 5.4, we get :

$$R^2(\hat{\sigma}) = d^2(c(\hat{\sigma}), c(\hat{\sigma}_p)) + R^2(\hat{\sigma}_p), \quad (5.8)$$

The set of points of $\mathbb{R}^d$ with equal weighted distances to the vertices of σ_p is an affine subspace we denote by $N(\hat{\sigma}_p)$. Writing $d(p, N(\hat{\sigma}_p))$ for the distance from p to $N(\hat{\sigma}_p)$, we have:

$$R^2(\hat{\sigma}) + w(p) = d^2(p, c(\hat{\sigma})) = d^2(p, N(\hat{\sigma}_p)) + (D(p, \sigma) - H(p, \hat{\sigma}))^2 \quad (5.9)$$

where $D(p, \sigma)$ is the altitude of p in σ and $H(p, \hat{\sigma}) = d(c(\hat{\sigma}), c(\hat{\sigma}_p))$ if p and $c(\hat{\sigma})$ are on the same side of $\text{aff}(\sigma_p)$ and $H(p, \hat{\sigma}) = -d(c(\hat{\sigma}), c(\hat{\sigma}_p))$ otherwise.

Using Equations 5.8 and 5.9 together, we get:

$$\begin{aligned} w(p) &= d^2(p, N(\hat{\sigma}_p)) + (D(p, \sigma) - H(p, \hat{\sigma}))^2 - R^2(\hat{\sigma}) \\ &= d^2(p, N(\hat{\sigma}_p)) + D^2(p, \sigma) - R^2(\hat{\sigma}_p) - 2D(p, \sigma)H(p, \hat{\sigma}). \end{aligned}$$

Writing $F(p, \hat{\sigma}) = d^2(p, N(\hat{\sigma}_p)) + D^2(p, \sigma) - R^2(\hat{\sigma}_p)$, we get:

$$w(p) = F(p, \hat{\sigma}) - 2D(p, \sigma)H(p, \hat{\sigma}).$$

Hence,

$$|w(p) - F(p, \hat{\sigma})| = 2D(p, \sigma) d(c(\hat{\sigma}), c(\hat{\sigma}_p)). \quad (5.10)$$




Figure 5.4: For the proof of Lemma 5.20.

Observe that $F(p, \hat{\sigma})$ depends on the locations of the vertices of σ and on the weights of the vertices of σ_p but does not depend on the weight $w(p)$ of p .

From Lemma 5.19, $D(p, \sigma) \leq \Theta_0 \frac{\Delta(\sigma)^2}{L(\sigma)}$, and, from Equation 5.8, we get that $d(c(\hat{\sigma}), c(\hat{\sigma}_p))$ is at most $R(\hat{\sigma})$. Therefore, if $R(\hat{\sigma}) \leq \varepsilon$, the weight $w(p)$ of p belongs to the interval $I(\sigma, p)$, centered at $F(p, \hat{\sigma})$, of measure

$$|I(\sigma, p)| \leq 4 \frac{\Delta(\sigma)^2}{L(\sigma)} \Theta_0 \varepsilon. \quad (5.11)$$

In the case where σ belongs to an $(\varepsilon, \bar{\eta})$ -net, we have $L(\sigma) \geq \bar{\eta} \varepsilon$ and, from Lemma 5.15, $\Delta(\sigma) \leq 2\sqrt{2}\varepsilon$. The weight range $|I(\sigma, p)|$ then satisfies

$$|I(\sigma, p)| \leq 32 \Theta_0 \frac{\varepsilon^2}{\bar{\eta}}.$$

□

5.3.4 Lovász local lemma

We will see in Section 5.3 how to compute, for a given input set of points P , a weighting scheme yielding to a thick weighted Delaunay triangulation. This construction relies on the constructive proof of the Lovász local lemma due to Moser and Tardos. In this section, we recall without proof these important results.

Lovász local lemma [1] is a powerful tool to prove the existence of combinatorial objects. Let $\mathcal{A}$ be a finite collection of “bad” events in some probability space. The lemma shows that the probability that none of these events occur is positive provided that the individual events occur with a sufficiently small probability and there is limited dependence among them. Here is the lemma in a simple form that will be sufficient for our purposes.

Theorem 5.21 (Lovász local lemma) *Let $\mathcal{A} = \{A_1, \dots, A_N\}$ be a finite set of events in some probability space. Suppose that each event A_i is independent of all but at most Γ of the other events and that $\Pr[A_i] \leq \varpi$ for all $1 \leq i \leq N$. If*

$$\varpi \leq \frac{1}{e(\Gamma + 1)} \quad (5.12)$$

where e denotes the base of the natural logarithm, then the probability that none of the events in $\mathcal{A}$ occurs is strictly positive.

Assume that the events are determined by a set of independent random variables. Each event is determined by a subset of those random variables. Two events are independent if the subsets of random variables determining each of them do not overlap. In such a case, Moser and Tardos gave a constructive proof of the Lovász lemma [109, 129]. The proof leads to a simple and natural algorithm that repeatedly checks whether some event in $\mathcal{A}$ occurs. In the affirmative, the algorithm picks an arbitrary occurring event, say A , and resamples A , where we call resampling of an event A the operation that consists in choosing new random values for the variables determining A .

Moser and Tardos proved that this simple algorithm quickly terminates, providing an assignment of the random variables that avoids all of the events in $\mathcal{A}$. It is important to note that the selection mechanism for picking an occurring event in the while loop is arbitrary and does not affect the correctness of the algorithm.

Algorithm 4 Moser Tardos algorithm

Input: a finite set $\mathcal{A}$ of events determined by a finite set $\mathcal{P}$ of independent random variables
for all $P \in \mathcal{P}$ **do**
 $v_P \leftarrow$ a random evaluation of $\mathcal{P}$
 while some event of $\mathcal{A}$ occurs **do**
 pick any such event $A \in \mathcal{A}$
 resample A
return $\{v_P, P \in \mathcal{P}\}$

Theorem 5.22 (Moser Tardos) *Assume that the same conditions as in Theorem 5.21 hold. Assume, in addition, that the events are determined by a set $\mathcal{P}$ of n independent random variables. Moser Tardos randomized algorithm computes an assignment of values for the variables in $\mathcal{P}$ such that no event in $\mathcal{A}$ occurs. The algorithm resamples each event at most $\frac{1}{\varpi}$ expected times and the expected total number of resampling steps is at most $O(n)$.*

5.3.5 Applying the Lovász local lemma to remove flakes

Let P be an $(\varepsilon, \bar{\eta})$ -net of $\mathbb{R}^d$ and $\tilde{w}_0$ a constant less than $1/4$. Now we take for the weights of the points in P independent random variables and we pick the weight of each point p uniformly at random in the interval $[0, \tilde{w}_0 L^2(p)]$. An event occurs when there exists a Θ_0 -flake $\sigma \subset P$ that has a weighted radius $R(\hat{\sigma})$ not greater than ε . Since we know from Lemma 5.15, that all simplices in $\text{Del}(\hat{P})$ have a weighted radius $R(\hat{\sigma})$ not greater than ε , removing all events will lead to a Θ_0 -thick complex. For convenience, we will often identify an event and the associated flake.

Lemma 5.23 *The probability that an event occurs is $\varpi \leq 32 \frac{\Theta_0}{\tilde{w}_0 \bar{\eta}^3}$.*

Proof Let σ be a Θ_0 -flake included in P and assume that the weights of all the vertices of σ except one, say p , have already been assigned. We know from Lemma 5.20 that if the weighted radius $R(\hat{\sigma})$ is not greater than ε , the weight $w(p)$ belongs to an interval of measure less than $|I| = 32 \Theta_0 \frac{\varepsilon^2}{\bar{\eta}}$. It

follows that the probability that an event occurs is at most

$$\frac{|I|}{\tilde{w}_0 L^2(p)} \leq 32 \frac{\Theta_0}{\tilde{w}_0 \bar{\eta}^3},$$

which implies the same bound on the probability of the event σ . $\square$

Lemma 5.24 *Each event overlaps at most Γ other events, where*

$$\Gamma + 1 \leq \left(\frac{13}{\bar{\eta}} \right)^{d(d+1)}.$$

Proof Let σ be a Θ_0 -flake with weighted center $c(\hat{\sigma})$ and weighted radius $R(\hat{\sigma}) \leq \varepsilon$. An event that overlaps σ is a Θ_0 -flake σ' with a weighted radius at most ε that shares a vertex p with σ . Let q be a vertex of σ' . We have:

$$\begin{aligned} d(c(\hat{\sigma}), q) &\leq d(c(\hat{\sigma}), p) + d(p, q) \\ &\leq \sqrt{R(\hat{\sigma})^2 + w(p)} + \Delta(\sigma') \end{aligned}$$

Since $R(\hat{\sigma}) \leq \varepsilon$, $w(p) \leq \tilde{w}_0 L(p)^2 \leq 4\tilde{w}_0 \varepsilon^2$ by Lemma 5.1, and $\Delta(\sigma) \leq 2\sqrt{2}\varepsilon$ by Lemma 5.15, we get:

$$d(c(\hat{\sigma}), q) \leq 3\sqrt{2}\varepsilon.$$

Therefore, any vertex of an event that overlap σ lies in the ball $B(c(\hat{\sigma}), r)$ of radius $r = \bar{r}\varepsilon = 3\sqrt{2}\varepsilon$. Since P is η -separated, this ball contains at most $J(r)$ points of P , with

$$\begin{aligned} J(r) &= \frac{(r + \frac{\eta}{2})^d}{(\frac{\eta}{2})^d} = \left(1 + \frac{2\bar{r}}{\bar{\eta}} \right)^d \leq \left(1 + \frac{6\sqrt{2}}{\bar{\eta}} \right)^d \\ &\leq \left(\frac{11}{\bar{\eta}} \right)^d. \end{aligned}$$

where the last inequality uses the fact that $\bar{\eta} \leq 2$.

Since the Θ_0 -flake σ also has its vertices within the ball $B(c(\hat{\sigma}), r)$, we can bound $\Gamma + 1$ by the number of simplices with vertices within $B(c(\hat{\sigma}), r)$, i.e.

$$\Gamma + 1 \leq \sum_{i=2}^{d+1} \binom{J(r)}{i} \leq (1 + J(r))^{d+1} \leq \left(1 + \left(\frac{11}{\bar{\eta}} \right)^d \right)^{d+1}.$$

Using $\bar{\eta} \leq 2$ and the fact that, for all $a \in \mathbb{R}^+$ and $b \in \mathbb{R}^+$, $a^d + b^d \leq (a+b)^d$, we get:

$$\Gamma + 1 \leq \left(\frac{13}{\bar{\eta}} \right)^{d(d+1)}.$$

□

We are now ready to apply the Lovász local lemma. We will assume that

$$\left(\frac{13}{\bar{\eta}} \right)^{d(d+1)} 32 \frac{\Theta_0}{\tilde{w}_0 \bar{\eta}^3} \leq \frac{1}{e}. \quad (5.13)$$

In view of Lemmas 5.23 and 5.24, this condition ensures that we have $\varpi \leq \frac{1}{e(\Gamma+1)}$, and therefore there exists a weight assignment on P with relative amplitude less than $\tilde{w}_0$, such that P includes no Θ_0 -flakes with a weighted radius less than ε . Since any simplex in the weighted Delaunay triangulation $\text{Del}(\hat{P})$ has a weighted radius less than ε (Lemma 5.15), we conclude that such an assignment yields a weighted Delaunay triangulation $\text{Del}(\hat{P})$ that has no Θ_0 -flakes.

Theorem 5.25 *Let P be an $(\varepsilon, \bar{\eta})$ -net of $\mathbb{T}^d$ and $\tilde{w}_0 \leq \frac{1}{4}$. If Equation 5.13 is satisfied, there is a weight assignment on P with relative amplitude less than $\tilde{w}_0$ such that the weighted Delaunay triangulation $\text{Del}(\hat{P})$ is Θ_0 -thick.*

Observe that the thickness Θ_0 that we can guarantee is bounded by Equation 5.13 that constrains Θ_0 to be small enough with respect to $\bar{\eta}$ and $\tilde{w}_0$. Note that the bound is very small and depends as 2^{-d^2} on d . Still it does not depend on the sampling density ε .

5.3.6 Algorithm

Based on the results of the previous subsections, we will apply Moser Tardos algorithm to our context so as to obtain a weighted Delaunay triangulation that is Θ_0 -thick. Algorithm 5 takes as input an $(\varepsilon, \bar{\eta})$ -net P of $\mathbb{T}^d$, a constant $\tilde{w}_0 < 1/4$ and a constant Θ_0 small enough to satisfy Equation 5.13. As proved below, the algorithm outputs a weighting scheme on P whose relative amplitude is smaller than $\tilde{w}_0$ and such that the weighted Delaunay triangulation $\text{Del}(\hat{P})$ is Θ_0 -thick. The algorithm maintains the weighted Delaunay triangulation $\text{Del}(\hat{P})$ while resampling the Θ_0 -flakes that occur in $\text{Del}(\hat{P})$

until they all disappear. Resampling a simplex σ consists in reassigned random weights to the vertices of σ . As already mentionned, the weights of the different vertices are picked independently and the weight $w(p)$ of vertex p is taken uniformly at random in the interval $[0, \tilde{w}_0 L^2(p)]$.

Algorithm 5 Thick weighted Delaunay triangulation

Input: $P, \tilde{w}_0, \Theta_0$

Initialize all weights to 0 and compute $\text{Del}(\hat{P}) = \text{Del}(P)$

while there are Θ_0 -flakes in $\text{Del}(\hat{P})$ **do**

 choose a Θ_0 -flake σ in $\text{Del}(\hat{P})$

 resample σ

 update $\text{Del}(\hat{P})$

Output: A weighting scheme on P and the corresponding weighted Delaunay triangulation which is granted to be Θ_0 -thick.

Theorem 5.26 *If P is an $(\varepsilon, \bar{\eta})$ -net of $\mathbb{T}^d$, $\tilde{w}_0$ a constant less than $1/4$ and Θ_0 a constant such that Equation 5.13 holds, Algorithm 5 outputs a weighting scheme $\hat{P}$ on P whose relative amplitude is smaller than $\tilde{w}_0$ and such that the weighted Delaunay triangulation $\text{Del}(\hat{P})$ is Θ_0 -thick. Its expected complexity is linear with respect to the size of P .*

Proof Algorithm 5 resamples the Θ_0 -flakes that occur in the weighted Delaunay triangulation $\text{Del}(\hat{P})$. Since the Θ_0 -flakes in $\text{Del}(\hat{P})$ have weighted radii at most ε , they are events as defined in Section 5.3.5. Therefore Algorithm 5 is a variant of Moser Tardos algorithm applied to the Θ_0 -flakes with small weighted radii. The main difference is that Algorithm 5 keeps only track of Θ_0 -flakes in the current $\text{Del}(\hat{P})$ and not of all possible Θ_0 -flakes with small weighted radii included in P . According to Theorem 5.22, Condition 5.13 ensures that Moser Tardos algorithm terminates whatever may be the order in which the events are resampled. Therefore the condition, a fortiori, guarantees the termination of Algorithm 5.

Each resampling involves the reweighting of at most $d + 1$ vertices. Since P is an $(\varepsilon, \bar{\eta})$ -net, and the weighting scheme has bounded relative amplitude, the weighted Delaunay complex $\text{Del}(\hat{P})$ can be updated in constant time. Thus the expected complexity of Algorithm 5 is proportional to the number of resampling which is $O(|P|)$ by Theorem 5.22. $\square$

5.4 Protection

We introduce now the notion of protection of a simplex. The notion of protection is stronger than the notion of thickness (see Lemma 5.27) and some positive protection can be obtained by perturbing the position of the points of P . This is another mean to ensure Delaunay triangulations to have positive thickness (other than weight assignment as discussed in the previous section). Protection will also be used in Section 6.2.

We say that a simplex $\sigma \subset \text{Del}(P)$ is δ -protected if there exists a point c_σ , called a δ -protection center of σ such that

$$\|c_\sigma - q\| > \|c_\sigma - p\| + \delta \quad \forall p \in \sigma \text{ and } \forall q \in P \setminus \sigma. \quad (5.14)$$

We will write $\bar{\delta} = \frac{\delta}{\varepsilon}$. Note that Equation 5.14 implies the following inequality

$$\|c_\sigma - q\|^2 > \|c_\sigma - p\|^2 + \delta^2 \quad \forall p \in \sigma \text{ and } \forall q \in P \setminus \sigma. \quad (5.15)$$

Lemma 5.27 (Separation and thickness from protection) *Let P be an $(\varepsilon, \bar{\eta})$ -net of a bounded domain $\Omega \subset \mathbb{R}^d$. Assume that $|P| > d + 1$ and that every d -simplex in $\text{Del}(P)$ is δ -protected. Then the separation ratio of P satisfies*

$$\bar{\eta} \geq \bar{\delta},$$

and the thickness of any simplex σ (of any dimension) of $\text{Del}(P)$ is at least

$$\Theta(\sigma) \geq \frac{\bar{\delta}^2}{8d}.$$

Before we prove the lemma, we state and prove two easy claims that will be useful to prove the second part of the lemma.

Claim 5.28 *Let $B = B(c, R)$ and $B' = B(c', R')$ be two n -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let H be the bisecting hyperplane of B and B' , i.e. the hyperplane that contains the $(n - 2)$ -sphere $S = \partial B \cap \partial B'$. Let θ be the angle of the cone (c, S) . Writing $\rho = \frac{R'}{R}$ and $\|c - c'\| = \lambda R$, we have*

$$\cos(\theta) = \frac{1 + \lambda^2 - \rho^2}{2\lambda}. \quad (5.16)$$

If $R \geq R'$, we have $\cos(\theta) \geq \frac{\lambda}{2}$.

Proof Let $q \in S$; applying the cosine rule to the triangle $[cc'q]$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos(\theta) = R'^2, \quad (5.17)$$

which proves Equation 5.16. If $R \geq R'$, then $\rho \leq 1$, and $\cos(\theta) \geq \lambda/2$ immediately follows from Equation 5.16. $\square$




Figure 5.5: Construction used in Claims 5.28 and 5.29.

If $B = B(c, R)$ is a d -ball, we denote by $B^{+\delta}$ the ball $B(c, \sqrt{R^2 + \delta^2})$.

Claim 5.29 Let $B = B(c, R)$ and $B' = B(c', R')$ be two n -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let $\tilde{\theta}$ be the angle of the cone $(c, \tilde{S})$ where $\tilde{S} = \partial B \cap \partial B'^{+\delta}$. Writing $\|c - c'\| = \lambda R$, we have

$$\cos(\tilde{\theta}) = \cos(\theta) - \frac{\delta^2}{2R^2\lambda}$$

Proof Let $\tilde{q} \in \tilde{S}$, applying the cosine rule to the triangle $[cc'\tilde{q}]$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos(\tilde{\theta}) = R'^2 + \delta^2.$$

Subtracting Equation 5.17 from the previous equality yields $\delta^2 = 2\lambda R^2(\cos(\theta) - \cos(\tilde{\theta}))$, which proves the lemma. $\square$

Proof of Lemma 5.27 1. Let p and q be two closest points of P . The edge $[pq]$ is an edge of $\text{Del}(P)$. We denote by σ a d -simplex of $\text{Del}(P)$ that

contains $[pq]$ and by σ' a d -simplex of $\text{Del}(P)$ that contains one of the two vertices of $[pq]$ but not the other. Since σ' is δ -protected, we must have $|e| \geq \delta$.

2. Since $\Theta(\tau) \geq \Theta(\sigma)$ for any simplex $\tau \subseteq \sigma$, it is sufficient to consider the case of a d -simplex σ . Since $|P| > d + 1$, there exists at least one other d -simplex σ' of $\text{Del}(P)$ that share a facet τ with σ . Let $B(\sigma) = B(c, R)$ and $B(\sigma') = B(c', R')$ be the circumscribing balls of σ and σ' respectively. The spheres ∂B and $\partial B'^{+\delta}$ intersect in a $(n - 2)$ -sphere $\tilde{S}$ which is contained in a hyperplane $\tilde{H}$ parallel to the hyperplane $H = \text{aff}(\tau)$. For any $\tilde{q} \in \tilde{S}$ we have

$$d(\tilde{H}, H) = d(\tilde{q}, H) = R(\cos(\theta) - \cos(\tilde{\theta})) = \frac{\delta^2}{2\|c - c'\|},$$

where the last equality follows from Claim 5.29 and $d(\tilde{H}, H)$ denotes the distance between the two parallel hyperplanes. See Figure 5.5 for an illustration. Since $p \in \partial B$, p belongs to $B(\sigma')^{+\delta}$ if and only if p lies in the strip bounded by H and $\tilde{H}$, which is equivalent to

$$d(p, H) = D(p, \sigma) < \frac{\delta^2}{2\|c - c'\|}.$$

We conclude that, if σ is δ -protected, the ball

$$B(\sigma)^{+\delta} = B(c_\sigma, \sqrt{R_\sigma^2 + \delta^2}) \subseteq B(c_\sigma, R_\sigma + \delta)$$

does not contain points of $P \setminus \sigma$, which implies from the inequality above that

$$d(p, H) = D(p, \sigma) \geq \frac{\delta^2}{2\|c - c'\|}.$$

For any simplex σ , we have $D(p, \sigma) \leq 2R(\sigma)$ for all $p \in \sigma$, where $R(\sigma)$ denotes the radius of the circumsphere of σ . For any σ in the triangulation of an ε -net, we have $R(\sigma) \leq \varepsilon$. Thus $D(p, \sigma) \leq 2\varepsilon$, and the inequality above yields $\|c - c'\| \geq \frac{\delta^2}{4\varepsilon}$.

We further obtain

$$\Theta(\sigma) = \frac{\min_{p \in \sigma} D(p, \sigma)}{d \Delta(\sigma)} \geq \frac{\delta^2}{8d\varepsilon^2} = \frac{\bar{\delta}^2}{8d}.$$

□

In Section 5.3.5, we have used a variant of Moser Tardos algorithm to obtain a thick *weighted* Delaunay complex. Instead of weighting the points, we

can perturb their positions so that the Delaunay triangulation of the perturbed point set P' is δ -protected for some small enough $\delta > 0$. Thanks to Lemma 5.27, the Delaunay triangulation of the perturbed points will have some positive thickness (see Exercise 5.12).

5.5 Exercises

Exercise 5.1 Let Ω be a subset of $\mathbb{R}^d$ and P a finite ε -dense set of points in Ω . Show that all Delaunay simplices of P with a vertex at distance greater than 2ε from the boundary of Ω belong to $\text{Del}_\Omega(P)$.

Exercise 5.2 (Hypergeometric distribution) Suppose a, b, c are positive integers, with $\max\{a, b\} \leq c$. The *hypergeometric distribution* with parameters a, b, c is the distribution of the random variable $X = |A \cap B|$, where A is a random sample of size a , from a universe C of size c , which has a subset B of size b . Prove that given $b \leq a$, the probability that the random sample A contains B , and is disjoint from another fixed set T with cardinality t , $T \cap B = \emptyset$, is at most $\left(\frac{a}{c}\right)^b \left(1 - \frac{a}{c}\right)^t$, for $a, c \gg b^2$.

Exercise 5.3 (Delaunay refinement) Adapt Algorithm 3 to add sample points in $\mathbb{T}^d$ so that the sample P is an ε -net of $\mathbb{T}^d$.

Exercise 5.4 (k -center clustering) Given is a set of points P and an integer k . The k -clustering problem consists in partitioning P into k clusters so as to minimize the maximum diameter of a cluster. Propose an algorithm that gives an approximation ratio of 2 for the k -clustering problem.

Exercise 5.5 Let Ω be an open subset of $\mathbb{R}^2$ and P an $(\varepsilon, \bar{\eta})$ -net of Ω in general position wrt circles. Show that all triangles of $\text{Del}_\Omega(P)$ have all their angles greater or equal to $\arcsin(\frac{\bar{\eta}}{2})$.

Exercise 5.6 (Max-min) Let P be a finite set of points in general position in the plane. To any triangulation T of P we attach the vector $V(T) = (\alpha_1, \dots, \alpha_{3t})$ where the $\alpha_i \in [0, \pi]$ are the angles of the t triangles of T , sorted by increasing values. Show that $\text{Del}(P)$ is, among all triangulations of P , the one that maximizes $V(T)$ for the lexicographic order. In particular, $\text{Del}(P)$ maximizes the smallest angle.

Exercise 5.7 Let U and V be two vector spaces of $\mathbb{R}^d$ with $\dim U \leq \dim V$. Show that $\angle(U, V) = \angle(V^\perp, U^\perp)$ where $U^\perp$ (resp., $V^\perp$) denotes the vector space normal to U (resp., V).

Exercise 5.8 Let h and h' be two affine spaces of the same dimension embedded in $\mathbb{R}^d$, and let u be a vector of $\mathbb{R}^d$. Show that $\angle(u, h') \leq \angle(u, h) + \angle(h, h')$.

Exercise 5.9 Let τ be a j -simplex. Show that $L(\tau) \geq D(p, \tau) \geq j! \Theta(\tau) \Delta(\tau)$ (see Lemma 5.19 for the notations).

Exercise 5.10 Bound from below the dihedral angles of a d -simplex of $\mathbb{R}^d$ as a function of its thickness.

Exercise 5.11 (Inheritance of protection) Let P be an $(\varepsilon, \bar{\eta})$ -net of a bounded domain $\Omega \subset \mathbb{R}^d$. We say that a simplex is δ -power protected if Equation 5.15 is satisfied, i.e. there exists a point c_σ such that

$$\|c_\sigma - q\|^2 > \|c_\sigma - p\|^2 + \delta^2 \quad \forall p \in \sigma \text{ and } \forall q \in L \setminus \sigma.$$

Show that if every d -simplex in $\text{Del}(P)$ is δ -power protected, then all simplices (of all dimensions) in $\text{Del}(P)$ are at least δ' -power protected where $\delta' = \frac{\delta}{d}$. deduce a similar result if one replaces the power protection by the protection as defined in Section 5.4. (Hint : use the lifting map introduced in the proofs of Theorems 4.3 and 4.6)

Exercise 5.12 (Protection via perturbation) Let P be an $(\varepsilon, \bar{\eta})$ -net P of the flat torus $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$. Propose an algorithm that perturbs the points of P so that the Delaunay triangulation of the perturbed point set P' is δ -protected for some a small enough $\delta > 0$ (and thus has some positive thickness by Lemma 5.27). If p is a point of P , the associated perturbed point is picked at random in the ball $B(p, \rho)$ for some $\rho > 0$. Use a variant of Moser Tardos algorithm .

5.6 Bibliographical notes

The farthest point insertion algorithm (Algorithm 3) has been popularized by Gonzales in the context of clustering data sets [92] and has found numerous applications in many fields. Gonzales proved Lemma 5.9 and Feder and

Greene [81] showed that no polynomial-time algorithm exists with a constant approximation ratio close to 2 unless $P = NP$. They also improved the $O(n^2)$ time complexity of Algorithm 3 to $O(n \log k)$ when the points live in Euclidean space. See also the paper by Har-Peled and Mendel [96] that contains further improvements and extensions, and, in particular, shows how to construct a hierarchical representation of a point set, called the *net-tree*, from the sequence of points provided by the farthest point insertion. The proof of the complexity of the incremental randomized construction of nets (Theorem 5.6) is due to Boissonnat, Devillers,

Dutta and Glisse [16]. Their paper contains also solutions to Exercises 5.2 and ??.

The notion of thick triangulations goes back to the early work on differential topology by Cairns [32], Whitehead [134], Whitney [135], Munkres [112] and others. Thick triangulations also play a central role in the work of Cheeger et al. [54] and Fu [85] on curvature measures. Since this notion appeared in different places and contexts, various names have been used, e.g. thickness, fullness or relative thickness. Our presentation follows the work of Boissonnat, Dyer and Ghosh [20]. Lemma 5.14 is due to Whitney [135].

More recently, thick triangulations have been found important in mesh generation where numerical simulations require meshes to be thick [74]. The notion of flake simplex introduced in this chapter is an extension of the notion of sliver introduced by Cheng et al. in the context of 3-dimensional mesh generation [55, 74]: a sliver is a flake with an upper bound on the ratio of its circumradius to the length of its shortest edge. Sliver removal in higher dimensions has been discussed in [100, 56].

Our weighting mechanism to remove flakes is inspired from the one used by Cheng et al. to remove slivers from 3-dimensional Delaunay triangulations. The weighting mechanism can be seen as a perturbation of the Euclidean metric. It is also possible to remove flakes and inconsistencies by perturbing the position of the points. This kind of perturbation may be preferred to the weighting mechanism in the context of mesh generation [101, 14, 77].

The notion of protection has been introduced by Boissonnat, Dyer and Ghosh to study the stability of Delaunay triangulations and the construction of Delaunay triangulations of manifolds [20, 19, 18]. A solution to Exercise 5.12 can be deduced from results in [19] and [21].

The Lovász local lemma, proved initially by Lovász and Erdős [1], is a cel-

celebrated result with a long history. The constructive proof of Moser and Tardos has been a breakthrough which is still the subject of intense research [109]. Its first introduction in Computational Geometry appears in [21].

Chapter 6

Filtrations

In this chapter, we introduce two simplicial complexes which have strong ties with Delaunay complexes. Their common point is that they allow to define filtrations. Filtrations, defined in Chapter 2, are sequences of nested complexes that allow to represent a set of points at various scales. They play an important role in persistent homology, a central tool in Topological Data Analysis, as will be demonstrated in Chapter 11.

We first define the alpha-complex, or α -complex, of a finite set of points P , which is a subcomplex of the Delaunay complex. Here α is a real parameter and varying α will lead to a filtration of the Delaunay complex. The definition extends to sets of weighted points : the weighted α -complex of a set of weighted points $\hat{P}$ is a subcomplex of the weighted Delaunay triangulation of $\hat{P}$, and varying α will lead to a filtration of the weighted Delaunay complex.

The α -complex of P is a combinatorial representation of the union of the balls $B(p, \alpha)$ of radius α centered at the points of P . It is thus closely related to the Čech complex introduced in Section 2.3. In fact, the α -complex shares with the Čech complex the property of having the same homotopy type as the union of the balls. A major difference between the two complexes is related to their size: the α -complex is usually much smaller than the Čech complex and can be computed more efficiently.

Owing to their capacity of representing union of balls and their topology, α -complexes play an important role in the description of proteins and macro molecules, and in drug design. In Geometric Inference, α -complexes also play an important role and we will see in Section 8.1 that they capture the homotopy type of well sampled manifolds.

Alpha shapes are constructed from the Delaunay complex and are therefore difficult to compute in high dimensions. In Section 6.2, we will introduce another complex, the so-called *witness complex*. The witness complex is defined from two point sets : L , which is the vertex set of the complex, and W that can be seen as an approximation of the space that contains L . In applications, L is usually a crude subset of W that can be extracted from W using, for example, the algorithms of Section 5.1. The witness complex can be defined and constructed in any discrete metric space and does not require the points to be embedded in a specific ambient metric space: we only need to know the pairwise distances between the points of L and W . This is a critical advantage in high dimensions over the Delaunay complex whose construction requires to evaluate the `in_ball` predicate whose

algebraic complexity depends on the dimension d of the ambient space (Exercise 4.5). In the case where the points live in Euclidean space, we will see that, under appropriate conditions, witness complexes and Delaunay complexes are identical. We will also introduce a variant called the relaxed witness which offers another filtration.

6.1 Alpha complexes

6.1.1 Definitions

Alpha complexes. Let P be a set of points in $\mathbb{R}^d$. From Lemma 4.1, we know that the simplices in the Delaunay complex $\text{Del}(P)$ are characterized by the *empty ball property* meaning that a simplex with vertices in P belong to $\text{Del}(P)$ iff it admits an *empty* circumscribing ball i.e. a circumscribing ball whose interior includes no point of P . We are interested here in sorting the simplices of $\text{Del}(P)$ according to the squared radius of their smallest empty circumscribing ball. For a simplex τ , we call this quantity the filtration value of τ and write it $\alpha(\tau)$.

Now, for any $\alpha \in \mathbb{R}$, we consider the subset $\mathcal{A}(P, \alpha)$ of the simplices in $\text{Del}(P)$ that have a filtration value at most α . Because a ball circumscribing a simplex circumscribes any face of this simplex, $\mathcal{A}(P, \alpha)$ is a subcomplex of $\text{Del}(P)$. It is called the α -complex of P .

Since the α -complex of P is a subcomplex of the Delaunay complex of P , it has a natural embedding if P is in general position wrt spheres (Theorem 4.3). The underlying space of the α -complex $\mathcal{A}(P, \alpha)$ is called the α -shape of P .

If P is in general position wrt spheres, the dimension of the α -complex $\mathcal{A}(P, \alpha)$ is at most the dimension d of the embedding space. It should be noted that the dimension of the α -complex may be strictly less than d and that the complex may not be pure, having some simplices which are not faces of simplices of maximal dimension.

The Delaunay filtration. The α -complex evolves when α increases, from the empty set for $\alpha < 0$, to the set of vertices of $\text{Del}(P)$ when $\alpha = 0$, and finally to the whole Delaunay triangulation $\text{Del}(P)$ when α is large enough.




Figure 6.1: The α -complex of a set of points in $\mathbb{R}^2$ for increasing value of the parameter α .

The parameter α defines an order on the simplices. This order is not total and some simplices may have the same value of the parameter. We define a total order by first sorting the simplices with a same α by increasing dimensions and then breaking ties arbitrarily. The total order we then obtain yields a filtration of the Delaunay complex.

We recall that a filtration of a simplicial complex K , as defined in Section 2.4, is a sequence of nested subcomplexes of K such that each subcomplex is obtained from the previous one by adding a simplex of K . The filtration of $\text{Del}(P)$ obtained by inserting the simplices in the order defined above is called a *Delaunay filtration* of P .

6.1.2 Computing alpha complexes and filtrations

Let P be a set of points in $\mathbb{R}^d$ and let $\text{Del}(P)$ be the corresponding Delaunay complex. For a d -simplex τ in $\mathbb{R}^d$, there is a unique ball of $\mathbb{R}^d$ circumscribing τ . We denote by $c(\tau)$ and $r(\tau)$ respectively the center and the radius of this ball. According to the definition, a d -simplex τ of $\text{Del}(P)$ belongs to the α -complex $\mathcal{A}(P, \alpha)$ iff $\alpha \geq r(\tau)^2$.

Things are a bit more complicated for simplices of dimension strictly less than d . Let τ be a simplex of dimension $k < d$. The centers of the d -balls circumscribing τ form a $(d - k)$ -flat of $\mathbb{R}^d$ we denote by $h(\tau)$. Consider the circumscribing ball of τ that has the smallest radius. The center of this smallest circumscribing ball, $c(\tau)$, is called the center of τ and its radius, $r(\tau)$, is called the circumradius of τ . The center $c(\tau)$ is the point where $h(\tau)$ intersects the k -flat $\text{aff}(\tau)$ spanned by τ .

Given a set of points P , we say that τ is a *Gabriel* simplex if its smallest circumball contains no point of P in its interior. Plainly, a Gabriel simplex is a simplex of $\text{Del}(P)$ but the converse is not necessarily true: a k -simplex τ of $\text{Del}(P)$ is a Gabriel simplex iff $c(\tau)$ belongs to the face $V(\tau)$ of the Voronoi diagram $\text{Vor}(P)$ dual to τ .

We conclude from the discussion that, if τ is a *Gabriel* simplex, its smallest enclosing ball is empty and it belongs to all α -complexes $\mathcal{A}(P, \alpha)$ for $\alpha \geq r(\tau)^2$. If τ a Delaunay simplex that is not Gabriel, the smallest empty ball that circumscribes τ circumscribes a coface of τ and therefore τ will have the same filtration value α as one of its cofaces. To prove this, consider the function $w(x)$ that associates to each point x in $h(\tau)$ its squared distance to the vertices in τ . Finding the filtration value $\alpha(\tau)$ of τ amounts to minimizing $w(x)$ under the condition that $x \in V(\tau)$. Since $w(x)$ is a convex function and $V(\tau)$ is a convex polyhedron, the minimum of $w(x)$ on $V(\tau)$ is reached at $c(\tau)$ iff $c(\tau) \in V(\tau)$ or on the boundary of $V(\tau)$ otherwise. In the first case, the simplex is a Gabriel simplex. In the last case, if we call σ the coface of higher dimension of τ whose dual Voronoi face contains $c(\tau)$ in its interior, we have $\alpha(\tau) = \alpha(\sigma)$.

For any $\tau \in \text{Del}(P)$ with dimension $k < d$, we denote by $U(\tau)$ the set of cofaces of τ in $\text{Del}(P)$ with dimension $k + 1$. Algorithm 6 computes for each simplex τ in $\text{Del}(P)$, the critical value $\alpha(\tau)$ at which τ enters the α -complex $\mathcal{A}(P, \alpha)$.

Algorithm 6 Computing Delaunay filtrations

Input: the set of points P in $\mathbb{R}^d$
 Compute the Delaunay complex $\text{Del}(P)$
for each d -simplex $\tau \in \text{Del}(P)$ **do**
 set $\alpha(\tau) = r(\tau)^2$ (the squared circumradius of τ)
for $k = d - 1, \dots, 0$ **do**
 for each d -simplex $\tau \in \text{Del}(P)$ **do**
 if τ is a Gabriel simplex **then**
 $\alpha(\tau) = r(\tau)^2$ (the squared smallest circumradius of τ)
 else
 $\alpha(\tau) = \min_{\sigma \in U(\tau)} \alpha(\sigma)$
Output: The critical α -value of each simplex in $\text{Del}(P)$ has been computed

6.1.3 Weighted alpha complex

The definition of α -complexes and α -shapes extend to the weighted case. Let $\hat{P}$ be a set of weighted points and let τ be a simplex with vertex set $\hat{P}_\tau \subset \hat{P}$. Let us recall that two weighted points are said to be orthogonal when their weighted distance is zero. A weighted point is said to be orthogonal to $\hat{P}_\tau$ when it is orthogonal to all weighted points in $\hat{P}_\tau$ and it is said to be free of any weighted point in $\hat{P}$ when it has a positive or null distance to any weighted point of $\hat{P}$. From Lemma 4.5, a simplex τ with vertex set $\hat{P}_\tau \subset \hat{P}$ belongs to the weighted Delaunay triangulation $\text{Del}(\hat{P})$ iff there is a weighted point orthogonal to $\hat{P}_\tau$ and free of weighted points in $\hat{P}$. For any value of $\alpha \in \mathbb{R}$, we consider the subset $\mathcal{A}(\hat{P}, \alpha)$ of simplices in $\text{Del}(\hat{P})$ for which there is a weighted point with weight at most α , orthogonal to $\hat{P}_\tau$ and free of weighted points in $\hat{P}$. The simplices in $\mathcal{A}(\hat{P}, \alpha)$ form a subcomplex of $\text{Del}(\hat{P})$ which is called the *weighted α -complex* of $\hat{P}$. Under the usual general position assumption, this complex naturally embeds in $\mathbb{R}^d$ and the underlying space of $\mathcal{A}(\hat{P}, \alpha)$ is called the *weighted α -shape* of the set $\hat{P}$ for the parameter value α . Notice that α -complexes and α -shapes are special cases of respectively weighted α -complexes and weighted α -shapes, obtained when all the weights of the considered weighted points are equal.

As before, we associate to each simplex τ of $\text{Del}(\hat{P})$ a filtration value $\alpha(\tau)$ that corresponds to the first time τ enters the filtration.

Algorithm 6 extends almost verbatim to the case of weighted points, pro-

vided that we replace circumballs by orthogonal weighted points and empty balls by free weighted points. Let $\hat{P}$ be a set of weighted points in $\mathbb{R}^d$, τ a simplex of $\text{Del}(\hat{P})$ and $\hat{P}_\tau$ the subset of $\hat{P}$ associated to the vertices of τ . We now say that τ is a Gabriel simplex iff the ball with smallest radius that is orthogonal to $\hat{P}_\tau$ is free of weighted points in $\hat{P}$. Likewise, the function $w(x)$ used to prove the correctness of Algorithm 6 is now the weighted distance to the vertices of τ , i. e. $w(x) = D(x, \hat{p}) = d(x, p)^2 - w_p$, if $\hat{p} = (p, w_p)$ is the weighted point associated to the vertex p of τ .

6.1.4 Application to union of balls

Lemma 6.1 *Let B be a finite set of balls in $\mathbb{R}^d$. The union $U(B)$ of balls in B is homotopy equivalent to the α -complex $\mathcal{A}(B, 0)$.*

See Figure 6.2 for an illustration of this fact.

Proof Each ball b in B may be regarded as a weighted point $(c(b), r^2(b))$ where $c(b)$ and $r^2(b)$ are respectively the center and the squared radius of b . We denote by $\text{Del}(B)$, $\text{Vor}(B)$ and $\mathcal{A}(B, \alpha)$ respectively the Delaunay complex, Voronoi diagram and α -shape of the balls in B . Let $V(b)$ be the cell of b in $\text{Vor}(B)$. We claim that $\{b \cap V(b), b \in B\}$ forms a finite convex cover of $U(B)$ as defined in Section 2.3. First, because b and $V(b)$ are both convex subsets of $\mathbb{R}^d$, each subset $b \cap V(b)$ is convex. We show next that $\{b \cap V(b), b \in B\}$ is a cover of $U(B)$, i.e.:

$$U(B) = \bigcup_{b \in B} b \cap V(b).$$

The inclusion $\bigcup_{b \in B} b \cap V(b) \subset U(B)$ is trivial. To show the reverse inclusion, let us consider a point p in $U(B)$. Point p belongs to at least one ball b_1 of B and let $b(p)$ be the ball in B whose Voronoi cell contains p . Since $p \in b_1$, the weighted distance $D(p, b_1)$ is negative and, since $b(p)$ minimizes the weighted distance to p , we have:

$$D(p, b(p)) \leq D(p, b_1) \leq 0,$$

which means that p belongs to $b(p)$ and therefore to $b(p) \cap V(b(p))$ and to $\bigcup_{b \in B} b \cap V(b)$. The claim is proved.

It follows from the claim and the Nerve theorem (Theorem 2.8) that the union $U(B)$ of balls in B is homotopy equivalent to the nerve of the cover

$\bigcup_{b \in B} b \cap V(b)$. We now show that the nerve of this cover is just the α -complex $\mathcal{A}(B, 0)$. Let $B' \subset B$ be a subset of B . The subset B' belongs to the nerve of the cover $\{b \cap V(b), b \in B\}$ iff the intersection $\bigcap_{b \in B'} b \cap V(b)$ is non empty. This in turn is equivalent to say that there exists a point x in $\bigcap_{b \in B'} b \cap V(b)$. Such a point x is at an equal negative weighted distance $w(x)$ to the balls in B' and at a greater weighted distance to any ball of $B \setminus B'$. In other words, the weighted point $\hat{x} = (x, w(x))$ is orthogonal to any ball in B' and has a positive weighted distance to any ball in $B \setminus B'$. Therefore B' belongs to $\text{Del}(B)$ and, since $w(x) \leq 0$, to the α -complex $\mathcal{A}(B, 0)$. $\square$

Let us consider the special case where all the balls in B have the same radius. Let $B(P, r)$ be the set of balls with radius r , centered at points of the set P . The union $U(B(P, r))$ of these balls is homotopy equivalent to the Čech complex $\check{\text{Cech}}(P, r)$ defined (in Section 2.5) as the nerve of the cover of $U(B(P, r))$ by balls in $B(P, r)$. Therefore, the Čech complex $\check{\text{Cech}}(P, r)$ and the α -complex $\mathcal{A}(B(P, r), 0)$ have the same homotopy type and capture both the homotopy type of the union of balls $U(B(P, r))$. However, the Čech complex $\check{\text{Cech}}(P, r)$ is often much bigger than the α -complex $\mathcal{A}(B(P, r), 0)$. In particular, the dimension of the Čech complex $\check{\text{Cech}}(P, r)$ may be larger than d , in fact it may be as large as the number of balls in $B(P, r)$. Accordingly the Čech complex $\check{\text{Cech}}(P, r)$ usually does not embed naturally in $\mathbb{R}^d$. Differently, the α -complex $\mathcal{A}(B(P, r), 0)$, being a subcomplex of $\text{Del}(B)$, embeds naturally in $\mathbb{R}^d$ under general position assumption.

6.2 Witness complexes

In this section, we introduce the witness complex, a variant of the Delaunay complex that can be defined using only distances (and not empty spheres). Hence the witness complex can be defined in any finite metric space where the input consists of the distance matrix of the data points that stores the pairwise distances (Section 5.1.3). Not every finite metric space can be isometrically embedded in a Euclidean space but if it is the case, we provide conditions under which the witness and the Delaunay complexes are identical. A practical situation, encountered for example in the context of sensor networks, is when the points come from some Euclidean space but their actual locations are not known.

The witness complex is defined from two sets of points L and W . The first




Figure 6.2: Homotopy between union of balls and α -shapes. The union $U(B)$ of balls in the set B is colored red, the weighted Voronoi diagram $\text{Vor}(B)$ is drawn in black, and the α -complex $\mathcal{A}(B, 0)$ is shown in blue.

one, called the set of *landmarks*, is finite. The other one, called the set of *witnesses*, serves as an approximation of the ambient space. A typical situation is when L is a subset of W , possibly a net of W extracted using one of the algorithms of Section 5.1.3. The witness complex $\text{Wit}(L, W)$ can be seen as a weak notion of Delaunay triangulation which is easy to compute, even in high dimensions, since it only involves comparisons of distances between input points.

Definition 6.2 (Witness of a simplex) *Let σ be a simplex with vertices in $L \subset \mathbb{R}^d$, and let w be a point of $W \subseteq \mathbb{R}^d$. We say that w is a witness of σ if*

$$\|w - p\| \leq \|w - q\| \quad \forall p \in \sigma \text{ and } \forall q \in L \setminus \sigma.$$

Definition 6.3 (Witness complex) *The witness complex $\text{Wit}(L, W)$ is the complex consisting of all simplexes σ such that any simplex $\tau \subseteq \sigma$ has a witness in W . In other words, $\text{Wit}(L, W)$ is the maximal simplicial complex with the property that all its simplices have a witness in W .*




Figure 6.3: A simplex and one of its witnesses.

In this section, we use the Euclidean distance to define witness complexes but the definition is general and extend to more general metric spaces and, in particular, to finite metric spaces where the only information we have about the input points is the distances between any two of them. In Euclidean space, the only predicates involved in the construction of $\text{Wit}(L, W)$ are (squared) distance comparisons, i.e. polynomials of degree 2 in the coordinates of the points. This is to be compared with the predicate that decides whether a point lies inside the ball circumscribing a d -simplex, whose degree depends on d (see Exercise 4.5).

6.2.1 Identity of witness and Delaunay complexes

When the points W and L live in Euclidean space, the witness complex can be seen as a weak Delaunay complex. The results below make this connection more precise. We first make the following easy observation.

Lemma 6.4 *If $W' \subseteq W$, then $\text{Wit}(L, W') \subseteq \text{Wit}(L, W)$.*

Let Ω be a subset of $\mathbb{R}^d$. As before (see Section 5.1.2), we write $\text{Del}_{|\Omega}(L)$ for the restriction of $\text{Del}(L)$ to Ω , i.e. the subcomplex of $\text{Del}(L)$ whose simplices have a circumcenter in Ω .

Lemma 6.5 *Let Ω be a subset of $\mathbb{R}^d$. $\text{Del}_{|\Omega}(L) \subseteq \text{Wit}(L, \Omega)$.*

Proof By definition, any simplex σ of $\text{Del}_\Omega(L)$ has an empty circumscribing ball whose center c belongs to Ω . This center is a witness of σ , and it is also a witness for all the faces of σ . $\square$

The following remarkable result provides a weak characterization of Delaunay complexes. It shows that Delaunay and witness complexes are identical when the set of witnesses cover the whole space $\mathbb{R}^d$.

Theorem 6.6 (Weak characterization) *For any convex $\Omega \subseteq \mathbb{R}^d$ and any finite point set $L \subset \Omega$, we have $\text{Wit}(L, \Omega) = \text{Del}_\Omega(L)$.*

Proof We have already proved that $\text{Del}_\Omega(L) \subseteq \text{Wit}(L, \Omega)$ (Lemma 6.5). We prove now the converse inclusion by an induction on the dimension k of the simplices. The claim holds for $k = 0$ since any vertex of $\text{Wit}(L, \Omega)$ is a point of L and thus a vertex of $\text{Del}_\Omega(L)$.




Figure 6.4: Proof of Lemma 6.6.

Assume now that any simplex of $\text{Wit}(L, \Omega)$ of dimension up to $k - 1$ is a simplex of $\text{Del}_\Omega(L)$ and let $\tau = [p_0, \dots, p_k]$ be a k -simplex of $\text{Wit}(L, \Omega)$ witnessed by a point w . We will say for convenience that a ball B witnesses τ if $B \cap L = \tau$. We denote by B_τ be the smallest ball centered at w that witnesses τ and by S_τ the sphere bounding B_τ .

If all the vertices of τ belong to S_τ , τ is a Delaunay simplex and we are done. Otherwise, we will show that one can find a new ball that witnesses τ such that its bounding sphere contains one more vertex of τ than S_τ . Refer




Figure 6.5: Two triangles that have a witness but not their common edge, even if $W = \mathbb{R}^d$. These two triangles are not Delaunay triangles.

to Figure 6.4. Write $\sigma = S_\tau \cap \tau$. By the induction hypothesis, σ is a simplex of $\text{Del}_\Omega(L)$ and therefore there exists an empty ball B_σ centered in Ω that circumscribes σ . Write c for its center. Consider the set of balls F centered on the line segment $s = [wc] \subset \Omega$ and circumscribing σ . Any ball in F is included in $B_\tau \cup B_\sigma$ and it circumscribes σ . Hence its interior contains no point of $L \setminus \tau$. Moreover, since the interior of B_σ is empty but the interior of B_τ is not, there exists a point z on s such that the ball of F centered at z , witnesses τ and contains $|\sigma| + 1$ points of τ on its boundary. Call this new ball B_τ . We can then carry on the induction and obtain a witness ball B_τ whose bounding sphere contains all the vertices of τ . Such a ball is thus a Delaunay ball and τ is a Delaunay simplex. $\square$

It is worth noticing that, for a simplex σ to belong to the witness complex, we required all the faces of σ to have a witness. As illustrated in Figure 6.5, this is mandatory for the theorem to hold.

We deduce from Lemma 6.4 and Theorem 6.6 the following corollary

Corollary 6.7 *Let Ω be a convex subset of $\mathbb{R}^d$ and let $W \subseteq \Omega$. We have $\text{Wit}(L, W) \subseteq \text{Del}_\Omega(L) \subseteq \text{Del}(L)$.*

If the points L are in general position with respect to spheres, we know that $\text{Del}(L)$ is embedded in $\mathbb{R}^d$ by Delaunay's theorem 4.3. It therefore follows from Corollary 6.7 that the same is true for $\text{Wit}(L, W)$. In particular, the dimension of $\text{Wit}(L, W)$ is at most d . When W is not the whole space $\mathbb{R}^d$ but a finite set of points, Theorem 6.6 no longer holds. However, the following lemma shows that both complexes are identical provided that W is dense enough and L is protected enough (refer to Section 5.4 for a definition of protection).

Lemma 6.8 (Identity from protection) *Let Ω be a convex subset of $\mathbb{R}^d$*

and let W and L be two finite sets of points in Ω . If W is ε -dense in Ω and if all simplices (of all dimensions) of $\text{Del}_\Omega(L)$ are δ -protected with $\delta \geq 2\varepsilon$, then $\text{Wit}(L, W) = \text{Del}_\Omega(L)$.

Proof By Corollary 6.7, we have $\text{Wit}(L, W) \subseteq \text{Del}_\Omega(L)$. We now prove the other inclusion. Let σ be a simplex in $\text{Del}_\Omega(L)$. By hypothesis, σ is δ -protected. Then there exists a point $c \in \Omega$ such that

$$\forall p \in \sigma, \forall q \in L \setminus \sigma, \quad \|c - p\| \leq \|c - q\| - \delta.$$

Since W is ε -dense in Ω , there exists a point $w \in W$ such that $\|w - c\| \leq \varepsilon$. Using the triangular inequality, we have for any $p \in \sigma$ and $q \in L \setminus \sigma$

$$\begin{aligned} \|w - p\| &\leq \|w - c\| + (\|c - q\| - \delta) \\ &\leq \|w - q\| + 2\|w - c\| - \delta \\ &\leq \|w - q\| + 2\varepsilon - \delta \end{aligned}$$

Hence, when $\delta \geq 2\varepsilon$, w is a witness for σ . $\square$

The above lemma requires the simplices of all dimensions to be δ -protected. In fact it is sufficient to check d -simplices only. Indeed, it can be proved that if the d -dimensional simplices are δ' -protected, for δ' slightly bigger than δ , then all simplices of all dimensions are also protected. See Exercise 5.11.

6.2.2 Computing witness complexes

Let L and W be two finite sets of points and $\text{Wit}(L, W)$ their witness complex. We assume, for convenience, that no two points of L are at the same distance from a point in W . We describe how to compute the k -skeleton of $\text{Wit}(L, W)$, denoted by $\text{Wit}^k(L, W)$, for any fixed k .

Let M be a matrix of size $|W| \times k$. The lines in M are associated to the elements of W and the line $M(w)$ associated to $w \in W$ stores the list of the k landmarks that are closest to w , sorted by increasing distance from w (M can be trivially computed in time $O(|W| \times |L| \log |L|)$ and, with more clever algorithms, in time $O(|W| (\log |W| + k))$ (see the bibliographical notes in Section 6.4). We write $M(w) = (M_1(w), \dots, M_k(w))$. Algorithm 7 below computes $\text{Wit}^k(L, W)$ from M . We assume without real loss of generality that $L \subset W$.

Under the general position assumption above, the number of i -simplices of $\text{Wit}(L, W)$ is at most $|W|$ for any $i \leq k$. Hence the total number of simplices of $\text{Wit}^k(L, W)$ is at most $k|W| = O(|W|)$.

Algorithm 7 Construction of the k -skeleton of a witness complex

Input: W a finite point set, $L \subset W$, k , M
 $\text{Wit}^k(L, W) := \emptyset$
 $W' := W$ {a set of active witnesses}
for $i = 0, \dots, k - 1$ **do**
 for each $w \in W'$ **do**
 if the i -simplex $\sigma(w) = [M_1(w), \dots, M_{i+1}(w)] \notin \text{Wit}^k(L, W)$ **then**
 if all the $(i - 1)$ -faces of $\sigma(w)$ are in $\text{Wit}^k(L, W)$ **then**
 add $\sigma(w)$ to $\text{Wit}^k(L, W)$
 else
 $W' := W' \setminus \{w\}$
Output: $\text{Wit}^k(L, W)$

The total number of for loops that are executed is $|\text{Wit}^k(L, W)| + |W| = O(|W|)$. Indeed, each loop either constructs a new simplex or removes a point from the set of active witnesses W' . A loop has to decide if a i -simplex as well as its i facets belong to the current complex $\text{Wit}^k(L, W)$. Each of these tests can be done in time $O(\log |L|)$ in one uses for example the simplex tree described in Exercise 2.1. Since $i \leq k$, the cost of a single loop is thus $O(k \log |L|)$. The overall complexity of the algorithm is therefore $O(k|W| \log |L|)$.

The algorithm is general and applies to any distance matrix M . In the case where L and W belong to $\mathbb{R}^d$, and if the points of L are in general position with respect to spheres, we know that $\text{Wit}(L, W)$ is a subcomplex of $\text{Del}(L)$ and thus embedded in $\mathbb{R}^d$ and of dimension at most d . Hence the entire witness complex is computed if one takes $k = d$.

6.2.3 Relaxed witness complexes

As before, W and L denote two sets of points in some subset $\Omega \subset \mathbb{R}^d$. L is finite.

Definition 6.9 (Relaxed witness) *Let σ be a simplex with vertices in L .*

We say that a point $w \in W$ is an α -witness of σ if

$$\|w - p\| \leq \|w - q\| + \alpha \quad \forall p \in \sigma \quad \text{and} \quad \forall q \in L \setminus \sigma.$$

Definition 6.10 (Relaxed witness complex) *The α -relaxed witness complex $\text{Wit}^\alpha(L, W)$ is the maximal simplicial complex with vertex set L whose simplices have an α -witness in W .*

For $\alpha = 0$, the relaxed witness complex is the standard witness complex. The parameter α defines a filtration on the witness complex, which can be used to compute persistent homology (Chapter 11.5).

Construction

We adapt Algorithm 7 above. At each step j , we insert, for each witness w , the j -dimensional simplices which are α -witnessed by w . Differently from the standard witness complex, there may be more than one j -simplex that is witnessed by a given witness $w \in W$. Consequently, we do not maintain a pointer from each active witness to the last inserted simplex it witnesses. We use simple top-down insertions from the root of the simplex tree 2.7.

Given a witness w and a dimension j , we generate all the j -dimensional simplices which are α -witnessed by w . For the ease of exposition, we suppose we are given the sorted list of nearest neighbors of w in L , noted $\{z_0 \cdots z_{|L|-1}\}$, and their distance to w , noted $m_i = d(w, z_i)$, with $m_0 \leq \cdots \leq m_{|L|-1}$, breaking ties arbitrarily. Note that if one wants to construct only the k -skeleton of the complex, it is sufficient to know the list of neighbors of w that are at distance at most $m_k + \alpha$ from w . We preprocess this list of neighbors for all witnesses. For $i \in \{0, \dots, |L| - 1\}$, we define the set A_i of landmarks z such that $m_i \leq d(w, z) \leq m_i + \alpha$. For $i \leq j + 1$, w α -witnesses all the j -simplices that contain $\{z_0, \dots, z_{i-1}\}$ and a $(j + 1 - i)$ -subset of A_i , provided $|A_i| \geq j + 1 - i$. We see that all j -simplices that are α -witnessed by w are obtained this way, and exactly once, when i ranges from 0 to $j + 1$.

For all $i \in \{0, \dots, j + 1\}$, we compute A_i and generate all the simplices which contain $\{z_0, \dots, z_{i-1}\}$ and a subset of A_i of size $(j + 1 - i)$. In order to easily update A_i when i is incremented, we maintain two pointers to the list of neighbors, one to z_i and the other to the end of A_i . We check in constant time if A_i contains more than $j + 1 - i$ vertices, and compute all the subsets of A_i of cardinality $j + 1 - i$ accordingly. See Figure 6.6.




Figure 6.6: Computation of the α -witnessed simplices σ of dimension 5. If z_3 is the first neighbor of w not in σ , then σ contains in particular $\{z_0, z_1, z_2\}$ and any 3-uplet of $A_3 = \{z_4, \dots, z_8\}$.

Restricted Delaunay complex and relaxed witness complex

We consider now the case where $\Omega \subseteq \mathbb{R}^d$ is not necessarily convex and extend Lemma 6.8. Specifically, we will show that, for a large enough relaxation, $\text{Del}_\Omega(L)$ is contained in the relaxed witness complex.

Lemma 6.11 *Let Ω be a subset of $\mathbb{R}^d$ (not necessarily convex). Assume that W is ε -dense in Ω and that all the simplices of $\text{Del}_\Omega(L)$ are δ -protected. If $\alpha \geq \max(0, 2\varepsilon - \delta)$, then $\text{Del}_\Omega(L) \subseteq \text{Wit}^\alpha(L, W)$.*

Proof Let σ be a d -simplex of $\text{Del}_\Omega(L)$ and write c_σ for its circumcenter. Since W is ε -dense in Ω , there exists a point w in W such that $\|c_\sigma - w\| \leq \varepsilon$. For any $p \in \sigma$ and $q \in L \setminus \sigma$, we then have

$$\begin{aligned}
 \|w - p\| &\leq \|c_\sigma - p\| + \|c_\sigma - w\| \\
 &\leq \|c_\sigma - q\| - \delta + \|c_\sigma - w\| \\
 &\leq \|w - q\| + 2\|c_\sigma - w\| - \delta \\
 &\leq \|w - q\| + 2\varepsilon - \delta
 \end{aligned}$$

which proves the lemma. $\square$

Note that if $\delta \geq 2\varepsilon$, the lemma gives $\text{Del}_\Omega(L) \subseteq \text{Wit}(L, W)$ (as in Lemma 6.8).

6.3 Exercises

Exercise 6.1 (Computing regularized α -complexes.) Let $\hat{P}$ be a set of weighted points in $\mathbb{R}^d$. Slightly modify Algorithm 6 so that it computes,

for each simplex τ of the Delaunay triangulation $\text{Del}(\hat{P})$, the value of the parameter α at which the simplex τ becomes part of the regularized α -complex $\mathcal{A}_r(\hat{P}, \alpha)$.

Exercise 6.2 (Full classification with respect to the α -complex.) Let $\hat{P}$ be a set of weighted points in $\mathbb{R}^d$. Each simplex τ of the Delaunay triangulation $\text{Del}(\hat{P})$ can be classified with respect to the α -complex $\mathcal{A}(\hat{P}, \alpha)$ as *external* if it does not belong to $\mathcal{A}(\hat{P}, \alpha)$, *singular* if it belongs to $\mathcal{A}(\hat{P}, \alpha)$ but none of its coface in $\text{Del}(\hat{P})$ does, *boundary* if it belongs to the boundary of α -complex and is not singular, which means that some of its cofaces belong to $\mathcal{A}(\hat{P}, \alpha)$ while others do not, and at last *internal* if it belongs to the interior of the α -complex, meaning that all its cofaces belong to $\mathcal{A}(\hat{P}, \alpha)$.

Modify Algorithm 6 so that it computes for any simplex τ of $\text{Del}(\hat{P})$, the at most three values of the parameter α where the status of the simplex changes from external to singular and then to boundary and interior.

Exercise 6.3 (Weighted witness complex) Show that the witness complex can be extended to weighted points and the weighted distance. Show that the identity results of Section 6.2.1 still hold when the Delaunay complex is replaced by its weighted counterpart.

Exercise 6.4 Show that, if L is a λ -sample of $\mathbb{R}^d$, the circumradius R_σ of any simplex σ in $\text{Wit}(L, W)$ is at most λ . (Hint : use Corollary 6.7).

Exercise 6.5 (Relaxed Delaunay complex) Let W and L be two finite sets of points in $\Omega \subset \mathbb{R}^d$. Let σ be a simplex with vertices in L . We say that a point $w \in W$ is an α -center of σ if

$$\|w - p\| \leq \|w - q\| + \alpha \quad \forall p \in \sigma \quad \text{and} \quad \forall q \in L.$$

The α -relaxed Delaunay complex $\text{Del}^\alpha(L, W)$ is the maximal simplicial complex with vertex set L whose d -simplices have an α -center in W . Show how to construct $\text{Del}^\alpha(L, W)$. Prove that $\text{Del}(L) \subseteq \text{Del}^\alpha(L, W)$ for $\alpha \geq 2\varepsilon$ if W is an ε -sample of Ω .

6.4 Bibliographical notes

Alpha-shapes were introduced by Edelsbrunner, Kirkpatrick and Seidel [76, 72]. Alpha-shapes are also widely used to represent union of balls [67] and to study the structure of macro molecules and various related problems like the docking of two molecules, see e.g. [70, 102, 69]

Witness complexes and relaxed Delaunay triangulations have been introduced in the seminal work of de Silva [63] who first proved Theorem 6.6 and several of its variants including the case of weighted points (Exercise 6.3). The proof presented in Section 6.2 is due to Attali et al. [4]. The 1-skeleton of the witness complex has been introduced earlier by Martinez and Schulten [107]. They showed that $\text{Del}(L)$ and $\text{Wit}(L, \mathbb{R}^d)$ have the same 1-skeleton and they proposed a dynamic algorithm for approximating the topology of a region of space, by a graph represented as a neural network.

The identity of witness and Delaunay complexes when the number of witnesses is finite is taken from [21]. The paper describes an algorithm to obtain $\text{Del}(L)$ from $\text{Wit}(L, W)$ using the algorithmic version of the local Lovász lemma (see Exercise 5.12 and Section 5.3.4). The paper also describes an algorithm to compute the relaxed Delaunay complex in time sublinear in the number of witnesses (Exercise 6.5).

Given a set of n points $P \subset \mathbb{R}^d$, one can construct a data structure called a well-separated pair decomposition. This data structure has many applications and can be used, in particular, to compute the k -nearest neighbors of all the points in P in time $O(n \log n + kn)$ [8]. See the book of Har-Peled for a recent account on well-separated decompositions [95].

Part III

Reconstruction of smooth submanifolds

Triangulating an object $\mathcal{M}$ entails computing a simplicial complex K which is homeomorphic to $\mathcal{M}$. This is a demanding quest and, in this part, we will make the strong assumption that $\mathcal{M}$ is a smooth submanifold of $\mathbb{R}^d$. We defer to the last part of the book other compromises between the generality of the shapes to be approximated and the quality of the approximation.

This part consists of three chapters. The first one provides conditions under which a simplicial complex is homeomorphic to a submanifold $\mathcal{M}$ of $\mathbb{R}^d$ and thus is a triangulation of $\mathcal{M}$ (Theorem 7.16). Prior to stating and proving the main theorem, we introduce the central notions of reach and ε -net on a manifold.

In Chapter 6.1, we introduce the alpha complex of a finite point set P and show how it can be constructed from the Delaunay triangulation of P . In the case where P is a dense enough sample of a submanifold $\mathcal{M} \subset \mathbb{R}^d$, the alpha complex has the same homotopy type as $\mathcal{M}$. However, the alpha complex of P is usually not homeomorphic to $\mathcal{M}$.

In Chapter 8, we study the manifold reconstruction problem. Given a finite set of points P on an unknown manifold $\mathcal{M}$, the goal is to compute a simplicial complex $\hat{\mathcal{M}}$ with vertex set P that triangulates $\mathcal{M}$. This problem is of primary importance when $\mathcal{M}$ is a surface of $\mathbb{R}^3$ (it is then known as the surface reconstruction problem). It also has applications in higher dimensions in the context of Data Analysis where data are considered as points in some Euclidean space, of usually high dimension.

A major difficulty, when considering *higher dimensional manifolds*, comes from the fact that triangulating high dimensional spaces requires exponential time and space (this phenomenon is called the curse of dimensionality). We therefore cannot afford to triangulate the ambient space as is commonly done when considering surface of $\mathbb{R}^3$. However, if the intrinsic dimension k of $\mathcal{M}$ is much smaller than the dimension d of the ambient space, we will see in Chapter 8 that we can walk around the curse of dimensionality and present an efficient algorithm that reconstructs a submanifold $\mathcal{M}$ from a finite sample P under appropriate conditions on P . The assumption that k is small even if d is large is common in Manifold Learning and Data Analysis. It reflects the fact that the data points which may live in a space of very high dimension are usually produced by a system with a limited number of degrees of freedom, and therefore lie on or close to a structure of small intrinsic dimension.

Chapter 7

Triangulation of submanifolds

Triangulating an object $\mathcal{M}$ entails computing a simplicial complex which is homeomorphic to $\mathcal{M}$. This is a demanding quest and, in this chapter, we will assume that $\mathcal{M}$ is a smooth and compact submanifold of $\mathbb{R}^d$ without boundary. The main goal of the chapter is to prove Theorem 7.16 that provides sufficient conditions under which a simplicial complex $\hat{\mathcal{M}}$ is a triangulation of a submanifold $\mathcal{M}$ of $\mathbb{R}^d$. These conditions require the simplices to be sufficiently small and thick and rely on the concepts of reach and of ε -net on a manifold. To prove the theorem, we will prove that the projection map that associates to a point of $\hat{\mathcal{M}}$ its closest point on $\mathcal{M}$ is a homeomorphism. The results of this chapter will be used in Chapter 8 to triangulate a submanifold $\mathcal{M}$ given only a finite point set on $\mathcal{M}$.

7.1 Reach and ε -nets on submanifolds

7.1.1 Submanifolds

Given an open set $U \subseteq \mathbb{R}^d$ and a non negative integer c , a map $\phi : U \rightarrow \mathbb{R}^d$ is said to be *c-differentiable*, or *of class $\mathcal{C}^c$* on U if its successive derivatives up to order c are well-defined and continuous on U . In particular, a 0-differentiable map is a continuous map. If moreover, $\phi : U \rightarrow V = \phi(U)$ is a bijection and $\phi^{-1} : V \rightarrow U$ is also c -differentiable, then ϕ is said to be a *c-differentiable diffeomorphism*.

Definition 7.1 (Submanifold) *A compact subset $\mathcal{M} \subset \mathbb{R}^d$ is a c-differentiable submanifold of dimension $k \leq d$, if for any $p \in \mathcal{M}$ there exist an open set $U \subset \mathbb{R}^d$ containing p , a c-differentiable diffeomorphism ϕ from U to an open set $V \subset \mathbb{R}^d$, and an affine k -dimensional subspace $A \subset \mathbb{R}^d$ such that*

$$\phi(U \cap \mathcal{M}) = A \cap V.$$

Intuitively, a submanifold of dimension k is a subset of $\mathbb{R}^d$ that is locally homeomorphic to an affine space of dimension k (see figure 7.1). A *curve* is a 1-dimensional submanifold of $\mathbb{R}^d$, and a *surface* is a 2-dimensional submanifold of $\mathbb{R}^d$. The submanifolds we will consider are differentiable and have no boundary, even if not explicitly mentionned.

The *tangent space* $T_p\mathcal{M}$ of a c -differentiable ($c > 0$) submanifold $\mathcal{M}$ at a point $p \in \mathcal{M}$ is the vector space spanned by the tangent vectors $\gamma'(0)$




Figure 7.1: A k -dimensional submanifold M of $\mathbb{R}^d$ (here $k = 1$ and $d = 2$) is a subset of $\mathbb{R}^d$ that is locally diffeomorphic to an open set of a k -dimensional affine subspace.

where $\gamma : (-1, 1) \rightarrow M \subset \mathbb{R}^d$ belongs to the set of differentiable curves contained in M such that $\gamma(0) = p$. If M is of dimension k , then $T_p M$ is a k -dimensional vector space. By an abuse of notation, we will also denote by $T_p M$ the affine subspace spanned by $T_p M$ and passing through p and we will denote by $N_p M$ the $(d - k)$ -dimensional affine subspace orthogonal to $T_p M$ (see Figure 7.2). In the sequel, to avoid heavy notation, when there is no ambiguity, we will drop the reference to M and denote $T_p M$ and $N_p M$ by T_p and N_p respectively.

In this chapter, we focus on submanifolds and do not consider manifolds defined in an intrinsic way, independently of any embedding in $\mathbb{R}^d$. However, the main results of the chapter and in particular the triangulation theorem (Theorem 7.16) can be extended to intrinsic smooth manifolds. See the bibliographical notes.

7.1.2 Projection map, medial axis and reach

Let M be a submanifold of $\mathbb{R}^d$. The *medial axis* of M is defined as the closure of the set of points $x \in \mathbb{R}^d$ that have more than one closest point on M (see Figure 7.3). We denote it by $\text{ax}(M)$.

We can associate to each point of $\text{ax}(M)$ a ball that is centered at that point, whose interior does not intersect M and that is maximal for inclusion. Such a ball will be called a *medial ball*.




Figure 7.2: The tangent and normal spaces to a 2-dimensional submanifold $\mathcal{M}$ of $\mathbb{R}^3$ at a point $p \in \mathcal{M}$.




Figure 7.3: The medial axis of a closed curve. We only show the component of the medial axis that is contained in the domain bounded by the curve. Various offsets of the curve are also shown as thin curves.

We define the *projection onto $\mathcal{M}$* as the mapping

$$\Pi : \mathbb{R}^d \setminus \text{ax}(\mathcal{M}) \rightarrow \mathcal{M}$$

that maps a point x to its (unique) closest point on $\mathcal{M}$.

The *reach* of $\mathcal{M}$, written $\text{rch}(\mathcal{M})$, is the infimum of the distance from a point of $\mathcal{M}$ to the medial axis $\text{ax}(\mathcal{M})$. As we will see, the reach encodes both local curvature considerations as well as global ones.

In this chapter, we will restrict our attention to the class of submanifolds of $\mathbb{R}^d$ with positive reach introduced by Federer [82]. This class includes all submanifolds of class C^2 and also some submanifolds whose principal curvatures may be discontinuous on subsets of measure 0. An example of such a submanifold is the *r-offset* of a solid cube, i.e. the set of points at distance at most r from the cube.

We now state some properties of submanifolds with positive reach.

Lemma 7.2 (Tubular neighborhood) *Let $\mathcal{M}$ be a manifold with positive reach $\text{rch}(\mathcal{M})$ and let $B_{N_p}(r)$ be the intersection of $B(p, r)$ with the normal space at p . If $r < \text{rch}(\mathcal{M})$, then, for every point $x \in B_{N_p}(r)$, $\Pi(x) = p$.*

The proof of this lemma for C^2 submanifolds follows from rather standard arguments in differential geometry. The result for submanifolds of positive reach is due to Federer [82]. From Lemma 7.2, we easily deduce the following lemma.

Lemma 7.3 *Let $\mathcal{M}$ be a submanifold of positive reach, and let $x \in \mathcal{M}$. Any open ball that is tangent to $\mathcal{M}$ at x and whose radius is at most $\text{rch}(\mathcal{M})$ does not intersect $\mathcal{M}$.*

Proof Let $B(c, r)$ be a ball tangent to $\mathcal{M}$ at x and assume that $r < \text{rch}(\mathcal{M})$. If the intersection of $\mathcal{M}$ and the open ball $B(c, r)$ is not empty, then $\Pi(c) \neq x$, contradicting Lemma 7.2. The result for $r = \text{rch}(\mathcal{M})$ now follows by taking the limit. $\square$

Lemma 7.4 *Let B be a closed ball that intersects $\mathcal{M}$. If $B \cap \mathcal{M}$ is not a topological ball, then B contains a point of the medial axis of $\mathcal{M}$.*




Figure 7.4: For the proof of Lemma 7.4.

Proof Write c for the center of B . The result is trivial when c belongs to the medial axis of $\mathcal{M}$. Therefore assume that $c \notin \text{ax}(\mathcal{M})$.

Let y be the (unique) point of $\mathcal{M}$ closest to c . We denote by B_y the closed ball centered at c with radius $\|c - y\|$ (see Figure 7.4). Plainly, the interior of B_y does not intersect $\mathcal{M}$ and $B_y \cap \mathcal{M} = \{y\}$, otherwise c would be a point of the medial axis. Hence y is an isolated critical point of the distance function from c , i.e. the function $d_c : \mathbb{R}^d \rightarrow \mathbb{R}$, $d_c(x) = \|x - c\|$. d_c is minimal at y and, for a small enough radius r strictly larger than $\|y - c\|$, $B(c, r) \cap \mathcal{M}$ is a topological ball of the same dimension as $\mathcal{M}$. Since $B \cap \mathcal{M}$ is not a topological ball, it follows from a basic result in Differential Topology [108, Theorem 3.1]¹ that there exists another critical point of d_c , say $z \in \mathcal{M}$, $z \neq y$, such that $r_c = \|c - z\| > \|c - y\|$ and the ball $B(c, r_c)$ is tangent to $\mathcal{M}$ at z . Consider the set $\mathcal{B}_z$ of closed balls that are tangent to $\mathcal{M}$ at z and are centered on the line segment $[cz]$. Note that $B(c, r_c)$ is the ball of $\mathcal{B}_z$ centered at c . Since the interior of $B(c, r_c)$ contains y and therefore intersects $\mathcal{M}$, there must exist a ball $B_z \in \mathcal{B}_z$ maximal for the inclusion whose interior does not intersect $\mathcal{M}$. The center of B_z belongs to $\text{ax}(\mathcal{M})$ and also to B since $B_z \subset B(c, r_c) \subset B$. The lemma is proved. $\square$

If $x \in \mathcal{M}$, $B(x, r)$ cannot intersect the medial axis of $\mathcal{M}$ for any $r < \text{rch}(x)$. Lemma 7.4 thus implies

¹Let $a = d_c(y)$ and $b = r$ and suppose that the set $d_c^{-1}(a, b)$, consisting of all $p \in \mathcal{M}$ with $a \leq d_c(p) \leq b$, contains no critical points of d_c (i.e. no point q of $\mathcal{M}$ where $B(c, \|c - q\|)$ is tangent to $\mathcal{M}$). Then $M^a = \{x \in \mathcal{M}, d_c \leq a\}$ is homeomorphic to $M^b = \{x \in \mathcal{M}, d_c \leq b\}$.

Corollary 7.5 *For any x of $\mathcal{M}$, and any $r < \text{rch}(x)$, the intersection of $\mathcal{M}$ with the ball $B(x, r)$ centered at x of radius r is a topological ball.*

Assume that we are given a finite set of points P on $\mathcal{M}$ and let $\text{Vor}(P)$ be the Voronoi diagram of P in the ambient space $\mathbb{R}^d$. The following lemma shows that the Voronoi cell of any $p \in P$ has a large extent in the normal space N_p at p .

Lemma 7.6 *Let $p \in P \subset \mathcal{M}$ and write N_p for the normal space of $\mathcal{M}$ at p . Then the ambient Voronoi cell $V(p)$ of p has a large extent in N_p . Specifically, $B(p, \text{rch}(\mathcal{M})) \cap N_p \subset V(p)$, where $B(c, r)$ denotes, as usual, the open ball centered at c of radius r .*

Proof Suppose $w \in (B(p, \text{rch}(\mathcal{M})) \cap N_p) \setminus V(p)$. Then the line segment $[pw]$ crosses the boundary of $V(p)$ at some point z and there exists some $u \in P \subset \mathcal{M}$, distinct from p , such that $\|z - u\| = \|z - p\|$. Since $z \in N_p$, p is the closest point to z on $\mathcal{M}$, and we have reached a contradiction with the definition of the reach. $\square$

7.1.3 ε -nets on a submanifold

We give a variant of the definition of ε -nets introduced in Section 5.1. This new definition is better adapted to the case of point samples on a submanifold and is sensitive to the reach of the manifold.

Definition 7.7 ($(\varepsilon, \bar{\eta})$ -net) *Let $\mathcal{M}$ be a submanifold of $\mathbb{R}^d$ of positive reach. A finite point set $P \subset \mathcal{M}$ is called an $(\varepsilon, \bar{\eta})$ -net of $\mathcal{M}$ if it is*

1. **ε -dense:** *any point x of $\mathcal{M}$ is at distance at most $\varepsilon \text{rch}(\mathcal{M})$ from a point of P (the distance is the Euclidean distance in $\mathbb{R}^d$),*
2. **η -separated:** *for any two points p, q of P , $\|p - q\| \geq \eta \text{rch}(\mathcal{M})$, where $\eta = \bar{\eta} \varepsilon$.*

We call ε the sampling radius of P and $\bar{\eta}$ the separation ratio of P . Note that, differently from Section 5.1, ε here is a dimensionless quantity.

This definition does not allow the sampling radius to vary over the submanifold. Hence, we are confined to uniform samples, which may be quite restrictive in practice. The results of this chapter can be extended to non-uniform ε -nets but, in order to keep the exposition simple and better outline the key ideas, we will restrict our attention to *uniform* ε -nets. We leave the extension to non-uniform ε -nets as an exercise (Exercise 7.4). See also the bibliographical notes.

7.2 Projection maps

In the rest of this chapter, $\mathcal{M}$ denotes a submanifold of $\mathbb{R}^d$ of positive reach $\text{rch}(\mathcal{M})$. The tangent space at $x \in \mathcal{M}$ is denoted by T_x . The angle between two vector subspaces U and V of $\mathbb{R}^d$ is denoted by $\angle(U, V)$ (see Section 5.2.2 for a definition).

If σ is a simplex, we denote by $\Delta(\sigma)$ its diameter (i.e. the length of its longest edge) and $\Theta(\sigma)$ its thickness. We further write for convenience $\Delta(\sigma) = \delta(\sigma) \text{rch}(\mathcal{M})$.

Lemma 7.8 *Let p and q be two points of $\mathcal{M}$. We have*

1. $\sin \angle(pq, T_p) \leq \frac{\|p-q\|}{2 \text{rch}(\mathcal{M})};$
2. *the distance from q to T_p is at most $\frac{\|p-q\|^2}{2 \text{rch}(\mathcal{M})}$.*

Proof 1. Let q' be the orthogonal projection of q onto T_p and let H be the plane $(pq q')$. Let in addition D be the open disk of H of radius $\text{rch}(\mathcal{M})$ that is tangent to $\mathcal{M}$ at p and whose center c is on the same side of $T_p \cap H$ as q (Refer to Figure 7.5). Since D is tangent to $\mathcal{M}$ and its radius is $\text{rch}(\mathcal{M})$, it follows from Lemma 7.3 that $d(c, \mathcal{M}) = \|c - p\|$ and that D does not intersect $\mathcal{M}$. Hence q does not belong to the interior of D (Lemma 7.3). Assume that the line segment $[pq]$ intersects the boundary of D in a point q'' distinct from p . We have

$$\|p - q\| \geq \|p - q''\| = 2 \text{rch}(\mathcal{M}) \sin \angle(pq, T_p)$$

which proves the first statement.




Figure 7.5: An illustration of the proof of Lemma 7.8 in the case where $\mathcal{M}$ is a curve of $\mathbb{R}^2$. $\theta = \angle(pq, T_p)$.

2. We have $\|q - q'\| = \|p - q\| \sin \angle(pq, T_p) \leq \frac{\|p - q\|^2}{2 \text{rch}(\mathcal{M})}$. $\square$

The previous lemma allows us to bound the Hausdorff distance between $\mathcal{M}$ and a simplex which has its vertices on $\mathcal{M}$.

Lemma 7.9 *Let σ be a simplex with its vertices on $\mathcal{M}$ and assume that $\delta(\sigma) < 1$. Then, for any point $x \in \sigma$, we have $d(x, \mathcal{M}) \leq 2\delta^2(\sigma) \text{rch}(\mathcal{M})$.*

Proof Write x' the point of $\mathcal{M}$ closest to x and note that $\|x - x'\| = d(x, \mathcal{M}) = d(x, T_{x'})$. For any vertex p of σ , we have

$$\|p - x'\| \leq \|p - x\| + \|x - x'\| \leq 2\|p - x\| \leq 2\Delta(\sigma) = 2\delta(\sigma) \text{rch}(\mathcal{M}).$$

Applying Lemma 7.8, we then get $d(p, T_{x'}) \leq 2\delta^2(\sigma) \text{rch}(\mathcal{M})$. This is true for all vertices p of σ and, since the function $d(\cdot, T_{x'})$ is affine on σ , it is also true for x and we have $d(x, T_{x'}) = \|x - x'\| \leq 2\delta^2(\sigma) \text{rch}(\mathcal{M})$. $\square$

The following lemma bounds the angle between two tangent spaces. Its proof relies on notions of differential geometry that go beyond the scope of this book. See the bibliographic notes (Section 7.5) for references.

Lemma 7.10 (Angle between tangent spaces) *Let $p, q \in \mathcal{M}$. Then $\sin \angle(T_p, T_q) \leq \|p - q\| / \text{rch}(\mathcal{M})$.*

We now introduce the important notion of distortion of a map.

Definition 7.11 (ξ -distortion map) A map $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a ξ -distortion map if for all $x, y \in U$ we have

$$|\|F(x) - F(y)\| - \|x - y\|| \leq \xi \|x - y\|,$$

or equivalently if

$$(1 - \xi) \|x - y\| \leq \|F(x) - F(y)\| \leq (1 + \xi) \|x - y\|.$$

We will need the next lemma.

Lemma 7.12 Let U be a subset of $\mathbb{R}^d$. A ξ -distortion map $F : U \rightarrow \mathbb{R}^d$ is an embedding (i.e. a homeomorphism onto its image) if $\xi < 1$. If it is smooth, then we have $|\|Ju\| - 1| \leq \xi$, where J denotes the Jacobian matrix of F and u is any unit vector of $\mathbb{R}^d$.

Proof Continuity and injectivity directly follow from Definition 7.11. The continuity of the inverse follows from Exercise 7.9. This proves that F is an embedding. The second part of the lemma follows from Definition 7.11. $\square$

We now study two maps that will play a crucial role in proving the main theorem of this chapter (Theorem 7.16). The first map has already been defined. It is the projection $\Pi : \mathbb{R}^d \setminus \text{ax}(\mathcal{M}) \rightarrow \mathcal{M}$ that maps a point to its closest point on $\mathcal{M}$.

We show now that Π is Lipschitz continuous in a neighborhood of $\mathcal{M}$, and that its restriction to a simplex with vertices on $\mathcal{M}$ has a *distortion* that is bounded as a function of the diameter and of the thickness of the simplex.

Lemma 7.13 (Distortion of Π) 1. Let x and y be two points of $\mathbb{R}^d \setminus \text{ax}(\mathcal{M})$. Write $x' = \Pi(x)$ and $y' = \Pi(y)$ for their (unique) projections onto $\mathcal{M}$. If $\|x - x'\| \leq \alpha \text{rch}(\mathcal{M})$ and $\|y - y'\| \leq \alpha \text{rch}(\mathcal{M})$ with $\alpha < 1$, then $\|x' - y'\| \leq \frac{1}{1-\alpha} \|x - y\|$.

2. Let σ be a simplex with its vertices on $\mathcal{M}$, and assume that $\delta(\sigma) \leq \frac{1}{2}$ and $\Theta(\sigma) \geq 5\delta(\sigma)$. Then the restriction of Π to σ is a ξ -distortion map where $\xi = 4\delta^2(\sigma) + \frac{16\delta^2(\sigma)}{\Theta^2(\sigma)} < 1$ and thus embeds σ in $\mathcal{M}$.




Figure 7.6: For the proof of Lemma 7.13 (1).

Proof Refer to Figure 7.6. Let $\theta_x = \angle(y' - x', T_{x'})$ and $\theta_y = \angle(y' - x', T_{y'})$. From Lemma 7.8, both $\sin \theta_x$ and $\sin \theta_y$ are bounded by

$$\frac{\|y' - x'\|}{2 \operatorname{rch}(\mathcal{M})}. \quad (7.1)$$

1. Let x'' and y'' be the orthogonal projections of x and y onto the line $(x'y')$. Let H_x be the hyperplane passing through x' and orthogonal to xx' , and write $\psi_x = \angle(x' - y', H_x)$ and observe that $\psi_x = \frac{\pi}{2} - \angle(x' - y', x - x') = \angle(x' - x, x - x'')$. By definition of angles, $\psi_x \leq \theta_x$ since $T_{x'} \subseteq H_x$. Using (7.1), we have

$$\|x' - x''\| = \|x - x'\| \sin \psi_x \leq \|x - x'\| \sin \theta_x \leq \alpha \operatorname{rch}(\mathcal{M}) \sin \theta_x \leq \frac{\alpha}{2} \|y' - x'\|.$$

Likewise : $\|y' - y''\| \leq \frac{\alpha}{2} \|y' - x'\|$. So

$$\|x - y\| \geq \|x'' - y''\| \geq \|y' - x'\| - \|x' - x''\| - \|y' - y''\| \geq (1 - \alpha) \|y' - x'\|.$$

The first part of the lemma follows.

2. Refer to Figure 7.7. Let x and y be two points of σ , $x' = \Pi(x)$ and $y' = \Pi(y)$. By Lemma 7.9, we know that $\|x - x'\|$ and $\|y - y'\|$ are at most $\alpha \operatorname{rch}(\mathcal{M})$ where $\alpha = 2\delta^2(\sigma) \leq 1/2$.




Figure 7.7: For the proof of Lemma 7.13 (2).

From the first part of the lemma, we have

$$\|x' - y'\| \leq \frac{1}{1 - \alpha} \|x - y\| \leq (1 + 2\alpha) \|x - y\| \leq 2 \|x - y\| \quad (7.2)$$

We now prove a lower bound on $\|x' - y'\|/\|x - y\|$.

Let, as before, H_x denote the hyperplane passing through x' and orthogonal to xx' . Let z be the projection of y onto H_x , z' the projection of y' onto H_x and write $\phi = \angle(x - y, x' - z) = \angle(x - y, H_x)$ (See Figure 7.7). We have

$$\|x' - y'\| \geq \|x' - z'\| \geq \|x' - z\| - \|z' - z\| = \|x - y\| \cos \phi - \|z' - z\| \quad (7.3)$$

Let us bound $\cos \phi$ and $\|z' - z\|$. We have

$$\begin{aligned} \phi = \angle(x - y, x' - z) &= \min_{v \in H_x} \angle(x - y, v) \\ &\leq \min_{v \in T_{x'}} \angle(x - y, v) \\ &\leq \max_{u \in \text{aff}(\sigma)} \min_{v \in T_{x'}} \angle(u, v) \\ &= \angle(\text{aff}(\sigma), T_{x'}). \end{aligned}$$

By the proof of Lemma 7.9, the vertices of σ are at distance at most $\alpha \text{rch}(\mathcal{M})$ from $T_{x'}$. We then deduce from Lemma 5.14

$$\sin \phi \leq \frac{2\alpha}{\delta(\sigma) \Theta(\sigma)} = \frac{4\delta(\sigma)}{\Theta(\sigma)} \quad \text{and so} \quad \cos \phi \geq 1 - \frac{16\delta^2(\sigma)}{\Theta^2(\sigma)} \quad (7.4)$$

Write $\chi = \angle(T_{y'}, H_x) \leq \angle(T_{y'}, T_{x'})$. By Lemma 7.10, we have

$$\sin \chi \leq \frac{\|x' - y'\|}{\text{rch}(\mathcal{M})}.$$

With $\|y - y'\| \leq \alpha \text{rch}(\mathcal{M})$ and (7.2), we deduce

$$\|z' - z\| = \|y - y'\| \sin t \leq \alpha \|x' - y'\| \leq 2\alpha \|x - y\|. \quad (7.5)$$

Using (7.3), (7.4) and (7.5) and $\alpha = 2\delta^2(\sigma)$, we get

$$\|x' - y'\| \geq \left(1 - 4\delta^2(\sigma) - \frac{16\delta^2(\sigma)}{\Theta^2(\sigma)}\right) \|x - y\| = (1 - \xi) \|x - y\|.$$

From (7.2) and $\alpha = 2\delta^2(\sigma) \leq 1/2$, we also have

$$\|x' - y'\| \leq (1 + 4\delta^2(\sigma)) \|x - y\| \leq (1 + \xi) \|x - y\|.$$

We conclude that

$$(1 - \xi) \|x - y\| \leq \|x' - y'\| \leq (1 + \xi) \|x - y\|.$$

Hence the restriction of Π to σ is a ξ -distortion map. It then follows from Lemma 7.12 that Π embeds σ in $\mathcal{M}$ if $\xi < 1$.

If $\Theta(\sigma) \geq 5\delta(\sigma)$, we have since $\Theta(\sigma) \leq 1$

$$\xi = 4\delta^2(\sigma) + \frac{16\delta^2(\sigma)}{\Theta^2(\sigma)} \leq \frac{20}{25} < 1.$$

□

We introduce now a second map $\Pi_p : \mathbb{R}^d \rightarrow T_p$ that maps points of $\mathbb{R}^d$ to its closest point on the tangent space T_p at p . We now bound the distortion of the restriction of Π_p to a neighbourhood of a p on $\mathcal{M}$ and of the restriction of Π_p to a simplex.

Lemma 7.14 (Distortion of Π_p) *Let $p \in \mathcal{M}$, and write $B(p, r)$ for the ball of $\mathbb{R}^d$ centered at p of radius r and $B_{\mathcal{M}}(p, r)$ for $B(p, r) \cap \mathcal{M}$.*

1. *Let ρ be a positive scalar such that $\rho < \frac{1}{2}$ and let $r = \rho \text{rch}(\mathcal{M})$. The restriction of Π_p to $B_{\mathcal{M}}(p, r)$ is a $4\rho^2$ -distortion map.*

2. Let σ be a simplex incident to p with its vertices on $\mathcal{M}$. If $\Theta(\sigma) > \delta(\sigma)$, then the restriction of Π_p to σ is a $\frac{\delta^2(\sigma)}{\Theta^2(\sigma)}$ -distortion map which embeds σ in T_p .

Proof 1. For all x and $y \in B_{\mathcal{M}}(p, r)$, we have from Lemmas 7.8 and 7.10

$$\sin \angle(xy, T_p) \leq \sin \angle(xy, T_x) + \sin \angle(T_x, T_p) \leq 2\rho < 1. \quad (7.6)$$

Write $x' = \Pi_p(x)$, $y' = \Pi_p(y)$ and $\theta = \angle(xy, T_p)$. Using (7.6), we get

$$\begin{aligned} \|x - y\| - \|x' - y'\| &= \|x - y\| (1 - \cos \theta) \\ &\leq \|x - y\| (1 - \sqrt{1 - 4\rho^2}) \\ &\leq 4\rho^2 \|x - y\|. \end{aligned}$$

Adding the fact that $\|x' - y'\| \leq \|x - y\|$, this shows that the restriction of Π_p to $B_{\mathcal{M}}(p, r)$ is a $4\rho^2$ -distortion map. It then follows from Lemma 7.12 that Π_p embeds $B_{\mathcal{M}}(p, r)$ in T_p . This ends the proof of the first part of the lemma.

2. Let $x, y \in \sigma$, $x' = \Pi_p(x)$, $y' = \Pi_p(y)$, and $\phi = \angle(x - y, x' - y') \leq \angle(\text{aff}(\sigma), T_p)$. By Lemma 7.8, the vertices of σ are at distance at most $h \text{rch}(\mathcal{M})$ from T_p , where $h = \frac{\delta^2(\sigma)}{2}$. It then follows from Lemma 5.14 that

$$\sin \phi \leq \frac{2h}{\delta(\sigma)\Theta(\sigma)} = \frac{\delta(\sigma)}{\Theta(\sigma)}. \quad (7.7)$$

Using (7.7), we get

$$\begin{aligned} \|x' - y'\| &= \|x - y\| \cos \phi \\ &\geq \|x - y\| \sqrt{1 - \frac{\delta^2(\sigma)}{\Theta^2(\sigma)}} \\ &\geq \|x - y\| \left(1 - \frac{\delta^2(\sigma)}{\Theta^2(\sigma)}\right). \end{aligned}$$

We deduce

$$\left(1 - \frac{\delta^2(\sigma)}{\Theta^2(\sigma)}\right) \|x - y\| \leq \|x' - y'\| \leq \|x - y\| \leq \left(1 + \frac{\delta^2(\sigma)}{\Theta^2(\sigma)}\right) \|x - y\|$$

The restriction of Π_p to σ is thus a $\frac{\delta^2(\sigma)}{\Theta^2(\sigma)}$ -distortion map and Lemma 7.12 ensures that Π_p embeds σ in T_p . This completes the proof of the lemma. $\square$

7.3 Triangulation of submanifolds

Definition 7.15 (Triangulation of a topological space) *A triangulation of a topological space X is a simplicial complex K and a homeomorphism $h : |K| \rightarrow X$.*

The following theorem provides sufficient conditions for a simplicial complex $\hat{\mathcal{M}}$ to be a triangulation of a submanifold $\mathcal{M}$ embedded in $\mathbb{R}^d$. This theorem will be used in Chapter 8 to reconstruct submanifolds from finite point sets. We recall that Π denotes the projection onto $\mathcal{M}$ and Π_p the orthogonal projection onto the tangent space T_p at $p \in \mathcal{M}$.

Theorem 7.16 (Triangulation of submanifolds) *Let $\mathcal{M}$ be a closed² k -submanifold of $\mathbb{R}^d$ of positive reach $\text{rch}(\mathcal{M})$ and let $\hat{\mathcal{M}}$ be a combinatorial k -manifold without boundary embedded in $\mathbb{R}^d$ that satisfies the following assumptions :*

- (a) *The vertices of $\hat{\mathcal{M}}$ belong to $\mathcal{M}$ and each connected component of $\mathcal{M}$ contains some vertices of $\hat{\mathcal{M}}$,*
- (b) *The diameter of any k -simplex of $\hat{\mathcal{M}}$ is less than $\Delta_0 = \delta_0 \text{rch}(\mathcal{M})$ where $\delta_0 \leq 1/5$,*
- (c) *The thickness of the k -simplices of $\hat{\mathcal{M}}$ is at least $\Theta_0 \geq 9\delta_0^2/\lambda_0$, where $L_0 = \lambda_0 \text{rch}(\mathcal{M})$ is a lower bound on the edge lengths of the simplices,*
- (d) *For any vertex $p \in \hat{\mathcal{M}}$, Π_p embeds $\text{star}(p, \hat{\mathcal{M}})$.*
- (e) *Let $r = \frac{7}{5} \delta_0 \text{rch}(\mathcal{M})$. Any vertex q of $\hat{\mathcal{M}}$ that belongs to $B(p, r)$ and is mapped by Π_p onto a point $q' \in \Pi_p(\text{star}(p))$ has to be a vertex of $\text{star}(p)$.*

Then the following facts hold :

1. *The restriction of Π to $\hat{\mathcal{M}}$ is a homeomorphism and thus $\hat{\mathcal{M}}$ is a triangulation of $\mathcal{M}$.*
2. *The Hausdorff distance between $\hat{\mathcal{M}}$ and $\mathcal{M}$ is at most $2\delta_0^2 \text{rch}(\mathcal{M})$.*

²i.e. compact and without boundary.

3. If σ is a k -simplex of $\hat{\mathcal{M}}$ and p one of its vertices, we have

$$\sin \angle(\text{aff}(\sigma), T_p) \leq \frac{\delta_0}{\Theta_0}.$$

The proof is given in the following sections. A more general version of the theorem holds. Indeed, we can remove the assumption that the simplicial complex $\hat{\mathcal{M}}$ is embedded in $\mathbb{R}^d$. The fact that the complex is naturally embedded in $\mathbb{R}^d$ is a consequence of this more general version of the theorem (see Exercise 7.12). In the sequel, we will denote by P the set of vertices of $\hat{\mathcal{M}}$.

7.3.1 Proof of statement 1

Write as before Π for the projection onto $\mathcal{M}$ and consider $\Pi|_{\hat{\mathcal{M}}} : \hat{\mathcal{M}} \rightarrow \mathcal{M}$, the restriction of Π to $\hat{\mathcal{M}}$. Note that $\Pi|_{\hat{\mathcal{M}}}$ is well defined since $\hat{\mathcal{M}}$ does not intersect the medial axis of $\mathcal{M}$ as $\delta_0 < 1$. We will prove that $\Pi|_{\hat{\mathcal{M}}}$ is a homeomorphism, which implies that $\hat{\mathcal{M}}$ is a triangulation of $\mathcal{M}$.

Here is an overview of the proof. We have already seen that Π embeds any simplex of $\hat{\mathcal{M}}$ into $\mathcal{M}$ (Lemma 7.13 (2)). We will extend this result and prove that $\Pi|_{\hat{\mathcal{M}}}$ is a local homeomorphism (Lemma 7.23). More specifically, we define an open cover of $\hat{\mathcal{M}}$ as follows. We attach to each vertex p of $\hat{\mathcal{M}}$ an open set V_p such that the union of the V_p , for all vertices p of $\hat{\mathcal{M}}$, covers $\hat{\mathcal{M}}$. Specifically W_p is defined as the set of points of $\text{star}(p) \subset \hat{\mathcal{M}}$ whose barycentric coordinate with respect to p is at least $\frac{1}{k+1}$, and V_p is an open set that contains W_p and is arbitrarily close to W_p . Since the barycentric coordinates in each k -simplex sum to 1, this ensures that the sets V_p cover $\hat{\mathcal{M}}$. We will show that Π embeds each V_p , which will prove that $\Pi|_{\hat{\mathcal{M}}}$ is a local homeomorphism.

We will then prove that $\Pi|_{\hat{\mathcal{M}}}$ is injective (Lemma 7.24) and surjective (Lemma 7.25). Furthermore, $\Pi|_{\hat{\mathcal{M}}}$ is continuous by Lemma 7.13(1) and its inverse is also continuous since $\hat{\mathcal{M}}$ and $\mathcal{M}$ are both k -manifolds without boundary (see Section 1.2.1). It will follow that $\Pi|_{\hat{\mathcal{M}}}$ is a homeomorphism and that $\hat{\mathcal{M}}$ and $\mathcal{M}$ are homeomorphic.




Figure 7.8: For the proof of Lemma 7.23.

7.3.2 Whitney's lemma

We write $\text{star}'(p)$ for $\Pi_p(\text{star}(p))$ and V'_p for $\Pi_p(V_p)$. With a slight abuse of notation, in the rest of the proof of Theorem 7.16, we write Π_p^{-1} for the inverse of the restriction of Π_p to $\text{star}'(p)$, which is well defined since Π_p embeds $\text{star}(p)$ by Hypothesis (d). We further define the map $F_p : \text{star}'(p) \subset T_p \rightarrow T_p$ as $F_p = \Pi_p \circ \Pi \circ \Pi_p^{-1}$. See Figure 7.8. By construction, F_p leaves the vertices of $\text{star}'(p)$ fixed: if q is a vertex of $\text{star}'(p)$, then $F_p(q) = q$. Our goal is to show that F_p embeds $V'_p = \Pi_p(V_p)$. To do so, we will use the following result due to Whitney³ that we state without proof [135, App. II, Lemma 15a].

Definition 7.17 *Let C be a k -simplicial complex embedded in $\mathbb{R}^k$. We say that a map $F : C \rightarrow \mathbb{R}^k$ is simplexwise positive if it is smooth and 1-1 on each k -simplex of C , and if the Jacobian of F is positive there.*

Lemma 7.18 (Whitney's lemma) *Assume C is a combinatorial k -manifold with boundary embedded in $\mathbb{R}^k$. Let $F : C \rightarrow \mathbb{R}^k$ be simplexwise positive in C . Then for any connected open subset Ω of $\mathbb{R}^k \setminus F(\partial C)$, any two points of Ω not in the image of the $(k-1)$ -skeleton of C are covered the same number of times. If this number is 1, then F , considered in the open subset $F^{-1}(\Omega)$ of C only, is one to one onto Ω .*

³Whitney proved a more general result which does not assume that C is embedded in $\mathbb{R}^k$.

7.3.3 Satisfying the conditions of Whitney's lemma

We will apply Whitney's lemma to $C = \text{star}'(p) = \Pi_p(\text{star}(p))$ and to $F = F_p$. This section contains the geometric arguments needed to prove that the conditions of Whitney's lemma are satisfied. After recalling that $\text{star}(p)$ is an oriented combinatorial k -manifold (Point 1), we show that F_p is simplexwise positive (Point 2). We then exhibit a point in $F_p(V_p) \setminus F_p(\partial \text{star}(p))$ that is covered only once (Point 3) and finally prove that $F_p(V_p)$ is included in a single component of $F_p(\text{star}(p)) \setminus F_p(\partial \text{star}(p))$ (Point 4). The proofs use extensively the bounds on the distortion of the two maps Π and Π_p of Section 7.2.

1. Since $\hat{\mathcal{M}}$ is a combinatorial k -manifold and Π_p embeds $\text{star}(p)$ (Hypothesis (d)), $\text{star}'(p)$ is a combinatorial k -manifold embedded in T_p . We can then assume that it is oriented.
2. We now prove that $F_p = \Pi_p \circ \Pi \circ \Pi_p^{-1}$ is simplexwise positive.

Lemma 7.19 (Distortion of F_p on a simplex) *Under the hypotheses of Theorem 7.16, the distortion of the restriction of F_p to any simplex of $\hat{\mathcal{M}}$ is less than $\frac{33\delta_0^2}{\Theta_0^2} \leq 1$.*

Proof Rename, for convenience, Π_p^{-1} as f_1 , the restriction of Π to $\text{star}(p)$ as f_2 , and the restriction of Π_p to $U_p = \Pi(\text{star}(p))$ as f_3 . We have $F_p = f_3 \circ f_2 \circ f_1$. Write ξ_i for the distortion of f_i , $i = 1, 2, 3$, and ξ for the distortion of F_p .

By hypothesis, the diameter of any simplex in $\text{star}(p)$ is at most δ_0 and its thickness is at least Θ_0 . Moreover, U_p is contained in a ball centered at p of radius $r = \delta_0 \text{rch}(\mathcal{M}) (1 + 2\delta_0) < \text{rch}(\mathcal{M})/3$ (Lemma 7.9).

Using the assumptions of Theorem 7.16 and Lemmas 7.14 and 7.13, we get

$$\xi_1 \leq \frac{\frac{\delta_0^2}{\Theta_0^2}}{1 - \frac{\delta_0^2}{\Theta_0^2}} \leq \frac{64}{63} \frac{\delta_0^2}{\Theta_0^2} \quad (\text{Exercise 7.9, Lemma 7.14(2) and } \Theta_0 \geq 8\delta_0)$$

$$\xi_2 \leq 4\delta_0^2 + \frac{16\delta_0^2}{\Theta_0^2} \leq 20 \frac{\delta_0^2}{\Theta_0^2} \quad (\text{Lemma 7.13 (2) and } \Theta_0 \leq 1)$$

$$\xi_3 \leq 4\delta_0^2(1 + 2\delta_0)^2 < 8\delta_0^2 \leq 8 \frac{\delta_0^2}{\Theta_0^2} \quad (\text{Lemma 7.14(1) and } \delta_0 \leq \frac{1}{5})$$

Using Exercise 7.10 and $\Theta_0 > 9\delta_0$, we obtain

$$\begin{aligned}\xi &\leq \xi_1 + \xi_2 + \xi_3 + \xi_1\xi_2 + \xi_2\xi_3 + \xi_3\xi_1 + \xi_1\xi_2\xi_3 \\ &< \frac{33\delta_0^2}{\Theta_0^2} < 1.\end{aligned}$$

□

It follows from Lemma 7.19 that F_p embeds any simplex σ of $\text{star}'(p)$ and that the Jacobian of F_p does not vanish on σ . To prove the last claim, let $\sigma \in \hat{\mathcal{M}}$ and, for $x \in \sigma$ and $v \in \text{aff}(\sigma)$ with $\|v\| = 1$, let $\alpha(t) = x + tv$. Since the restriction of F_p to σ is a ξ -distortion map (Lemma 7.19), and rewriting $F_p = F$ for convenience, we have

$$(1 - \xi)\|(x + tv) - x\| \leq \|F(x + tv) - F(x)\| \leq (1 + \xi)\|(x + tv) - x\|,$$

so for all $t \neq 0$,

$$1 - \xi \leq \frac{\|F(x + tv) - F(x)\|}{|t|} \leq 1 + \xi.$$

Since F is differentiable,

$$\lim_{t \rightarrow 0} \frac{\|F(x + tv) - F(x)\|}{|t|} = \left\| \lim_{t \rightarrow 0} \frac{F(x + tv) - F(x)}{t} \right\| = \|dF_x(v)\|,$$

so

$$1 - \xi \leq \|dF_x(v)\| \leq 1 + \xi,$$

which yields the claimed result.

We have shown that the Jacobian of F_p does not vanish on σ . It remains to prove that it is positive. To do so, observe that the differential of the restriction of F_p to σ at the point $p \in \sigma$ is the identity. This follows from the observation that the differential of Π at p restricted to T_p is the identity, since p lies on the manifold.

3. We will now prove that there is a point in $F_p(V_p) \setminus F_p(\partial \text{star}(p))$ and not in the image of the $(k - 1)$ -skeleton that is covered once. For this purpose, we choose the image of the barycenter of a k -simplex of $\text{star}'(p)$.

Lemma 7.20 *Assume that the conditions of Theorem 7.16 hold and let b' be the barycenter of a k -simplex σ in $\text{star}'(p)$. Then the image of b' , $F_p(b')$, is covered once, i.e. $F_p^{-1}(F_p(b')) = \{b'\}$.*

Proof We denote as usual by $\Delta(\sigma) = \delta(\sigma) \operatorname{rch}(\mathcal{M})$ the diameter of σ and by $D(p, \sigma)$ the altitude of p in σ . The distance from the barycenter b' to the boundary of σ is

$$\min_{p \in \operatorname{vert}(\sigma)} \frac{D(p, \sigma)}{k+1} = \frac{k\Theta(\sigma)\Delta(\sigma)}{k+1} \geq \frac{\Theta_0\Delta_0}{2} \quad (7.8)$$

using $k \geq 1$. Since the distortion of the restriction of F_p to σ is less than 1 (Lemma 7.19), it is an embedding (Lemma 7.12) and therefore no point $z' \in \sigma$ can verify $F_p(z') = F_p(b')$. Let us consider now the case of a point z' of $\operatorname{star}'(p) \setminus \sigma$. Write $b = \Pi_p^{-1}(b')$ and $z = \Pi_p^{-1}(z')$, and observe that

$$\|F_p(b') - b'\| = \|\Pi_p(\Pi(b) - b)\| \leq 2\delta_0^2 \operatorname{rch}(\mathcal{M}) \quad (\text{By Lemma 7.9})$$

and similarly for $\|F_p(z') - z'\|$.

Hence,

$$\begin{aligned} \|F_p(b') - F_p(z')\| &\geq \|b' - z'\| - \|F_p(b') - b'\| - \|F_p(z') - z'\| \\ &\geq d(b', \partial\sigma) - 4\delta_0^2 \operatorname{rch}(\mathcal{M}) \\ &\geq \frac{\Theta_0 \delta_0 \operatorname{rch}(\mathcal{M})}{2} - 4\delta_0^2 \operatorname{rch}(\mathcal{M}) \quad (\text{By Equation 7.8}) \end{aligned}$$

It follows that $F_p(b') \neq F_p(z')$ since we assumed $\Theta_0 > 8\delta_0$. $\square$

Since $b' \in V'_p$, the previous lemma implies that there exists a point in $F_p(V'_p)$ that is covered exactly once.

4. The next lemma shows that $F_p(V'_p) \cap F_p(\partial \operatorname{star}'(p)) = \emptyset$.

Lemma 7.21 *Assume that the conditions of Theorem 7.16 hold and let $x \in \sigma \subset \operatorname{star}'(p)$ and $y \in T_p \setminus \operatorname{star}'(p)$. If the barycentric coordinate of x with respect to p is at least $\frac{1}{k+1}$, then $F_p(x) \neq F_p(y)$.*

Proof Arguing as in the proof of Lemma 7.20, we get

$$\begin{aligned} \|F_p(x) - F_p(y)\| &\geq \|x - y\| - \|F_p(x) - x\| - \|F_p(y) - y\| \\ &\geq d(x, \partial \operatorname{star}'(p)) - 4\delta_0^2 \operatorname{rch}(\mathcal{M}) \\ &\geq \frac{k}{k+1} \lambda_0 \Theta_0 \operatorname{rch}(\mathcal{M}) - 4\delta_0^2 \operatorname{rch}(\mathcal{M}) \quad (\text{By Lemma 7.22}) \\ &\geq \frac{1}{2} \lambda_0 \Theta_0 \operatorname{rch}(\mathcal{M}) - 4\delta_0^2 \operatorname{rch}(\mathcal{M}) \quad (k > 1) \end{aligned}$$

Hence, $F_p(x) \neq F_p(y)$ if $\Theta_0 > 8\delta_0^2/\lambda_0$. $\square$

Lemma 7.22 *Let σ be a k -simplex of $\text{star}'(p)$ and let x be a point of σ whose barycentric coordinate that is associated to p in σ satisfies $\lambda_p(x) \geq \alpha$. Then $d(x, \partial\text{star}'(p)) \geq \alpha k L_0 \Theta_0$, where $L_0 = \lambda_0 \text{rch}(\mathcal{M})$ is a lower bound on the edge lengths of the simplices and Θ_0 is a lower bound on their thickness.*

Proof Let γ_σ be a line segment joining two points x_σ and y_σ of σ and denote by $\lambda_p(x_\sigma)$ and $\lambda_p(y_\sigma)$ respectively their p -th barycentric coordinates. Then $|\gamma_\sigma| \geq |\lambda_p(x_\sigma) - \lambda_p(y_\sigma)| D(p, \sigma)$, where $D(p, \sigma)$ denotes the altitude of p in σ and we indicate the length of a curve by $|\cdot|$.

Let $\text{star}'_\alpha(p) = \{z \in \text{star}(p), \lambda_p(z) \geq \alpha\}$. Equivalently, $\text{star}'_\alpha(p)$ is the image of $\text{star}'(p)$ by the homothety of center p and ratio $1 - \alpha$. Consider a shortest line segment γ connecting a point of $\partial\text{star}'(p)$ to a point of $\partial\text{star}'_\alpha(p)$. Plainly, γ intersects $\text{star}'(p)$ and $\text{star}'_\alpha(p)$ only at its endpoints and the intersection of γ with any simplex σ of $\text{star}'(p)$, if non empty, is a line segment we denote by γ_σ . If we now note that the barycentric coordinates coincide on a face that is shared by two simplices we see that:

$$\begin{aligned} |\gamma| = \sum_{\sigma} |\gamma_\sigma| &\geq \sum_{\sigma} |\lambda_p(x_\sigma) - \lambda_p(y_\sigma)| D(p, \sigma) \\ &\geq \sum_{\sigma} |\lambda_p(x_\sigma) - \lambda_p(y_\sigma)| k \Delta(\sigma) \Theta_0 \\ &\geq \alpha k L_0 \Theta_0 \end{aligned}$$

where x_σ is the point where γ enters σ and y_σ where it leaves. This yields $d_H(\partial\text{star}'(p), \partial\text{star}'_\alpha(p)) \geq \alpha k L_0 \Theta_0$, where d_H denotes the Hausdorff distance. $\square$

7.3.4 Local homeomorphism

Lemma 7.23 ($\Pi|_{\hat{\mathcal{M}}}$ is a local homeomorphism) *Assume that the conditions of Theorem 7.16 and that the conditions of Whitney's lemma are fulfilled. Then Π embeds V_p in $\mathcal{M}$. The restriction of Π to $\hat{\mathcal{M}}$ is thus a local homeomorphism.*

Proof Since the conditions of Whitney's lemma are fulfilled, F_p embeds V'_p in T_p . Consequently, since $F_p = \Pi_p \circ \Pi \circ \Pi_p^{-1}$, Π must embed V_p into $\mathcal{M}$. Moreover, since the set of all V_p for all vertices of $\hat{\mathcal{M}}$ cover $\hat{\mathcal{M}}$, we have proved that the map Π restricted to $\hat{\mathcal{M}}$ is a local homeomorphism on $\mathcal{M}$. This completes the proof of Lemma 7.23. $\square$

7.3.5 The restriction of Π to $\hat{\mathcal{M}}$ is injective

Lemma 7.24 (Injectivity of $\Pi|_{\hat{\mathcal{M}}}$) *Under the conditions of Theorem 7.16, Π is injective on $\hat{\mathcal{M}}$.*

Proof Let x be a point of $\Pi(\hat{\mathcal{M}})$. $\Pi|_{\hat{\mathcal{M}}}^{-1}(x)$ is non empty and finite since Π embeds each simplex of $\hat{\mathcal{M}}$ (Lemma 7.13) and there are only finitely many simplexes in $\hat{\mathcal{M}}$. For each point $y \in \Pi|_{\hat{\mathcal{M}}}^{-1}(x) \cap \hat{\mathcal{M}}$, we choose a sufficiently small open neighborhood $U(y)$ of y such that $U(y)$ is homeomorphic to a k -ball and $U(y)$ is contained in some V_p . This is possible since the V_p , $p \in P$, constitute an open cover of $\hat{\mathcal{M}}$. Since, as already noticed, Π embeds any V_p on $\mathcal{M}$, it also embeds each $U(y)$ in $\mathcal{M}$. Hence the preimage under $\Pi|_{\hat{\mathcal{M}}}$ of any sufficiently small open neighborhood $U(x)$ of a point x of $\mathcal{M}$ is a union of disjoint open sets, each of which is contained in some $V_p \subset \hat{\mathcal{M}}$ and mapped homeomorphically onto $U(x)$ by Π . (In topological terms, $(\hat{\mathcal{M}}, \Pi)$ is a covering space of $\Pi(\hat{\mathcal{M}})$.)

We now show that the cardinality of $\Pi|_{\hat{\mathcal{M}}}^{-1}(x)$ is constant over each connected component of $\mathcal{M}$. Indeed, consider the function G that associates to a point x of $\mathcal{M}$ the cardinality of $\Pi|_{\hat{\mathcal{M}}}^{-1}(x)$. From the above discussion, G is locally constant, which implies that it is constant on each connected component of $\mathcal{M}$. Hence $\Pi|_{\hat{\mathcal{M}}}$ covers all the points of a connected component of its image $\Pi(\hat{\mathcal{M}})$ the same number of times. This number is 1 as is shown next.

Assume, for a contradiction, that there exists a vertex q and a simplex σ of $\hat{\mathcal{M}}$ such that q belongs to the image $\Pi(\sigma)$ but q is not a vertex of σ . Specifically, let $x \in \sigma$ be such that $\Pi(x) = q$. Recall that $F_p = \Pi_p \circ \Pi \circ \Pi_p^{-1}$ and observe that $F_p(x) = F_p(q)$ for any vertex p of σ . In the rest of the proof, we denote by p a vertex of σ such that the associated barycentric coordinate of x satisfies $\lambda_p(x) \geq \frac{1}{k+1}$ (which must exist).

We write for convenience $x' = \Pi_p(x)$ and $q' = \Pi_p(q)$. Plainly, x' belongs to $\Pi_p(\sigma)$ but q' does not belong to $\text{star}'(p) = \Pi_p(\text{star}(p))$ by Hypothesis (e). We conclude that $F_p(x') \neq F_p(q')$ from Lemma 7.21. $\square$

7.3.6 The restriction of Π to $\hat{\mathcal{M}}$ is surjective

Lastly we prove that the restriction of Π to $\hat{\mathcal{M}}$ is surjective.

Lemma 7.25 (Surjectivity of $\Pi|_{\hat{\mathcal{M}}}$) *Under the conditions of Theorem 7.16, $\Pi|_{\hat{\mathcal{M}}}$ is surjective on $\mathcal{M}$.*

Proof Since $\hat{\mathcal{M}}$ has no boundary, the same is true for $\Pi(\hat{\mathcal{M}})$ by the invariance of domain theorem,⁴. Hence any connected component of $\mathcal{M}$ that intersects $\Pi(\hat{\mathcal{M}})$ has to be totally included in $\Pi(\hat{\mathcal{M}})$. Since by Hypothesis 1 of Theorem 7.16, every connected component of $\mathcal{M}$ intersects $\Pi(\hat{\mathcal{M}})$, $\mathcal{M}$ is included in $\Pi(\hat{\mathcal{M}})$. $\square$

7.3.7 End of proof of statement 1

We have shown that $\Pi|_{\hat{\mathcal{M}}}$ is a bijection from $\hat{\mathcal{M}}$ to $\mathcal{M}$. Furthermore, $\Pi|_{\hat{\mathcal{M}}}$ is continuous by Lemma 7.13(1) and its inverse is also continuous since $\hat{\mathcal{M}}$ is compact and $\mathcal{M}$ is a metric space (see Section 1.2.1). It follows that $\Pi|_{\hat{\mathcal{M}}}$ is a homeomorphism and that $\hat{\mathcal{M}}$ and $\mathcal{M}$ are homeomorphic, which is the first statement in Theorem 7.16.

7.3.8 Proof of statements 2 and 3, and of Theorem 7.16

To prove the second statement, let x be a point of $\mathcal{M}$ and let y be the point of $\hat{\mathcal{M}}$ such that $y = \Pi(x)$. Such a point y exists by the first statement. It then follows from Lemma 7.9 that $\|x - y\| \leq 2\delta_0^2 \text{rch}(\mathcal{M})$.

The third statement of the theorem follows from Lemma 5.14 as in the proof of Lemma 7.14 (Equation (7.7)).

⁴ The invariance of domain theorem states that, given an open subset $U \subseteq \mathbb{R}^d$ and an injective and continuous function $f : U \rightarrow \mathbb{R}^d$ then f is a homeomorphism between U and $f(U)$. The theorem and its proof are due to L. E. J. Brouwer [113, Theorem 36.5].

We have thus proved the three statements of Theorem 7.16.

7.4 Exercises

Exercise 7.1 (ε -nets on submanifolds) Let P be an ε -dense sample of a submanifold $\mathcal{M}$. Let $P' := \emptyset$ and apply the following procedure : while there exists a point p of $P \setminus P'$ whose distance to the current set P' is greater than $\varepsilon \text{rch}(\mathcal{M})$, insert p in P' . Show that P' is an $(\varepsilon \text{rch}(\mathcal{M}), 1)$ -net of P and a $(2\varepsilon, \frac{1}{2})$ -net of $\mathcal{M}$. Adapt the algorithms of Section 5.1 to the case of submanifolds.

Exercise 7.2 (Size of nets) Show that the size of an $(\varepsilon, \bar{\eta})$ -net of a k -submanifold with a bounded sampling ratio $\bar{\eta}$ depends exponentially on k .

Exercise 7.3 (Computing $\text{rch}(\mathcal{M})$) Let $\mathcal{M}$ be a smooth submanifold and write $\text{ax}(\mathcal{M})$ for the medial axis of $\mathcal{M}$. Show that $\text{rch}(\mathcal{M})$ is the radius of a ball $B(m)$ centered at a point $m \in \text{ax}(\mathcal{M})$ and tangent to $\mathcal{M}$. Show that if $B(m)$ has only one contact point, $B(m)$ is osculating $\mathcal{M}$ at the contact point. If $B(m)$ has two distinct contact points, the two contact points are the endpoints of a diameter of $B(m)$. Propose an algorithm to compute $\text{rch}(\mathcal{M})$ when $\mathcal{M}$ is a hypersurface of $\mathbb{R}^d$ implicitly defined as $f(x) = 0$ where f is a differentiable function defined over $\mathbb{R}^d$ for which 0 is a regular value.

Exercise 7.4 (Non uniform nets) We define the local feature size $\text{lfs}(x)$ of a point x of a submanifold $\mathcal{M}$ as the distance from x to the medial axis of $\mathcal{M}$, $\text{ax}(\mathcal{M})$. We say that a finite set of points $P \subset \mathcal{M}$ is a *non uniform* $(\varepsilon, \bar{\eta})$ -net if it satisfies the following two properties:

1. Any point x of $\mathcal{M}$ is at distance at most $\varepsilon \text{lfs}(x)$ from a point of P .
2. For any two points p, q of P , $\|p - q\| \geq \eta \min(\text{lfs}(p), \text{lfs}(q))$.

Show that lfs is a 1-Lipschitz function, and, using this fact, extend the results of this chapter to sufficiently dense non uniform ε -nets.

Exercise 7.5 (Size of non-uniform nets) Provide upper and lower bounds on the size of a non-uniform $(\varepsilon, \bar{\eta})$ -net of a k -submanifold as a function of k and the integral over $\mathcal{M}$ of $\text{lfs}(x)$, the local feature size of $\mathcal{M}$ at x .

Exercise 7.6 (Geodesic distance) Let any two points x and y be two points on $\mathcal{M}$ such that $\|x - y\| \leq \alpha \text{rch}(\mathcal{M})$, $\alpha < 1$. Show that the length of a shortest path on $\mathcal{M}$ joining x to y is at most $\frac{\|x-y\|}{1-\alpha}$.

Exercise 7.7 (Angle between tangent spaces) Prove Lemma 7.10 (with possibly a bigger constant) using elementary arguments.

Exercise 7.8 Let x be a point of $B_{\mathcal{M}} = B(p, r) \cap \mathcal{M}$ where $r = \rho \text{rch}(\mathcal{M})$, $\rho < 1$. Show that there exists a point $y \in T_p$ such that $\Pi(y) = x$, where Π denotes the projection on $\mathcal{M}$, and $\|x - y\| \leq \frac{r}{1-\rho}$.

Exercise 7.9 (Distortion) Let $F : U \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a ξ -distortion map. Show that F^{-1} is well defined and that it is a $\frac{\xi}{1-\xi}$ -distortion map.

Exercise 7.10 If F_1 is a ξ_1 -distortion map and F_2 is a ξ_2 -distortion map, then $F_1 \circ F_2$ is a $(\xi_1 + \xi_2 + \xi_1 \xi_2)$ -distortion map.

Exercise 7.11 (Ambient isotopy) Assume that the conditions of Theorem 7.16 are satisfied. Show that the restriction of Π to $\hat{\mathcal{M}}$ induces an ambient isotopy

$$\Phi^* : \mathbb{R}^d \times [0, 1] \longrightarrow \mathbb{R}^d$$

such that the map $\Phi^*(\cdot, 0)$ restricted to $\hat{\mathcal{M}}$ is the identity map on $\hat{\mathcal{M}}$ and $\Phi^*(\hat{\mathcal{M}}, 1) = \mathcal{M}$. The isotopy does not move the points by more than $O(\delta_0^2 \text{rch}(\mathcal{M}))$.

Hint: Let

$$\Phi : \hat{\mathcal{M}} \times [0, 1] \longrightarrow \mathbb{R}^d, \quad (x, t) \mapsto x + t(\Pi(x) - x)$$

Note that $\Phi(\cdot, 0)$ is an identity map on $\hat{\mathcal{M}}$ and $\Phi(\cdot, 1) = \Pi|_{\hat{\mathcal{M}}}$. The map Φ is an isotopy because the maps

$$\Phi_t : \hat{\mathcal{M}} \longrightarrow \mathbb{R}^d, \quad x \mapsto \Phi(x, t)$$

are homeomorphisms between $\hat{\mathcal{M}}$ and $\Phi_t(\hat{\mathcal{M}})$.

Isotopy Φ can be extended to an ambient isotopy $\Phi^* : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ such that $\Phi^*(\cdot, 0)|_{\hat{\mathcal{M}}} = \Phi(\cdot, 0)$ and $\Phi^*(\cdot, 1)|_{\hat{\mathcal{M}}} = \Phi(\cdot, 1)$.

Exercise 7.12 (Extension of Theorem 7.16) Remove in Theorem 7.16 the assumption that the simplicial complex is embedded in $\mathbb{R}^d$ and assume instead that we are given an abstract finite simplicial complex K without boundary. We then consider the piecewise linear map $\iota : K \hookrightarrow \mathbb{R}^d$ defined by the coordinates of the vertices of K and define $\hat{\mathcal{M}}$ as $\iota(K)$. Using essentially the same proof as in Sections 7.3.1-7.3.8, prove that $\Pi \circ \iota(K)$ is homeomorphic to the underlying space of $|K|$ and that $\hat{\mathcal{M}}$ is embedded in $\mathbb{R}^d$.

7.5 Bibliographical notes

The medial axis was introduced as a tool in image analysis. It is by now widely used to represent and analyze shapes. For mathematical and algorithmic properties, we refer the reader to the work of Lieutier [104] and to the survey paper [3]. Lieutier proved that any bounded open subset of $\mathbb{R}^d$ is homotopy equivalent to its medial axis [104]. The notion of local feature size (Exercise 7.4) was introduced by Amenta and Bern in their seminal paper on surface reconstruction [2]. See also [25, 91]. The related notion of reach has been introduced earlier by Federer [82] who proved Lemmas 7.2, 7.8 and the first part of Lemma 7.13. Many other results on sets of positive reach can also be found in this paper. The link between the reach and the principal curvatures of a submanifold has been established by Niyogi, Smale and Weinberger [115] who also proved a weaker form of Lemma 7.10. The improved bound is due to M. Wintraecken [136].

Proving that any smooth manifold can be triangulated has been the subject of many important developments in the mathematical community by Cairns [32], Whitehead [134], Whitney [135], Munkres [112] and others. Our proof uses Lemma 7.18 which is due to Whitney. More recently, the Delaunay triangulation turned out to be a useful tool in this context leading to efficient algorithms. Variants of Theorem 7.16 have been proved by Cheng and al. [57], Boissonnat and Ghosh [22] and Dyer et al. [66]. The proof presented here is based on M. Wintraecken's Ph.D. thesis and on unpublished notes by R. Dyer. A proof of Exercise 7.11 can be found in [98].

Chapter 8

Reconstruction of submanifolds

Let $\mathcal{M}$ be an unknown manifold $\mathcal{M}$ and let $P \subset \mathcal{M}$ be a known finite sample. The reconstruction problem is to recover from P alone the topological type of $\mathcal{M}$ (manifold reconstruction) or the homotopy type of $\mathcal{M}$ (homotopy reconstruction). More concretely, we look for a simplicial complex $\hat{\mathcal{M}}$ which approximates $\mathcal{M}$, and, in particular, is provably homeomorphic or homotopy equivalent to $\mathcal{M}$.

In Section 8.1, we will see that, under appropriate conditions on P , $\mathcal{M}$ has the same homotopy type as the union of balls centered on the points of P . It follows that the alpha-shape of $\mathcal{M}$, for an appropriate value of α has the same homotopy type as $\mathcal{M}$. This method however has some limitations. The complex captures the homotopy type but is not in general a triangulation of $\mathcal{M}$ and computing the alpha-shape is limited to low dimensional spaces.

In the following sections, we overcome these limitations and show how to reconstruct a simplicial complex with vertex set P that is homeomorphic to $\mathcal{M}$. We have to face two main difficulties. First, even when a point set is a dense and sparse sampling on a manifold $\mathcal{M}$, the Delaunay simplexes of dimension higher than 2 are not guaranteed to be thick (see Section 5.2) and therefore not guaranteed to approximate the tangent bundle of $\mathcal{M}$ (see Lemma 5.14). This is an issue since the main theorem of Chapter 7, Theorem 7.16, that provides conditions to triangulate a submanifold, requires simplices to be thick. To be able to reconstruct manifolds of dimension greater than 2, we will need to explicitly take care of non-thick simplices using techniques similar to what has been done in Chapter 5 in the Euclidean case.

The second major difficulty comes from the so-called curse of dimensionality. We have seen (Theorems 4.4 and 5.4) that the size of the Delaunay triangulation of n points grows exponentially with the dimension d of the embedding space. As a consequence, when d is large, we cannot afford to compute the d -dimensional Delaunay triangulation $\text{Del}(P)$ or any other subdivision of $\mathbb{R}^d$. Instead, we will introduce a subcomplex of $\text{Del}(P)$, called the tangential Delaunay complex, whose complexity depends on the intrinsic dimension of $\mathcal{M}$ and not on the ambient dimension d . This complex is defined locally and the various local triangulations are glued together so as to constitute a manifold complex that is embedded in $\mathbb{R}^d$ and triangulates $\mathcal{M}$.

8.1 Homotopy reconstruction using alpha-shapes

Unions of balls play a central role in manifold reconstruction. Indeed, while Chapter 6 shows how to capture the homotopy type of a union of balls with a simplicial complex, this section shows that the homotopy type of a sampled manifold can be obtained from the homotopy type of a union of balls. More precisely, we show here that the homotopy type, hence the homology groups, of a manifold $\mathcal{M}$ with positive reach can be obtained from union of balls centered at the points of a sample of $\mathcal{M}$ provided that the sample is dense enough with respect to the reach.

Theorem 8.1 *Let $\mathcal{M}$ be a manifold with positive reach $\text{rch}(\mathcal{M})$, and $P \subset \mathcal{M}$ be a point sample of $\mathcal{M}$ with sampling radius $\frac{\varepsilon}{2}\text{rch}(\mathcal{M})$, meaning that any point x of $\mathcal{M}$ is at distance less than $\frac{\varepsilon}{2}\text{rch}(\mathcal{M})$ from the closest sample point. If $\varepsilon < \sqrt{\frac{3}{5}}$, the union of balls with radius $\varepsilon \text{rch}(\mathcal{M})$ centered on P , is homotopy equivalent to $\mathcal{M}$.*

Proof In the following, we denote by $B(p, r)$ the ball with radius r centered on p and write for short U for the union of balls with radius $\varepsilon \text{rch}(\mathcal{M})$ centered on points of P :

$$U = \bigcup_{p \in P} B(p, \varepsilon \text{rch}(\mathcal{M})).$$

Obviously, $\mathcal{M}$ is included in U . To prove the homotopy equivalence we prove below that U deformation retracts to $\mathcal{M}$. For all $x \in U$ and $t \in [0, 1]$, we define

$$F(x, t) = (1 - t)x + t\Pi(x),$$

where $\Pi(x)$ is the projection on $\mathcal{M}$. F is continuous from Lemma 7.13. For all $x \in U$, $F(x, 0) = x$, and $F(x, 1) = \Pi(x)$ is in $\mathcal{M}$, and for all $x \in \mathcal{M}$ and $t \in [0, 1]$, $F(x, t) = x$. Therefore F is a deformation retracts from U to $\mathcal{M}$, provided that $F(x, t)$ belongs to U for any $(x, t) \in U \times [0, 1]$, which is proved now.

Let us consider Π_U , the restriction of Π to U . The preimage $\Pi_U^{-1}(y)$ of a point $y \in \mathcal{M}$ is

$$\Pi_U^{-1}(y) = N_y \cap U \cap B(y, \text{rch}(\mathcal{M})), \quad (8.1)$$

where N_y is the normal subspace of $\mathcal{M}$ at y . The ball with radius $\text{rch}(\mathcal{M})$ centered at y , $B(y, \text{rch}(\mathcal{M}))$, appears in Equation 8.1 to remove orphan

components of $N_y \cap U$, i. e. components which do not contain y , and may arise from the fact that $\mathcal{M}$ is curved. Therefore,

$$\Pi_U^{-1}(y) = \bigcup_{p \in P} B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap B(y, \text{rch}(\mathcal{M})).$$

We also consider the subset $st(y)$ defined as

$$st(y) = \bigcup_{p \in P \cap B(y, \varepsilon \text{rch}(\mathcal{M}))} B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap B(y, \text{rch}(\mathcal{M})).$$

Obviously, $st(y) \subset \Pi_U^{-1}(y)$. Then Lemma 8.2 below proves that $st(y)$ is star shaped with respect to y and Lemma 8.3 proves that $st(y) = \Pi_U^{-1}(y)$. It follows that for any $y \in \mathcal{M}$, $\Pi_U^{-1}(y)$ is star shaped with respect to y and that for any $(x, t) \in U \times [0, 1]$, $F(x, t)$ belongs to U . $\square$

Lemma 8.2 *The subset $st(y)$ is star shaped with respect to y .*

Proof Let z be an arbitrary point in $st(y)$. Then $z \in B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap B(y, \text{rch}(\mathcal{M}))$ for some $p \in P \cap B(y, \varepsilon \text{rch}(\mathcal{M}))$. Since $p \in B(y, \varepsilon \text{rch}(\mathcal{M}))$, $y \in B(p, \varepsilon \text{rch}(\mathcal{M}))$. Since z and y are both in $B(p, \varepsilon \text{rch}(\mathcal{M}))$, the segment zy is entirely contained in $B(p, \varepsilon \text{rch}(\mathcal{M}))$. At the same time, zy is entirely contained in N_y and in $B(y, \text{rch}(\mathcal{M}))$ and therefore in $st(y)$. $\square$

Lemma 8.3 *The subset $st(y)$ coincides with the preimage $\Pi_U^{-1}(y)$.*

Proof We are left to show that $\Pi_U^{-1}(y) \subset st(y)$. Let z be a point in $B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap B(y, \text{rch}(\mathcal{M}))$ where p is a point of P such that $p \notin B(y, \varepsilon \text{rch}(\mathcal{M}))$. Lemma 8.4 shows that the distance from z to y is at most $\varepsilon^2 \text{rch}(\mathcal{M})$ and Lemma 8.5 shows that if P is $(\frac{\varepsilon}{2} \text{rch}(\mathcal{M}))$ -dense in $\mathcal{M}$ with $\varepsilon < \frac{3}{5}$, there is some point $q \in P \cap B(y, \varepsilon \text{rch}(\mathcal{M}))$ such that $z \in B(q, \varepsilon \text{rch}(\mathcal{M}))$ which achieves the proof. $\square$

Lemma 8.4 *Let z be a point in $B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap B(y, \text{rch}(\mathcal{M}))$ where p is a point of P such that $p \notin B(y, \varepsilon \text{rch}(\mathcal{M}))$. The distance $d(y, z)$ from z to y is at most $\varepsilon^2 \text{rch}(\mathcal{M})$.*




Figure 8.1: For the proof of Lemma 8.3.

Proof Let us consider the plane H through y, z and p . See Figure 8.1. In the plane H , we call n_y the line $H \cap N_y$, and h_y the line orthogonal to n_y that passes through y . Point z belongs to n_y and point p lies anywhere outside of the two balls with radius $\text{rch}(\mathcal{M})$ tangent to h_y at y . Since the distance $d(y, z)$ is constrained by the fact that $d(p, z) \leq \varepsilon \text{rch}(\mathcal{M})$, the maximum of the distance $d(y, z)$ occurs in the configuration where points p and z lie on the same side of h_y , as shown in Figure 8.1. We denote by r_p the radius of the circle tangent to $\mathcal{M}$ in y and going through p , and by θ_p the angle $\angle(h_y, yp)$ between h_y and yp .

We have $r_p \geq \text{rch}(\mathcal{M})$ and $d(y, p) = 2r_p \sin \theta_p$. Furthermore we have (see Figure 8.1) :

$$\begin{aligned}
 d(y, z) &= d(y, p) \sin \theta_p + \sqrt{d(z, p)^2 - d(y, p)^2 \cos^2 \theta_p} \\
 &= 2r_p \sin^2 \theta_p + \sqrt{d(z, p)^2 - r_p^2 \sin^2 2\theta_p} \\
 &\leq 2r_p \sin^2 \theta_p + \sqrt{\varepsilon^2 \text{rch}(\mathcal{M})^2 - r_p^2 \sin^2 2\theta_p} \\
 &\stackrel{\text{def}}{=} f(r_p, \theta_p)
 \end{aligned} \tag{8.2}$$




Figure 8.2: For the proof of Lemma 8.4.

We have:

$$\begin{aligned} \frac{df}{d\theta_p} &= 2r_p \sin 2\theta_p - \frac{2r_p^2 \sin 2\theta_p \cos 2\theta_p}{\sqrt{\varepsilon^2 \text{rch}(\mathcal{M})^2 - r_p^2 \sin^2 2\theta_p}} \\ &= 2r_p \sin 2\theta_p \left(1 - \frac{r_p \cos 2\theta_p}{\sqrt{\varepsilon^2 \text{rch}(\mathcal{M})^2 - r_p^2 \sin^2 2\theta_p}} \right). \end{aligned}$$

Since $\varepsilon < \sqrt{\frac{3}{5}} < 1$ and $\text{rch}(\mathcal{M}) \leq r_p$, the function f is monotonically decreasing with respect to θ_p and reaches its maximum when θ_p is minimum, i.e. when $d(y, p) = 2r_p \sin \theta_p = \varepsilon \text{rch}(\mathcal{M})$. We have then

$$\begin{aligned} f(r_p, \theta_p) &= \varepsilon \text{rch}(\mathcal{M}) \sin \theta_p + \sqrt{\varepsilon^2 \text{rch}(\mathcal{M})^2 - \varepsilon^2 \text{rch}(\mathcal{M})^2 \cos^2 \theta_p} \\ &= 2\varepsilon \text{rch}(\mathcal{M}) \sin \theta_p = \frac{\varepsilon^2 \text{rch}(\mathcal{M})^2}{r_p} \leq \varepsilon^2 \text{rch}(\mathcal{M}), \end{aligned}$$

which, together with Equation 8.2, achieves the proof. $\square$

Lemma 8.5 *Let z be a point in $B(p, \varepsilon \text{rch}(\mathcal{M})) \cap N_y \cap b_{\text{rch}(\mathcal{M})}(y)$ where p is a point of P such that $p \notin B(y, \varepsilon \text{rch}(\mathcal{M}))$. If P is $(\frac{\varepsilon}{2} \text{rch}(\mathcal{M}))$ -dense in $\mathcal{M}$ with $\varepsilon < \sqrt{\frac{3}{5}} \text{rch}(\mathcal{M})$, there is some point $q \in P \cap b_\varepsilon(y)$ such that $z \in b_\varepsilon(q) \cap N_y$.*

Proof Because P is $(\frac{\varepsilon}{2}\text{rch}(\mathcal{M}))$ -dense in $\mathcal{M}$, there is a point $q \in P$ at distance at most $\frac{\varepsilon}{2}\text{rch}(\mathcal{M})$ from y . We prove that $z \in B(q, \varepsilon)$. We consider now the plane H' through y, z and q and write n'_y for the line $H' \cap N_y$, and h'_y for the line orthogonal to H' that passes through y . Figure 8.2, drawn in the plane H' , shows the worst situation for the distance $d(z, q)$. Denoting now by θ_q the angle $\angle(h'_y, yq)$ between h'_y and yq , we have

$$\begin{aligned} d(z, q)^2 &\leq d(y, q)^2 \cos^2 \theta_q + (d(y, q) \sin \theta_q + d(y, z))^2 \\ &\leq d(y, q)^2 + 2d(y, q)d(y, z) \sin \theta_q + d(y, z)^2. \end{aligned}$$

Since $d(y, q) \leq \frac{\varepsilon}{2}\text{rch}(\mathcal{M})$ and $d(y, z) \leq \varepsilon^2\text{rch}(\mathcal{M})$ (Lemma 8.4), we get

$$d(z, q)^2 \leq \frac{\varepsilon^2}{4}\text{rch}(\mathcal{M})^2 + \varepsilon^3\text{rch}(\mathcal{M})^2 \sin \theta_q + \varepsilon^4\text{rch}(\mathcal{M})^2.$$

We have $d(y, q) = 2r_q \sin \theta_q \leq \frac{\varepsilon}{2}\text{rch}(\mathcal{M})$ where r_q is the radius of the circle tangent to $\mathcal{M}$ at y and going through q . Since point q lies outside the two balls with radii $\text{rch}(\mathcal{M})$ that are tangent to T_y at y , r_q is greater than $\text{rch}(\mathcal{M})$ and therefore $\sin \theta_q \leq \frac{\varepsilon}{4}$. Thus,

$$\begin{aligned} d(z, q)^2 &\leq \frac{\varepsilon^2}{4}\text{rch}(\mathcal{M})^2 + \frac{\varepsilon^4}{4}\text{rch}(\mathcal{M})^2 + \varepsilon^4\text{rch}(\mathcal{M})^2 \\ &\leq \varepsilon^2\text{rch}(\mathcal{M})^2 \left(\frac{1}{4} + \frac{5}{4}\varepsilon^2 \right), \end{aligned}$$

which is not greater than $\varepsilon^2\text{rch}(\mathcal{M})^2$ if $\varepsilon^2 \leq \frac{3}{5}$. $\square$

8.2 Tangential Delaunay complex

We introduce in this section a data structure, named the tangential Delaunay complex. Let $\mathcal{M}$ be a k -submanifold of $\mathbb{R}^d$. The only knowledge we have on $\mathcal{M}$ is a finite set of points $P \in \mathcal{M}$ and the tangent spaces at each point of P . The tangential Delaunay complex $\text{Del}_{T\mathcal{M}}(P)$ is a k -dimensional subcomplex of the d -dimensional Delaunay complex $\text{Del}(P)$. An important property is that $\text{Del}_{T\mathcal{M}}(P)$ can be constructed without computing any data structure of dimension higher than k , and in particular without computing the full Delaunay complex. We will see in Section 8.3 that $\text{Del}_{T\mathcal{M}}(P)$ can be used to reconstruct a triangulation of $\mathcal{M}$.

8.2.1 Definition

Let P be a finite set of $n > k + 1$ points on $\mathcal{M}$. Let $\text{Del}(P)$ be the Delaunay complex of P , i.e. the collection of all the simplices with vertices in P that admit an empty circumscribing d -dimensional ball. A ball of $\mathbb{R}^d$ is called *empty* if its interior contains no point of P . For $p \in P$, we denote by $\text{Del}_p(P)$ the Delaunay complex of P restricted to the tangent space T_p , i.e. the subcomplex of $\text{Del}(P)$ formed by all the simplices with vertices in P that admit an empty circumscribing ball centered on T_p . Equivalently, the simplices of $\text{Del}_p(P)$ are the simplices of $\text{Del}(P)$ whose Voronoi dual face intersect T_p .

In the rest of this chapter, we assume that the points of P are in general position wrt spheres so that the Delaunay complex $\text{Del}(P)$ naturally embeds in $\mathbb{R}^d$ as a triangulation of P (see Section 4.3). We will further assume that P satisfies the following *transversality condition*: T_p contains no point that is equidistant from more than $k + 1$ points of P . The transversality condition implies that the restricted Delaunay complex $\text{Del}_p(P)$ is a subcomplex of dimension at most k . Furthermore, if T_p intersects a $(d - k)$ -face f of $\text{Vor}(P)$, the intersection consists of a single point c . The unique point c is the center of an empty ball circumscribing the k -simplex of $\text{Del}(P)$ dual to f . It is easy to see that applying an infinitesimal perturbation to P is sufficient to ensure the transversality condition.

We write $\text{star}(p)$ for the closed star of p in $\text{Del}_p(P)$, i.e. the subcomplex of $\text{Del}_p(P)$ consisting of the simplices of $\text{Del}_p(P)$ that are incident to p together with their faces (see Figure 8.3). In the following, we will simply call star a closed star. For a k -simplex σ in $\text{star}(p)$, we write $B_p(\sigma)$ for the ball centered on T_p that circumscribes σ , $c_p(\sigma)$ for its center and $R_p(\sigma)$ for its radius. Observe that $R_p(\sigma) \geq R(\sigma)$, where $R(\sigma)$ is the radius of the smallest ball circumscribing σ .

Definition 8.6 (Tangential Delaunay complex) *We call tangential Delaunay complex, or tangential complex for short, the simplicial complex $\text{Del}_{T\mathcal{M}}(P) = \{\sigma, \sigma \in \text{star}(p), p \in P\}$.*

Plainly, $\text{Del}_{T\mathcal{M}}(P)$ is a subcomplex of $\text{Del}(P)$ and is therefore a simplicial complex embedded in $\mathbb{R}^d$ if the points of P are in general position wrt spheres (see Section 4.3). The following lemma is crucial since it shows




Figure 8.3: $\mathcal{M}$ is in red. The sample P is the set of blue circles. The green line is the tangent space T_p at p . The Voronoi diagram of the sample is in black. The Delaunay triangulation $\text{Del}(P)$ is drawn in blue with, in bold, $\text{star}(p)$.

that computing the tangential complex reduces to computing n weighted Delaunay triangulations in the k -dimensional flats T_p , $p \in P$.

We define a map $\Psi_p : P \rightarrow T_p \times \mathbb{R}$ that associates to each point $p_i \in P$ a weighted point in T_p . Specifically, $\Psi_p(p_i) = (p'_i, p''_i) \in T_p \times \mathbb{R}$, where p'_i is the orthogonal projection of p_i onto T_p and $p''_i = -\|p_i - p'_i\|^2$. Observe that, under the transversality assumption, Ψ_p is 1-1.

It is known that the d -dimensional Voronoi diagram $\text{Vor}(P)$ intersects T_p along the weighted k -dimensional Voronoi diagram $\text{Vor}(\Psi_p(P))$ (see Exercise 4.11). Accordingly $\text{Del}_p(P)$, the restriction of the d -dimensional Delaunay complex $\text{Del}(P)$ to T_p , is isomorphic to the weighted Delaunay complex $\text{Del}(\Psi_p(P))$. Note that the transversality condition implies the fact that the set $(\Psi_p(P))$ is in general position wrt to spheres in T_p . Therefore, the simplicial complex $\text{Del}(\Psi_p(P))$ is naturally embedded in T_p . Moreover, the simplices of $\text{Del}(\Psi_p(P))$ are obtained by projecting onto T_p the simplices of $\text{Del}_p(P)$. Conversely, the simplices of $\text{Del}_p(P)$ can be deduced from the simplices of $\text{Del}(\Psi_p(P))$ by a piecewise linear map that we call the *lifting map*. Specifically, the lifting map lifts each weighted point (p', p'') associated to a vertex of $\text{Del}(\Psi_p(P))$ to the unique point $p \in P$ such that $(p', p'') = \Psi_p(p)$. The lift of a simplex σ_p of $\text{Del}_p(P)$ is then the geometric simplex σ whose vertices are the lifts of the vertices of σ_p . $\text{Del}_p(P)$ is then the image of $\text{Del}(\Psi_p(P))$ by the *lifting map*. We summarize our discussion in the following lemma :

Lemma 8.7 *If the points of P are in general position and satisfy the transversality condition, $\text{Del}_p(P)$ is the lift of $\text{Del}(\Psi_p(P))$, the k -dimensional weighted Delaunay triangulation of $\Psi_p(P)$ in T_p .*


We deduce from the lemma an efficient algorithm to compute $\text{star}(p) = \text{star}(p, \text{Del}_p(P))$: project P onto T_p , compute $\text{star}_{T_p}(p) = \text{star}(p, \text{Del}(\Psi_p(P)))$, the star of p in $\text{Del}(\Psi_p(P))$, and then lift $\text{star}_{T_p}(p)$ to $\text{star}(p)$. Apart from the projection of the points onto T_p , this algorithm involves only operations in the k -dimensional flat T_p . If P is an $(\varepsilon, \bar{\eta})$ -net of $\mathcal{M}$, we can even restrict our attention to the subset of P inside the ball of radius 2ε centered at p . The transversality condition with respect to the tangent plane T_p is therefore only required for points in that ball.

8.2.2 Inconsistent simplices

In general, the tangential complex is *not* a combinatorial manifold (see Definition 2.17). This is due to the presence of so-called *inconsistent simplices*.

Definition 8.8 (Inconsistent simplex) *A simplex $\sigma \in \text{Del}_{T\mathcal{M}}(P)$ is called inconsistent if σ does not belong to the stars of all its vertices. Let σ be an inconsistent simplex and let p_i and p_j be two vertices of σ so that σ is in $\text{star}(p_i)$ but not in $\text{star}(p_j)$. We say that the pair of vertices (p_i, p_j) witnesses the inconsistent simplex σ .*

Refer to Figure 8.4. Let σ be an inconsistent k -simplex witnessed by the pair of vertices (p_i, p_j) . The simplex σ belongs to the star of p_i but not to the star of p_j . Equivalently, the Voronoi $(d - k)$ -dimensional face $\text{Vor}(\sigma)$ dual to σ intersects T_{p_i} (at a point $c_{p_i}(\sigma)$) but does not intersect T_{p_j} . Observe that $c_{p_i}(\sigma)$ is the center of an empty d -dimensional ball $B_{p_i}(\sigma)$ circumscribing σ . Let $c_{p_j}(\sigma)$ denote the intersection of $\text{aff}(\text{Vor}(\sigma))$ with T_{p_j} . Differently from $B_{p_i}(\sigma)$, the d -dimensional ball $B_{p_j}(\sigma)$ centered at $c_{p_j}(\sigma)$ that circumscribes σ contains a subset $P_j(\sigma)$ of points of P in its interior. Therefore, the line segment $[c_{p_i}(\sigma) c_{p_j}(\sigma)]$ intersects the interior of some Voronoi cells (among which are the cells of the points of $P_j(\sigma)$). Let p_l be the point of $P \setminus \sigma$ whose Voronoi cell is hit first by the segment $[c_{p_i}(\sigma) c_{p_j}(\sigma)]$, when oriented from $c_{p_i}(\sigma)$ to $c_{p_j}(\sigma)$. We write σ^I for the $(k + 1)$ -simplex $\text{conv}(\sigma, p_l)$ and $i(\sigma^I)$ for the first point of the oriented segment $[c_{p_i}(\sigma) c_{p_j}(\sigma)]$ that belongs to $\text{Vor}(p_l)$. Observe that since the point $i(\sigma^I)$ belongs to $\text{Vor}(\sigma) \cap \text{Vor}(p_l)$, σ^I is



Definition 8.9 (Inconsistency trigger) *An inconsistency trigger of a k -simplex σ of $\text{Del}_{\mathcal{T}\mathcal{M}}(P)$ is a $(k+1)$ -simplex of $\text{Del}(P)$ that is the inconsistency trigger of σ for some ordered pair of vertices of σ that witnesses the inconsistency.*

Since an inconsistency trigger σ^I is a $(k+1)$ -simplex of $\text{Del}(P)$, we will use the same notations for σ^I as for any other simplex, e.g. $R(\sigma^I)$ for the radius of the smallest circumscribing ball or $\Theta(\sigma^I)$ for its thickness. We write $\text{Inc}(p)$




Figure 8.5: In Figure (a), $\mathcal{M}$ is in red, the sample P is the set of blue points, green lines denote tangent spaces, the Voronoi diagram of the sample is in black, and $\text{Del}_{T\mathcal{M}}(P)$ is drawn in blue. The dashed lines indicate the inconsistent edges. In Figure (b), the blue lines denote $\text{Del}_{T\mathcal{M}}(P)$ and the grey triangles denote the inconsistency triggers.

for the set of inconsistency triggers incident to p and $\text{Inc}(P) = \cup_{p \in P} \text{Inc}(p)$. We also define the *completed tangential complex* $\text{Del}_{T\mathcal{M}}^{+I}(P) = \text{Del}_{T\mathcal{M}}(P) \cup \text{Inc}(P)$. See Figure 8.5.

Calculating $\text{Del}_{T\mathcal{M}}^{+I}(P)$ is easy once we know $\text{Del}(P)$. Indeed, it suffices to detect the inconsistent k -simplices that do not appear in the stars of all their vertices. Let σ be an inconsistent simplex witnessed by the pair (p_i, p_j) . To compute the associated inconsistency trigger, we need to identify the point p_l . This can be done by computing the restriction of $\text{Vor}(P)$ to the line $L_{ij} = (c_{p_i}(\sigma), c_{p_j}(\sigma))$. This in turn can be done by projecting the points of P onto L_{ij} and computing a 1-dimensional weighted Voronoi diagram (see Exercise 4.11).

8.2.3 Geometric properties of inconsistency triggers

We give some simple geometric lemmas that, in particular, bound the thickness of the inconsistency triggers. We will use angles between affine spaces as defined in Section 5.2.2.

Lemma 8.10 *Let P be an ε -dense sample¹ of a submanifold $\mathcal{M}$. Let $\alpha_0 = \frac{1}{8}(1 + 2\varepsilon - \sqrt{1 - 12\varepsilon + 4\varepsilon^2}) \approx \varepsilon$ and assume that $\alpha_0 < 1/4$. Then*

1. $\text{Vor}(p) \cap T_p \subseteq B(p, \alpha_0 \text{rch}(\mathcal{M}))$.
2. for any k -simplex $\sigma \in \text{star}(p)$, $R(\sigma) \leq R_p(\sigma) \leq \alpha_0 \text{rch}(\mathcal{M})$ (recall that $R_p(\sigma)$ is the radius of the circumscribing ball of σ centered on T_p .)
3. for any $\sigma \in \text{Del}_{T\mathcal{M}}(P)$, $\Delta(\sigma) \leq 2\alpha_0 \text{rch}(\mathcal{M})$.

Proof We prove (i). The other statements easily follow. Let $x \in \text{Vor}(p) \cap T_p$ and write $\|p - x\| = \alpha \text{rch}(\mathcal{M})$. Let x' be the point of $\mathcal{M}$ closest to x and let $x'' = \Pi_p(x')$. We have $\|x - x'\| \leq \|x - p\|$ and by the triangular inequality, $\|p - x'\| \leq 2\|x - p\| = 2\alpha \text{rch}(\mathcal{M})$. It then follows from Lemma 7.8 that

$$\|x' - x''\| \leq \frac{\|p - x'\|^2}{2\text{rch}(\mathcal{M})} \leq 2\alpha^2 \text{rch}(\mathcal{M}).$$

Now observe that $\|x' - x''\| = \|x - x'\| \cos \phi$, where $\phi = \angle(T_{x'}, T_p)$. Assuming $\alpha \leq \frac{1}{4}$, we have, by Lemma 7.10, $\cos \phi \geq 1 - 8\alpha^2$. We conclude that

$$\|x - x'\| \leq \frac{2\alpha^2 \text{rch}(\mathcal{M})}{1 - 8\alpha^2}.$$

Since P is ε -dense in $\mathcal{M}$, there exists a point $q \in P$, such that $\|x' - q\| \leq \varepsilon \text{rch}(\mathcal{M})$. Together with $x \in \text{Vor}(p)$, this implies

$$\|x - p\| \leq \|x - q\| \leq \|x - x'\| + \|x' - q\| \leq \left(\frac{2\alpha^2}{1 - 8\alpha^2} + \varepsilon \right) \text{rch}(\mathcal{M}) \quad (8.3)$$

We thus have

$$\alpha \leq \frac{2\alpha^2}{1 - 8\alpha^2} + \varepsilon \leq \frac{2\alpha^2}{1 - 2\alpha} + \varepsilon, \quad (8.4)$$

where the last inequality is obtained by taking $\alpha \leq 1/4$.

Let $\alpha_0 = \frac{1}{8}(1 + 2\varepsilon - \sqrt{1 - 12\varepsilon + 4\varepsilon^2})$ and $\alpha_1 = \frac{1}{8}(1 + 2\varepsilon + \sqrt{1 - 12\varepsilon + 4\varepsilon^2})$ and note that $\alpha_0 \approx \varepsilon$ and $\varepsilon \leq \alpha_0 < \frac{1}{4} \leq \alpha_1$. The inequality above is satisfied either when $\alpha \leq \alpha_0$ or when $\alpha \geq \alpha_1$. Hence $\text{Vor}(p)$ is contained in the union of $B(p, \alpha_0)$ and $\mathbb{R}^d \setminus B(p, \alpha_1)$. However, since $\text{Vor}(p)$ is connected and contains p , $\text{Vor}(p)$ must be entirely contained in $B(p, \alpha_0 \text{rch}(\mathcal{M}))$. $\square$

¹See Definition 7.7.

The following lemmas bound the size and shape of inconsistency triggers.

In particular the lemmas state that if an inconsistent simplex is small, thick and has a small circumradius, its inconsistency triggers have small circumradii (Lemma 8.11) and cannot be thick (Lemma 8.12).

Lemma 8.11 *Let σ be an inconsistent k -simplex and let σ^I be a $(k+1)$ -simplex that triggers the inconsistency of σ . Writing $\theta = \max_{p \in \sigma} (\angle(\text{aff}(\sigma), T_p))$, we have*

$$\sin \theta \leq \frac{\Delta(\sigma)}{\Theta(\sigma) \text{rch}(\mathcal{M})} \quad \text{and} \quad R(\sigma^I) \leq \frac{R(\sigma)}{\cos \theta}.$$

Proof We assume that the vertex pair (p_i, p_j) is the witness pair of σ associated to the trigger σ^I and use the same notations as in Subsection 8.2.2. Therefore σ belongs to $\text{Del}_{p_i}(P)$ but not to $\text{Del}_{p_j}(P)$, and the vertices of σ^I lie in the closure of $B_{ij} = B_{p_j}(\sigma) \setminus B_{p_i}(\sigma)$. Observe also that c_{p_i}, c_{p_j} and $i(\sigma^I)$ lie in the $(d-k)$ -flat that contains $c(\sigma)$ and is perpendicular to $\text{aff}(\sigma)$. Hence the orthogonal projection of these four points onto $\text{aff}(\sigma)$ is $c(\sigma)$.

We now bound $\theta = \angle(\text{aff}(\sigma), T_p)$ for any $p \in \sigma$. By Lemma 7.8 (2), we have for any $p, q \in \sigma$, $d(q, T_p) \leq \frac{\Delta^2(\sigma)}{2\text{rch}(\mathcal{M})}$ and by Lemma 5.14

$$\sin \theta \leq \frac{2 \frac{\Delta^2(\sigma)}{2\text{rch}(\mathcal{M})}}{\Theta(\sigma) \Delta(\sigma)} = \frac{\Delta(\sigma)}{\Theta(\sigma) \text{rch}(\mathcal{M})}.$$

Since the orthogonal projection of c_{p_i} onto $\text{aff}(\sigma)$ is $c(\sigma)$, $\omega = \angle(p_i - c_{p_i}, p_i - c(\sigma)) = \min_{u \in \text{aff}(\sigma)} \angle(p_i - c_{p_i}, u) \leq \theta$. We thus have

$$R_{p_i}(\sigma) = \|p_i - c_{p_i}\| = \frac{R(\sigma)}{\cos \omega} \leq \frac{R(\sigma)}{\cos \theta}.$$

and we get the same bound if p_i is replaced by p_j or by any other vertex p in the above inequality. Since $i(\sigma^I) \in [c_{p_i} c_{p_j}]$, we also have $R(\sigma^I) \leq \|i(\sigma^I) - p_i\| \leq R(\sigma)/\cos \theta$. $\square$

We deduce from the lemma a bound on the thickness of inconsistency triggers. This is a crucial property to be used later to remove inconsistencies.

Lemma 8.12 *The thickness of the $(k+1)$ -simplex σ^I that triggers the inconsistency of a k -simplex σ satisfies :*

$$\Theta(\sigma^I) \leq \frac{\Delta(\sigma^I)}{2(k+1)\text{rch}(\mathcal{M})} \left(1 + \frac{2}{\Theta(\sigma)}\right)$$

Proof Let p_l be the vertex of σ^I that is not a vertex of σ . We bound the altitude $D(p_l, \sigma^I)$ of σ^I . Let $q \in \sigma$. We deduce

$$\begin{aligned} D(p_l, \sigma^I) &= \|p_l - q\| \sin \angle(p_l - q, \text{aff}(\sigma)) \\ &\leq \Delta(\sigma^I) (\sin \angle(p_l - q, T_q) + \sin \angle(T_q, \text{aff}(\sigma))) \\ &\leq \Delta(\sigma^I) \left(\frac{\Delta(\sigma^I)}{2\text{rch}(\mathcal{M})} + \frac{\Delta(\sigma)}{\Theta(\sigma)\text{rch}(\mathcal{M})} \right) \quad (\text{Lemmas 7.8 \& 8.11}) \\ &\leq \frac{\Delta^2(\sigma^I)}{2\text{rch}(\mathcal{M})} \left(1 + \frac{2}{\Theta(\sigma)}\right) \end{aligned}$$

The bound on the thickness then follows from the definition of thickness (see Section 5.2). $\square$

Corollary 8.13 *Let P be an ε -dense sample of $\mathcal{M}$ and let σ^I be a $(k+1)$ -simplex that triggers the inconsistency of a k -simplex σ . Assume that σ is Θ_0 -thick and that Θ_0 satisfies*

$$\Theta_0^{k+1} > \frac{\alpha_0}{k+1} \left(\cos \arcsin \frac{2\alpha_0}{\Theta_0^k} \right)^{-1} \left(1 + \frac{2}{\Theta_0^k}\right) \quad (8.5)$$

where α_0 is defined in Lemma 8.10. Then σ^I is not Θ_0 -thick.

Proof We need to prove that $\Theta(\sigma^I) < \Theta_0^{k+1}$. From Lemma 8.12 we have

$$\Theta(\sigma^I) \leq \frac{\Delta(\sigma^I)}{2(k+1)\text{rch}(\mathcal{M})} \left(1 + \frac{2}{\Theta(\sigma)}\right) \leq \frac{\Delta(\sigma^I)}{2(k+1)\text{rch}(\mathcal{M})} \left(1 + \frac{2}{\Theta_0^k}\right).$$

We now bound $\Delta(\sigma^I)$:

$$\begin{aligned} \Delta(\sigma^I) &\leq 2R(\sigma^I) \\ &\leq 2R(\sigma) \left(\cos \arcsin \frac{\Delta(\sigma)}{\Theta(\sigma)\text{rch}(\mathcal{M})} \right)^{-1} \quad (\text{Lemma 8.11}) \\ &\leq 2\alpha_0 \text{rch}(\mathcal{M}) \left(\cos \arcsin \frac{2\alpha_0}{\Theta_0^k} \right)^{-1} \quad (\text{Lemma 8.10}) \end{aligned}$$

Hence

$$\Theta(\sigma^I) \leq \frac{\alpha_0}{k+1} \left(\cos \arcsin \frac{2\alpha_0}{\Theta_0^k} \right)^{-1} \left(1 + \frac{2}{\Theta_0^k} \right).$$

It follows that if Equation 8.5 holds, $\Theta(\sigma^I) < \Theta_0^{k+1}$, which means that σ^I is not Θ_0 -good. $\square$

Note that the condition in the corollary is satisfied when the thickness of the inconsistent simplexes satisfy $\Theta(\sigma) > \Theta_0^k = \Omega(\varepsilon^{1/(k+1)})$.

8.3 Submanifold reconstruction

Let $\mathcal{M}$ denote a submanifold of $\mathbb{R}^d$ that is compact, closed, differentiable, and whose reach is positive. The only knowledge we have about $\mathcal{M}$ is its dimension k together with a finite point sample $P \subset \mathcal{M}$ and the tangent spaces at those points. From that knowledge, we want to construct a triangulation of $\mathcal{M}$. We assume that P is an $(\varepsilon, \bar{\eta})$ -net of $\mathcal{M}$ (See Definition 7.7). The parameters ε and $\bar{\eta}$ need not to be known but must satisfy some conditions to be made explicit in the analysis of the algorithm. In particular, we need ε to be sufficiently small.

The main idea of the proposed reconstruction algorithm is to rely on the tangential complex $\text{Del}_{T\mathcal{M}}(P)$. However, we know from the previous section that $\text{Del}_{T\mathcal{M}}(P)$ may not be a combinatorial manifold, and our first goal is to obtain a complex that is a k -combinatorial manifold. This will be achieved by assigning weights to the points of P and resorting to the weighted tangential complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ which is the natural counterpart of the tangential complex when the point set P is replaced by a set $\hat{P}$ of weighted points. Properties of tangential complexes extend to weighted tangential complexes and, in particular, we will show that the occurrence of inconsistencies in $\text{Del}_{T\mathcal{M}}(\hat{P})$ is triggered by the existence of non-thick simplices of dimension at most $k+1$ in $\text{Del}(\hat{P})$. Then, since it is possible to assign weights to the points of P to get a thick weighted Delaunay complex (Section 5.3), it is a fortiori possible to get a weighted tangential complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ with no inconsistency. Finally, an algorithm is derived from Moser Tardos algorithm (Section 5.2). The algorithm uses data structures of dimension at most k . Its output is a k -dimensional combinatorial manifold that we will show to be homeomorphic to $\mathcal{M}$ and close to $\mathcal{M}$.

8.3.1 Weight assignment

We recall the definition of a *weighting scheme* introduced in Section 5.3. Given a point set $P = \{p_1, \dots, p_n\} \subseteq \mathbb{R}^d$, a weighting scheme on P is a function w that assigns to each point $p_i \in P$ a non-negative real weight $w(p_i)$. We write $\hat{p}_i = (p_i, w(p_i))$ and $\hat{P} = \{\hat{p}_1, \dots, \hat{p}_n\}$.

We recall the definition of the *relative amplitude* of w as $\tilde{w} = \max_{p \in P} \frac{w(p)}{L^2(p)}$ where $L(p) = \min_{q \in P \setminus \{p\}} \|p - q\|$. Note that $2\varepsilon \text{rch}(M) \geq L(p) \geq \bar{\eta} \varepsilon \text{rch}(\mathcal{M})$ since P is an $(\varepsilon, \bar{\eta})$ -net.

From this point onwards, we assume that all weights are non negative and that $\tilde{w} \leq \tilde{w}_0$, for some constant $\tilde{w}_0 < 1/4$. Hence, for any point $p \in P$, $w(p) \in [0, \tilde{w}_0 L^2(p)] \subset [0, 4\tilde{w}_0 \varepsilon^2 \text{rch}^2(M)]$.

The condition on $\tilde{w}$ implies in particular that the weighted points in $\hat{P}$ have pairwise positive weighted distances which ensures that any point in P appears as a vertex of the weighted Delaunay complex $\text{Del}(\hat{P})$ (see Section 4.4.3).

We can extend the definition of the tangential complex to the case of a set $\hat{P}$ of weighted points. We simply need to replace in Definition 8.6 the Delaunay complex $\text{Del}(P)$ by the weighted Delaunay complex $\text{Del}(\hat{P})$. The role played above by Delaunay balls will be played by Delaunay weighted points, i.e. weighted points orthogonal to weighted Delaunay simplices and free of weighted points in $\hat{P}$. Lemma 8.7 remains valid provided that the mapping Ψ_p is extended to weighted points as follows. If $\hat{p}_i = (p_i, w_i) \in \mathbb{R}^d \times \mathbb{R}$ is a weighted point, we define $\Psi_p(\hat{p}_i) = (p_i, w_i - \|p_i - p'_i\|^2)$, where p'_i is the orthogonal projection of p_i on the tangent space T_p .

We now extend to the weighted case the properties of the inconsistency triggers shown in Section 8.2.

Let P be an ε -dense sample of a submanifold $\mathcal{M}$ and assume that the weighting scheme on P has a relative amplitude $\tilde{w} \leq \tilde{w}_0 < 1/4$. We denote by $\text{Del}_p(\hat{P})$ the weighted Delaunay complex restricted to T_p and by $\text{star}(\hat{p})$ the star of p in $\text{Del}_p(\hat{P})$. We further write $\text{Vor}(\hat{p})$ for the cell of p in the weighted Voronoi diagram $\text{Vor}(\hat{P})$. For any k -simplex σ in $\text{Del}_p(\hat{P})$, we denote respectively by $R(\sigma)$ and $R(\hat{\sigma})$ the circumradius and weighted radius of σ , by $R_p(\sigma)$ the radius of the ball centered on T_p that circumscribes σ and by $R_p(\hat{\sigma})$ the radius of the ball centered on T_p that is orthogonal to $\hat{\sigma}$.

Lemma 8.14 *Let*

$$\begin{aligned}\hat{\varepsilon} &= \varepsilon(1 + 2\sqrt{\tilde{w}_0}), \\ \hat{\alpha}_0 &= \frac{1}{8}(1 + 2\hat{\varepsilon} - \sqrt{1 - 12\hat{\varepsilon} + 4\hat{\varepsilon}^2}) \approx \varepsilon(1 + 2\sqrt{\tilde{w}_0}), \\ \hat{\alpha}_1 &= \hat{\alpha}_0 + 2\varepsilon\sqrt{\tilde{w}_0} \approx \varepsilon(1 + 4\sqrt{\tilde{w}_0})\end{aligned}$$

and assume $\hat{\alpha}_0 \leq \frac{1}{4}$. Then we have:

- (i) $\text{Vor}(\hat{p}) \cap T_p \subseteq B(p, \hat{\alpha}_0 \text{rch}(\mathcal{M}))$,
- (ii) for any k -simplex $\sigma \in \text{star}(\hat{p})$, $R(\hat{\sigma}) \leq R_p(\hat{\sigma}) \leq \hat{\alpha}_0 \text{rch}(\mathcal{M})$.
- (iii) for any $\sigma \in \text{Del}_{T\mathcal{M}}(\hat{P})$, $\Delta(\sigma) \leq 2\hat{\alpha}_1 \text{rch}(\mathcal{M})$.

Proof The proof of (i) is an easy adaptation of the proof of Lemma 8.10. We define x , x' and x'' and α as in the proof of Lemma 8.10 and still get

$$\|x - x'\| \leq \frac{2\alpha^2 \text{rch}(\mathcal{M})}{1 - 8\alpha^2}.$$

Since P is ε -dense in $\mathcal{M}$, there exists a point $q \in P$, such that $\|x' - q\| \leq \varepsilon \text{rch}(\mathcal{M})$. Since x is in $\text{Vor}(\hat{p})$, this implies that

$$\begin{aligned}\|x - p\|^2 + w(p) &\leq \|x - q\|^2 + w(q) \\ \|x - p\|^2 &\leq (\|x - x'\| + \|x' - q\|)^2 + w(q),\end{aligned}$$

which yields

$$\alpha^2 \leq \left(\frac{2\alpha^2}{1 - 8\alpha^2} + \varepsilon \right)^2 + 4\tilde{w}_0\varepsilon^2$$

Hence,

$$\alpha \leq \frac{2\alpha^2}{1 - 8\alpha^2} + \varepsilon(1 + 2\sqrt{\tilde{w}_0}).$$

which is just Equation 8.4 where ε has been replaced by $\hat{\varepsilon}$ so that we can end the proof as in Lemma 8.10.

Let $c_p(\hat{\sigma})$ be the point of T_p that is orthogonal to σ . Property (ii) follows easily since $c_p(\hat{\sigma})$ belongs to $\text{Vor}(\hat{p}) \cap T_p$ and $d(c_p(\hat{\sigma}), p)^2 = R_p(\hat{\sigma})^2 + w(p)$. Hence $R_p(\hat{\sigma}) \leq d(c_p(\hat{\sigma}), p) \leq \hat{\alpha}_0 \text{rch}(\mathcal{M})$.

To prove Property (iii), we bound the Euclidean distance $d(c_p(\hat{\sigma}), q)$ between $c_p(\hat{\sigma})$ and a vertex q of σ as follows:

$$\begin{aligned} d(c_p(\hat{\sigma}), q)^2 &= R_p(\hat{\sigma})^2 + w(q) \\ &\leq \hat{\alpha}_0^2 \text{rch}(\mathcal{M})^2 + 4\tilde{w}_0 \varepsilon^2 \text{rch}(\mathcal{M})^2. \end{aligned}$$

Hence,

$$d(c_p(\hat{\sigma}), q) \leq \left(\hat{\alpha}_0 + 2\varepsilon\sqrt{\tilde{w}_0} \right) \text{rch}(\mathcal{M}) = \hat{\alpha}_1 \text{rch}(\mathcal{M})$$

Since this bound holds for any vertex q of σ , the bound given in (iii) for the diameter $\Delta(\sigma)$ follows. $\square$

Lemma 8.15 *Let σ be an inconsistent k -simplex of $\text{Del}_{T\mathcal{M}}(\hat{P})$ and let σ^I be a $(k+1)$ -simplex that triggers the inconsistency of σ . If θ is defined by $\theta = \max_{p \in \sigma} (\angle(\text{aff}(\sigma), T_p))$, we have*

$$\sin \theta \leq \frac{\Delta(\sigma)}{\Theta(\sigma) \text{rch}(\mathcal{M})} \quad \text{and} \quad R(\hat{\sigma}^I) \leq \frac{R(\hat{\sigma})}{\cos \theta}.$$

Proof The proof is the same as the proof of Lemma 8.11 $\square$

Lemma 8.12 applies verbatim to the weighted case and finally an easy adaptation of the proof of Corollary 8.13 yields the following weighted version.

Corollary 8.16 *Let P be an ε -dense sample of $\mathcal{M}$ and let w be a weighting scheme on P with a relative amplitude $\tilde{w} \leq \tilde{w}_0 < 1/4$. Let σ be an inconsistent k -simplex and let σ^I be a $(k+1)$ -simplex that triggers the inconsistency of σ . Assume that σ is Θ_0 -thick and that Θ_0 satisfies*

$$\Theta_0^{k+1} > \frac{\hat{\alpha}_1}{k+1} \left(\cos \arcsin \frac{2\hat{\alpha}_0}{\Theta_0^k} \right)^{-1} \left(1 + \frac{2}{\Theta_0^k} \right) \quad (8.6)$$

where $\hat{\alpha}_0$ and $\hat{\alpha}_1$ are defined in Lemma 8.14. Then σ^I is not Θ_0 -thick.

Proof We need to prove that $\Theta(\sigma^I) < \Theta_0^{k+1}$. From Lemma 8.12 we have

$$\Theta(\sigma^I) \leq \frac{\Delta(\sigma^I)}{2(k+1)\text{rch}(\mathcal{M})} \left(1 + \frac{2}{\Theta(\tau)} \right) \quad (8.7)$$

We now bound $\Delta(\sigma^I)$. Let $c(\hat{\sigma}^I)$ be the weighted center of σ^I and let q be any vertex of σ^I . We have

$$\begin{aligned} d(c(\hat{\sigma}^I), q)^2 &= R(\hat{\sigma}^I)^2 + w(q) \\ &\leq \frac{R(\hat{\sigma})^2}{\cos \theta^2} + 4\tilde{w}_0 \varepsilon^2 \text{rch}(\mathcal{M})^2 \\ &\leq \frac{\hat{\alpha}_0^2 + 4\tilde{w}_0 \varepsilon^2}{\cos \theta^2} \text{rch}(\mathcal{M})^2 \leq \frac{\hat{\alpha}_1^2}{\cos \theta^2} \text{rch}(\mathcal{M})^2 \end{aligned}$$

where θ is given by Lemma 8.15 and $\hat{\alpha}_1$ is defined in Lemma 8.14. Since this is true for any vertex q of σ^I , $\Delta(\sigma^I)$ can be bound by $2\frac{\hat{\alpha}_1}{\cos \theta} \text{rch}(\mathcal{M})$. Plugging this bound in Equation 8.7 yields that $\Theta(\sigma^I) < \Theta_0^{k+1}$ if σ is Θ_0 -thick and Θ_0 satisfies Equation 8.6. $\square$

Note that Equation 8.6 in Corollary 8.16 is satisfied if $\Theta_0 = \Omega(\varepsilon^{\frac{1}{k(k+1)}})$.

8.3.2 Reconstruction algorithm

We conclude from the previous section that if P is an $(\varepsilon, \bar{\eta})$ -net of $\mathcal{M}$, the occurrence of inconsistencies in the tangential complex $\text{Del}_{T\mathcal{M}}(P)$ is due to the occurrence of non-thick simplices of dimension up to $k+1$ in the Delaunay complex $\text{Del}(P)$. This property still holds in the weighted case provided that the weighting scheme has a small relative amplitude. We can thus proceed in a way similar to what has been done in Section 5.3 to remove non-thick simplices from Delaunay complexes using a variant of Moser-Tardos algorithm. Here we want to find a weighting scheme on P such that the simplices of the tangential complex as well as the inconsistency triggers (all are simplices of $\text{Del}(\hat{P})$) are thick. As a consequence, the tangential complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ won't contain any inconsistency.

The algorithm depends on two parameters $\tilde{w}_0$ and Θ_0 . Here as in Section 5.3, we call resampling the operation which consists in reassigning the weights of the vertices of a simplex. The weights are taken independently. The weight of a vertex p is taken uniformly at random in $[0, \tilde{w}_0 L^2(p)]$. The algorithm maintains the weighted tangential complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ and *resamples* the Θ_0 -flakes that may appear in $\text{Del}_{T\mathcal{M}}(\hat{P})$ or in the $(k+1)$ simplices that trigger an inconsistency. Note that the algorithm does not compute the full weighted Delaunay complex $\text{Del}(\hat{P})$ but only a subcomplex of dimension $k+1$.

Algorithm 8 Tangential complex with no inconsistency

Input: $P, \{T_p, p \in P\}, \tilde{w}_0, \Theta_0$ Initialize all weights to 0 and compute $\text{Del}_{T\mathcal{M}}(\hat{P})$ **while** there are Θ_0 -flakes or inconsistencies in $\text{Del}_{T\mathcal{M}}(\hat{P})$ **do** **while** there is a Θ_0 -flake σ in $\text{Del}_{T\mathcal{M}}(\hat{P})$ **do** resample σ update $\text{Del}_{T\mathcal{M}}(\hat{P})$ **if** there is an inconsistent simplex σ in $\text{Del}_{T\mathcal{M}}(\hat{P})$ **then** compute a trigger simplex σ^I associated to σ resample the flake $\sigma \subset \sigma^I$ update $\text{Del}_{T\mathcal{M}}(\hat{P})$ **Output:** A weighting scheme on P and the corresponding weighted tangential Delaunay complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ which is granted to be Θ_0 -thick and to have no inconsistency.

Theorem 8.17 *Let $\mathcal{M}$ be a submanifold of positive reach and let P be an $(\varepsilon, \bar{\eta})$ -net of $\mathcal{M}$ for a sufficiently small ε . Let in addition $\tilde{w}_0 < 1/4$ and Θ_0 be two constants satisfying Equation 8.6 and the following Equation*

$$\left[\frac{6}{\bar{\eta}} \frac{(1 + 4\sqrt{\tilde{w}_0})}{\cos \theta} \right]^{d(d+1)} \frac{16\Theta_0}{\tilde{w}_0 \bar{\eta}^3} \frac{(1 + 2\sqrt{\tilde{w}_0})(1 + 4\sqrt{\tilde{w}_0})^2}{\cos^3 \theta} \leq \frac{1}{e} \quad (8.8)$$

where

$$\theta = \arcsin \frac{2\hat{\alpha}_1}{\Theta_0^k} \approx \arcsin \frac{2\varepsilon(1 + 4\sqrt{\tilde{w}_0})}{\Theta_0^k}.$$

Then Algorithm 8 terminates and outputs a weight assignment on the points of P such that the weighted tangential complex $\text{Del}_{T\mathcal{M}}(\hat{P})$ is free of inconsistencies. The expected time complexity of the algorithm is $O(|P|)$.

Proof First note that the algorithm removed inconsistencies only when there are no Θ_0 -flakes in $\text{Del}_{T\mathcal{M}}(\hat{P})$. Then Equation 8.6 ensures that the trigger of an inconsistency is not Θ_0 -thick (Lemma 8.16) and therefore includes a Θ_0 -flake that can be resampled.

The resampled simplices are Θ_0 -flakes of $\text{Del}(\hat{P})$ that either belong to $\text{Del}_{T\mathcal{M}}(\hat{P})$ or are included in an inconsistency trigger. Let σ be any resampled flake. From Lemmas 8.14 and 8.15, the weighted radius $R(\hat{\sigma})$ is at most

$$R(\hat{\sigma}) \leq \frac{\hat{\alpha}_0}{\cos \theta} \text{rch}(\mathcal{M}),$$

and, from Lemma 8.14 and the proof of Corollary 8.16, its diameter is bounded by

$$\Delta(\sigma) \leq \frac{2\hat{\alpha}_1}{\cos \theta} \text{rch}(\mathcal{M}).$$

Therefore Algorithm 8 is a variant of Moser-Tardos algorithm similar to the algorithm of Section 5.3.5 that removes from $\text{Del}(\hat{P})$ all Θ_0 -flakes with small weighted radii. Then, arguing as in Section 5.3.5, we show that Algorithm 8 terminates provided that the condition expressed in Equation 8.8 (similar to Equation 5.13) is satisfied. See Exercise 8.3 for details. The expected number of resampled simplices is $O(|P|)$. Since each resampling can be performed in constant time, the expected complexity of Algorithm 8 is also $O(|P|)$.

Assume that $\bar{\eta}$ and $\tilde{w}_0 < \frac{1}{4}$ are fixed. Then Equation 8.6 is satisfied if Θ_0 is sufficiently large with respect to ε and Algorithm 8 terminates provided that Θ_0 is sufficiently small with respect to $\bar{\eta}$ and $\tilde{w}_0$ and is sufficiently large with respect to ε . Those conditions can be satisfied if the sampling radius ε is sufficiently small. $\square$

8.3.3 Guarantees on the reconstruction

The simplicial complex $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}(\hat{P})$ output by the algorithm is free of inconsistencies. In addition, it is a good approximation of $\mathcal{M}$ as stated in the following theorem.

Theorem 8.18 (Guarantees) *Under the same hypotheses as in Theorem 8.17, the Delaunay tangential complex $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}(\hat{P})$ output by the algorithm satisfies the following properties:*

1. *All the simplices in $\hat{\mathcal{M}}$ are Θ_0 -thick.*
2. *$\hat{\mathcal{M}}$ is a piecewise linear k -submanifold without boundary;*
3. *$\hat{\mathcal{M}}$ is homeomorphic to $\mathcal{M}$;*
4. *The Hausdorff distance between $\hat{\mathcal{M}}$ and $\mathcal{M}$ is at most $4\varepsilon^2 \text{rch}(\mathcal{M})$;*
5. *If σ is a k -simplex of $\hat{\mathcal{M}}$ and p is a vertex of σ , we have*

$$\sin \angle(\text{aff}(\sigma), T_p) \leq \frac{2\varepsilon}{\Theta_0}.$$

Proof The first statement directly follows from the algorithm. Proving the second statement reduces to proving that the link of any vertex of $\hat{\mathcal{M}}$ is a topological $(k-1)$ -sphere. We first observe that, since $\hat{\mathcal{M}}$ contains no inconsistencies, the star of any vertex p in $\hat{\mathcal{M}}$ is identical to $\text{star}(\hat{p})$, the star of p in $\text{Del}_p(\hat{P})$. Hence, to prove the second statement, it is enough to prove that the link of p in $\text{Del}_p(\hat{P})$ is a topological $(k-1)$ -sphere, which is done in the next lemma.

Lemma 8.19 *$\hat{\mathcal{M}}$ is a simplicial combinatorial manifold.*

Proof It is sufficient to prove that, for any $p \in P$, the link of p in $\hat{\mathcal{M}}$ is a topological $(k-1)$ -sphere. By Lemma 8.7, $\text{star}(\hat{p})$ is isomorphic to $\text{star}_p(\hat{p})$, the star of p in $\text{Del}(\Psi_p(\hat{P}))$. Since $\text{star}_p(\hat{p})$ is a k -dimensional triangulated topological ball under the general position and transversality assumptions, the same is true for $\text{star}(\hat{p})$. To prove the lemma, it is then sufficient to show that p cannot belong to the boundary of $\text{star}_p(\hat{p})$. Consider the dual cell V of $p = \Psi_p(p)$ in the weighted Voronoi diagram $\text{Vor}(\Psi_p(\hat{P}))$. V is the intersection of the Voronoi cell of p with T_p , i.e. $V = \text{Vor}(p) \cap T_p$. By Lemma 8.10, V is bounded, which implies that p cannot belong to the boundary of $\text{star}_p(\hat{p})$. It follows that p cannot belong to the boundary of $\text{star}(\hat{p})$. $\square$

Statements 3 – 5 in Theorem 8.18 then follow from Theorem 7.16. We need to check that the five hypotheses of Theorem 7.16 are satisfied.

Hypotheses (a) – (c) of the theorem are satisfied provided that ε is small enough (see Exercises 8.2 and 8.3). Hypothesis (d) is satisfied thanks to Lemma 8.7. We now prove that Hypothesis (e) holds also.

Lemma 8.20 *Hypothesis (e) of Theorem 7.16 is satisfied.*

Proof Let q be a vertex of $\hat{\mathcal{M}}$ that q belongs to $U_p = B(p, r) \cap \mathcal{M}$ where $r = \varepsilon \text{rch}(M) (1+2\varepsilon)$. We write $q' = \Pi_p(q)$ for the projection of q onto T_p and $\text{star}'(p) = \Pi_p(\text{star}(p))$ for the projection of $\text{star}(p)$ onto T_p . To prove that Hypothesis (e) is satisfied, it is sufficient to prove that q' is a vertex of the k -dimensional weighted Delaunay triangulation $\text{Del}(\Psi_p(P))$ (see Lemma 8.7), which is equivalent to proving the following claim.

Claim *The Voronoi cell of q intersects T_p .*

Proof of the claim. To prove the claim, we will make use of Lemma 7.6 that states that $V(q)$ has a large extent in N_q and therefore intersects T_p . Denote by q'' the (unique) intersection point of N_q and T_p , and by q' the projection of q onto T_p . Since $q \in U_p$, we have $\|p - q\| < \varepsilon(1 + 2\varepsilon) \text{rch}(\mathcal{M}) = \varepsilon' \text{rch}(\mathcal{M})$, and it follows from Lemma 7.8 that $\|q - q'\| < \frac{\varepsilon'^2 \text{rch}(\mathcal{M})}{2}$. Write now $\theta = \angle(T_p, T_q)$ and recall that $\theta < 2\varepsilon'$ (Lemma 7.10). It follows that

$$\|q - q''\| < \frac{\|q - q'\|}{\cos \theta} < \text{rch}(\mathcal{M})$$

by our assumption on ε . It then follows from Lemma 7.6 that $q'' \in V(q) \cap T_p$, which proves the claim. $\square$

The claim implies that q' has a non empty cell in the weighted Voronoi diagram $\text{Vor}(\Psi_p(P))$ embedded in T_p (see Section 8.7). Equivalently, q is a vertex of $\text{Del}_p(P)$ and q' is a vertex of $\text{Del}(\Psi_p(P))$. Since $\text{Del}(\Psi_p(P))$ is a triangulation embedded in T_p (under the transversality condition of Section 8.2.1), q' cannot belong to $\text{star}'(p)$ without being a vertex of $\text{star}'(p)$ and, accordingly, q must be a vertex of $\text{star}(P)$. Therefore, Hypothesis 4' holds and the lemma is proved. $\square$

This ends the proof of Lemma 8.20. All hypotheses of Theorem 7.16 are satisfied. This ends the proof of Theorem 8.18. $\square$

8.4 Exercises

Exercise 8.1 Given is an $(\varepsilon, \bar{\eta})$ -net P of a differentiable (unknown) submanifold $\mathcal{M} \in \mathbb{R}^d$. Propose a method to approximate the tangent space T_p of $\mathcal{M}$ at $p \in P$.

Exercise 8.2 (Distance between components) Let P be an ε -dense sample of a submanifold $\mathcal{M}$ of positive reach $\text{rch}(\mathcal{M})$. Prove that each connected component of $\mathcal{M}$ contains at least one point of P and therefore at least one vertex of $\text{Del}_{T\mathcal{M}}(P)$. (Hint : show that the distance between any two connected components of $\mathcal{M}$ is at least $2\text{rch}(\mathcal{M})$).

Exercise 8.3 Show that Algorithm 8 terminates if the condition expressed as Equation 8.8 is satisfied. Observe that, for given values of $\bar{\eta}$ and $\tilde{w}_0$, the condition is satisfied if Θ_0 is sufficiently small and ε is sufficiently small compared to Θ_0^k .

Sketch of the proof. Equation 8.8 is the same as Equation 5.12 of Theorem 5.21 applied to the case of Algorithm 8. We use the notations of Section 8.3.2.

Using the bounds on the weighted radius and the diameter of a resampled flake σ (Lemma 8.14), we first bound, as in Section 5.3, the measure of the weight range $I(\sigma, p)$ of a vertex p of σ :

$$\begin{aligned} |I(\sigma, p)| &\leq 4\Theta_0 \frac{\Delta(\sigma)^2}{L(\sigma)} R(\hat{\sigma}) \\ &\leq \frac{4\Theta_0}{\varepsilon\bar{\eta}} \frac{4\hat{\alpha}_1^2 \hat{\alpha}_0}{\cos^3 \theta} \text{rch}(\mathcal{M})^2 \\ &\leq \frac{16\Theta_0}{\bar{\eta}} \frac{(1 + 2\sqrt{\tilde{w}_0})(1 + 4\sqrt{\tilde{w}_0})^2}{\cos^3 \theta} \varepsilon^2 \text{rch}(\mathcal{M})^2. \end{aligned}$$

Then, the probability ϖ that such a flake occurs is at most $\frac{|I(\sigma, p)|}{\tilde{w}_0 L(p)^2}$ (Lemma 5.23), where $L(p)$ is the length of the shortest edge incident to p in $\text{Del}(\hat{P})$. Thus

$$\varpi \leq \frac{16\Theta_0}{\tilde{w}_0 \bar{\eta}^3} \frac{(1 + 2\sqrt{\tilde{w}_0})(1 + 4\sqrt{\tilde{w}_0})^2}{\cos^3 \theta}.$$

We then bound $\Gamma + 1$ where Γ is the number of events overlapping a given flake σ . Overlapping events are flakes that share at least one vertex. Since resampled flakes have their diameter bounded by $\Delta_m = \frac{2\hat{\alpha}_1 \text{rch}(\mathcal{M})}{\cos^3 \theta}$, flakes overlapping a given flake σ have their vertices within the ball $B(c(\sigma), \frac{3}{2}\Delta_m)$ where $c(\sigma)$ is the circumcenter of σ . Using then the fact that two vertices are at least $\bar{\eta}\varepsilon \text{rch}(\mathcal{M})$ apart, we bound the number of vertices in this ball and then $\Gamma + 1$ using a now standard volume argument:

$$\begin{aligned} \Gamma + 1 &\leq \left(\frac{\frac{3}{2}\Delta_m}{\frac{1}{2}\bar{\eta}\varepsilon \text{rch}(\mathcal{M})} \right)^{d(d+1)} \\ &\leq \left(\frac{6\hat{\alpha}_1}{\bar{\eta} \cos \theta} \right)^{d(d+1)} \approx \left(\frac{6(1 + 4\sqrt{\tilde{w}_0})}{\bar{\eta} \cos \theta} \right)^{d(d+1)} \end{aligned}$$

The condition for the termination of Algorithm 8 is then obtained by plugging the bounds on ϖ and $\Gamma + 1$ into Equation 5.12 of Theorem 5.21.

8.5 Bibliographical notes

Alpha-shapes and weighted alpha-shapes are among the first tools introduced in the area of shape reconstruction where one seeks to construct an approximation of the shape of a three-dimensional object from a set of points measured on the boundary of the object. References are provided in the bibliographical notes of Chapter 6. Theorem 8.1 that states that an α -complex of a finite sample of a manifold has the same homotopy type as the manifold is due to Niyogi, Smale and Weinberger [115].

The tangential complex has been independently defined by Freedman [83] and by Boissonnat and Flottoto [12]. Boissonnat and Ghosh later showed how to remove inconsistencies in the tangential complex by star stitching [13] and proved Theorems 8.17 and 7.16. The tangential complex can be seen as a light variant of the cocone introduced by Cheng, Dey and Ramos [56].

The approach followed in this chapter that defines local triangulations and remove inconsistencies among the local triangulations has been pioneered by Shewchuk to maintain triangulations of moving points [126] and by Boissonnat, Wormser and Yvinec to generate anisotropic meshes [14]. The central question behind this approach is the stability of Delaunay triangulations and the existence and construction of Delaunay triangulations on manifolds [20, 19, 18].

In this chapter, we have assumed that the dimension of the submanifold is known and that the tangent space can be computed at any data point. Giesen and Wagner have shown how to estimate the dimension [91]. Estimating the tangent space can be done using principal component analysis (PCA) [99].

An implementation of the tangential complex can be found in the Gudhi library [130]. Fig. 8.6 shows a projection in $\mathbb{R}^3$ of the reconstruction of a Riemann surface embedded in $\mathbb{R}^8$.




Figure 8.6: A Riemann surface embedded in $\mathbb{R}^8$ reconstructed using the tangential Delaunay complex. For visualization purposes, the surface has been projected in $\mathbb{R}^3$.

Part IV

Distance-based inference

Chapter 9

Stability of distance functions

In the sequel, all the considered shapes and their approximations are represented by compact subsets of an Euclidean space $\mathbb{R}^d$. We use indifferently the words shape and compact set. In this chapter, we address the general problem consisting in recovering the topology of a shape from an approximation.

In general, the topological and geometric invariants of a shape cannot be directly extracted from the corresponding invariants of an approximating shape. In particular this is always the case for a (continuous) shape K , e.g. a surface in $\mathbb{R}^3$, approximated by a finite point cloud data set K' : point clouds in themselves do not carry any non trivial topological or geometric structure. Moreover, the occurrence of some features may depend on a “scale” at which the data and the shape are considered: for example, viewed with human eyes, the surface of a real world object may look very regular but at a microscopic scale it appears as a much more complicated surface with many holes and tunnels. It is thus necessary to “build” some scale-dependant geometric structure on top of such point clouds to recover informations about the shapes they approximate. For that purpose, the approach we adopt in this chapter consists in considering the distance functions to compact sets and to compare the topology of the sublevel sets (i.e. the offsets) of close compact sets. The underlying intuition is that “at some scales” (i.e. for some range values of the offsets), two close compact sets should have the same offset topology as illustrated on figure 9.1. The goal of this chapter is to turn this intuition into a formal framework with rigorous statements. This requires to proceed in two steps. First, one needs to understand how the topology of the offsets K^r of a given compact set K evolve with the parameter r . The answer to this question is given by the theory of critical points for distance functions. Second, it is necessary to compare the topology of the offsets of two close (for the Hausdorff distance) compact sets. This leads to stability results and sampling conditions necessary to ensure correct geometric inference.

The general mathematical framework introduced in this chapter allows to generalize and extend the results of Chapter 6, Section 8.1 and Chapter 8 to a wide class of non smooth shapes encountered in practical applications.

9.1 Distance function and Hausdorff distance

Given a compact subset $K \subset \mathbb{R}^d$, the *distance function* d_K to K is the non-negative function defined by

$$d_K(x) = \inf_{y \in K} d(x, y) \quad \text{for all } x \in \mathbb{R}^d$$

where $d(x, y) = \|x - y\|$ is the euclidean distance between x and y in $\mathbb{R}^d$. The distance function to K is continuous and indeed 1-Lipschitz: for all $x, x' \in \mathbb{R}^d$, $|d_K(x) - d_K(x')| \leq \|x - x'\|$. Moreover, K is completely characterized by d_K since $K = d_K^{-1}(0)$.

For any non-negative real number r , the *r-offset* K^r of K is the r -sublevel set of d_K defined by

$$K^r = d_K^{-1}([0, r]) = \{x \in \mathbb{R}^d : d_K(x) \leq r\}.$$

Recall from Section 1.2.2 that if K and K' are two compact subset of $\mathbb{R}^d$, then the Hausdorff distance $d_H(K, K')$ is the infimum of the sets of non negative numbers $r \geq 0$ such that $K' \subset K^r$ and $K \subset K'^r$. Indeed, the Hausdorff distance can be expressed in various equivalent ways in terms of distance functions:

Proposition 9.1 *Let $K, K' \subset \mathbb{R}^d$ be two compact sets. The Hausdorff distance $d_H(K, K')$ between K and K' is defined by any of the following equivalent assertions:*

- $d_H(K, K')$ is the smallest number r such that $K \subset K'^r$ and $K' \subset K^r$.
- $d_H(K, K') = \max(\sup_{x \in K} d_{K'}(x), \sup_{x \in K'} d_K(x))$.
- $d_H(K, K') = \|d_K - d_{K'}\| := \sup_{x \in \mathbb{R}^d} |d_K(x) - d_{K'}(x)|$.

9.2 Critical points of distance functions

Given a compact set $K \subset \mathbb{R}^d$, the distance function d_K is usually not differentiable. For example, if K is the union of the four sides of a square in the plane, d_K is not differentiable along the diagonals of K . Nevertheless,




Figure 9.1: Various offsets of a point cloud data set sampled around a torus in $\mathbb{R}^3$.

it is possible to define a generalized gradient vector field $\nabla_K : \mathbb{R}^d \rightarrow \mathbb{R}^d$ for d_K that coincides with the classical gradient at the points where d_K is differentiable.

For any point $x \in \mathbb{R}^d$ we denote by $\Gamma_K(x)$ the set of points in K closest to x :

$$\Gamma_K(x) = \{y \in K : d(x, y) = d_K(x)\} \subset \mathbb{R}^d.$$

This is a non empty compact subset of K .

Let $B_K(x)$ be the smallest closed ball enclosing $\Gamma_K(x)$ and let $c_K(x)$ be its center and $F_K(x)$ its radius (see figure 9.2). For $x \in \mathbb{R}^d \setminus K$, the generalized gradient $\nabla d_K(x)$ is defined by

$$\nabla d_K(x) = \frac{x - c_K(x)}{d_K(x)}$$

and for $x \in K$, $\nabla d_K(x) = 0$.

The norm of the gradient is given by

$$\|\nabla d_K(x)\|^2 = 1 - \frac{F_K^2(x)}{d_K^2(x)}. \quad (9.1)$$




Figure 9.2: The gradient of the distance function.

Equivalently, the norm of $\nabla d_K(x)$ is the cosine of the half angle of the smallest circular cone with apex x that contains $\Gamma_K(x)$. Intuitively, the direction of $\nabla d_K(x)$ is the one along which the directional derivative of d_K is the largest or, in other words, the one in which the slope of the graph $\{(y, d_K(y)) : y \in \mathbb{R}^d\} \subset \mathbb{R}^{d+1}$ is the largest at the point $(x, d_K(x))$ (see Figure 9.3).

The map $x \in \mathbb{R}^d \rightarrow \nabla d_K(x) \in \mathbb{R}^d$ is in general not continuous. In other words, ∇d_K is a discontinuous vector field. Nevertheless it is possible to show [103, 120] that $x \rightarrow \|\nabla d_K(x)\|$ is a lower semi-continuous function, i.e. for any $a \in \mathbb{R}$, $\|\nabla d_K\|^{-1}((-\infty, a])$ is a closed subset of $\mathbb{R}^d$. Moreover, ∇d_K is integrable in the following sense.

Proposition 9.2 *There exists a continuous map $\mathfrak{C} : \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that for any $x \in \mathbb{R}^d$, the map $t \rightarrow \mathfrak{C}(t, x)$, called a trajectory of ∇d_K is a*




Figure 9.3: The graph of the distance to a square (in red) in the plane. Along the diagonal of the square, the direction of the gradient is given by the diagonals and its norm is the slope of the edges of the graph above the diagonals.

solution of the differential equation

$$\frac{dX}{dt} = \nabla d_K(X(t))$$

satisfying $X(0) = x$. Moreover this continuous trajectory can be parametrized by arc length $s \rightarrow \mathfrak{C}(t(s), x)$ and one has

$$d_K(\mathfrak{C}(t(l), x)) = d_K(x) + \int_0^l \|\nabla d_K(\mathfrak{C}(t(s), x))\| ds. \quad (9.2)$$

The above equation implies that d_K is non decreasing along the trajectories of ∇d_K . It can also be shown [103] that F_K is also non decreasing along the trajectories of ∇d_K .

The gradient ∇d_K allows to define the notion of critical point for d_K in the same way as for differentiable functions.

Definition 9.3 (Critical point) *A point x is a critical point of d_K if $\nabla d_K(x) = 0$. A real $c \geq 0$ is a critical value of d_K if there exists a critical point $x \in \mathbb{R}^d$ such that $d_K(x) = c$. A regular value of d_K is a value which is not critical.*

When there is no risk of confusion, we make the small abuse of language consisting in calling a critical (resp. regular) point of d_K a critical (resp. regular) point of K .

9.3 Topology of the offsets

Using the notion of critical point defined in the previous section, it appears that some properties of distance functions are similar to the ones of differentiable functions. In particular, the sublevel sets of d_K are topological submanifolds of $\mathbb{R}^d$ and their topology can change only at critical points. These properties are formalized in the following two theorems.

Theorem 9.4 *Let $K \subset \mathbb{R}^d$ be a compact set and let r be a regular value of d_K . The level set $d_K^{-1}(r)$ is a $(d-1)$ -dimensional topological submanifold of $\mathbb{R}^d$.*

Theorem 9.5 (Isotopy Lemma) *Let $K \subset \mathbb{R}^d$ be a compact set and let $r_1 < r_2$ be two real numbers such that $[r_1, r_2]$ does not contain any critical value of d_K . Then all the level sets $d_K^{-1}(r), r \in [r_1, r_2]$ are homeomorphic (and even isotopic) and the set $A(r_1, r_2) = \{x \in \mathbb{R}^d : r_1 \leq d_K(x) \leq r_2\}$ is homeomorphic to $d_K^{-1}(r_1) \times [r_1, r_2]$.*

An immediate consequence of these two results is that the topology of the offsets of K can only change at critical values and for any regular value r of d_K , the offset K^r is a d -dimensional topological manifold with boundary. In particular, the topology of the small offsets $K^r, r > 0$, cannot change while r is smaller than the smallest positive critical value of d_K (if it exists). This leads to the notion of *weak feature size*.

Definition 9.6 (Weak feature size) *Let $K \subset \mathbb{R}^d$ be a compact set. The weak feature size $\text{wfs}(K)$ of K is the infimum of the positive critical values of d_K . If d_K does not have critical values, $\text{wfs}(K) = +\infty$.*

It follows from the Isotopy Lemma 9.5 that if $0 < \alpha \leq \beta < \text{wfs}(K)$, then K^α and K^β are isotopic. In a more intuitive way, the knowledge of K at precision, or scale, α gives the same information for any choice of $0 < \alpha <$

$\text{wfs}(K)$. Moreover, the following result allows to compare the topology of the offsets of two close compact sets with positive weak feature sizes.

Theorem 9.7 *Let $K, K' \subset \mathbb{R}^d$ and $\varepsilon > 0$ be such that $d_H(K, K') < \varepsilon$, $\text{wfs}(K) > 2\varepsilon$ and $\text{wfs}(K') > 2\varepsilon$. Then for any $0 < \alpha \leq 2\varepsilon$, K^α and K'^α are homotopy equivalent.*

Proof Let $\delta > 0$ be such that $\text{wfs}(K) > 2\varepsilon + \delta$ and $\text{wfs}(K') > 2\varepsilon + \delta$. It is enough to prove that $K^{\delta+2\varepsilon}$ and $K'^{\delta+2\varepsilon}$ are homotopy equivalent. Since $d_H(K, K') < \varepsilon$, the following diagram, where each map is an inclusion, is commutative.

$$\begin{array}{ccccc}
 K^\delta & \xrightarrow{a_0} & K^{\delta+\varepsilon} & \xrightarrow{a_1} & K^{\delta+2\varepsilon} \\
 & \searrow d_0 & & \searrow d_1 & \\
 & c_0 & & c_1 & \\
 K'^\delta & \xrightarrow{b_0} & K'^{\delta+\varepsilon} & \xrightarrow{b_1} & K'^{\delta+2\varepsilon}
 \end{array}$$

The Isotopy Lemma 9.5 implies that the inclusions a_0, a_1, b_0 and b_1 are homotopy equivalences. Let s_0, s_1, r_0 and r_1 be the homotopic inverses of a_0, a_1, b_0 and b_1 respectively.

The following computation, where $\cong$ denotes the homotopy equivalence relation, shows that c_1 is an homotopy equivalence between $K'^{\delta+\varepsilon}$ and $K^{\delta+2\varepsilon}$ with homotopic inverse $r_1 \circ d_1 \circ s_1$:

$$\begin{aligned}
 c_1 \circ r_1 \circ d_1 \circ s_1 &\cong c_1 \circ (r_1 \circ b_1) \circ d_0 \circ s_0 \circ s_1 \\
 &\cong (c_1 \circ d_0) \circ s_0 \circ s_1 \\
 &\cong a_1 \circ a_0 \circ s_0 \circ s_1 \cong id_{K^{\delta+2\varepsilon}}
 \end{aligned}$$

Similarly, we get $r_1 \circ d_1 \circ s_1 \circ c_1 \cong id_{K'^{\delta+\varepsilon}}$

□

The previous theorem shows that the compact sets with positive weak feature size provide a class of compact sets with interesting topological stability properties. Moreover, it is possible to show that this class is large enough to include most of the shapes encountered in practical applications. In particular, smooth manifolds, polyhedra, polyhyedral sets, semi-algebraic sets and more generally the so-called subanalytic compact sets (i.e. obtained from

analytic equations and inequalities) all have positive weak feature size. Nevertheless, the previous theorem suffers from a weakness that prevents it to be really useful in practice. Indeed, the assumption involving both $\text{wfs}(K)$ and $\text{wfs}(K')$ is hardly satisfied in practical situations. For example if K is a continuous shape approximated by a finite point cloud K' , $\text{wfs}(K')$ is equal to half of the distance between the two closest points of K' which is usually smaller than $2d_H(K, K')$ as illustrated on figure 9.4. As a consequence, even if the weak feature size of K is large, whatever the quality of the approximation of K by K' the assumptions of the theorem 9.7 may not be satisfied. This phenomenon can also be interpreted as a lack of continuity of the map $K \rightarrow \text{wfs}(K)$ or as an instability property of the critical points of distance functions.




Figure 9.4: A segment K approximated by a point cloud K' (in red). The weak feature size of K' is obviously smaller than two times the Hausdorff distance between K and K' .

9.4 Stability of critical points

Since the topology of the offsets of a compact set can only change for critical values of d_K , it is natural to study the stability of these critical points when K is replaced by a close compact set K' . Unfortunately, it appears that the critical points are unstable, as illustrated on figure 9.5.

To overcome this stability problem we introduce a parametrized notion of critical point.

Definition 9.8 (μ -critical points) *Let $K \subset \mathbb{R}^d$ be a compact set and let $0 \leq \mu \leq 1$. A point $x \in \mathbb{R}^d$ is μ -critical for d_K if $\|\nabla d_K(x)\| \leq \mu$.*

Note that the notion of 0-critical point coincide with the notion of critical point of Definition 9.3 . The family of μ -critical points satisfies the following fundamental stability property.




Figure 9.5: When K is a rectangle, the set of critical points of d_K is the whole segment $[AB]$. This segment collapses to a single critical point A' as one stretches the bottom side of K to obtain K' . Nevertheless, along the previously critical segment, the norm of the gradient of $d_{K'}$ remains small.

Theorem 9.9 (Critical point stability theorem) *Let K and K' be two compact subsets of $\mathbb{R}^d$ and $d_H(K, K') \leq \varepsilon$. For any μ -critical point x of K , there is a $(2\sqrt{\varepsilon/d_K(x)} + \mu)$ -critical point of K' at distance at most $2\sqrt{\varepsilon d_K(x)}$ from x .*

The proof of this theorem follows from two technical lemmas. The first one shows that the function d_K cannot grow too fast in a neighborhood of a μ -critical point.

Lemma 9.10 *Let $K \subset \mathbb{R}^d$ be a compact set and x one of its μ -critical points. For any $y \in \mathbb{R}^d$, we have:*

$$d_K^2(y) \leq d_K^2(x) + 2\mu d_K(x) \|x - y\| + \|x - y\|^2.$$

Proof Let $\Gamma = \Gamma_K(x)$ be the set of points closest to x on K , and let S be the sphere with center x and radius $d_K(x)$. Let also $c = c_K(x)$ be the center of the minimal enclosing ball of Γ , and $\alpha = \arccos^{-1}(\mu)$ (see Figure 9.6).




Figure 9.6: Proof of lemma 9.10

For any $x' \in \Gamma$ we have

$$\begin{aligned}
 d_K^2(y) \leq \|y - x'\|^2 &= ((y - x) + (x - x')) \cdot (y - x) + (x - x') \cdot (y - x) \\
 &= \|y - x\|^2 + \|x - x'\|^2 + 2((y - x) \cdot (x - x')) \\
 &= d_K^2(x) + 2d_K(x)\|x - y\| \cos(y - x, x - x') + \|x - y\|^2
 \end{aligned}$$

To prove the lemma it is thus sufficient to prove the following claim.

Claim: there exists a point $x' \in \Gamma$ such that the angle between $(y - x)$ and $(x' - x)$ is not smaller than α .

We distinguish between two cases.

Case 1: $\mu \neq 0$

Assume that the claim is not satisfied. Then for any $x' \in \Gamma$ the angle between $(x' - x)$ and $(y - x)$ is smaller than α . Since Γ is compact, there exists $\alpha' < \alpha$ such that Γ is contained in the “circular” cone with apex x and axis the half-line directed by $y - x$ and apex angle α' . This cone intersects S along a $(d - 2)$ -sphere with center c' and radius $R' = d_K(x) \sin \alpha'$. Since Γ is also contained in S , Γ is contained in the ball of center c' and radius $R' = d_K(x) \sin \alpha' < d_K(x) \sin \alpha = F_K(x)$: a contradiction.

Case 2: $\mu = 0$

$\alpha = \frac{\pi}{2}$ and Γ cannot be enclosed in any hemisphere of S . So there is at least one point $x' \in \Gamma$ such that the angle between $(y - x)$ and $(x' - x)$ is not smaller than $\frac{\pi}{2}$. $\square$

The next lemma allows to study the behavior of the μ -critical points when K is replaced by a close approximation K' .

Lemma 9.11 *Let K and K' be two compact subsets of $\mathbb{R}^d$ and $d_H(K, K') \leq \varepsilon$. For any μ -critical point x of K and any $\rho > 0$, there is a μ' -critical point of K' at distance at most ρ from x , with:*

$$\mu' \leq \frac{\rho}{2d_K(x)} + \frac{2\varepsilon}{\rho} + \mu$$

Proof Let us consider the trajectory $s \rightarrow C(s)$ of the vector field $\nabla d_{K'}$ parameterized by arc length and starting at x . If C reaches a critical point of K' before $s = \rho$, the lemma holds. Assume this is not the case. Letting $y = C(\rho)$, we have:

$$d_{K'}(y) - d_{K'}(x) = \int_0^\rho \|\nabla d_{K'}(C(s))\| ds$$

Therefore, there must exist a point p on the curve C between $s = 0$ and $s = \rho$ such that:

$$\|\nabla d_{K'}(p)\| \leq \frac{d_{K'}(y) - d_{K'}(x)}{\rho} \quad (9.3)$$

The curve C being parametrized by arc length, note that $\|p - x\| \leq \rho$. Now Lemma 9.10 applied to x , y , and K reads:

$$d_K(y) \leq \sqrt{d_K^2(x) + 2\mu d_K(x)\|x - y\| + \|x - y\|^2}$$

Also, since $\varepsilon = d_H(K, K')$, we have that for all $z \in \mathbb{R}^d$, $|d_K(z) - d_{K'}(z)| \leq \varepsilon$. Hence:

$$\begin{aligned} d_{K'}(y) - d_{K'}(x) &\leq \sqrt{d_K^2(x) + 2\mu d_K(x)\|x - y\| + \|x - y\|^2} \\ &\quad - d_K(x) + 2\varepsilon \\ &\leq d_K(x) \left[\sqrt{1 + \frac{2\mu\|x - y\|}{d_K(x)} + \frac{\|x - y\|^2}{d_K^2(x)}} - 1 \right] \\ &\quad + 2\varepsilon \\ &\leq \mu\|x - y\| + \frac{\|x - y\|^2}{2d_K(x)} + 2\varepsilon \end{aligned}$$

the last inequality coming from the fact that $\sqrt{1+u} \leq 1 + \frac{u}{2}$ for $u \geq 0$. Noticing that $\|x - y\| \leq \rho$, dividing by ρ , and applying equation (9.3) shows that p satisfies the desired requirements. $\square$

Proof [of Theorem 9.9] The bound of the previous lemma can be optimized by choosing $\rho = 2\sqrt{\varepsilon d_K(x)}$. It then becomes equal to $2\sqrt{\varepsilon/d_K(x)} + \mu$. The theorem follows immediately. $\square$

Remark 9.12 *Note that since $d_{K'}$ is increasing along the trajectories of $\nabla d_{K'}$ (see Equation 9.2), the μ' -critical point p for $d_{K'}$ of Lemma 9.11 can be chosen such that $d_{K'}(p) \geq d_{K'}(x)$.*

9.5 The critical function of a compact set

The critical point stability Theorem 9.9 plays a fundamental role to get topological stability results. It allows to introduce a general framework for inferring the topology and the geometry of a large class of (non-smooth) compact sets. For that purpose, we first introduce a one variable real-valued function that encodes the criticality of the level sets of d_K .

Definition 9.13 *The critical function of a compact set $K \subset \mathbb{R}^d$, $\chi_K : (0, +\infty) \rightarrow \mathbb{R}_+$ is defined by*

$$\chi_K(r) = \inf_{x \in d_K^{-1}(r)} \|\nabla d_K(x)\|$$

An example of a critical function is given in Figure 9.7. Note that from the isotopy Lemma 9.5 the zeros of the critical functions correspond to the changes in the topology of the offsets of K . As we will see later, the main interest of the critical function χ_K is to provide information about the topological stability of some offsets of the compact sets contained in a neighborhood of K . In particular, whether a compact set K is a Hausdorff approximation of a simple compact set or not can be directly read from its critical function.

Using the critical points stability Theorem 9.9, we easily get the following stability result for the critical function.




Figure 9.7: The critical function of a square embedded in $\mathbb{R}^3$ with edge length equal to 50 (top) and the critical function of a point cloud sampling this square (bottom).

Theorem 9.14 (Critical function stability theorem) *Let K and K' be two compact subsets of $\mathbb{R}^d$ such that $d_H(K, K') \leq \varepsilon$. For all $r \geq 0$, we have:*

$$\inf\{\chi_{K'}(u) \mid u \in I(r, \varepsilon)\} \leq \chi_K(r) + 2\sqrt{\frac{\varepsilon}{r}}$$

where $I(r, \varepsilon) = [r - \varepsilon, r + 2\chi_K(r)\sqrt{\varepsilon r} + 3\varepsilon]$

This result shows that if the critical function of K' is not smaller than some value α on the interval $I(d, \varepsilon)$ then the critical function of K at the point r

cannot be smaller than $\alpha - 2\sqrt{\frac{\varepsilon}{r}}$. In particular, if $\alpha > 2\sqrt{\frac{\varepsilon}{r}}$ then r cannot be a critical value of d_K . Since the topology of the offsets of K can only change at critical values, it is thus possible to locate intervals on which the topology of the offsets of K does not change. The Figures 9.7 and 9.8 illustrate this property.

From an algorithmic point of view, it is not difficult to see that when K is a finite point cloud, the critical function of K can be easily computed from the Voronoi diagram of K - see Exercise 9.11.

Proof Let $r \geq 0$ and let $x \in d_K^{-1}(r)$ be such that $\|\nabla d_K(x)\| = \chi_K(r)$. The existence of such a point x comes from the fact that the infimum involved in the definition of χ_K is indeed a minimum. This follows from the lower semi-continuity of $\|\nabla d_K\|$ and the compactness of $d_K^{-1}(r)$. The critical point stability Theorem 9.9 implies that there exists a point p which is $(2\sqrt{\frac{\varepsilon}{r}} + \chi_K(r))$ -critical for $d_{K'}$ at distance at most $2\sqrt{\varepsilon r}$ from x . Applying Lemma 9.10 to x , p and K we get

$$\begin{aligned} d_K(p) &\leq \sqrt{r^2 + 4\chi_K(r)r\sqrt{\varepsilon r} + 4\varepsilon r} \\ &\leq r\sqrt{1 + 4\chi_K(r)\sqrt{\varepsilon/r} + 4\varepsilon/r} \\ &\leq r + 2\chi_K(r)\sqrt{\varepsilon r} + 2\varepsilon \end{aligned}$$

Now, according to Remark 9.12, p can be chosen such that $d_{K'}(p) \geq d_{K'}(x)$. Using that $|d_{K'}(p) - d_K(p)| \leq \varepsilon$, the theorem follows from the above inequality. $\square$

The example of Figure 9.8 illustrates the critical function Theorem 9.14. The critical function of a point cloud sampling a torus shape reveals three intervals with stable topology for K : the first one corresponds to offsets having the topology of a torus (bottom left), the second one corresponds to solid torus with a hole homeomorphic to a ball inside (bottom middle - not visible from outside) and the third one is unbounded and correspond to offsets that have the topology of a ball (bottom right).

9.6 Sampling conditions and μ -reach

In this section we introduce the μ -reach, a stronger regularity condition than the weak feature size, that allows to give stronger reconstruction results than Theorem 9.7.




Figure 9.8: A points set (left) sampled around a torus shape in $\mathbb{R}^3$ and its critical function (the upper curve). The lowest curve represents the lower bound for the critical function of any shape K at distance less than some fixed threshold (here 0.001, the diameter of the torus being 10) from the point cloud.

Definition 9.15 For $0 < \mu \leq 1$, the μ -reach $r_\mu(K)$ of a compact set $K \subset \mathbb{R}^d$ is defined as

$$r_\mu(K) = \inf\{r > 0 : \chi_K(r) < \mu\}$$

By analogy with the wfs, the μ -reach is the infimum of the μ -critical values of d_K . When $\mu = 1$, $r_\mu(K)$ is known as the *reach*. When K is a compact smooth submanifold of $\mathbb{R}^d$, it coincides with the reach defined in Chapter 7. The function $\mu \rightarrow r_\mu(K)$ is non increasing and we have

$$\lim_{\mu \rightarrow 0^+} r_\mu(K) \leq \text{wfs}(K)$$

Note that the above inequality can be strict (Exercise 9.10).

It follows from the critical point stability Theorem 9.9 that the positiveness of the μ -reach of a compact set K' implies some constraints on the location

of the critical points of any close enough approximation K of K' . More precisely we have the following result (Exercise 9.12).

Theorem 9.16 (Critical values separation theorem) *Let K and K' be two compact subsets of $\mathbb{R}^d$, ε the Hausdorff distance between K and K' , and μ a non-negative number. The distance function d_K has no critical values in the interval $(4\varepsilon/\mu^2, r_\mu(K') - 3\varepsilon)$. Besides, for any $\mu' < \mu$, χ_K is larger than μ' on the interval*

$$\left(\frac{4\varepsilon}{(\mu - \mu')^2}, r_\mu(K') - 3\sqrt{\varepsilon r_\mu(K')}\right).$$

Note that taking μ too small does not give any information on the critical values, since the lower bound then exceeds the upper bound. It is also possible to build examples showing that the bounds of the interval in the above theorem are tight.

The notion of μ -reach allows to introduce the following sampling condition.

Definition 9.17 *Given two positive real numbers κ and μ , one says that a compact set $K \subset \mathbb{R}^d$ is a (κ, μ) -approximation of a compact set $K' \subset \mathbb{R}^d$ if*

$$d_H(K, K') \leq \kappa r_\mu(K').$$

The notion of (κ, μ) -approximation generalizes the notion of $(\varepsilon, \bar{\eta})$ -net introduced for smooth submanifolds in Chapter 7.

9.7 Offset reconstruction

Equipped with the notion of (κ, μ) -approximation and the stability properties of critical points of distance functions proved in the previous section, we are now able to get easily offsets reconstruction results from approximations.

Theorem 9.18 (Isotopic reconstruction theorem) *Let $K' \subset \mathbb{R}^d$ be a compact set such that $r_\mu = r_\mu(K') > 0$ for some $\mu > 0$. Let K be a (κ, μ) -approximation of K' where*

$$\kappa < \min\left(\frac{\sqrt{5}}{2} - 1, \frac{\mu^2}{16 + 2\mu^2}\right).$$

Assume that r, r' are such that

$$0 < r' < \text{wfs}(K') \quad \text{and} \quad \frac{4\kappa r_\mu}{\mu^2} \leq r < r_\mu(K') - 3\kappa r_\mu.$$

Then the offset K^r and level set $d_K^{-1}(r)$ are isotopic to $K^{r'}$ and $d_{K'}^{-1}(r')$ respectively.

The proof of the isotopy between the offsets is beyond the scope of this book. We prove here the following weaker version.

Theorem 9.19 (Homotopic reconstruction theorem) *Let $K \subset \mathbb{R}^d$ be a (κ, μ) -approximation of a compact set $K' \subset \mathbb{R}^d$. If*

$$\kappa < \frac{\mu^2}{5\mu^2 + 12}$$

then the complement of K^α is homotopy equivalent to the complement of K' and K^α is homotopy equivalent to K'^η as soon as

$$0 < \eta < \text{wfs}(K') \quad \text{and} \quad \frac{4d_H(K, K')}{\mu^2} \leq \alpha < r_\mu(K') - 3d_H(K, K')$$

Proof The critical values separation Theorem 9.16 applied to K and K' ensures that d_K does not have any critical value in the interval $(4\varepsilon/\mu^2, r_\mu(K') - 3\varepsilon)$ where $\varepsilon = d_H(K, K')$. It follows from the isotopy Lemma 9.5 that all the offsets of K corresponding to the values contained in this interval are isotopic. It is thus sufficient to prove the theorem for $\alpha = 4\varepsilon/\mu^2$. Since the critical functions of K and K^α are related by the relation

$$\chi_{K^\alpha}(r) = \chi_K(r + \alpha)$$

(see Exercise 9.9), we have $\text{wfs}(K^\alpha) \geq r_\mu(K') - 3\varepsilon - 4\varepsilon/\mu^2$. We also have

$$d_H(K^\alpha, K') \leq \varepsilon + \frac{4\varepsilon}{\mu^2}$$

According to Theorem 9.7, the theorem holds as soon as

$$d_H(K^\alpha, K') \leq \frac{1}{2} \min(\text{wfs}(K^\alpha), \text{wfs}(K'))$$

An easy computation shows that this inequality holds when $\kappa < \frac{\mu^2}{5\mu^2 + 12}$. $\square$

9.8 Exercises

Exercise 9.1 Prove that d_K is 1-Lipschitz.

Exercise 9.2 Prove Proposition 9.1.

Exercise 9.3 Prove that $\sigma_K(x)$ the smallest enclosing ball containing $\Gamma_K(x)$ exists and is unique.

Exercise 9.4 Show that the map from $\mathbb{R}^d$ to the space of compact subsets of $\mathbb{R}^d$ is semi-continuous, i.e.

$$\forall x, \forall r > 0, \exists \alpha > 0, \|y - x\| \leq \alpha \Rightarrow \Gamma_K(y) \subset \{z : d(z, \Gamma_K(x)) \leq r\}$$

Exercise 9.5 Show that for any $x \in \mathbb{R}^d$, $c_K(x)$ is the point on the convex hull of $\Gamma_K(x)$ closest to x .

Exercise 9.6 Prove that for any $x \in \mathbb{R}^d$ one has the following equivalence: $(x \text{ is a critical point of } d_K) \Leftrightarrow (x \text{ lies in the convex hull of } \Gamma_K(x))$.

Exercise 9.7 Let P be a finite set of points in $\mathbb{R}^2$ in general position. Prove that a point $x \in \mathbb{R}^2$ is a critical point of d_P if and only if it satisfies one of the following conditions:

- $x \in P$,
- x is the intersection point between a Voronoi edge and its dual Delaunay edge,
- x is a Voronoi vertex contained in its dual Delaunay triangle.

How does this result generalize for finite point clouds in higher dimensions?

Exercise 9.8 Let $K = \{p_1, \dots, p_n\} \subset \mathbb{R}^d$ be a finite point set. Prove that $\text{wfs}(K) = \frac{1}{2} \min_{i \neq j} \|p_i - p_j\|$.

Exercise 9.9 Show that for any compact set $K \subset \mathbb{R}^d$ and any $\alpha \geq 0$,

$$\chi_{K^\alpha}(r) = \chi_K(r + \alpha) \quad \text{for all } r \geq 0.$$

Hint: first prove the same kind of relation between d_{K^α} and d_K .

Exercise 9.10 Give an example of a compact set (e.g. a compact subset of $\mathbb{R}^2$) K such that $\lim_{\mu \rightarrow 0^+} r_\mu(K) \neq \text{wfs}(K)$.

Exercise 9.11 Let $P \subset \mathbb{R}^d$ be a finite set of points. We denote $\text{Vor}(P)$ and $\text{Del}(P)$ the Voronoi diagram and Delaunay triangulation of P . We use the notations of Section 9.2.

1. What can be said about the restriction of F_P to a Voronoi cell of $\text{Vor}(P)$?
2. For each Delaunay simplex τ , let V_τ its dual Voronoi cell and let $f_\tau : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ the function defined by

$$f_\tau(r) = \sqrt{1 - \frac{F_P^2(\tau)}{r^2}} \quad \text{if } r \in d_P(V_\tau)$$

and $f_\tau(r) = 1$ otherwise. Show that the critical function of P is equal to the minimum, taken over all the Delaunay simplices of $\text{Del}(P)$, of the functions f_τ .

3. Deduce from the previous question an algorithm that takes $\text{Del}(P)$ as input and output the critical function of P .

Exercise 9.12 Prove the critical values separation Theorem 9.16. Hint: this is a consequence of Theorem 9.14 - see also [37].

9.9 Bibliographical notes

The distance functions framework for geometric inference has been introduced in [37] and the proof of the isotopic reconstruction Theorem 9.18 is given in [47]. It is not restricted to topological inference and shape reconstruction but has been extended to other geometric inference problems. For example, it has been used to prove stability results for normals [47] and curvatures [39] estimation of compact sets with positive μ -reach. It has also been used to prove that some smoothing operations involving offsets of shapes in Computer Aided Geometric Design are theoretically well-founded [38].

Distance functions have been widely studied and used in Riemannian geometry [53, 93] and non-smooth analysis [58]. The notion of critical point introduced in this chapter coincides with the notion of critical point for distance function used in Riemannian geometry and non-smooth analysis where the notion of Clarke gradient [58] is closely related to the above defined gradient. The general properties of the gradient of d_K and of its trajectories given in Section 9.2, are established in [103] or, in a more general setting, in [120].

Distance functions are also particular cases of the so-called semiconcave functions. Many of the results presented in this section can be deduced from general properties of semiconcave functions [120]. This allows in particular to extend most of the results given in this section to compact subsets of Riemannian manifolds.

When K is a finite set of points, several variants of the gradient flow $\mathfrak{C}$ defined in this chapter have been previously and independently considered in the literature [68, 90, 34].

The notion of weak feature size has been introduced in [43, 44, 45]. The positiveness of the weak feature size of large classes of compact sets is discussed in [86, p. 1045] and [44], proposition 3.6. The notion of reach for compact subsets of $\mathbb{R}^d$ has been introduced by H. Federer [82] in the setting of geometric measure theory.

Chapter 10

Distance to probability measures

The problem of outliers As we have seen in Chapter 9, the use of distance functions provides an interesting approach for the robust estimation of the topological and geometric properties of shapes in $\mathbb{R}^d$ from approximating data. Here, approximation is meant with respect to the Hausdorff distance that requires all the data points to be located in a close neighborhood of the shape. However, in many practical applications the data come with outliers, i.e. observations (points) that are not located close to the approximated shape. For such data the Hausdorff distance is no longer relevant to formalize the notion of approximation: just adding one point p at distance R from a given data set K makes the Hausdorff distance between K and $K \cup \{p\}$ equal to R (see Figure 10.1). As a consequence, the distance-based approach of Chapter 9 fails for data corrupted by noise and outliers as illustrated on Figure 10.2. To overcome this issue, in this chapter, we adapt the distance-based framework for geometric inference to the general framework of data carrying noise and outliers.




Figure 10.1: Adding just one point at distance R to a point cloud sampling a circle changes the Hausdorff distance between the shape and the sample by R .

10.1 Extending the sampling theory for compact sets

All the inference results of Chapter 9 follow from only three fundamental properties of distance functions:




Figure 10.2: In addition to a set of points densely sampled on the surface S of a tangle cube we consider a single “outlier” located away from S (left). When considering the offsets, the new added point creates a connected component that makes the estimation of the topology of S from the offsets of the data impossible (right): e.g. the estimated number of connected components (two) is clearly wrong.

- Stability of the map $K \rightarrow d_K$: for any compact subsets K, K' of $\mathbb{R}^d$ we have

$$\|d_K - d_{K'}\|_\infty = d_H(K, K')$$

where $\|d_K - d_{K'}\|_\infty = \sup_{x \in \mathbb{R}^d} |d_K(x) - d_{K'}(x)|$.

- For any compact set $K \subset \mathbb{R}^d$, the distance function d_K is 1-Lipschitz: for any $x, x' \in \mathbb{R}^d$, $|d_K(x) - d_K(x')| \leq \|x - x'\|$.
- For any compact set $K \subset \mathbb{R}^d$, the distance function d_K^2 is 1-semiconcave: $x \rightarrow \|x\|^2 - d_K^2(x)$ is convex (see Exercise 10.1).

The first property is an obvious condition to ensure that the offsets of two close compact sets are close to each other. The second and third properties are the fundamental ingredients to prove the existence and integrability of the gradient of d_K (Section 9.2) and the isotopy lemma of Section 9.3. These results still hold for general proper semiconcave functions, motivating the following definition of functions that are of particular interest for geometric inference.

Definition 10.1 A non-negative function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^+$ is a distance-like if

- i) ϕ is 1-Lipschitz,
- ii) ϕ^2 is 1-semiconcave, i.e. $x \rightarrow \|x\|^2 - \phi^2(x)$ is convex
- iii) ϕ is proper i.e., for any compact set $K \subset \mathbb{R}$, $\phi^{-1}(K)$ is compact.

Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a distance-like function. The 1-semiconcavity of ϕ^2 allows to define a notion of gradient vector field $\nabla\phi(x)$ for ϕ , defined everywhere and satisfying $\|\nabla\phi(x)\| \leq 1$. Although not continuous, the vector field $\nabla\phi$ is sufficiently regular to be integrated in a continuous locally Lipschitz flow $\Phi^t : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $t \geq 0$. The flow Φ^t integrates the gradient $\nabla\phi$ in the sense that for every $x \in \mathbb{R}^d$, the curve $\gamma : t \mapsto \Phi^t(x)$ is right-differentiable, and for every $t > 0$, $\left. \frac{d\gamma}{dt} \right|_{t-} = \nabla\phi(\gamma(t))$. Moreover, for any integral curve $\gamma : [a, b] \rightarrow \mathbb{R}^d$ parametrized by arc-length, one has:

$$\phi(\gamma(b)) = \phi(\gamma(a)) + \int_a^b \|\nabla\phi(\gamma(t))\| dt.$$

Definition 10.2 Let ϕ be a distance-like function. We denote by $\phi^r = \phi^{-1}([0, r])$ the r sublevel set of ϕ .

1. A point $x \in \mathbb{R}^d$ will be called α -critical (with $\alpha \in [0, 1]$) if the inequality $\phi^2(x+h) \leq \phi^2(x) + 2\alpha\|h\|\phi(x) + \|h\|^2$ is true for all $h \in \mathbb{R}^d$. A 0-critical point is simply called a critical point. It follows from the 1-semiconcavity of ϕ^2 that $\|\nabla\phi(x)\|$ is the infimum of the $\alpha \geq 0$ such that x is α -critical.
2. The weak feature size of ϕ at r is the maximum $r' > 0$ such that ϕ doesn't have any critical value between r and $r+r'$. We denote it by $\text{wfs}_\phi(r)$. For any $0 < \alpha < 1$, the α -reach of ϕ is the maximum r such that $\phi^{-1}((0, r])$ does not contain any α -critical point.

Notice that the α -reach is always a lower bound for the weak-feature size, with $r = 0$.

The Isotopy Lemma 9.5 extends to distance-like functions.

Theorem 10.3 (Extended isotopy lemma) Let ϕ be a distance-like function and $r_1 < r_2$ be two positive numbers such that ϕ has no critical points in the subset $\phi^{-1}([r_1, r_2])$. Then all the sublevel sets $\phi^{-1}([0, r])$ are isotopic for $r \in [r_1, r_2]$.

The proof of the following theorem, showing that the offsets of two uniformly close distance-like functions with large weak feature size have the same homotopy type, relies on Theorem 10.3 and is almost verbatim the same as the one of Theorem 9.7.

Proposition 10.4 *Let ϕ and ψ be two distance-like functions, such that $\|\phi - \psi\|_\infty \leq \varepsilon$. Suppose moreover that $\text{wfs}_\phi(r) > 2\varepsilon$ and $\text{wfs}_\psi(r) > 2\varepsilon$. Then, for every $0 < \eta \leq 2\varepsilon$, $\phi^{r+\eta}$ and $\psi^{r+\eta}$ have the same homotopy type.*

The Critical Point Stability Theorem 9.9 also holds for distance-like functions.

Theorem 10.5 *Let ϕ and ψ be two distance-like functions with $\|\phi - \psi\|_\infty \leq \varepsilon$. For any α -critical point x of ϕ , there exists a α' -critical point x' of ψ with $\|x - x'\| \leq 2\sqrt{\varepsilon\phi(x)}$ and $\alpha' \leq \alpha + 2\sqrt{\varepsilon/\phi(x)}$.*

Proof The proof is almost verbatim the same as the proof of Theorem 9.9 $\square$

Corollary 10.6 *Let ϕ and ψ be two ε -close distance-like functions, and suppose that $\text{reach}_\alpha(\phi) \geq R$ for some $\alpha > 0$. Then, ψ has no critical value in the interval $]4\varepsilon/\alpha^2, R - 3\varepsilon[$.*

Proof The proof is almost verbatim the same as the proof of Theorem 9.16. $\square$

Theorem 10.7 (Extended reconstruction theorem) *Let ϕ, ψ be two ε -close distance-like functions, with $\text{reach}_\alpha(\phi) \geq R$ for some positive α . Then, for any $r \in [4\varepsilon/\alpha^2, R - 3\varepsilon]$, and for $0 < \eta < R$, the sublevel sets ψ^r and ϕ^η are homotopy equivalent, as soon as*

$$\varepsilon \leq \frac{R}{5 + 4/\alpha^2}$$

Proof By the extended isotopy Lemma 10.3, all the sublevel sets ψ^r have the same homotopy type, for r in the given range. Let us choose $r = 4\varepsilon/\alpha^2$. We have:

$$\text{wfs}_\phi(r) \geq R - 4\varepsilon/\alpha^2 \text{ and } \text{wfs}_\psi(r) \geq R - 3\varepsilon - 4\varepsilon/\alpha^2$$

By Proposition 10.4, the sublevel sets ϕ^r and ψ^r have the same homotopy type as soon as the uniform distance ε between ϕ and ψ is smaller than $\frac{1}{2}\text{wfs}_\phi(r)$ and $\frac{1}{2}\text{wfs}_\psi(r)$. This is true, provided that $2\varepsilon \leq R - \varepsilon(3 + 4/\alpha^2)$. The theorem follows. $\square$

Remark that in the Definition 10.2 the notion of α -reach could be made dependent on a parameter r , i.e. the (r, α) -reach of ϕ could be defined as the maximum r' such that the set $\phi^{-1}((r, r + r'])$ does not contain any α -critical value. A reconstruction theorem similar to Theorem 10.7 would still hold under the weaker condition that the (r, α) -reach of ϕ is positive.

10.2 Measures and the Wasserstein distance W_2

To overcome the problem of outliers, a first idea is to consider geometric objects as mass distributions, i.e. measures, instead of purely geometric compact sets. Considering probability measures as the new class of studied objects leads to a much better adapted framework to cope with noise and outliers.

10.2.1 Replacing compact sets by measures

The definition of measure relies on the notion of σ -algebra. All the measures considered in this book will be defined on the so-called Borel σ -algebra whose definition is given below.

Definition 10.8 *A σ -algebra on a set X is a collection Σ of subsets of X such that:*

- i) $\emptyset \in \Sigma$,
 - ii) if $A \in \Sigma$, then $A^c \in \Sigma$,
 - iii) if $(A_n)_{n \in \mathbb{N}}$ is a countable family of elements of Σ , then $\cup_{n \in \mathbb{N}} A_n \in \Sigma$.
- The set of Borel sets of $\mathbb{R}^d$ is the smallest σ -algebra containing all the open sets of $\mathbb{R}^d$ (and thus all the closed sets).*

Definition 10.9 *A Borel measure or, for short in this book, a measure, μ on $\mathbb{R}^d$ is a map from the set of Borel subsets B of $\mathbb{R}^d$ to the set of non-negative real numbers such that whenever (B_i) is a countable family of dis-*

joint Borel subsets of $\mathbb{R}^d$, $\mu(\cup_{i \in \mathbb{N}} B_i) = \sum_i \mu(B_i)$. A probability measure is a measure whose total mass $\mu(\mathbb{R}^d)$ is equal to 1.

Definition 10.10 The support of a measure μ is the smallest, with respect to inclusion, closed set K on which the mass of μ is concentrated, i.e. such that $\mu(\mathbb{R}^d \setminus K) = 0$.

For some compact sets such as point clouds, submanifolds of $\mathbb{R}^d$ or some more general shapes there exist natural ways to associate probability measures whose support are these compact sets as shown in the following examples.

Given a point $x \in \mathbb{R}^d$ the Dirac measure δ_x at x is defined as $\delta_x(B) = 1$ if $x \in B$ and $\delta_x(B) = 0$ otherwise. Given a set of n points C , the empirical or uniform measure μ_C , associated to C , is defined as $\mu_C(B) = \frac{1}{n}|B \cap C|$. It is the sum of n Dirac masses of weight $1/n$, centered at each point of C .

Given a compact k -dimensional manifold $M \subseteq \mathbb{R}^d$, let vol_M be the k -dimensional volume measure on M . As M is compact, $\text{vol}_M(M)$ is finite and we define a probability measure μ_M supported on M by $\mu_M(B) = \text{vol}_M(B \cap M) / \text{vol}_M(M)$, for any Borel set $B \subseteq \mathbb{R}^d$. For example, if M is a curve, $\mu_M(B)$ is the fraction of the total length of M that is contained in B ; similarly, if M is a surface, $\mu_M(B)$ is the fraction of the total area of M that is contained in B . Notice that if M is a finite union of submanifolds $M_1, \dots, M_k$, then we can define probability measures on M just by considering weighted sums of the measures μ_{M_i} .

10.2.2 The Wasserstein distance W_2

There exist a whole family of Wasserstein distances W_p ($p \geq 1$) between probability measures in $\mathbb{R}^d$. Their definition relies on the notion of transport plan between measures. Although some of the results of this chapter can be stated for any distance W_p , for technical reasons that become clear in the following we only consider the W_2 distance.

A transport plan between two probability measures μ and ν on $\mathbb{R}^d$ is a probability measure π on $\mathbb{R}^d \times \mathbb{R}^d$ such that for every Borel sets $A, B \subseteq \mathbb{R}^d$, $\pi(A \times \mathbb{R}^d) = \mu(A)$ and $\pi(\mathbb{R}^d \times B) = \nu(B)$. Intuitively $\pi(A \times B)$ corresponds to the amount of mass of μ contained in A that will be transported to B by

the transport plan. The cost of such a transport plan π is given by

$$\mathcal{C}(\pi) = \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\pi(x, y) \right)^{1/2}$$

As an example, consider two probability measures with finite supports μ and ν :

$$\mu = \sum_{i=1}^m c_i \delta_{x_i} \quad \text{and} \quad \nu = \sum_{j=1}^n d_j \delta_{y_j}$$

with $\sum_{i=1}^m c_i = 1$ and $\sum_{j=1}^n d_j = 1$. A transport plan between μ and ν is then a $n \times m$ matrix $\Pi = (\pi_{i,j})$ with non negative entries such that

$$\sum_{j=1}^n \pi_{i,j} = c_i \quad \text{and} \quad \sum_{i=1}^m \pi_{i,j} = d_j.$$

The coefficient $\pi_{i,j}$ can be seen as the amount of the mass of μ located at x_j that is transported to y_i . The cost of such a transport plan is then given by

$$\mathcal{C}(\Pi) = \left(\sum_{i=1}^n \sum_{j=1}^m \pi_{i,j} \|x_j - y_i\|^2 \right)^{1/2}.$$

Definition 10.11 *The Wasserstein distance of order 2 between two probability measures μ and ν on $\mathbb{R}^d$ is the minimum cost $\mathcal{C}(\pi)$ of a transport plan π between μ and ν . It is denoted by $W_2(\mu, \nu)$.*

Notice that $W_2(\mu, \nu)$ may be infinite. However, if μ and ν have finite moment of order 2, i.e. $\int_{\mathbb{R}^d} \|x\|^2 d\mu(x) < +\infty$ and $\int_{\mathbb{R}^d} \|x\|^2 d\nu(x) < +\infty$, then $W_2(\mu, \nu)$ is finite and the space of probability measures with finite moment of order 2 endowed with W_2 is a metric space [132].

Even for measures with finite support, the computation of the Wasserstein distance is very expensive. However, it provides an interesting notion to quantify the resilience to outliers. To illustrate this, consider a set $C = \{x_1, x_2, \dots, x_N\}$ of N points in $\mathbb{R}^d$ and a noisy version C' obtained by replacing the first n points in C by points y_i such that $d_C(y_i) = R > 0$ for $i = 1, \dots, n$. If we denote by $\mu = \frac{1}{N} \sum_{p \in C} \delta_p$ and $\nu = \frac{1}{N} \sum_{q \in C'} \delta_q$ the empirical measures associated to C and C' respectively then one has

$$W_2(\mu, \nu) = \sqrt{\frac{n}{N}} (R + \text{diam}(C))$$

while the Hausdorff distance between C and C' is at least R . To prove this inequality, consider the transport plan Π from ν to μ that moves the outliers back to their original position and leave the other points fixed. The matrix Π of this transport plan (see the above example) is defined by $\pi_{i,j} = 1/N$ if $i = j$ and 0 otherwise. Since $\|x_i - y_i\| \leq R + \text{diam}(C)$ for $i = 0 \cdots n$ and $x_i = y_i$ for $i > n$, we immediately deduce that the cost of this transport plan is upper bounded by $\sqrt{n/N}(R + \text{diam}(C))$. As a consequence, replacing a small amount of points ($n \ll N$) of C by outliers results in a new measure that remains close to the original one.

From a geometric inference point of view, since, in practice, we are working with point sets sampled according to some unknown probability distribution μ , the question of the convergence of the empirical measure μ_N to μ with respect to the Wasserstein distance is of fundamental importance. This question is beyond the scope of this book but has been a subject of study in probability and statistics for a long time. For example, if μ is supported on a compact set, then μ_N converges almost surely to μ in the W_2 distance. However all the stability and inference results stated in this chapter only rely on the Wasserstein distance between the considered measures and are independent of any convergence property of empirical measures.

10.3 Distance function to a probability measure

In this section, we associate to any probability measure in $\mathbb{R}^d$ a family of real valued functions that are both distance-like and robust with respect to perturbations of the probability measure.

10.3.1 Definition

The distance function to a compact set K evaluated at $x \in \mathbb{R}^d$ is defined as the smallest radius r such that the closed ball centered at x of radius r contains at least a point of K . A natural idea to adapt this definition when K is replaced by a measure μ is to consider the smallest radius r such that the ball with center x and radius r contains a given fraction m of the total mass of μ .

Definition 10.12 *Let μ be a probability measure on $\mathbb{R}^d$ and $0 \leq m < 1$ a*

given parameter. We denote by $\delta_{\mu,m}$ the function defined by

$$\delta_{\mu,m} : x \in \mathbb{R}^d \mapsto \inf\{r > 0; \mu(\bar{B}(x,r)) > m\}$$

where $\bar{B}(x,r)$ denotes the closed ball with center x and radius r .

Notice that for $m = 0$, the definition coincides with the (usual) distance function to the support of the measure μ . Moreover for any $m \in [0, 1]$, $\delta_{\mu,m}$ is 1-Lipschitz.

Unfortunately $\delta_{\mu,m}$ is not robust with respect to perturbations of the measure μ . Indeed, the map $\mu \rightarrow \delta_{\mu,m}$ is not continuous as shown by the following example. Let $\mu_\varepsilon = (\frac{1}{2} - \varepsilon)\delta_0 + (\frac{1}{2} + \varepsilon)\delta_1$ be the weighted sum of two Dirac measures at 0 and 1 in $\mathbb{R}$ and let $m = 1/2$. Then, for $\varepsilon \geq 0$ one has $\delta_{\mu_\varepsilon,1/2}(t) = |1 - t|$ for $t < 0$ while if $\varepsilon < 0$, one obtains $\delta_{\mu_\varepsilon,1/2}(t) = |t|$ which means that $\varepsilon \mapsto \delta_{\mu_\varepsilon,1/2}$ is not continuous at $\varepsilon = 0$.

To overcome this issue we define the distance function associated to μ as a L^2 average of the pseudo-distances $\delta_{\mu,m}$ for a range $[0, m_0]$ of parameters m :

Definition 10.13 (Distance-to-measure) *Let μ be a probability measure on $\mathbb{R}^d$, and m_0 be a positive mass parameter $0 < m_0 \leq 1$. The distance function to μ with parameter m_0 is the function $d_{\mu,m_0} : \mathbb{R}^d \rightarrow \mathbb{R}_+$ defined by :*

$$d_{\mu,m_0}^2 : \mathbb{R}^d \rightarrow \mathbb{R}_+, \quad x \mapsto \frac{1}{m_0} \int_0^{m_0} \delta_{\mu,m}^2(x) dm$$

10.3.2 Distance function to empirical measures

An interesting property of the above defined functions is that they have a simple expression in terms of nearest neighbors. More precisely, let C be a point cloud with n points in $\mathbb{R}^d$, and μ_C be the uniform measure on it: $\mu_C = \frac{1}{n} \sum_{p \in C} \delta_p$. For $0 < m \leq 1$, the function $\delta_{\mu_C,m}$ evaluated at a given point $x \in \mathbb{R}^d$ is by definition equal to the distance between x and its k -th nearest neighbor in C , where k is the smallest integer larger than mn . Hence the function $m \mapsto \delta_{\mu_C,m}(x)$ is constant and equal to the distance from x to its k -th nearest neighbor in C on each interval $[\frac{k-1}{n}, \frac{k}{n})$. Integrating the square of this piecewise constant functions gives the following expression for d_{μ,m_0}^2 , where $m_0 = k_0/n$:

$$\begin{aligned}
d_{\mu, m_0}^2(x) &= \frac{1}{m_0} \int_0^{m_0} \delta_{\mu, m}^2(x) dm = \frac{1}{m_0} \sum_{k=1}^{k_0} \frac{1}{n} \delta_{\mu, k/n}^2(x) \\
&= \frac{1}{k_0} \sum_{p \in \text{NN}_C^{k_0}(x)} \|p - x\|^2
\end{aligned}$$

where $\text{NN}_C^{k_0}(x)$ denote the set of the first k_0 nearest neighbors of x in C . As a consequence the pointwise evaluation of $d_{\mu_C, k_0/n}^2(x)$ reduces to a k -nearest neighbor query in C .




Figure 10.3: Computation of the distance to the empirical measure associated to a point set C ($x_C(k)$ denotes the k -th nearest neighbor of x).

10.3.3 Equivalent formulation

In this paragraph, we provide another characterization of the distance function to a measure d_{μ, m_0} showing that it is in fact the distance function to a closed set, but in the non Euclidean space of probability measures endowed with the W_2 metric (see figure 10.4). This equivalent formulation will be used to deduce that $\mu \rightarrow d_{\mu, m_0}$ is Lipschitz and $x \rightarrow d_{\mu, m_0}^2(x)$ is semiconcave.

Definition 10.14 *A measure ν is a submeasure of another measure μ if for every Borel subset B of $\mathbb{R}^d$, $\nu(B) \leq \mu(B)$. The set of all submeasures of a given measure is denoted by $\text{Sub}(\mu)$, while the set of submeasures of μ with a prescribed total mass $m_0 > 0$ is denoted by $\text{Sub}_{m_0}(\mu)$.*




Figure 10.4: The distance function to a measure as a usual distance function in an infinite dimensional space.

Proposition 10.15 *For any probability measure μ on $\mathbb{R}^d$, the distance function to μ evaluated at x is defined as:*

$$d_{\mu, m_0}(x) = \min \{m_0^{-1/2} W_2(m_0 \delta_x, \nu) ; \nu \in \text{Sub}_{m_0}(\mu)\} \quad (10.1)$$

Moreover, for any measure μ_{x, m_0} that realizes the above minimum one has:

$$d_{\mu, m_0}(x) = \left(\frac{1}{m_0} \int_{\mathbb{R}^d} \|x - h\|^2 d\mu_{x, m_0}(h) \right)^{1/2}$$

Said otherwise, the distance d_{μ, m_0} evaluated at a point $x \in \mathbb{R}^d$ is the minimal Wasserstein distance between the Dirac mass $m_0 \delta_x$ and the set of submeasures of μ with total mass m_0 .

The set $\mathcal{R}_{\mu, m_0}(x)$ of submeasures minimizing the above expression corresponds to the nearest neighbors of the Dirac measure $m_0 \delta_x$ on the set of submeasures $\text{Sub}_{m_0}(\mu)$. It is not empty but it might not be reduced to a

single element. Indeed, it coincides with the set of submeasures μ_{x,m_0} of total mass m_0 whose support is contained in the closed ball $\bar{B}(x, \delta_{\mu,m_0}(x))$, and whose restriction to the open ball $B(x, \delta_{\mu,m_0}(x))$ coincides with μ .

10.3.4 Stability of the distance function to a measure

The characterization of d_{μ,m_0} given in Proposition 10.15 provides a rather easy way to prove the stability of $\mu \mapsto d_{\mu,m_0}$.

Theorem 10.16 (Distance function stability) *If μ and μ' are two probability measures on $\mathbb{R}^d$ and $m_0 > 0$, then $\|d_{\mu,m_0} - d_{\mu',m_0}\|_\infty \leq \frac{1}{\sqrt{m_0}} W_2(\mu, \mu')$.*

The proof of theorem 10.16 follows from the following proposition.

Proposition 10.17 *Let μ and μ' be two probability measures on $\mathbb{R}^d$. Then,*

$$d_H(\text{Sub}_{m_0}(\mu), \text{Sub}_{m_0}(\mu')) \leq W_2(\mu, \mu')$$

where $d_H(.,.)$ is the Hausdorff distance in the space of probability measures endowed with the W_2 metric.

Proof (sketch of) Let ε be the Wasserstein distance of order 2 between μ and μ' , and π be a corresponding optimal transport plan, i.e. a transport plan between μ and μ' such that $\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\pi(x, y) = \varepsilon^2$. Given a submeasure ν of μ , one can find a submeasure π' of π that transports ν to a submeasure ν' of μ' (notice that this later claim is not completely obvious and its formal proof is beyond the scope of this book. It can be proven using the Radon-Nykodim theorem). Then,

$$W_2(\nu, \nu')^2 \leq \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\pi'(x, y) \leq \varepsilon^2$$

This shows that $\text{dist}(\nu, \text{Sub}_{m_0}(\mu')) \leq \varepsilon$ for every submeasure $\nu \in \text{Sub}_{m_0}(\mu)$. The same holds by exchanging the roles of μ and μ' , thus proving the bound on the Hausdorff distance. $\square$

Proof [of Theorem 10.16] The following sequence of equalities and inequalities, that follows from Propositions 10.15 and 10.17, proves the theo-

rem.:

$$\begin{aligned}
 d_{\mu, m_0}(x) &= \frac{1}{\sqrt{m_0}} \text{dist}_{W_2}(m_0 \delta_x, \text{Sub}_{m_0}(\mu)) \\
 &\leq \frac{1}{\sqrt{m_0}} (d_H(\text{Sub}_{m_0}(\mu), \text{Sub}_{m_0}(\mu')) + \text{dist}_{W_2}(m_0 \delta_x, \text{Sub}_{m_0}(\mu'))) \\
 &\leq \frac{1}{\sqrt{m_0}} W_2(\mu, \mu') + d_{\mu', m_0}(x)
 \end{aligned}$$

□

10.3.5 The distance to a measure is a distance-like function.

The subdifferential of a function $f : \Omega \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$ at a point x , is the set of vectors v of $\mathbb{R}^d$, denoted by $\partial_x f$, such that for all small enough vector h , $f(x+h) \geq f(x) + h.v$. This gives a characterization of convexity: a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex if and only if its subdifferential $\partial_x f$ is non-empty for every point x . If this is the case, then f admits a derivative at a point x if and only if the subdifferential $\partial_x f$ is a singleton, in which case the gradient $\nabla_x f$ coincides with its unique element [58].

Proposition 10.18 *The function $v_{\mu, m_0} : x \in \mathbb{R}^d \mapsto \|x\|^2 - d_{\mu, m_0}^2$ is convex, and its subdifferential at a point $x \in \mathbb{R}^d$ is given by*

$$\partial_x v_{\mu, m_0} = \left\{ 2x - \frac{2}{m_0} \int_{h \in \mathbb{R}^d} (x-h) d\mu_{x, m_0}(h) ; \mu_{x, m_0} \in \mathcal{R}_{\mu, m_0}(x) \right\}$$

Proof For any two points x and y of $\mathbb{R}^d$, let μ_{x, m_0} and μ_{y, m_0} be in $\mathcal{R}_{\mu, m_0}(x)$ and $\mathcal{R}_{\mu, m_0}(y)$ respectively. Thanks to Proposition 10.15, we have the following sequence of equalities and inequalities:

$$\begin{aligned}
 d_{\mu, m_0}^2(y) &= \frac{1}{m_0} \int_{h \in \mathbb{R}^d} \|y-h\|^2 d\mu_{y, m_0}(h) \\
 &\leq \frac{1}{m_0} \int_{h \in \mathbb{R}^d} \|y-h\|^2 d\mu_{x, m_0}(h) \\
 &\leq \frac{1}{m_0} \int_{h \in \mathbb{R}^d} \|x-h\|^2 + 2(x-h).(y-x) + \|y-x\|^2 d\mu_{x, m_0}(h) \\
 &\leq d_{\mu, m_0}^2(x) + \|y-x\|^2 + v.(y-x)
 \end{aligned}$$

where v is the vector defined by

$$v = \frac{2}{m_0} \int_{h \in \mathbb{R}^d} (x - h) d\mu_{x,m_0}(h).$$

The inequality can be rewritten as:

$$(\|y\|^2 - d_{\mu,m_0}^2(y)) - (\|x\|^2 - d_{\mu,m_0}^2(x)) \geq (2x - v) \cdot (y - x)$$

which shows that the vector $(2x - v)$ belongs to the subdifferential of v at x . By the characterization of convex functions by that we recalled above, one deduces that v_{μ,m_0} is convex.

The proof of the reverse inclusion is slightly more technical and beyond the scope of the book. $\square$

Corollary 10.19 *The function d_{μ,m_0}^2 is 1-semiconcave. Moreover,*

(i) d_{μ,m_0}^2 is differentiable almost everywhere in $\mathbb{R}^d$, with gradient defined by

$$\nabla_x d_{\mu,m_0}^2 = \frac{2}{m_0} \int_{h \in \mathbb{R}^d} [x - h] d\mu_{x,m_0}(h)$$

where μ_{x,m_0} is the only measure in $\mathcal{R}_{\mu,m_0}(x)$.

(ii) the function $x \in \mathbb{R}^d \mapsto d_{\mu,m_0}(x)$ is 1-Lipschitz.

Proof (i). It follows from the fact that a convex function is differentiable at almost every point, at which its gradient is the only element of the subdifferential at that point.

(ii). The gradient of d_{μ,m_0} can be written as:

$$\nabla_x d_{\mu,m_0} = \frac{\nabla_x d_{\mu,m_0}^2}{2d_{\mu,m_0}} = \frac{1}{\sqrt{m_0}} \frac{\int_{h \in \mathbb{R}^d} [x - h] d\mu_{x,m_0}(h)}{(\int_{h \in \mathbb{R}^d} \|x - h\|^2 d\mu_{x,m_0}(h))^{1/2}}$$

Using the Cauchy-Schwartz inequality we find $\|\nabla_x d_{\mu,m_0}\| \leq 1$ which proves the statement. $\square$

10.4 Applications to geometric inference

Reconstruction from point clouds in presence of outliers was the main motivation for introducing the distance function to a measure. In this section, we show how the extended reconstruction Theorem 10.7 can be applied to distance to measure functions. It is also possible to adapt most of the topological and geometric inference results of Chapter 9 in a similar way.

10.4.1 Distance to a measure vs. distance to its support

In this section, we compare the distance functions d_{μ, m_0} to a measure μ and the distance function d_S to its support S , and study the convergence properties of d_{μ, m_0} to d_S as the mass parameter m_0 converges to zero. Remark that the function δ_{μ, m_0} (and hence the distance d_{μ, m_0}) is always larger than the distance function d_S , i.e. for any $x \in \mathbb{R}^d$, $d_S(x) \leq d_{\mu, m_0}(x)$. As a consequence, to obtain a convergence result of d_{μ, m_0} to d_S as m_0 goes to zero, we just need to upper bound $d_{\mu, m_0} - d_S$ by a function converging to 0 as m_0 goes to 0. It turns out that the convergence speed of d_{μ, m_0} to d_S depends on the way the mass of μ contained within any ball $B(p, r)$ centered at a point p of the support increases with r . Let us define:

- (i) We say that a non-decreasing positive function $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a *uniform lower bound on the growth of μ* if for every point p in the support of μ and every $\varepsilon > 0$, $\mu(B(p, \varepsilon)) \geq f(\varepsilon)$;
- (ii) The measure μ has *dimension at most k* if there is a constant $C(\mu)$ such that $f(\varepsilon) = C(\mu)\varepsilon^k$ is a uniform lower bound on the growth of μ , for ε small enough.

Lemma 10.20 *Let μ be a probability measure and f be a uniform lower bound on the growth of μ . Then $\|d_{\mu, m_0} - d_S\|_\infty < \varepsilon$ as soon as $m_0 < f(\varepsilon)$.*

Proof Let ε and m_0 be such that $m_0 < f(\varepsilon)$ and let x be a point in $\mathbb{R}^d$, p a projection of x on S , i.e. a point p such that $\|x - p\| = d_S(x)$. By assumption, $\mu(B(x, d_S(x) + \varepsilon)) \geq \mu(B(p, \varepsilon)) \geq m_0$. Hence, $\delta_{\mu, m_0}(x) \leq d_S(x) + \varepsilon$. The function $m \mapsto \delta_{\mu, m}(x)$ being non-decreasing, we get: $m_0 d_S^2(x) \leq \int_0^{m_0} \delta_{\mu, m}^2(x) dm \leq m_0 (d_S(x) + \varepsilon)^2$. Taking the square root of this expression proves the lemma. $\square$

Corollary 10.21 (i) If the support S of μ is compact, then d_S is the uniform limit of d_{μ, m_0} as m_0 converges to 0:

$$\|d_{\mu, m_0} - d_S\|_\infty = \sup_{x \in \mathbb{R}^d} |d_{\mu, m_0}(x) - d_S(x)| \xrightarrow{m_0 \rightarrow 0} 0$$

(ii) If the measure μ has dimension at most $k > 0$, then

$$\|d_{\mu, m_0} - d_S\|_\infty \leq C(\mu)^{-1/k} m_0^{1/k}$$

Proof (i) If S is compact, there exists a sequence $x_1, x_2, \dots$ of points in S such that for any $\varepsilon > 0$, $S \subseteq \cup_{i=1}^n B(x_i, \varepsilon/2)$ for some $n = n(\varepsilon)$. By definition of the support of a measure, $\eta(\varepsilon) = \min_{i=1 \dots n} \mu(B(x_i, \varepsilon/2))$ is positive. Now, for any point $x \in S$, there is a x_i such that $\|x - x_i\| \leq \varepsilon/2$. Hence, $B(x_i, \varepsilon/2) \subseteq B(x, \varepsilon)$, which means that $\mu(B(x, \varepsilon)) \geq \eta(\varepsilon)$. (ii) Follows straightforwardly from the lemma. $\square$

For example, the uniform probability measure on a k -dimensional compact submanifold S has dimension at most k . The following proposition gives a more precise convergence speed estimate based on curvature.

Proposition 10.22 Let S be a smooth k -dimensional submanifold of $\mathbb{R}^d$ whose curvature radii are lower bounded by R , and μ the uniform probability measure on S , then

$$\|d_S - d_{\mu, m_0}\|_\infty \leq C(S)^{-1/k} m_0^{1/k}$$

for m_0 small enough and $C(S) = (2/\pi)^k \beta_k / \text{vol}^k(S)$ where β_k is the volume of the unit ball in $\mathbb{R}^k$.

Notice in particular that the convergence speed of d_{μ, m_0} to d_S depends only on the *intrinsic* dimension k of the submanifold S , and not on the ambient dimension d . The proof of this result is beyond the scope of this book and relies on the so-called Günther-Bishop theorem (cf [89, Section 3.101]).

10.4.2 Shape reconstruction from noisy data

The previous results lead to shape reconstruction theorems that work for noisy data with outliers. To fit in our framework we consider shapes that

are defined as supports of probability measures. Let μ be a probability measure of dimension at most $k > 0$ with compact support $K \subset \mathbb{R}^d$ and let $d_K : \mathbb{R}^d \rightarrow \mathbb{R}_+$ be the (Euclidean) distance function to K . If μ' is another probability measure (e.g. the empirical measure given by a point cloud sampled according to μ), one has

$$\|d_K - d_{\mu', m_0}\|_\infty \leq \|d_K - d_{\mu, m_0}\|_\infty + \|d_{\mu, m_0} - d_{\mu', m_0}\|_\infty \quad (10.2)$$

$$\leq C(\mu)^{-1/k} m_0^{1/k} + \frac{1}{\sqrt{m_0}} W_2(\mu, \mu') \quad (10.3)$$

This inequality insuring the closeness of d_{μ', m_0} to the distance function d_K for the sup-norm follows immediately from the stability Theorem 10.16 and Corollary 10.21. As expected, the choice of m_0 is a trade-off: small m_0 lead to better approximation of the distance function to the support, while large values of m_0 make the distance functions to measures more stable. Eq. (10.3) leads to the following corollary of Theorem 10.7:

Corollary 10.23 *Let μ be a measure and K its support. Suppose that μ has dimension at most k and that $\text{reach}_\alpha(d_K) \geq R$ for some $R > 0$. Let μ' be another measure, and ε be an upper bound on the uniform distance between d_K and d_{μ', m_0} . Then, for any $r \in [4\varepsilon/\alpha^2, R - 3\varepsilon]$, the r -sublevel sets of d_{μ, m_0} and the offsets K^η , for $0 < \eta < R$ are homotopy equivalent as soon as:*

$$W_2(\mu, \mu') \leq \frac{R\sqrt{m_0}}{5 + 4/\alpha^2} - C(\mu)^{-1/k} m_0^{1/k+1/2}$$

Figure 10.5 illustrates Theorem 10.7 on a sampled mechanical part with 10% of outliers. In this case μ' is the normalized sum of the Dirac measures centered on the data points and the (unknown) measure μ is the uniform measure on the mechanical part.

10.5 Exercises

Exercise 10.1 Let $K \subset \mathbb{R}^d$ be a compact set. Show that the map $x \rightarrow \|x\|^2 - d_K^2(x)$ is convex.

Hint: Recall that the supremum of any family of convex functions is convex and show that $x \rightarrow \|x\|^2 - d_K^2(x)$ is the supremum of a set of affine functions.




Figure 10.5: On the left, a point cloud sampled on a mechanical part to which 10% of outliers (uniformly sampled in a box enclosing the model) have been added. On the right, the reconstruction of an isosurface of the distance function d_{μ_C, m_0} to the uniform probability measure on this point cloud.

Exercise 10.2 Let $\mu_1, \dots, \mu_k$ be measures on $\mathbb{R}^d$ and let $\lambda_1, \dots, \lambda_k \in \mathbb{R}$. Show that if all the λ_i are non negative then $\mu = \sum_{i=1}^k \lambda_i \mu_i$ is a measure. Show that the set of probability measures on $\mathbb{R}^d$ is convex.

Exercise 10.3 Let μ and ν be two probability measures with finite supports:

$$\mu = \sum_{j=1}^m c_j \delta_{x_j} \quad \text{and} \quad \nu = \sum_{i=1}^n d_i \delta_{y_i}$$

where $x_1, \dots, x_m$ and $y_1, \dots, y_n$ are points in $\mathbb{R}^d$ and $\sum_{j=1}^m c_j = \sum_{i=1}^n d_i = 1$. Show that any transport plan between μ and ν can be represented as a $n \times m$ matrix $\Pi = (\pi_{i,j})$ with non negative entries such that

$$\sum_{i=1}^n \pi_{i,j} = c_j \quad \text{and} \quad \sum_{j=1}^m \pi_{i,j} = d_i.$$

Exercise 10.4 Let μ be a probability measure on $\mathbb{R}^d$ and let $m \in [0, 1)$. Show that $\delta_{\mu, m}$ is 1-Lipschitz:

$$\forall x, y \in \mathbb{R}^d, |\delta_{\mu, m}(x) - \delta_{\mu, m}(y)| \leq \|x - y\|.$$

Exercise 10.5 Let μ be a probability measure on $\mathbb{R}^d$ and let $0 < m_0 < 1$ and $x \in \mathbb{R}^d$. Recall that we denote by $\mathcal{R}_{\mu, m_0}(x)$ the set of submeasures minimizing the right hand term of Equation (10.1). Give an example of a probability measure μ such that $\mathcal{R}_{\mu, m_0}(x)$ contains only one element and an example such that $\mathcal{R}_{\mu, m_0}(x)$ contains an infinite number of elements.

(Hint: for $\mathcal{R}_{\mu, m_0}(x)$ to contain more than one submeasure, the measure $\mu(S(x, \delta_{\mu, m_0}))$ of the sphere of center x and radius δ_{μ, m_0} must be positive.

10.6 Bibliographical notes

Most of the chapter comes from [40] that introduces and studies stability properties of distance functions to a probability measure. Approximation and computation of the distance to measure function and its connections with power distances has been studied in [94]. Extensions of the distance-to-measure framework to general metric spaces has been considered in [29]. Statistical aspects of distance-to-measure functions in relation with density estimation and deconvolution have been considered in [10, 31, 49, 51].

Wasserstein distances are closely related to the theory of optimal transportation (see e.g. [132]). The distance W_1 is also known as the earth-mover distance, and has been used in various domains such as, e.g., vision [119] or image retrieval [123].

General results about semiconcave functions can be found in [120, 33].

The convergence properties of empirical measure with respect to the Wasserstein metric have been widely studied and quantitative results can be found in [28].

The complete proofs of Propositions 10.15, 10.18 and 10.4 are given in [40].

The Günther-Bishop Theorem is stated in [89, Section 3.101] and the proof of Proposition 10.22 can be found in [40].

Chapter 11

Homology inference

Although the distance-based approach introduced in the two previous chapters provides a powerful mathematical framework for shape reconstruction, it is not always possible, nor desirable, to fully reconstruct the approximated shapes from data. This chapter focuses on weaker topological invariants, homology, Betti numbers and persistent homology, that turn out to be easier to infer and that are widely used in applied topology and topological data analysis. The introduction of homology is restricted to the minimum that is necessary to understand the basic ideas of homology inference and persistent homology and its usage in topological data analysis.

11.1 Simplicial homology and Betti numbers

In this section we introduce the basic notions of simplicial homology that are necessary to define and study topological persistence. To avoid minor technical discussions about the orientation of simplices, we restrict to the homology with coefficients in the finite field $\mathbb{Z}/2\mathbb{Z} = \{0, 1\}$. In the sequel of this chapter, K denotes a finite d -dimensional simplicial complex.

11.1.1 The space of k -chains

For any non negative integer k , the space of k -chains is the vector space of all the formal sums (with coefficient in $\mathbb{Z}/2\mathbb{Z}$) of k -dimensional simplices of K . More precisely, if $\{\sigma_1, \dots, \sigma_p\}$ is the set of k -simplices of K any k -chain c can be uniquely written

$$c = \sum_{i=1}^p \varepsilon_i \sigma_i \quad \text{with } \varepsilon_i \in \mathbb{Z}/2\mathbb{Z}$$

If $c' = \sum_{i=1}^p \varepsilon'_i \sigma_i$ is another k -chain, the sum of two k -chains and the product of a chain by a scalar are defined by

$$c + c' = \sum_{i=1}^p (\varepsilon_i + \varepsilon'_i) \sigma_i \quad \text{and} \quad \lambda.c = \sum_{i=1}^p (\lambda \varepsilon_i) \sigma_i$$

where the sums $\varepsilon_i + \varepsilon'_i$ and the products $\lambda \varepsilon_i$ are modulo 2.

Definition 11.1 *The space of k -chains is the set $C_k(K)$ of the simplicial k -chains of K with the two operations defined above. This is a $\mathbb{Z}/2\mathbb{Z}$ -vector space whose zero element is the empty chain $0 = \sum_{i=1}^p 0.\sigma_i$.*

Notice that the set of k -simplices of K is a basis of $C_k(K)$. For example, for the simplicial complex K of Figure 11.1, $C_1(K)$ is the $\mathbb{Z}/2\mathbb{Z}$ -vector space generated by the edges $e_1 = [a, b]$, $e_2 = [b, c]$, $e_3 = [c, a]$, $e_4 = [c, d]$, i.e.

$$C_1(K) = \{0, e_1, e_2, e_3, e_4, e_1 + e_2, e_1 + e_3, e_1 + e_4, e_2 + e_3, e_2 + e_4, e_3 + e_4, e_1 + e_2 + e_3, \dots\}$$

Summing $e_1 + e_2$ with $e_2 + e_3 + e_4$ gives $e_1 + e_3 + e_4$.




Figure 11.1: A very simple simplicial complex made of 4 vertices, 4 edges and 1 triangle.

Chains with coefficient in $\mathbb{Z}/2\mathbb{Z}$ have an obvious geometric interpretation: since any k -chain can be uniquely written as $c = \sigma_{i_1} + \sigma_{i_2} + \dots + \sigma_{i_m}$ where the σ_{i_j} are k -simplices, c can be considered as the union of the simplices σ_{i_j} . The sum of two k -chains is equal to their symmetric difference.

11.1.2 The boundary operator and homology groups

Definition 11.2 (Boundary of a simplex) *The boundary $\partial(\sigma)$ of a k -simplex σ is the sum of its $(k-1)$ -faces. This is a $(k-1)$ -chain.*

If $\sigma = [v_0, \dots, v_k]$ is a k -simplex, then

$$\partial(\sigma) = \sum_{i=0}^k [v_0 \cdots \hat{v}_i \cdots v_k]$$

where $[v_0 \cdots \hat{v}_i \cdots v_k]$ is the $(k-1)$ -simplex spanned by the sets of all the vertices of σ except v_i .

Remark 11.3 Notice that in the general case where the coefficient of the chains are taken in another field than $\mathbb{Z}/2\mathbb{Z}$ it is important to take into account the ordering of the vertices in σ and the boundary of σ has to be defined as $\partial(\sigma) = \sum_{i=0}^k (-1)^i [v_0 \cdots \hat{v}_i \cdots v_k]$ - see [113].

The boundary operator, defined on the simplices of K , extends linearly to $C_k(K)$.

Definition 11.4 *The boundary operator is the linear map defined by*

$$\begin{aligned} \partial : C_k(K) &\rightarrow C_{k-1}(K) \\ c &\rightarrow \partial c = \sum_{\sigma \in c} \partial(\sigma) \end{aligned}$$

Notice that one should denote ∂_k the above defined operator but to avoid heavy notations one usually omits the index in the notations.

Proposition 11.5 *The boundary of the boundary of a chain is always zero:*

$$\partial\partial := \partial \circ \partial = 0.$$

Proof Since the boundary operator is linear, it is sufficient to check the property for a simplex. Let $\sigma = [v_0 \cdots v_k]$ be a k -simplex.

$$\begin{aligned} \partial\partial\sigma &= \partial \left(\sum_{i=0}^k [v_0 \cdots \hat{v}_i \cdots v_k] \right) \\ &= \sum_{i=0}^k \partial[v_0 \cdots \hat{v}_i \cdots v_k] \\ &= \sum_{j < i} [v_0 \cdots \hat{v}_j \cdots \hat{v}_i \cdots v_k] + \sum_{j > i} [v_0 \cdots \hat{v}_i \cdots \hat{v}_j \cdots v_k] \\ &= 0 \end{aligned}$$

□

The boundary operator defines a sequence of linear maps between the spaces of chains.

Definition 11.6 (Chain complex) *The chain complex associated to a complex K of dimension d is the following sequence of linear operators*

$$\{0\} \rightarrow C_d(K) \xrightarrow{\partial} C_{d-1}(K) \xrightarrow{\partial} \cdots \xrightarrow{\partial} C_{k+1}(K) \xrightarrow{\partial} C_k(K) \xrightarrow{\partial} \cdots \xrightarrow{\partial} C_0(K) \xrightarrow{\partial} \{0\}$$

For $k \in \{0, \dots, d\}$, the set $Z_k(K)$ of k -cycles of K is the kernel of $\partial : C_k \rightarrow C_{k-1}$:

$$Z_k(K) := \ker(\partial : C_k \rightarrow C_{k-1}) = \{c \in C_k : \partial c = 0\}$$

The image $B_k(K)$ of $\partial : C_{k+1} \rightarrow C_k$ is the set of k -chains bounding a $(k+1)$ -chain:

$$B_k(K) := \text{im}(\partial : C_{k+1} \rightarrow C_k) = \{c \in C_k : \exists c' \in C_{k+1}, c = \partial c'\}$$

B_k and Z_k are subspaces of C_k and according to Proposition 11.5, one has

$$B_k(K) \subset Z_k(K) \subset C_k(K).$$

Examples of chains, cycles and boundaries are given in Figure 11.2.




Figure 11.2: Examples of chains, cycles and boundaries: c_1 is a cycle which is not a boundary, c_2 is a boundary and c_3 is a chain that is not a cycle.

Definition 11.7 (Homology groups) The k^{th} homology group of K is the quotient vector space

$$H_k(K) = Z_k/B_k$$

$H_k(K)$ is a vector space and its elements are the homology classes of K . The dimension $\beta_k(K)$ of $H_k(K)$ is the k^{th} Betti number of K .

The homology class of a cycle $c \in Z_k(K)$ is the set $c + B_k(K) = \{c + b : b \in B_k(K)\}$. Two cycles c, c' that are in the same homology class are said to be *homologous*.




Figure 11.3: Examples of Betti numbers for simple simplicial complexes: from left to right, an edge, the boundary of a triangle, a triangle and the boundary of a tetrahedron.

11.2 An algorithm to compute Betti numbers

Let K be a finite simplicial complex of dimension d and

$$\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K$$

a filtration of K such that for any $i = 0, \dots, m-1$,

$$K^{i+1} = K^i \cup \sigma^{i+1} \quad \text{where } \sigma^{i+1} \text{ is a simplex.}$$

By considering the evolution of the Betti numbers of the filtration as we add the simplices σ^i we get the following algorithm.

To prove the correctness of the algorithm, one has to understand how the topology of the filtration evolves each time we add a simplex. Let assume that the Betti numbers of K^{i-1} have been computed and add the simplex σ^i of dimension $k+1$ to get K^i . Remark that according to the definition of filtration, σ^i cannot be part of the boundary of any $(k+2)$ -simplex of K^i . As a consequence if σ^i is contained in a $(k+1)$ -cycle in K^i , this cycle is not the boundary of a $(k+2)$ -chain in K^i . Let consider the two alternatives of the algorithm that are illustrated in Figure 11.4:

Case 1: assume that σ^i is contained in a $(k+1)$ -cycle c in K^i . Then c cannot be homologous to any $(k+1)$ -cycle in K^{i-1} . Indeed, otherwise there

Algorithm 9 Betti numbers computation

Input: A filtration of a d -dimensional simplicial complex K containing m simplices.

$\beta_0 \leftarrow 0; \beta_1 \leftarrow 0; \dots; \beta_d \leftarrow 0$

for $i = 1$ to m **do**

$k = \dim \sigma^i - 1$

if σ^i is contained in a $(k+1)$ -cycle in K^i **then**

$\beta_{k+1} \leftarrow \beta_{k+1} + 1$

else

$\beta_k \leftarrow \beta_k - 1$

Output: The Betti numbers $\beta_0, \beta_1, \dots, \beta_d$ of K .

would exist a cycle c' in K^{i-1} such that $c + c'$ is the boundary of a $(k+2)$ -chain d . But since σ^i cannot be contained in c' , it has to be contained in $c + c' = \partial d$ contradicting the above remark. So, c creates a new homology class which is linearly independent of the classes created by the cycles in K^{i-1} . As a consequence, $\beta_{k+1}(K^i) \geq \beta_{k+1}(K^{i-1}) + 1$. To conclude this first case, it is sufficient to remark that adding the $(k+1)$ -simplex σ^i to K^{i-1} cannot increase the dimension of the k^{th} homology group by more than one: if c and c' are two $(k+1)$ -cycles containing σ^i , then $c + c'$ is a $(k+1)$ -cycle in K^{i-1} implying that c' is contained in the linear subspace spanned by $Z_{k+1}(K^{i-1})$ and c . It follows that $\dim Z_{k+1}(K^i) \leq \dim Z_{k+1}(K^{i-1}) + 1$ and since $B_{k+1}(K^{i-1}) \subset B_{k+1}(K^i)$, $\beta_{k+1}(K^i) \leq \beta_{k+1}(K^{i-1}) + 1$.

Case 2: assume that σ^i is not contained in any $(k+1)$ -cycle in K^i . Then the k -cycle $\partial\sigma^i$ is not a boundary in K^{i-1} . Indeed, otherwise there would exist a chain c in K^{i-1} such that $\partial c = \partial\sigma^i$ or equivalently $\partial(c + \sigma^i) = 0$. Thus $c + \sigma^i$ is a $(k+1)$ -cycle in K^i containing σ^i : a contradiction. As a consequence, since the k -cycle $\partial\sigma^i$ which is not a boundary in K^{i-1} becomes a boundary in K^i , one has $\beta_k(K^i) \leq \beta_k(K^{i-1}) - 1$. One proves as in Case 1 that this inequality is indeed an equality.

The above discussion suggests to distinguish between the two types of simplices in the filtration of K that will play an important role in the definition of topological persistence.

Definition 11.8 Let K be a d -dimensional simplicial complex and let

$$\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$$




Figure 11.4: Examples of insertion of a 1-simplex σ^i to a subcomplex K^{i-1} . When σ^i is not contained in a 1-cycle (left), its insertion results in the connection of two connected components of K^{i-1} . When σ^i is contained in a 1-cycle (right), its insertion create a new 1-cycle that is not homologous to any previously existing in K^{i-1} .

be a filtration of K . A simplex σ^i is called *positive* if it is contained in a $(k+1)$ -cycle in K^i (which is necessarily not a boundary in K^i according to the remark at the beginning of the proof of the correctness of the algorithm) and *negative* otherwise.

With the above definition the k^{th} Betti number of K is equal to the difference between the number of positive k -simplices (which are creating k -cycles) and the number of negative $(k+1)$ -simplices (which are “killing” k -cycles).

As an example, if one considers the simplicial complex K of figure 11.1 with the filtration defined by the simplices ordering $\emptyset, a, b, c, ab, bc, d, ac, cd, abc$, then the positive simplices are a, b, c, d and ac . The Betti numbers of K are $\beta_0 = 1, \beta_1 = 0$ and $\beta_2 = 0$.

It is important to notice that the above algorithm needs to be able to decide whether a simplex is positive or negative. This is not, a priori, an obvious question but an answer will be given in Section 11.5. It is also important to notice that the algorithm not only computes the Betti numbers of K but also the Betti numbers of all the subcomplexes K^i of the filtration.

11.3 Singular homology and topological invariance

The homology groups and Betti numbers are topological invariants: if K and K' are two simplicial complexes with homeomorphic geometric realizations then their homology groups are isomorphic and their Betti numbers are equal. This result is still true if the geometric realizations of K and K' are just homotopy equivalent. The proof of this invariance property is a classical, but not obvious, result in algebraic topology. It is beyond the scope of this book (see [113, 97] for details).

Singular homology is another notion of homology that allows to consider general spaces that are not necessarily homeomorphic to simplicial complexes. The definition of singular homology is similar to the one of simplicial homology except that it relies on the notion of singular simplex. Let Δ_k be the standard k -dimensional simplex in $\mathbb{R}^{k+1}$, i.e. the geometric simplex spanned by the vertices x_i , $i = 1, \dots, k+1$, whose all coordinates are 0 except the i^{th} one which is equal to 1. Given a topological space X , a *singular k -simplex* σ is a continuous map $\sigma : \Delta_k \rightarrow X$. As in the case of simplicial homology, the space of singular k -chains is the vector space of formal linear combinations of singular k -simplices. The boundary $\partial\sigma$ of a singular k -simplex is the sum of the restriction of σ to each of the $(k-1)$ -faces of Δ_k . Proposition 11.5 still holds for the (singular) boundary operator and the k^{th} singular homology group of X is defined similarly as the quotient of the space of cycles by the space of boundaries.

A remarkable fact is that simplicial and singular homology are related in the following way: if X is a topological space homeomorphic to the support of a simplicial complex K , then the singular homology groups of X are isomorphic to the simplicial homology groups of K . For example, if X is a surface and if K and K' are two triangulations of X , then the homology groups $H_k(K)$ and $H_k(K')$ are isomorphic. Thus they have the same Betti numbers that are, indeed, the ones of X . As a consequence, in the sequel of the Chapter, we will consider indifferently simplicial or singular homology.

Another important property of singular (and thus simplicial) homology is that continuous maps between topological spaces canonically induce homomorphisms between their homology groups. Indeed, if $f : X \rightarrow Y$ is a continuous map between two topological spaces and if $\sigma : \Delta_k \rightarrow X$ is a singular simplex in X , then $f \circ \sigma : \Delta_k \rightarrow Y$ is a singular simplex in Y . So, f induces a linear map between the spaces of chains on X and Y that preserves

cycles and boundaries. As a consequence, f also induces an homomorphism $f_* : H_k(X) \rightarrow H_k(Y)$. Moreover if f is an homeomorphism, f_* is an isomorphism and $(f_*)^{-1} = (f^{-1})_*$. Similarly, if f is an homotopy equivalence with homotopic inverse $g : Y \rightarrow X$, then f_* is an isomorphism with inverse g_* . As a consequence, two spaces that are homotopy equivalent have the same Betti numbers. Notice that, when X is not homotopy equivalent to a finite simplicial complex, its Betti numbers might not be finite.

11.4 Betti numbers inference

Singular homology allows to consider Betti numbers of compact sets in $\mathbb{R}^d$ and of their offsets. Using its connexion to simplicial homology and the distance functions framework of Chapter 9, we derive explicit methods to infer the Betti numbers of compact subsets with positive weak feature size.

Let $K \subset \mathbb{R}^d$ be a compact set with $\text{wfs}(K) > 0$ and let $P \in \mathbb{R}^d$ be a finite set of points such that $d_H(K, P) < \varepsilon$ for some given $\varepsilon > 0$. Recall that, from the isotopy Lemma 9.5, all the r -offsets K^r of K , for $0 < r < \text{wfs}(K)$, are homeomorphic and thus have isomorphic homology groups. The goal of this section is to provide an effective method to compute the Betti numbers $\beta_k(K^r)$, $0 < r < \text{wfs}(K)$, from P .

Theorem 11.9 *Let $K \subset \mathbb{R}^d$ be a compact set with $\text{wfs}(K) > 0$ and let $P \in \mathbb{R}^d$ be a finite set of points such that $d_H(K, P) < \varepsilon$ for some given $\varepsilon > 0$. Assume that $\text{wfs}(K) > 4\varepsilon$. For $\alpha > 0$ such that $4\varepsilon + \alpha < \text{wfs}(K)$, let $i : P^{\alpha+\varepsilon} \hookrightarrow P^{\alpha+3\varepsilon}$ be the canonical inclusion. Then for any non negative integer k and any $0 < r < \text{wfs}(K)$,*

$$H_k(K^r) \cong \text{im} (i_* : H_k(P^{\alpha+\varepsilon}) \rightarrow H_k(P^{\alpha+3\varepsilon}))$$

where im denotes the image of the homomorphism and $\cong$ means that the two groups are isomorphic.

Proof Since $d_H(K, P) < \varepsilon$, we have the following sequence of inclusion maps

$$K^\alpha \subseteq P^{\alpha+\varepsilon} \subseteq K^{\alpha+2\varepsilon} \subseteq P^{\alpha+3\varepsilon} \subseteq K^{\alpha+4\varepsilon}$$

that induces the following sequence of homomorphisms (the one induced by the canonical inclusion maps) at the homology level

$$H_k(K^\alpha) \rightarrow H_k(P^{\alpha+\varepsilon}) \rightarrow H_k(K^{\alpha+2\varepsilon}) \rightarrow H_k(P^{\alpha+3\varepsilon}) \rightarrow H_k(K^{\alpha+4\varepsilon}).$$

Since $\text{wfs}(K) > \alpha + 4\varepsilon$, it follows from the isotopy Lemma 9.5 that the homomorphisms $H_k(K^\alpha) \rightarrow H_k(K^{\alpha+2\varepsilon})$ and $H_k(K^{\alpha+2\varepsilon}) \rightarrow H_k(K^{\alpha+4\varepsilon})$ induced by the inclusion maps are indeed isomorphisms. Notice that the above sequence implies that the rank of these isomorphisms is finite (see Exercise 11.5). It immediately follows that the rank of $H_k(P^{\alpha+\varepsilon}) \rightarrow H_k(P^{\alpha+3\varepsilon})$ is equal to the rank of these isomorphisms which is equal to $\beta_k(K^\alpha)$. $\square$

Theorem 11.9 shows that the Betti numbers of the offsets of K can be deduced from the offsets of P . However, the direct computation of the homology groups of a union of balls, which is a continuous object and not a finite simplicial complex, is not obvious. To overcome this issue, recall that the Nerve Theorem 2.8 implies that for any $r \geq 0$, P^r is homotopy equivalent to $\check{\text{Cech}}(P, r)$. As a consequence $H_k(P^r)$ and $H_k(\check{\text{Cech}}(P, r))$ are isomorphic. Moreover, one can show that the isomorphisms can be chosen to commute with the ones induced by inclusions maps [46], making the following diagram commutative

$$\begin{array}{ccc} H_k(P^r) & \rightarrow & H_k(P^{r'}) \\ \uparrow & & \uparrow \\ H_k(\check{\text{Cech}}(P, r)) & \rightarrow & H_k(\check{\text{Cech}}(P, r')) \end{array} \quad (11.1)$$

We immediately obtain the following result.

Proposition 11.10 *Assume that $\text{wfs}(K) > 4\varepsilon$. For $\alpha > 0$ such that $4\varepsilon + \alpha < \text{wfs}(K)$, let $i : \check{\text{Cech}}(P, \alpha + \varepsilon) \hookrightarrow \check{\text{Cech}}(P, \alpha + 3\varepsilon)$ be the canonical inclusion. Then for any non negative integer k and any $0 < r < \text{wfs}(K)$,*

$$H_k(K^r) \cong \text{im}(i_* : H_k(\check{\text{Cech}}(P, \alpha + \varepsilon)) \rightarrow H_k(\check{\text{Cech}}(P, \alpha + 3\varepsilon))).$$

Thanks to the previous proposition, inferring the Betti numbers of K^r now boils down to homology computation on finite Čech complexes. However, as already noticed in Chapter 2.5, computing Čech complexes require to determine if finite sets of balls intersect, which quickly becomes prohibitive as d and the cardinality of P increase. Using the interleaving property between Čech and Vietoris-Rips filtrations established in Lemma 2.13, we obtain the following theorem.

Theorem 11.11 *Assume that $\text{wfs}(K) > 9\varepsilon$. For any $2\varepsilon \leq \alpha \leq \frac{1}{4}(\text{wfs}(K) - \varepsilon)$ and any $0 < r < \text{wfs}(K)$ we have*

$$\beta_k(K^r) = \text{rk}(H_k(\text{Rips}(P, \alpha) \rightarrow H_k(\text{Rips}(P, 4\alpha)))$$

where $\text{rk}(H_k(\text{Rips}(P, \alpha) \rightarrow H_k(\text{Rips}(P, 4\alpha)))$ denotes the rank of the homomorphism induced by the inclusion $\text{Rips}(P, \alpha) \subseteq \text{Rips}(P, 4\alpha)$.

This last result raises two questions. The first one is about how to compute the rank of the homomorphisms induced by the inclusion maps between the homology groups of the subcomplexes of the *Vietoris–Rips* complex. This question will be answered in next Section when the persistence algorithm will be presented. The second, and more tricky, question is about the effective choice of α when K and $\text{wfs}(K)$ are not known. This is an ill-posed problem because $\text{wfs}(K)$ does not depend continuously of K - see Section 9.3. However, it is possible for the user to try to guess a good choice of α using the following algorithm.

Algorithm 10 Betti numbers inference

Input: $P \subset \mathbb{R}^d$ a finite set.

Let $L := \emptyset$, $\alpha := +\infty$;

while $L \subsetneq P$ **do**

Let $p := \arg\max_{w \in P} \min_{v \in L} \|w - v\|$; // pick arbitrary p if $L = \emptyset$

$L \leftarrow L \cup \{p\}$;

$\alpha \leftarrow \max_{w \in P} \min_{v \in L} \|w - v\|$;

Update $\text{Rips}(L, 4\alpha)$ and $\text{Rips}(L, 16\alpha)$;

Compute $\beta_{k, 4\alpha}^{16\alpha} = \text{rk}(i_* : H_k(\text{Rips}(L, 4\alpha)) \rightarrow H_k(\text{Rips}(L, 16\alpha)))$;

Output: diagram showing the evolution of persistent Betti numbers, i.e. the ranks of i_* versus α .

When applied to a point cloud sampled around a compact subset of $\mathbb{R}^d$ with positive weak feature size, the algorithm provides the diagrams of rank numbers $\beta_{k, 4\alpha}^{16\alpha}$ that are constant on some intervals of values α as illustrated on Figure 11.5. Identifying these intervals allows the user to determine the scales at which the topological features of K^r can be inferred. Notice that intervals on which the persistent Betti numbers are constant can appear at different scales, reflecting multiscale topological features of the offsets of K - see Figure 11.6.

The previous algorithm comes with the following theoretical guarantees justifying the existence of intervals of constant persistent Betti numbers [46].

Theorem 11.12 *Let $K \subset \mathbb{R}^d$ be a compact set with $\text{wfs}(K) > 0$ and let $P \in \mathbb{R}^d$ be a finite set of points such that $d_H(K, P) < \varepsilon$ for some given*




Figure 11.5: An example of persistent Betti numbers curves obtained from a point cloud sampled on a torus. They are plotted as functions of $1/\alpha$. These numbers are constant on an interval containing $[18/\text{wfs}(K), 1/\epsilon]$ and correspond to the three first Betti numbers of the torus: $\beta_0 = 1$, $\beta_1 = 2$ and $\beta_2 = 1$.

$\epsilon > 0$. Assume that $\text{wfs}(K) > 18\epsilon$. Then at each iteration of the algorithm such that $\epsilon < \alpha < \frac{1}{18}\text{wfs}(K)$,

$$\beta_k(K^r) = \beta_{k,4\alpha}^{16\alpha}$$

for any $r \in (0, \text{wfs}(K))$ and any non negative integer k .

The example of Figure 11.6 where P is sampled around a smooth 2-dimensional torus in $\mathbb{R}^{1000}$ illustrates this property: it would not have been possible to do the computation if the complexity was exponential in $d = 1000$.

11.5 Persistent homology

The algorithm of Section 11.2 to compute the Betti numbers of a filtered simplicial complex also provides the Betti numbers of all the subcomplexes of the filtration. Intuitively, the goal of persistent homology is to keep track of all this information and to pair the creation and destruction time of homology classes appearing during the process.




Figure 11.6: The persistent Betti numbers curves obtained from a point cloud sampled on a closed curve spiraling on a torus which was itself non linearly embedded into $\mathbb{R}^{1000}$. They are plotted as functions of $1/\delta$. We distinguish two intervals on which the computed ranks are constant. The right-most one (corresponding to the smaller range of δ) exhibits the Betti numbers of a circle $(1, 1, 0)$, while left most one (intuitively corresponding to a larger scale) exhibits the Betti numbers of the torus $(1, 2, 0)$.

11.5.1 A simple example

Before formally introducing persistent homology, we first consider a very simple example. Let $f : [0, 1] \rightarrow \mathbb{R}$ be the function whose graph is represented on Figure 11.7. We are interested in studying the evolution of the topology of the sublevel sets $f^{-1}((-\infty, t])$ as t increases. The topology of the sublevel sets changes when t crosses the critical values a, b, c, d and e . Passing through the critical value a creates a connected component and for $a \leq t < b$, $f^{-1}((-\infty, t])$ is a connected set (an interval). When t passes through the critical value b a second connected component appears and for $b \leq t < c$, $f^{-1}((-\infty, t])$ has two connected components. When t reaches the value c , the two connected components are merged: the most recently created component, when t passed through b , is merged into the older one. One then pairs the two values b and c corresponding to the “birth” and “death” of the component. In the persistent homology framework, this pairing is either represented by the interval on the left of the graph of f on Figure 11.7 or by the point with coordinates (b, c) in the plane on the right of Figure 11.7. The length $c - b$ of the interval (b, c) represents the lifespan of the component created at b . Intuitively, the larger the interval is, the more

relevant the corresponding component is. Now, continuing to increase t , a new connected component is created when one reaches d which is merged at $t = e$ giving rise to a new persistence interval (d, e) . Notice that a is not paired to any (finite) value since the first created component is never merged into another one. As a consequence it is paired with $+\infty$ and represented by the interval $(a, +\infty)$. At the end, the pairs are either represented as a set of intervals (called a *barcode*) or as a diagram in the plane (called the *persistence diagram* - see Figure 11.7 on the right). For technical reasons that will become clear later in this chapter the diagonal $\{y = x\}$ is added to the diagram.




Figure 11.7: The persistence diagram of a real valued function.

When one considers functions f defined over higher dimensional spaces, passing through critical values may not only change the connectedness of the sublevel sets but also other topological features: creation/destruction of cycles, voids, etc... All these events corresponds to a change in the corresponding homology groups (H_0 for connected components, H_1 for cycles, H_2 for voids,...). In the sequel of this section we show that we can define pairs and persistence diagrams for each dimension.

Now, if we replace the function f by a function g on the Figure 11.8 which is close to f , we observe that the number of pairs of g is much larger than the one of f . However most of these pairs correspond to intervals with short length (points close to the diagonal) while the pairs corresponding to long interval are close to the ones of f . In other words, the topological features having a large persistence with respect to the size of the perturbation are preserved while the topological features created by the perturbation have a small persistence. We will see that this is a general phenomenon: two

close functions have close persistence diagrams. The stability properties of persistence diagrams are of fundamental importance to formalize and quantify the notion of topological noise and to handle noisy data.




Figure 11.8: An approximation g of f and its persistence diagram.

11.5.2 Persistent homology of a filtration

We first define the notion of persistence for a filtration of a simplicial complex. Its goal is to study the evolution of the homology of the subcomplexes of the filtration.

Let K be a d -dimensional simplicial complex and let

$$\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$$

be a filtration of K such that, for any $i = 0, \dots, m-1$, $K^{i+1} = K^i \cup \sigma^{i+1}$ where σ^{i+1} is a simplex.

For any $0 \leq n \leq m$, we denote by C_k^n the set of k -chains (with coefficients in $\mathbb{Z}/2\mathbb{Z}$) of K^n . Notice that the restriction of the boundary operator to C_k^n has its image contained in C_{k-1}^{n-1} . We denote by Z_k^n and B_k^n the sets of k -cycles and k -boundaries of K^n respectively. The k -th homology group of K^n is thus

$$H_k^n = Z_k^n / B_k^n$$

With these notations, we have the following inclusions

$$Z_k^0 \subset Z_k^1 \subset \dots \subset Z_k^n \subset \dots \subset Z_k^m = Z_k(K)$$

$$B_k^0 \subset B_k^1 \subset \dots \subset B_k^n \subset \dots \subset B_k^m = B_k(K)$$

Definition 11.13 (Persistent Betti numbers) For $p \in \{0, \dots, m\}$ et $l \in \{0, \dots, m - p\}$, the k -th persistent Betti number of K^l is the dimension of the vector space $H_k^{l,p} = Z_k^l / (B_k^{l+p} \cap Z_k^l)$.

The k -th persistent Betti number of K^l represents the number of independent homology classes of k -cycles in K^l that are not boundaries in K^{l+p} . Intuitively, a k -cycle in K^l generating a non-zero element in $H_k^{l,p}$ is a cycle that has appeared in the filtration before the step $l + 1$ and that is still not a boundary at step $l + p$. We have seen in Section 11.2 that a homology class is created when a positive simplex is added in the filtration and that a homology class is destroyed when a negative simplex is added. Persistent homology provides a natural way to pair positive and negative simplices such that whenever a positive simplex is added to the filtration it creates a homology class and a corresponding cycle that becomes a boundary when the negative simplex to which it is paired is added.

Cycle associated to a positive simplex

Lemma 11.14 Let $\sigma = \sigma^i$ be a positive k -simplex in the filtration of K . There exists a unique k -cycle c that is not a boundary in K^i , that contains σ and that does not contain any other positive k -simplex.

Proof The lemma is proven by induction on the order of the positive k -simplices in the filtration. Assume that for any positive k -simplex added to the filtration before σ there exists a k -cycle, that is not a boundary and that contains σ but no other positive k -simplex. Since σ is positive, there exists a k -cycle d that is not a boundary in K^i and that contains σ . Let σ_{i_j} , $j = 1, \dots, p$ be the positive k -simplices different from σ contained in d and let c_j be the cycles which are not boundaries containing them and not containing any other positive simplices. Then

$$c = d + c_1 + \dots + c_p$$

is a k -cycle in which σ is the only positive simplex. Since $\sigma = \sigma^i$ is the last simplex added in K^i there does not exist any $(k + 1)$ -simplex in K^i containing σ in its boundary. As a consequence, c cannot be a boundary cycle.

The uniqueness of c is proven in a similar way. $\square$

Persistent homology basis and persistent pairs

The k -cycles associated to the positive k -simplices in Lemma 11.14 allow to maintain a basis of the k -dimensional homology groups of the subcomplexes of the filtration. At the beginning, the basis of $H_k(K^0)$ is empty. Bases of the $H_k(K^i)$ are built inductively in the following way. Assume that the basis of H_k^{i-1} has been built and that the i -th simplex σ^i is positive and of dimension k . We add to the basis of H_k^{i-1} the homology class of the cycle c^i associated to σ^i to obtain a basis of H_k^i . Indeed, since c^i is the sum of σ^i and negative simplices, it is not homologous to any linear combination of cycles defining the basis of H_k^{i-1} . Since $\dim H_k^i = \dim H_k^{i-1} + 1$ we thus obtain a basis of H_k^i .

Now assume that the basis of H_k^{j-1} is built and the j -th simplex σ^j is negative and of dimension $k + 1$. Let $c^{i_1}, \dots, c^{i_p}$ be the cycles associated to the positive simplices $\sigma^{i_1}, \dots, \sigma^{i_p}$ whose homology classes form a basis of H_k^{j-1} . The boundary $d = \partial\sigma^j$ of σ^j is a k -cycle in K^{j-1} which is not a boundary in K^{j-1} but is a boundary in K^j (see the proof of the algorithm to compute Betti numbers in the previous section). So it can be written in a unique way as

$$d = \partial\sigma^j = \sum_{k=1}^p \varepsilon_k c^{i_k} + b$$

where $\varepsilon_k \in \{0, 1\}$ and b is a boundary. We then denote $l(j) = \max\{i_k : \varepsilon_k = 1\}$ and we remove the homology class of $c^{l(j)}$ from the basis of H_k^{j-1} .

Claim: We obtain a basis of H_k^j .

Since $\dim H_k^{j-1} = \dim H_k^j + 1$ we just need to prove that $c^{l(j)}$ is, up to a boundary, a linear combination of the cycles c^{i_k} in Z_k^j , $i_k \neq l(j)$, which is equivalent to the above decomposition of d .

Definition 11.15 (Persistent pairs) *The pairs of simplices $(\sigma^{l(j)}, \sigma^j)$ are called the persistence pairs of the filtration of K .*

Intuitively, the homology class created by $\sigma^{l(j)}$ in $K^{l(j)}$ is destroyed by σ^j in K^j . The persistence of this pair is $j - l(j) - 1$. From the above discussion we deduce a first algorithm to compute the persistent pairs of a filtration of a simplicial complex K of dimension d with m simplices.

Algorithm 11 Persistent pairs computation

Input: A filtration of a d -dimensional simplicial complex K containing m simplices.

$$L_0 = L_1 = \dots = L_{d-1} = \emptyset$$

for $j = 0$ to m **do**

$$k = \dim \sigma^j - 1$$

if σ^j is negative **then**

$l(j) =$ the largest index of the positive k -simplices associated to $\partial\sigma^j$;

$$L_k \leftarrow L_k \cup \{(\sigma^{l(j)}, \sigma^j)\};$$

Output: Return the persistent pairs in each dimension $L_0, L_1, \dots, L_{d-1}$;

Notice that, as for the algorithm to compute the Betti numbers, the main issue with this algorithm is to determine $l(j)$. We overcome it by considering a matrix version of the above algorithm.

Persistence algorithm: the matrix version

We now present an easy to implement version of the persistence algorithm. It relies on a simple reduction of the matrix of the boundary operator.

Let K be a simplicial complex of dimension d and

$$\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$$

be a filtration of K such that for any $i = 0, \dots, m-1$,

$$K^{i+1} = K^i \cup \sigma^{i+1} \quad \text{where } \sigma^{i+1} \text{ is a simplex.}$$

Let $M = (m_{i,j})_{i,j=1,\dots,m}$ be the matrix with coefficients in $\mathbb{Z}/2\mathbb{Z}$ of the boundary operator defined by:

$$m_{i,j} = 1 \quad \text{if } \sigma^i \text{ is a face of } \sigma^j \quad \text{and} \quad m_{i,j} = 0 \quad \text{otherwise.}$$

Hence, if σ^j is a $(k+1)$ -simplex the j -th column of M represents the set of the k -dimensional faces of the boundary of σ^j . Since the simplices of K are ordered according to the filtration, the matrix M is upper triangular. For any column C_j of M , we denote by $l(j)$ the index of the lowest line of M containing a non-zero term in the column C_j :

$$(i = l(j)) \Leftrightarrow (m_{i,j} = 1 \quad \text{and} \quad m_{i',j} = 0 \quad \forall i' > i)$$

Notice that $l(j)$ is not defined when the column C_j does not contain any non zero term. We then have the following very simple algorithm to compute the persistent pairs.

Algorithm 12 Persistent computation - Matrix version

Input: A filtration of a d -dimensional simplicial complex K containing m simplices and the matrix M of the boundary operator.

for $j = 0$ to m **do**

while there exists $j' < j$ such that $l(j') == l(j)$ **do**

$C_j = C_j + C_{j'} \pmod{2}$;

Output: Return the pairs $(l(j), j)$;

Proposition 11.16 *The previous algorithm computes the persistent pairs of the filtration $\emptyset = K^0 \subset K^1 \subset \dots \subset K^m = K$.*

Proof The result immediately follows from a sequence of elementary facts.

Fact 1 : At each step of the algorithm, the column C_j represents a chain of the following form

$$\partial \left(\sigma^j + \sum_{i < j} \varepsilon_i \sigma^i \right) \quad \text{with } \varepsilon_i \in \{0, 1\}.$$

This is proven by an immediate induction.

Fact 2 : At the end of the algorithm, if j is such that $l(j)$ is defined, then $\sigma^{l(j)}$ is a positive simplex.

Indeed, the column C_j represents a chain of the form

$$\sigma^{l(j)} + \sum_{p < l(j)} \eta_p \sigma^p \quad \text{with } \eta_p \in \{0, 1\},$$

but according to Fact 1, C_j also represents a boundary in K^j . So the previous chain is a cycle (since $\partial \circ \partial = 0$) in $K^{l(j)}$ containing $\sigma^{l(j)}$. The simplex $\sigma^{l(j)}$ is thus positive.

Fact 3 : If at the end of the algorithm the column C_j only contains zero terms, then σ^j is positive.

Indeed, according to Fact 1, we have

$$\partial \left(\sigma^j + \sum_{i < j} \varepsilon_i \sigma^i \right) = 0,$$

so σ^j is contained in a cycle of K^j .

Fact 4 : If at the end of the algorithm, the column C_j contains non zero terms, then $(\sigma^{l(j)}, \sigma^j)$ is a persistence pair.

Combining Facts 1 and 2, the boundary of σ^j can be written as

$$\partial \sigma^j = \sigma^{l(j)} + \sum_{p < l(j)} \eta_p \sigma^p + \partial \left(\sum_{i < j} \varepsilon_i \sigma^i \right)$$

Moreover $\sigma^{l(j)}$ is positive and thus was added to the persistent homology basis at time $l(j)$ and has not been paired before time j . Remarking that at the end of the algorithm a line of the matrix cannot contain more than one lowest non-zero term of a column, we deduce that $(\sigma^{l(j)}, \sigma^j)$ is a persistence pair. $\square$

Remark 11.17 Notice that the time complexity of the above algorithm is $O(m^3)$ in the worst case. However in practical applications it usually happens to be much faster ($O(m)$ or $O(m \log m)$).

11.5.3 Persistence diagrams and bottleneck distance

For a fixed k , the persistent pairs of simplices of respective dimensions k and $k + 1$ are conveniently represented as a diagram in the plane $\mathbb{R}^2$: each pair $(\sigma^{l(j)}, \sigma^j)$ is represented by the point of coordinates $(l(j), j)$. For each positive simplex σ^i which is not paired to any negative simplex in the filtration, we associate the pair $(\sigma^i, +\infty)$. For technical reasons that will become clear in the next section, we add to this finite set the diagonal $\{y = x\}$ of $\mathbb{R}^2$ to get the k -dimensional persistence diagram of the filtration. More generally, if the filtration is indexed by a non decreasing sequence of real numbers, as, e.g., in the case of a filtration associated to the sublevel sets of a function,

$$\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \dots \subset K^{\alpha_m} = K \quad \text{with} \quad \alpha_0 \leq \alpha_1 \leq \dots \leq \alpha_m$$

a persistent pair of simplices (σ_i, σ_j) is represented in the diagram as the point of coordinates (α_i, α_j) . In this later case, we have to take care that, since the sequence (α_i) is non decreasing, several pairs can be associated to the same point in the plane. A persistence diagram is thus a multiset and the *multiplicity* of a point is defined as the number of pairs associated to this point. By convention, the points on the diagonal have infinite multiplicity. By convention, if a simplex σ_i is not paired, we then add the point of coordinates $(\alpha_i, +\infty)$ to the persistence diagram. Notice that, as a consequence, the persistence diagram is a multiset in $\overline{\mathbb{R}}^2$ where $\overline{\mathbb{R}} = \mathbb{R} \cup +\infty$.

Persistence diagrams can be compared using a matching distance called the *bottleneck distance*.

Definition 11.18 (Bottleneck distance) *Let D_1 and D_2 be two persistence diagrams. The bottleneck distance between D_1 and D_2 is defined as*

$$d_B(D_1, D_2) = \inf_{\gamma} \sup_{p \in D_1} \|p - \gamma(p)\|_{\infty}$$

where γ is the set of bijections between the multi-sets D_1 and D_2 (a point with multiplicity $m > 1$ is considered as m disjoint copies) and $\|p - q\|_{\infty} = \max(|x_p - x_q|, |y_p - y_q|)$. By convention, if $y_p = y_q = +\infty$, then $\|p - q\|_{\infty} = |x_p - x_q|$.

The above definition motivates the inclusion of the diagonal of $\mathbb{R}^2$ in the definition of persistence diagram: it allows to compare diagrams that do not have the same numbers of off-diagonal points by matching them with points on the diagonal (see Figure 11.9).

11.5.4 Persistence modules, interleaving and stability

Persistence can be defined in a purely algebraic way that turns out to be particularly useful and powerful in many settings. The proofs of the results presented in this section are beyond the scope of this book but the algebraic framework introduced in this section allows to efficiently prove the results of the next sections.

The notion of persistence can be extended to more general sequences of vector spaces in the following way.




Figure 11.9: The bottleneck distance between two diagrams

Definition 11.19 (Persistence modules) A persistence module $\mathbb{V}$ over a subset I of the real numbers $\mathbb{R}$ is an indexed family of vector spaces $(V_a \mid a \in I)$ and a doubly-indexed family of linear maps $(v_a^b : V_a \rightarrow V_b \mid a \leq b)$ which satisfy the composition law $v_b^c \circ v_a^b = v_a^c$ whenever $a \leq b \leq c$, and where v_a^a is the identity map on V_a .

Definition 11.20 (q-tameness) The persistence module $\mathbb{V}$ is said to be q -tame if $\text{rk}(v_a^b) < +\infty$ whenever $a < b$.

The sequence of homology groups of the filtration of K considered in the previous section together with the homomorphisms induced by the inclusion maps is a persistence module indexed over the set $I = \{0, 1, \dots, m\}$. It can be shown that the persistence diagram of a filtration $\{\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \dots \subset K^{\alpha_m} = K\}$ is completely determined by the rank of the homomorphisms $H_k(K^{\alpha_i}) \rightarrow H_k(K^{\alpha_j})$ for any $i < j$. This property extends to q -tame modules.

Theorem 11.21 *If a persistence module $\mathbb{V}$ is q -tame, then it has a well-defined persistence diagram $\mathbf{dgm}(\mathbb{V}) \subset \overline{\mathbb{R}}^2$. When $\mathbb{V}$ is the persistence module defined by a filtration of a finite simplicial complex, this diagram coincides with the one defined in Section 11.5.3.*

To avoid technical difficulties in the sequel of this section, we assume that all the considered persistence modules are indexed by $\mathbb{R}$. This is not a restrictive assumption in our setting: if $\mathbb{V}$ is a persistence module indexed by a finite set $\alpha_1 < \alpha_2 < \dots < \alpha_m$ it can be extended to a piecewise constant module indexed by $\mathbb{R}$ by defining $V_\alpha = V_{\alpha_j}$ for $\alpha \in [\alpha_j, \alpha_{j+1})$ with the convention $\alpha_0 = -\infty$ and $\alpha_{m+1} = +\infty$. The linear maps are then defined in an obvious way by letting $v_a^b = Id$ whenever $\alpha_j \leq a \leq b < \alpha_{j+1}$.

Let $\mathbb{U}, \mathbb{V}$ be persistence modules over $\mathbb{R}$, and let ε be any real number. A homomorphism of degree ε is a collection Φ of linear maps

$$(\phi_a : U_a \rightarrow V_{a+\varepsilon} \mid a \in \mathbb{R})$$

such that $v_{a+\varepsilon}^{b+\varepsilon} \circ \phi_a = \phi_b \circ u_a^b$ for all $a \leq b$. We write

$$\begin{aligned} \text{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}) &= \{\text{homomorphisms } \mathbb{U} \rightarrow \mathbb{V} \text{ of degree } \varepsilon\}, \\ \text{End}^\varepsilon(\mathbb{V}) &= \{\text{homomorphisms } \mathbb{V} \rightarrow \mathbb{V} \text{ of degree } \varepsilon\}. \end{aligned}$$

Composition is defined in the obvious way. For $\varepsilon \geq 0$, the most important degree- ε endomorphism is the shift map

$$1_{\mathbb{V}}^\varepsilon \in \text{End}^\varepsilon(\mathbb{V}),$$

which is the collection of maps $(v_a^{a+\varepsilon})$ from the persistence structure on $\mathbb{V}$.

If Φ is a homomorphism $\mathbb{U} \rightarrow \mathbb{V}$ of any degree, then by definition $\Phi 1_{\mathbb{U}}^\varepsilon = 1_{\mathbb{V}}^\varepsilon \Phi$ for all $\varepsilon \geq 0$.

Definition 11.22 (Interleaving of persistence modules) *Two persistence modules $\mathbb{U}, \mathbb{V}$ are said to be ε -interleaved if there are maps*

$$\Phi \in \text{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}), \quad \Psi \in \text{Hom}^\varepsilon(\mathbb{V}, \mathbb{U})$$

such that $\Psi\Phi = 1_{\mathbb{U}}^{2\varepsilon}$ and $\Phi\Psi = 1_{\mathbb{V}}^{2\varepsilon}$.

The notion of interleaving allows to state the fundamental stability theorem for persistence diagrams.

Theorem 11.23 (Stability of persistence diagrams) *Let $\mathbb{U}, \mathbb{V}$ be two q -tame persistent modules that are ε -interleaved for some $\varepsilon \geq 0$. Denoting by $\text{dgm}(\mathbb{U})$ and $\text{dgm}(\mathbb{V})$ their persistence diagrams, we have*

$$d_B(\text{dgm}(\mathbb{U}), \text{dgm}(\mathbb{V})) \leq \varepsilon.$$

In the next sections, we apply the stability to different settings.

11.5.5 Persistence stability for functions

Let $f : X \rightarrow \mathbb{R}$ be a real-valued function defined on a topological space X . Let consider the *sublevel set filtration* $\{F_\alpha = f^{-1}((-\infty, \alpha])\}_{\alpha \in \mathbb{R}}$ and consider the (singular) homology groups $H_k(F_\alpha)$ of these sublevel sets. Notice that the canonical inclusion $F_\alpha \subseteq F_\beta$ whenever $\alpha \leq \beta$ induces an homeomorphism $H_k(F_\alpha) \rightarrow H_k(F_\beta)$. So, the sublevel sets filtration of f induces a persistence module $\mathbb{F}_k$.

Definition 11.24 *A real-valued function f defined on a topological space X is said to be q -tame if its associated persistence modules $\mathbb{F}_k$ are q -tame for any non negative integer k .*

The following result provides sufficient conditions for f and g to be q -tame.

Proposition 11.25 *If X is homeomorphic to a finite simplicial complex and $f : X \rightarrow \mathbb{R}$ is continuous, then f is q -tame. In particular, $\text{dgm}(\mathbb{F}_k)$ is well-defined.*

Proposition 11.26 *Let $f, g : X \rightarrow \mathbb{R}$ be two functions defined on a topological space X such that $\|f - g\|_\infty = \sup_{x \in X} |f(x) - g(x)| < \varepsilon$. Then the persistence modules $\mathbb{F}_k$ and $\mathbb{G}_k$ induced by the sublevel sets filtrations of f and g are ε -interleaved.*

Proof Since $\|f - g\|_\infty < \varepsilon$, we have, for any $\alpha \in \mathbb{R}$, $F_\alpha \subseteq G_{\alpha+\varepsilon} \subseteq F_{\alpha+2\varepsilon} \subseteq G_{\alpha+3\varepsilon} \subseteq \dots$. These inclusions induce homomorphisms $H_k(F_\alpha) \rightarrow H_k(G_{\alpha+\varepsilon})$ and $H_k(G_\alpha) \rightarrow H_k(F_{\alpha+\varepsilon})$ for all $\alpha \in \mathbb{R}$. The sets of these homomorphisms define an ε -interleaving between $\mathbb{F}_k$ and $\mathbb{G}_k$. $\square$

In the sequel, when there is no ambiguity, the notation $\text{dgm}(f)$ denotes the persistence diagram of $\mathbb{F}_k$, for any k . Applying the Persistence Stability Theorem 11.23, we immediately obtain the following stability result for functions.

Theorem 11.27 *Let X be a topological space and let $f, g : X \rightarrow \mathbb{R}$ be two q -tame functions. Then*

$$d_B(\text{dgm}(\mathbb{F}_k), \text{dgm}(\mathbb{G}_k)) \leq \|f - g\|_\infty.$$

From a practical point of view, the above theorem provides a rigorous way to approximate the persistence diagrams of continuous functions defined on a triangulated space. For example, assume that X is a triangulated manifold in $\mathbb{R}^d$ where the diameter of each simplex is upper bounded by some $\delta > 0$ and assume $f : X \rightarrow \mathbb{R}$ to be c -Lipschitz for some $c > 0$, i.e. $|f(x) - f(x')| \leq c\|x - x'\|$. Then, given a non negative integer k , one can easily check that the bottleneck distance between the persistence diagram of f and the persistence diagram of the filtered complex induced by the values of f at the vertices of the triangulation is upper bounded by $c\delta$ (see Exercise 11.8).

11.5.6 Persistence stability for compact sets and complexes built on top of point clouds

Proposition 11.28 *Let $X \subset \mathbb{R}^d$ be a compact set. The distance function $d_X : \mathbb{R}^d \rightarrow \mathbb{R}$ is q -tame.*

Proof Given $0 \leq \alpha < \beta$ and a non negative integer k , we just need to prove that the homomorphism $H_k(X^\alpha) \rightarrow H_k(X^\beta)$ induced by the inclusion of the offsets, $X^\alpha \subset X^\beta$, has finite rank. Denoting $\varepsilon = (\beta - \alpha)/2 > 0$, since X is compact, there exists a finite subset $P \subseteq X$ of X such that $d_H(X, P) < \varepsilon$. As a consequence we have the following inclusion

$$X^\alpha \subseteq P^{\alpha+\varepsilon} \subseteq X^\beta$$

that induces the following sequence of homomorphisms

$$H_k(X^\alpha) \rightarrow H_k(P^{\alpha+\varepsilon}) \rightarrow H_k(X^\beta).$$




Figure 11.10: Comparing the persistence diagrams of two close functions defined on a segment

Now, it follows from the nerve Lemma that $P^{\alpha+\varepsilon}$ is a finite union of balls homotopy equivalent to $\check{\text{Cech}}(P, \alpha + \varepsilon)$ which is a finite simplicial complex. As a consequence $\dim H_k(P^{\alpha+\varepsilon}) < +\infty$ and $rk(H_k(X^\alpha) \rightarrow H_k(X^\beta)) \leq \dim H_k(P^{\alpha+\varepsilon}) < +\infty$. $\square$

The above proposition implies that distance functions to compact subsets of $\mathbb{R}^d$ have well-defined persistence diagrams. Recalling that if $X, Y \subset \mathbb{R}^d$ are compact, then $d_H(X, Y) = \|d_X - d_Y\|_\infty$, we immediately obtain the following corollary.

Corollary 11.29 *Let $X, Y \subset \mathbb{R}^d$ be compact. Then*

$$d_B(\text{dgm}(d_X), \text{dgm}(d_Y)) \leq d_H(X, Y).$$

In particular, if $P, Q \subset \mathbb{R}^d$ are finite point clouds, then for any non negative integer k ,

$$d_B(\text{dgm}(H_k(\check{\text{Cech}}(P))), \text{dgm}(H_k(\check{\text{Cech}}(Q)))) \leq d_H(P, Q)$$

where $\check{\text{Cech}}(.)$ denotes the Čech filtration.

The second part of the corollary follows from the nerve Lemma 2.8 and the paragraph before Proposition 11.10 showing that P^α and $\check{\text{Cech}}(P, \alpha)$ are homotopy equivalent for any α and the homotopy equivalences can be chosen to commute with the inclusion maps at the homology level. As a consequence the persistence modules induced by the sublevel sets of d_P and the Čech filtration are 0-interleaved. Using that P^α is also homotopy equivalent to the alpha-complex $\mathcal{A}(P, \alpha)$ the same result also holds when $\check{\text{Cech}}(P)$ is replaced by the alpha-complex filtration $\mathcal{A}(P)$.

Application: topological signatures for shapes. The same kind of result can be also established for the Vietoris-Rips complex and, thanks to these stability properties, the obtained persistence diagrams can be considered as robust multiscale topological signatures associated to point clouds. They can thus be used to compare the topological structure of points clouds sampled from different shapes. Notice that if the finite point cloud P is transformed by an isometry of $\mathbb{R}^d$ into another point cloud P' , then the Čech, the Vietoris-Rips and alpha-shape filtrations of P and P' are the same. However, $d_H(P', Q)$ can become much larger than $d_H(P, Q)$ while the bottleneck distance between persistence diagrams remains unchanged, making the second inequality of Corollary 11.29 less interesting. To overcome this issue one can consider the *Gromov-Hausdorff distance* defined in the following way.

Definition 11.30 Let $X, Y \subset \mathbb{R}^d$ be two compact sets and let $\varepsilon \geq 0$. An ε -correspondence between X and Y is a subset $C \subseteq X \times Y$ such that

- (i) for any $x \in X$, there exists $y \in Y$ such that $(x, y) \in C$;
- (ii) for any $y \in Y$, there exists $x \in X$ such that $(x, y) \in C$;
- (iii) for any $(x, y), (x', y') \in C$, $|d(x, x') - d(y, y')| \leq \varepsilon$, where $d(x, x') = \|x - x'\|$ is the Euclidean distance.

The Gromov-Hausdorff distance between X and Y is defined by

$$d_{GH}(X, Y) = \inf\{\varepsilon \geq 0 : \text{there exists an } \varepsilon\text{-correspondence between } X \text{ and } Y\}.$$

Notice that the above definition can be extended verbatim to any pair of compact metric spaces. Indeed the Gromov-Hausdorff distance allows to compare compact metric spaces, up to isometry, independently of any embedding. Coming back to the point clouds P, P' and Q where P' is the image

of P by an ambient isometry of $\mathbb{R}^d$, we have that $d_{GH}(P, Q) = d_{GH}(P', Q)$. Moreover, Corollary 11.29 has the following generalization.

Theorem 11.31 *Let $P, Q \subset \mathbb{R}^d$ be finite point clouds and let $\text{Filt}(\cdot)$ be any of the Čech, Vietoris – Rips or alpha-shape filtered complexes. Then for any non negative integer k ,*

$$d_B(\text{dgm}(H_k(\text{Filt}(P))), \text{dgm}(H_k(\text{Filt}(Q)))) \leq d_{GH}(P, Q).$$

This theorem can be extended to point clouds in non Euclidean metric spaces, except for the alpha-shape filtration which is no longer defined. In particular, to define the Vietoris-Rips complex, one just needs to know the pairwise distances between the points. As the computation of the Gromov-Hausdorff distance is usually intractable in practice, we can thus use the persistence diagrams of the Vietoris-Rips filtrations to compare the topological structure of finite data sets coming with pairwise distance information. Thanks to Theorem 11.31, the bottleneck distance provides a discriminative comparison tool: if the bottleneck distance between the diagrams is large, the two corresponding sets are far away from each other with respect to d_{GH} . The reverse is not true.

11.6 Exercises

Exercise 11.1 Let K be a finite simplicial complex. Prove that $\beta_0(K)$ is equal to the number of connected components of K .

Hint: use the result of Exercise 2.1.

Exercise 11.2 Compute the Betti numbers of the simplicial complexes of Figure 11.3.

Exercise 11.3 (Difficult) Let P be a finite set of points in $\mathbb{R}^2$. Prove that for any $\alpha \geq 0$

$$\text{Rips}(P, \alpha) \subseteq \check{\text{Cech}}(P, \alpha \sqrt{\frac{d}{2(d+1)}}) \subseteq \text{Rips}(P, 2\alpha \sqrt{\frac{d}{2(d+1)}})$$

Hint: see [62], Theorem 2.5.

Exercise 11.4 Chains with coefficient in $\mathbb{Z}/2\mathbb{Z}$ have an obvious geometric interpretation: since any k -chain can be uniquely written as $c = \sigma_{i_1} + \sigma_{i_2} + \cdots + \sigma_{i_m}$ where the σ_{i_j} are k -simplices, c can be considered as the union of the simplices σ_{i_j} . Show that the sum of two k -chains is equal to their symmetric difference¹.

Exercise 11.5 Let $P \subset \mathbb{R}^d$ be a finite set of points. Prove that, for any $r \geq 0$, the Betti numbers $\beta_k(P^r)$ of the r -offset P^r are finite.

Hint: use the Nerve theorem.

Exercise 11.6 Let $\mathbb{F}$ be a filtration of a simplicial complex K . Prove that all the vertices of K are positive and that an edge σ^i is positive if and only if the two ends (vertices) of σ^i are in the same connected component of K^{i-1} .

Exercise 11.7 Let $\mathbb{F}$ be a filtration of a simplicial complex K .

1. Prove that any cycle in K contains at least one positive simplex.
2. Prove that the cycle associated to a positive simplex in lemma 11.14 is uniquely defined.

Exercise 11.8 Let X be a (finitely) triangulated subset of $\mathbb{R}^d$ and let $f : X \rightarrow \mathbb{R}$ be a c -Lipschitz function, $c > 0$. Let K_f be the filtration induced by f on the triangulation of X . Denoting by $\delta > 0$ the largest diameter of the simplices of the triangulation of X , prove that,

$$d_B(\text{dgm}(K_f), \text{dgm}(f)) \leq c\delta.$$

11.7 Bibliographical notes

A detailed introduction to algebraic topology and simplicial and singular homology can be found [113, 97].

The results of Section 11.4 are derived from [46] where complete proofs are given.

Persistent homology has been independently introduced by different authors [71, 84, 121] and has known important developments during the last decade.

¹The symmetric difference of two sets A and B is defined by $A \Delta B = (A \cup B) \setminus (A \cap B)$.

The stability of persistence diagrams has been initially proven by [60] for tame continuous functions defined on triangulable spaces. It was then extended and generalized by [35] and [41] to a more algebraic framework that appeared of fundamental importance in topological data analysis. An introductory course to computational topology is provided in [75] and a recent and general presentation of persistent homology theory is given in [117].

Theorem 11.31 is a particular case of a result proven in [42] and it has found various applications in shape classification [36] and in statistical analysis of data - see, e.g., [50, 80, 11, 48].

Bibliography

- [1] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, New York, USA, 2nd edition, 2008.
- [2] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.
- [3] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of the medial axes. a state-of-the-art report. In T. Möller, B. Hamann, and R. Russell, editors, *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, volume 555 of *Mathematics and Visualization*, pages 109–125. Springer-Verlag, 2009.
- [4] D. Attali, Herbert Edelsbrunner, and Yuriy Mileyko. Weak witnesses for Delaunay triangulations of submanifolds. In *Proc. ACM Sympos. Solid and Physical Modeling*, pages 143–150, 2007.
- [5] Dominique Attali, André Lieutier, and David Salinas. Efficient data structure for representing and simplifying simplicial complexes in high dimensions. *Int. J. Comput. Geometry Appl.*, 22(4):279–304, 2012.
- [6] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16:78–96, 1987.
- [7] Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013.
- [8] Callahan Paul B. and Kosaraju Sambasiva Rao. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *Journal of the ACM*, 42(1):67–90, 1995.
- [9] M. Berger. *Géométrie (vols. 1-5)*. Fernand Nathan, Paris, 1977.

- [10] Gérard Biau, Frédéric Chazal, David Cohen-Steiner, Luc Devroye, Carlos Rodriguez, et al. A weighted k-nearest neighbor density estimate for geometric inference. *Electronic Journal of Statistics*, 5:204–237, 2011.
- [11] AndrewJ. Blumberg, Itamar Gal, MichaelA. Mandell, and Matthew Pancia. Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Foundations of Computational Mathematics*, 14(4):745–789, 2014.
- [12] J-D. Boissonnat and J. Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004.
- [13] J-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential delaunay complexes. In *Proc. 26th Annual Symposium on Computational Geometry*, 2010.
- [14] J-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. ACM Symp. on Computational Geometry*, pages 270–277, 2008.
- [15] Jean-Daniel Boissonnat, Karthik C. S., and Sébastien Tavenas. Building Efficient and Compact Data Structures for Simplicial Complexes. In *International Symposium on Computational Geometry 2015*, Eindhoven, Netherlands, June 2015.
- [16] Jean-Daniel Boissonnat, Olivier Devillers, Kunal Dutta, and Marc Glisse. Delaunay triangulation of a random sample of a good sample has linear size. Technical report, 2017.
- [17] Jean-Daniel Boissonnat, Olivier Devillers, Sylvain Pion, Monique Teilaud, and Mariette Yvinec. Triangulations in CGAL. *Comput. Geom. Theory Appl.*, 22:5–19, 2002.
- [18] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay triangulation of manifolds. Rapport de recherche RR-8389, INRIA, October 2013.
- [19] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay stability via perturbations. *Int. J. Comput. Geometry Appl.*, 24(2):125–152, 2014.

- [20] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. The Stability of Delaunay Triangulations. *Int. J. on Comput. Geometry Appl.*, 23(4 & 5):303–333, 2014.
- [21] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. A probabilistic approach to reducing algebraic complexity of delaunay triangulations. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 595–606, 2015.
- [22] Jean-Daniel Boissonnat and Arijit Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete and computational Geometry*, (November), 2013.
- [23] Jean-Daniel Boissonnat and Clément Maria. The Simplex Tree: An Efficient Data Structure for General Simplicial Complexes. *Algorithmica*, 70(3):406–427, November 2014.
- [24] Jean-Daniel Boissonnat, Frank Nielsen, and Richard Nock. Bregman Voronoi diagrams. *Discrete and Computational Geometry*, 44(2), 2010.
- [25] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67:405–451, 2005.
- [26] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Curved Voronoi diagrams. In Jean-Daniel Boissonnat and Monique Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, pages 67–116. Springer-Verlag, Mathematics and Visualization, 2006.
- [27] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.
- [28] F. Bolley, A. Guillin, and C. Villani. Quantitative Concentration Inequalities for Empirical Measures on Non-compact Spaces. *Probab. Theory Rel.*, 137(3):541–593, 2007.
- [29] Mickaël Buchet, Frédéric Chazal, Steve Y Oudot, and Donald R Sheehy. Efficient and robust persistent homology for measures. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 168–180. SIAM, 2015.

- [30] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33.
- [31] Claire Caillerie, Frédéric Chazal, Jérôme Dedecker, Bertrand Michel, et al. Deconvolution for the wasserstein metric and geometric inference. *Electronic Journal of Statistics*, 5:1394–1423, 2011.
- [32] S. S. Cairns. A simple triangulation method for smooth manifolds. *Bull. Amer. Math. Soc.*, 67:380–390, 1961.
- [33] P. Cannarsa and C. Cinestrari. *Semiconcave Functions, Hamilton-Jacobi Equations, and Optimal Control*, volume 58. Birkhauser, 2004.
- [34] R. Chaine. A geometric convection approach of 3d-reconstruction. In *1st Symposium on Geometry Processing*, pages 218–229, 2003.
- [35] F. Chazal, D. Cohen-Steiner, L. J. Guibas, M. Glisse, and S. Y. Oudot. Proximity of persistence modules and their diagrams. In *Proc. 25th ACM Sympos. Comput. Geom.*, 2009.
- [36] F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Memoli, and S. Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum*, pages 1393–1403, 2009.
- [37] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in euclidean space. *Discrete and Computational Geometry*, 41(3):461–479, 2009.
- [38] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert. Shape smoothing using double offsets. In *Proc. ACM solid and Physical Modeling*, 2007.
- [39] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert. Stability of Curvature Measures. *Computer Graphics Forum (proc. SGP 2009)*, pages 1485–1496, 2009.
- [40] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.
- [41] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. The structure and stability of persistence modules. arXiv:1207.3674 [math.AT], 2012.
- [42] F. Chazal, V. de Silva, and S. Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014.

- [43] F. Chazal and A. Lieutier. The λ -medial axis. *Graphical Models*, 67(4):304–331, July 2005.
- [44] F. Chazal and A. Lieutier. Weak feature size and persistent homology: Computing homology of solids in $\mathbb{R}^n$ from noisy data samples. In *Proc. 21st Annual ACM Symposium on Computational Geometry*, pages 255–262, 2005.
- [45] F. Chazal and A. Lieutier. Stability and computation of topological invariants of solids in $\mathbb{R}^n$. *Discrete Comput. Geom.*, 37(4):601–617, 2007.
- [46] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 232–241, 2008.
- [47] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. Normal cone approximation and offset shape isotopy. *Comput. Geom. Theory Appl.*, 42(6-7):566–581, 2009.
- [48] Frederic Chazal, Brittany Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Subsampling methods for persistent homology. In David Blei and Francis Bach, editors, *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2143–2151. JMLR Workshop and Conference Proceedings, 2015.
- [49] Frédéric Chazal, Brittany T Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Robust topological inference: Distance to a measure and kernel distance. *arXiv preprint arXiv:1412.7197*, 2014.
- [50] Frédéric Chazal, Marc Glisse, Catherine Labruère, and Bertrand Michel. Convergence rates for persistence diagram estimation in topological data analysis. *Journal of Machine Learning Research*, 16:3603–3635, 2015.
- [51] Frédéric Chazal, Pascal Massart, and Bertrand Michel. Rates of convergence for robust geometric inference. *arXiv preprint arXiv:1505.07602*, 2015.
- [52] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.

- [53] J. Cheeger. Critical points of distance functions and applications to geometry. In *Geometric Topology: recent developments, Montecani Terme*, volume 1504 of *Lecture Notes in Math.*, pages 1–38, 1990.
- [54] J. Cheeger, W. Müller, and R. Schrader. On the curvature of piecewise flat spaces. *Comm. Mth. Phys.*, 92:405–454, 1984.
- [55] S-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S-H. Teng. Sliver Exudation. *Journal of ACM*, 47:883–904, 2000.
- [56] S-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold Reconstruction from Point Samples. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pages 1018–1027, 2005.
- [57] Siu-Wing Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *Proc. ACM-SIAM Sympos. Discrete Algorithms*, pages 1018–1027, 2005.
- [58] F. H. Clarke. *Optimization and Nonsmooth Analysis*. Wiley-Interscience, New-York, 1983.
- [59] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [60] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37(1):103–120, 2007.
- [61] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.
- [62] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7:339–358, 2007.
- [63] Vin de Silva. A weak characterisation of the delaunay triangulation. *Geometriae Dedicata*, 135(1):39–64, 2008.
- [64] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934.
- [65] T. K. Dey. *Curve and Surface Reconstruction : Algorithms with Mathematical Analysis*. Cambridge University Press, 2007.

- [66] Ramsay Dyer, Gert Vegter, and Mathijs Wintraecken. Riemannian simplices and triangulations. In *31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands*, pages 255–269, 2015.
- [67] H. Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13(1):415–440, 1995.
- [68] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. In *B. Aronov, S. Basu, J. Pach, M. Sharir (Eds.), Ricky Pollack and Eli Goodman Festschrift*, pages 379–404. Springer-Verlag, Berlin, 2003.
- [69] H. Edelsbrunner, M. Facello, P. Fu, and J. Liang. Measuring proteins and voids in proteins. In *Proc. 28th Ann. Hawaii Internat. Conf. System Sciences*, volume V of *Biotechnology Computing*, pages 256–264, 1995.
- [70] H. Edelsbrunner, M. Facello, and J. Liang. On the definition and the construction of pockets in macromolecules. *Discrete Applied Mathematics*, 88(1-3):83–102, 1998.
- [71] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
- [72] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. In *Proceedings of the 1992 workshop on Volume visualization*, pages 75–82. ACM, 1992.
- [73] Herbert Edelsbrunner. Algorithms in combinatorial geometry, volume 10 of *eatcs monographs on theoretical computer science*, 1987.
- [74] Herbert Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, 2001.
- [75] Herbert Edelsbrunner and John L Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.
- [76] Herbert Edelsbrunner, David Kirkpatrick, and Raimund Seidel. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29(4):551–559, 1983.
- [77] Herbert Edelsbrunner, Xiang-Yang Li, Gary L. Miller, Andreas Stathopoulos, Dafna Talmor, Shang-Hua Teng, Alper Üngör, and Noel

- Walkington. Smoothing and cleaning up slivers. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 273–277, 2000.
- [78] Herbert Edelsbrunner and Hubert Wagner. Topological data analysis with bregman divergences. *CoRR*, abs/1607.06274, 2016.
- [79] David Eppstein, Maarten Löffler, and Darren Strash. Listing all maximal cliques in sparse graphs in near-optimal time. In *Algorithms and Computation - 21st International Symposium, ISAAC 2010, Jeju Island, Korea, December 15-17, 2010, Proceedings, Part I*, pages 403–414, 2010.
- [80] B. T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, S. Balakrishnan, A. Singh, et al. Confidence sets for persistence diagrams. *The Annals of Statistics*, 42(6):2301–2339, 2014.
- [81] Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 434–444, New York, NY, USA, 1988. ACM.
- [82] Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.
- [83] D Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002.
- [84] P. Frosini and C. Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9:596–603, 1999.
- [85] J. H. Fu. Convergence of curvatures in secant approximations. *J. Differential Geometry*, 37:177–190, 1993.
- [86] J. H. G. Fu. Tubular neighborhoods in Euclidean spaces. *Duke Mathematical Journal*, 52(4):1025–1046, 1985.
- [87] W. Fulton. *Algebraic Topology: a First Course*. Springer-Verlag, 1995.
- [88] Jean Gallier. Notes on Convex Sets, Polytopes, Polyhedra Combinatorial Topology, Voronoi Diagrams and Delaunay Triangulations. Research Report RR-6379, 2007.

- [89] S. Gallot, D. Hulin, and J. Lafontaine. *Riemannian Geometry*. Springer-Verlag, 1990.
- [90] J. Giesen and M. John. The flow complex: a data structure for geometric modeling. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 285–294, 2003.
- [91] Joachim Giesen and Uli Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete & Computational Geometry*, 32:245–267, 2004.
- [92] T.F Gonzales. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38(2-3):293–306, 1985.
- [93] K. Grove. Critical point theory for distance functions. In *Proc. of Symposia in Pure Mathematics*, volume 54, 1993. Part 3.
- [94] Leonidas Guibas, Dmitriy Morozov, and Quentin Mérigot. Witnessed k-distance. *Discrete & Computational Geometry*, 49(1):22–45, 2013.
- [95] Sariel Har-Peled. *Geometric approximation algorithms*, volume 173. American Mathematical Society, 2011.
- [96] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [97] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [98] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, NY, 1976.
- [99] I. T. Jolliffe. *Principal component analysis*. Springer, 2002.
- [100] X-Y. Li. Generating Well-Shaped d -dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.
- [101] X-Y. Li. Generating Well-Shaped d -dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.
- [102] J. Liang, H. Edelsbrunner, and C. Woodward. Anatomy of protein pockets and cavities: measurement of binding site geometry and implications for ligand design. *Protein Science*, 7:1884–1897, 1998.

- [103] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, September 2004.
- [104] André Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.
- [105] Y. Ma, P. Niyogi, G. Sapiro, and R. Vidal Ed. Dimensionality reduction via subspace and submanifold learning. *IEEE Signal Processing Magazine*, 28(2), 2011.
- [106] Clément Maria, Jean-Daniel Boissonnat, Marc Glisse, and Mariette Yvinec. The Gudhi Library: Simplicial Complexes and Persistent Homology. In *The 4th International Congress on Mathematical Software (ICMS)*, Hanyang University, Seoul, Korea, France, August 2014.
- [107] Thomas Martinetz and Klaus Schulten. Topology representing networks. *Neural Networks*, 7(3):507–522, 1994.
- [108] J. Milnor. *Morse Theory*. Cambridge, 2006.
- [109] Robin A. Moser and Gábor Tardos. A constructive proof of the generalized Lovász lemma. *Journal of the ACM*, (2), 2010.
- [110] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.
- [111] Ketan Mulmuley. *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice Hall, 1994.
- [112] J. R. Munkres. *Elementary differential topology*. Princeton University Press, 1966.
- [113] J. R. Munkres. *Elements of algebraic topology*. Addison-Wesley, Redwood City, California, 1984.
- [114] James Munkres. *Topology*. Prentice Hall, Inc, 2000.
- [115] P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete and Computational Geometry*, 39(1):419–441, 2008.

- [116] Atsuyuki Okabe, Barry Boots, and Kokichi Sugihara. *Spatial Tesselations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Chichester, UK, 1992.
- [117] Steve Y. Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*, volume 209 of *AMS Mathematical Surveys and Monographs*. American Mathematical Society, 2015.
- [118] D. Pedoe. *Geometry, a comprehensive course*. Dover Publications, New York, 1970.
- [119] S. Peleg, M. Werman, and H. Rom. A unified approach to the change of resolution: space and gray-level. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(7):739–742, 1989.
- [120] A. Petrunin. Semiconcave functions in alexandrov’s geometry. In *Surveys in Differential Geometry: Metric and Comparison Geometry, volume XI*. International Press of Boston, 2007.
- [121] V. Robins. Toxard computing homology from finite approximations. *Topology Proceedings*, 24:503–532, 1999.
- [122] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, New Jersey, 1970.
- [123] Y. Rubner, C. Tomasi, and L.J. Guibas. The Earth Mover’s Distance as a Metric for Image Retrieval. *Int. J. Comput. Vision*, 40(2):99–121, 2000.
- [124] R. Seidel. The upper bound theorem for polytopes: an easy proof of its asymptotic version. *Comput. Geom. Theory Appl.*, 5:115–116, 1995.
- [125] Raimund Seidel. A convex hull algorithm optimal for point sets in even dimensions. 1981.
- [126] J. Shewchuk. Star splaying: An algorithm for repairing delaunay triangulations and convex hulls. In *Proc. 21st Annual Symposium on Computational Geometry*, pages 237–246, 2005.
- [127] R. Sibson. A vector identity for the Dirichlet tessellation. *Math. Proc. Camb. Phil. Soc.*, 87:151–155, 1980.

- [128] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnet, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.
- [129] J. Spencer. Robin Moser makes Lovász Local Lemma Algorithmic! Technical report.
- [130] The GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015.
- [131] L.N. Trefethen and D. Bau. *Numerical linear algebra*. Society for Industrial Mathematics, 1997.
- [132] C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.
- [133] Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In H. Maurer, editor, *New Results and New Trends in Computer Science*, volume 555 of *Lecture Notes Comput. Sci.*, pages 359–370. Springer-Verlag, 1991.
- [134] J. H. C. Whitehead. On C^1 -complexes. *Ann. of Math.*, 41:809–824, 1940.
- [135] H. Whitney. *Geometric integration theory*. Princeton University Press, 1957.
- [136] Mathijs Wintraecken. Bounds on the angle between tangent spaces and the metric distortion for C^2 manifolds with given positive reach. In *European Workshop on Computational Geometry (EUROCG2017)*, 2017.
- [137] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, 1994.

Index

- ε -correspondence, 282
- (κ, μ) -approximation, 229
- adjacency graph, 38, 39
- affine diagrams, 86
- affinely independent, 30
- alpha complex, 137, 185
- altitude, 112
- Betti number, 259
 - persistent, 271
- boundary, 20
- boundary operator, 258
- Bregman
 - complex, 91
 - diagram, 88, 90
 - divergence, 88
- Čech complex, 35
- cell complex, 49, 74
- chain complex, 258
- closure, 20
- combinatorial manifold, 38
- compact space, 21
- connected space, 21
- contractible space, 24
- convex polyhedron, 48
- convex polytope, 46
 - boundary complex, 49
 - face, 47
 - supporting hyperplane, 46
- cover, 33
 - good cover, 34
- critical function, 225
- critical point
 - α -critical point, 238
 - μ -critical point, 221
 - of distance function, 218
- cycle, 259
- deformation retract, 25
- Delaunay
 - ball, 77
 - complex, 76
 - filtration, 138
 - tangential complex, 190
 - triangulation, 77, 81
- dense sample, 102
- diameter, 112
- diffeomorphism, 158
- distance
 - bottleneck, 276
 - Euclidean distance, 21
 - Gromov-Hausdorff, 282
 - Hausdorff, 26
 - Wasserstein, 242
 - weighted distance, 81
- distance function, 215
- distance-like function, 237
- distance-to-measure, 244
- distortion, 166
- dual complexes, 52
- embedding, 23
 - natural embedding, 33

- empty ball property, 76, 137
- Euclidean space, 21
- filtration, 34
- flake, 119
- Gabriel simplex, 139
- general position, 48
- general position wrt spheres, 77
- geometric realization, 32
- Hasse diagram, 39
- homeomorphism, 22
- homologous, 260
- homology group
 - simplicial, 259
 - singular, 263
- homotopy, 24
- homotopy equivalence, 24
- isomorphism
 - of simplicial complexes, 32
- isotopy
 - ambient, 23
 - lemma, 219, 238
- Kullback-Leibler
 - divergence, 89
- Legendre duality, 92
- link, 37
- lower envelope, 72
- lower hull, 51
- maximization diagram, 73, 82
- measure
 - Borel, 240
 - Dirac, 241
 - empirical, 241
 - probability, 241
 - submeasure, 245
 - support, 241
 - uniform, 241
- medial axis, 159
- medial ball, 159
- metric space, 20
 - discrete metric space, 108
 - Euclidean space, 21
- minimization diagram, 72, 82
- Moser Tardos algorithm, 198
- nerve, 33
- Nerve Theorem, 34
- net, 103, 163
 - $(\varepsilon, \bar{\eta})$ -net, 103
- offset, 25, 215
- persistence diagram, 275
- persistence module, 277
 - homomorphism, 278
 - interleaving, 278
- persistent pair, 272
- polytope, 46
- protection, 128
 - δ -protected, 128
 - δ -protection center, 128
- pseudo-inverse, 112
- q-tameness, 277
 - function, 279
- reach, 161, 228
 - α -reach, 238
 - μ -reach, 228
- reconstruction
 - distance-like functions, 239
 - homotopy, 185
 - isotopic, 229
 - submanifold, 198
- restriction, 104
- sampling radius, 102

- semiconcave, 237
- separation ratio, 102
- simplex, 30
 - adjacent, 39
 - boundary, 257
 - incident, 39
 - positive, negative, 262
 - tree, 39
- simplicial chain, 256
- simplicial complex, 31
 - abstract, 32
 - boundary complex, 37
 - dimension of, 31
 - pure, 37
- simplicial polytope, 48
- singular value, 112
- skeleton, 31
- space
 - discrete metric space, 108
 - compact, 21
 - connected, 21
 - contractible, 24
 - path connected, 21
 - underlying, 31
- star, 37
- subcomplex, 31
- submanifold, 158
- tangent space, 158
- tangential complex, 189, 190
- thickness, 112
- topological
 - ball, 22
 - invariant, 25
 - space, 20
- topological space
 - triangulation, 38, 171
- topology, 20
 - induced, 20
 - metric, 21
- transport plan, 241
- triangulation
 - of a point set, 38
 - of a topological space, 38, 171
- tubular neighborhood
 - see offset, 25
- upper bound theorem, 53
- upper envelope, 52, 73, 82
- vertex, 30
 - scheme, 32
- Vietoris-Rips complex, 36
- Voronoi
 - k -order diagram, 86
 - cell, 73
 - diagram, 73
- weak feature size, 219, 238
- weighted
 - center, 118
 - distance, 81
 - orthogonal weighted points, 117
 - point, 81
 - radius, 118
- weighted α -complex, 140
- weighted α -shape, 140
- weighted Delaunay
 - complex, 83
 - triangulation, 85
- weighted Voronoi
 - cell, 82
 - diagram, 82
- weighting scheme, 117, 199
- witness complex, 143
 - relaxed witness complex, 148