Computational Geometry and Topology for Data Analysis

Jean-Daniel Boissonnat **Frédéric Chazal**
Mariette Yvinec

January 7, 2016

Contents

Introduction

Context and challenges. During the last decade, the wide availability of measurement devices and simulation tools has led to an explosion in the amount of available data in almost all domains of science, industry, economy and even everyday life. Often these data come as point clouds embedded in Euclidean space, or in possibly more general metric spaces, usually of high dimensions. Those points are usually not uniformly distributed in the embedding space but lie close to some *low-dimensional structure*, which reflects the fact that the physical system that produced the data has a moderate number of degrees of freedom. Understanding the geometry of the underlying structure of the data is key to understanding the underlying system.

In many applications, these structures are highly nonlinear and have a non trivial topology : a precise understanding of their geometry is out of reach of current techniques. Let us illustrate our motivation and objectives through the paradigmatic example of *energy landscapes of molecules*. Understanding energy landscapes is a major challenge in chemistry and biology but, despite a lot of efforts and a wide variety of approaches, little is understood about the actual structures of these landscapes. The case of cyclo-octane C_8H_{16}, a cyclic alkane used in manufacture of plastics, is instructive. This relatively simple molecule has been studied in chemistry for over 40 years but it is only very recently that its conformational space has been fully understood. By analyzing a dataset of 1M points in $\mathbb{R}^{72}$, each describing a cyclo-octane conformation, it has been shown that the conformation space of cyclo-octane has the unexpected geometry of a multi-sheeted 2-dimensional surface composed of a sphere and a Klein bottle, intersecting in two rings [**?**]. Besides its fundamental interest, such a discovery opens new avenues for understanding the energy landscape of cyclo-octane. Extending this type of analysis to large molecules, in particular to proteins, would have tremendous implications. Many other examples with high potential remain to be solved in domains as varied as neurosciences, medical imaging, speech recognition and astrophysics.

Topological Data Analysis. The crux is to have access to robust and efficient data structures and algorithms to represent and analyze the possibly highly nonlinear underlying geometric structure of data. This is the object of study of the emerging field of Topological Data Analysis. The field finds its root in computational geometry and topology, and in several areas of mathematics like algebraic topology, non smooth analysis and geometric

measure theory. Reconstructing a surface of 3-space from a point sample is a classical problem in topological data analysis that has been widely studied in the last decades and is by now well understood. This book intends to provide mathematical and algorithmic foundations to extend the topological approach to higher-dimensional data. The first major difficulty comes from the fact that the complexity of data structures and algorithms used to approximate shapes rapidly grows as the dimensionality increases, which makes them intractable in high dimensions. This phenomenon, referred to as the *curse of dimensionality*, prevents, in particular, subdividing the ambient space, as is usually done in 3-space, since the size of any such subdivision depends exponentially on the ambient dimension. Instead, any practical method must be *sensitive to the intrinsic dimension* (usually unknown) of the shape under analysis. The second main difficulty comes from the fact that high-dimensional data often suffer from significant *defects*, including sparsity, noise, and outliers that may hide the intrinsic dimension of the underlying structure. A fundamental objective of Topological Data Analysis is to guarantee the infered geometry and topology to be close to the original ones even when the data are corrupted by various types of noise and outliers.

This book. 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 faithfull 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 or the Rips complexes will also be considered.

The book contains a first chapter that recall some background material and is then divided into three main parts.

Part I presents the basic geometric data structures that are useful for geometric inference. The first chapter (chapter 2) introduces grid data structures, namely quadtrees and their higher dimensional counterparts. This chapter reviews the known data structures to search nearest neighbors and approximate nearest neighbors in high dimensional spaces. This is a fundamental tool in geometric inference that is required, for instance, to compute the distance from a query point to a sample set or to find the neighbors of a given point (either the k-nearest neighbors or all neighbors within a certain distance). Searching nearest neighbors is a typical example of a geometric problem that suffers from the curse of dimensionality. All efficient solutions that have been proposed in 2D or 3D become intractable when the dimension of the ambient space gets large, and new approaches are mandatory.

Chapter 3 introduces potypes. Polytopes appear in disguise in Voronoi diagrams and Delaunay complexes and their combinatorial and algorithmic study play a central role in this book.

In Chapter 4, we introduce simplicial complexes.

Part II introduces Delaunay complexes and their variants. Delaunay complexes are defined from Voronoi diagrams which are natural space partitions induced by the distance function to a sample. Delaunay triangulations appear as the underlying basic data structure for manifold reconstruction. The extensions of Voronoi diagrams and Delaunay triangulations to weighted distances, affine diagrams and regular triangulations are then presented in Chapter 6 together with their relevant applications to quadratic distance functions, kth-nearest neighbor searching and Bregman divergences which are used in information theory, image processing and statistical analysis (chapter 6).

The last chapter of Part II show how to remove badly shaped simplices from a Delaunay triangulation and how to extend the notion of Delaunay triangulation to finite metric spaces.

Part III is devoted to reconstructing a submanifold $\mathbb{M}$ from a finite point sample $P \in \mathbb{M}$. The ultimate goal is to construct a triangulation of $\mathbb{M}$, i.e. a simplicial complex K that is homeomorphic to $\mathbb{M}$. This is a demanding quest and, in this part, we will restrict our attention to the case where $\mathbb{M}$ is a smooth submanifold of $\mathbb{R}^d$. We defer to Part IV other compromises between the generality of the shapes to be approximated and the precision of the approximation.

Part III consists of four chapters. In the first one, we introdude the basic concepts and results and state a theorem that provides conditions for a simplicial complex $\hat{\mathbb{M}}$ with vertex set $P \subset \mathbb{M}$ to be both a triangulation of $\mathbb{M}$ and a good geometric approximation of $\mathbb{M}$.

The second chapter introduces alpha-complexes. We show how to construct alpha-complexes from Delaunay triangulations and weighted Delaunay triangulations. We also show conditions under which the alpha-complex of $P \subset \mathbb{M}$ has the same homotopy type as $\mathbb{M}$. Alpha-complexes provide natural filtrations of Delaunay and weighted Delaunay complexes and can also be used as the input of algorithms computing persistent homology to be described in Chapter 13.

Although Delaunay triangulations have many beautiful properties, their simplices, in dimension greater than 2, may be arbitrarily flat 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 many works on the triangulation of manifolds, in numerical simulations to ensure the convergence of numerical methods solving partial differential equations, and will play a central role in Chapter 10. In the third chapter of Part III, we introduce random perturbation techniques to get rid of flat simplices.

The last chapter of Part III is devoted to the problem of reconstructing submanifolds from point samples. This problem is of primary importance when $\mathbb{M}$ is a surface of $\mathbb{R}^3$ (it is then known as the surface reconstruction problem). It also founds 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. We therefore 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 $\mathbb{M}$ is small, even if the dimension d of the ambient space may be very large. Chapter 10 takes advantage of this assumption and presents a reconstruction algorithm whose complexity is linear in d and only exponential ion k.

In many applications, the simplifying assumptions made in Part III are not satisfied, the geometric structures of the data are not manifolds and triangulating such structures from noisy point clouds is usually out of reach. Part IV addresses more general situations and presents methods to recover topological and geometric properties of more general shapes. Chapter 11 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 11 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 12 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 11 and 12 provide strong results on the topology of the sublevel sets of distance functions, computing and manipulating such sublevel sets are 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 13 introduces persistent homology and provides tools to robustly infer the homology of the sampled shape.

Target audience. This book is intended to students, researchers and engineers working in Computational Geometry, Computational Topology and Data Analysis. The material in this book originates from a course that we gave at MPRI (*Master Parisien de Recherche en Informatique, Parisian Master of Research in Computer Science).*

Acknowlegments.

Chapter 1

Background material on 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, compact sets and submanifolds

A *topology* on a set X is a family $\mathcal{O}$ of subsets of X that satisfies the three following conditions:
i) the empty set $\emptyset$ and X are elements of $\mathcal{O}$,
ii) any union of elements of $\mathcal{O}$ is an element of $\mathcal{O}$,
iii) 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*.

A map $f : X \to X'$ 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.

Recall that a *metric (or distance)* on X is a map $d : X \times X \to [0, +\infty)$ such that:
i) for any $x, y \in X$, $d(x, y) = d(y, x)$,
ii) for any $x, y \in X$, $d(x, y) = 0$ if and only if $x = y$,
iii) 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. For example, the standard topology in an Euclidean space is the one induced by the metric defined by the norm: $d(x, y) = \|x - y\|$.

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, compacity is characterized sequentially: a metric space X

is compact if and only if any sequence in X has a convergent subsequence. In the Euclidean case, 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).

In this book, most of the considered shapes will be represented as compact subsets of an Euclidean space $\mathbb{R}^d$, endowed with the metric topology. Continuous and differential maps in Euclidean spaces are defined in the standard way. Recall that a k-differentiable diffeomorphism $\phi : U \to V$ from an open set $U \subset \mathbb{R}^d$ to an open set $V \subset R^d$, is an homeomorphism (see definition 1.2) such that ϕ and ϕ^{-1} are k-differentiable on U and V respectively.

Definition 1.1 (Submanifold) *A compact subset $M \subset \mathbb{R}^d$ is a k-differentiable submanifold of dimension $m < d$, if for any $p \in M$ there exist an open set $U \subset R^d$ containing p, a k-differentiable diffeomorphism ϕ from U to an open set $V \subset R^d$, and an affine m-dimensional subspace $A \subset \mathbb{R}^d$ such that*

$$\phi(U \cap M) = A \cap V.$$

Intuitively, a submanifold of dimension m is a subset of $\mathbb{R}^d$ that looks locally like an open set of an affine space of dimension m. More generally, manifolds are defined in an intrinsic way, independently of any embedding in $\mathbb{R}^d$. Submanifolds as defined above are manifolds and we slightly abuse terminology by using indifferently the terms manifolds or submanifolds to denote a submanifold of an Euclidean space. We also call *curve* a 1-dimensional submanifold of $\mathbb{R}^d$, and *surface* a 2-dimensional submanifold of $\mathbb{R}^d$.

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

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

Definition 1.2 (Homeomorphism) *Two topological spaces X and Y are homeomorphic if there exists a continuous bijective map $h : X \to Y$ such that its inverse h^{-1} is also continuous. The map h is called an homeomorphism.*

When the space X is compact and the space Y is Hausdorff (i.e. for any pair of points $(y, y') \in Y$, $y \neq y'$ there exist two disjoint open sets $U, U' \subset Y$ such that $y \in U$ and $y' \in U'$), the continuity of h^{-1} does not need to be checked to prove that h is an homeomorphism. Indeed, any continuous bijection between X and Y is an homeomorphism.

As an example, a circle and a simple closed polygonal curve [1] are homeomorphic. On the other hand, a circle and a segment are not homeomorphic. If X is homeomorphic to the standard unit ball of $\mathbb{R}^d$, X is called a *topological ball*.

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

Definition 1.3 (Isotopy) *An isotopy between X and Y is a continuous application $F : X \times [0, 1] \to \mathbb{R}^d$ such that $F(., 0)$ is the identity map on X, $F(X, 1) = Y$ and for any $t \in [0, 1]$, F is an homeomorphism between X and $F(X, t)$. An ambient isotopy between X and Y is a map $F : \mathbb{R}^d \times [0, 1] \to \mathbb{R}^d$ 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 an 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. Note that restricting an ambient isotopy to $X \times [0, 1]$ gives an isotopy. In fact, a classical result of differential topology states that the converse is also true under some quite general conditions: if X and Y are isotopic, then they are ambient isotopic (see [81]). As a consequence, in this book we do not pay attention to the considered notion of isotopy.

The notion of isotopy is stronger than homeomorphy in the sense that if X and Y are isotopic, then they are obviously homeomorphic. In general two homeomorphic subspaces of $\mathbb{R}^d$ may not be isotopic. This is the case for a

[1] A continuous polygonal curve P with consecutive edges $e_1 = [p_1, p_2], e_2 = [p_2, p_3], \cdots e_n = [p_n, p_{n+1}]$ is simple and closed if and only if $e_i \cap e_j = \emptyset$ whenever $2 \leq |i - j|$ mod (n), $e_i \cap e_{i+1} = p_{i+1}$ for $i = 1, \cdots n - 1$ and $e_n \cap e_1 = p_1$

knotted and an unknotted torus embedded in $\mathbb{R}^3$ as the ones on figure 1.1. Note that proving these two surfaces are not isotopic is a non obvious task that require some background in algebraic topology.

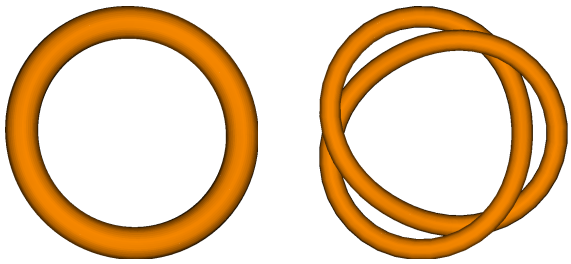

Figure 1.1: Two surfaces embedded in $\mathbb{R}^3$ homeomorphic to a torus that are not isotopic.

In general deciding whether two spaces are homeomorphic or not 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 $f_0, f_1 : X \to Y$ are *homotopic* if there exists a continuous map $H : [0,1] \times X \to Y$ 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 (see [72], pp. 171-172 or [94] p. 108 for more details).

Definition 1.4 *Two topological spaces X and Y have the same homotopy type (or are homotopy equivalent) if there exist two continuous maps $f : X \to Y$ and $g : Y \to X$ 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.*

Definition 1.5 *A contractible space is a space that has the same homotopy type as a single point.*

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

- $\forall x \in X$, $H(0,x) = x$
- $\forall x \in X$, $H(1,x) \in Y$
- $\forall y \in Y$ *and* $\forall t \in [0,1]$, $H(t,y) \in Y$

then X and Y are homotopy equivalent. If the last property of H is replaced by the stronger one $\forall y \in Y$ and $\forall t \in [0,1]$, $H(t,y) = y$, then, by definition, H defines 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*. They 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: a point and a segment are homotopy equivalent but they are not homeomorphic. 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 and Fréchet distances

Given a compact set X of $\mathbb{R}^d$, define X^ε to be the *tubular neighborhood* or *offset* of X of radius ε, i.e. the set of all points at distance at most ε from X

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

Definition 1.7 *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,*

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

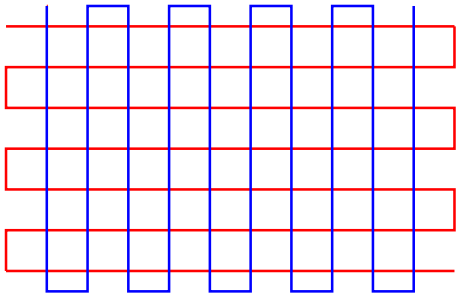

Figure 1.2: Two curves whose Haudorff distance is small compared to their Fréchet distance.

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

In some cases, the Hausdorff distance is not always a good measure of the similarity of two shapes. This is illustrated in Figure **??** where two curves are close for the Hausdorff distance but look quite different.

A more satisfactory measure of the similarity of two shapes is the so-called Fréchet distance.

Definition 1.8 *The* Fréchet distance *between two subsets X and Y of $\mathbb{R}^d$ is*

$$d_F(X,Y) = \inf_h \sup_{p\in X} d(p,h(p)),$$

where h ranges over all homeomorphisms from X to Y.

Observe that the Fréchet distance is finer than the Hausdorff distance to compare shapes ($d_F(X,Y) \geq d_H(X,Y)$) but it is more restrictive since it only allows to compare shapes that are homeomorphic.

Note that if $F : X \times [0,1] \to \mathbb{R}^d$ is an isotopy between X and Y, then $\sup_{x\in X} \|F(x,0) - F(x,1)\|$ is an upper bound on the Fréchet distance between X and Y.

1.3 Exercises

Exercise 1.1 *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.*

Exercise 1.2 *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 $\gamma : [0,1] \to C$ 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 1.3 *Give examples of simplicial complexes in $\mathbb{R}^3$ that are homeomorphic to a ball, a sphere, and a torus.*

Exercise 1.4 *Prove that any abstract simplicial complex of dimension d can be realized as an isomorphic geometric simplicial complex in $\mathbb{R}^{2d+1}$.*

1.4 Bibliographical notes

Our presentation of simplicial complexes follows the one in Munkres [93]. The nerve theorem and its variants are classical results in algebraic topology. A proof is given in Hatcher [80], Section 4G.

Part I

Geometric data structures

Chapter 2

Grids and approximate nearest neighbor search

2.1 Nearest neighbor queries

Nearest neighbor queries are ubiquitous in geometric data analysis.

Definition 2.1 (Nearest neighbor) *Given a set of points P in space $\mathbb{R}^d$ and a query point q, the nearest neighbor of q in P is the point of P that achieves the smallest distance to q.*

The nearest neighbor of q in P is denoted by $\mathrm{nn(q, P)}$ or simply by $\mathrm{nn(q)}$ when there is no ambiguity about the set P. Thus, if $d(x, y)$ denotes the distance between two points x and y,

$$d\left(q, \mathrm{nn(q, P)}\right) = \min_{p \in P} d(q, p) = d(q, P).$$

Nearest neighbor queries of course highly depend on the definition of the distance function $d(x, y)$ between two points. We assume here that the distance is the Euclidean distance, i.e.: $d(x, y) = \|x - y\|$.

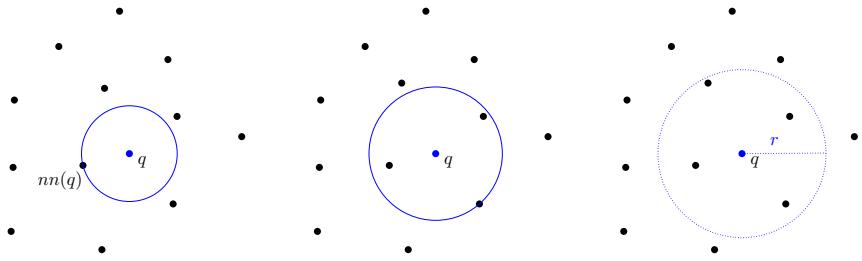

Figure 2.1: Nearest neighbor and related queries. Left: the nearest neighbor of q, middle: the k-nearest neighbors of q, right: the set of points at distance at most r from q.

Related queries are the k-nearest neighbors queries, asking for the subset $P_k(q)$ of k points in P that are nearest to q, i.e.

$$d(q, p) \leq d(q, p'),\ \forall (p, p') \text{ with } p \in P_k(q) \text{ and } p' \in P \setminus P_k(q).$$

A neighborhood range query is another variant asking for the subset $P(q.r) \subset P$ of points in P at distance at most r from the query q. See Figure 2.1. In many applications, several nearest neighbor queries are issued for the same set of points P, and it might be worth to prepocess the set P into a data structure yielding fast answers to nearest neighbor queries. The goal is then

Storage space	Query time	
$O\left(n^{2^{d+1}}\right)$	$O\left(2^d \log n\right)$	Dobkin, Lipton [55]
$O\left(n^{\lceil \frac{d}{2} \rceil (1+\delta)}\right)$	$O\left(2^{O(d \log d} \log n\right)$	Clarkson [50]
$O\left(n^{d+\delta}\right)$	$O\left(d^5 \log n\right)$	Meiser 93 [90]
$O\left(m\right)$	$O\left(n/m^{1/\lceil d/2 \rceil}\right)$	Agarwal, Erickson 99 [1]

Table 2.1: Storage requirements and query time for a few nearest neighbor algorithms

to optimize the trade-off between the memory footprint of the data structure and the query time. One should keep in mind, that the naive solution, which consists in computing the distances $d(p, q)$ from q to all the points in P, yields a linear query time $O(dn)$ and a data structure with linear size $O(dn)$, where n is the cardinal of P.

2.2 The curse of dimensionality and approximate nearest neighbor queries

Common data structures to handle repetitive nearest neighbor queries on a given point set P are grids or derived from grids. The main idea is to store the points of P in the grid cells. A nearest neighbor query is then answered by locating the query point in the grid and searching for the nearest neighbor in the cell containing q and in surrounding cells. Obviously the method is not efficient unless the point set is known to have a uniform spatial distribution or if the data structure use spatial partitions that adapt to the local density of the point set. Nevertheless, unless point sets are known to have particularly smooth spatial distributions, all grid based data structures are bound to have a size that is exponential with respect to the dimension of the ambiant space. More generally, the nearest neighbor problem is a typical case of problem suffering from the *curse of dimensionality*. It seems that any data structure yielding a sublinear query time requires a storage space whose size grows exponentially with the space dimension. See for example table 2.1 where the storage space and query time of a few typical nearest neighbor algorithms are shown.

A popular way to escape the curse of dimensionality consists in dropping the requirement to get the exact nearest neighbor of a query point. Many

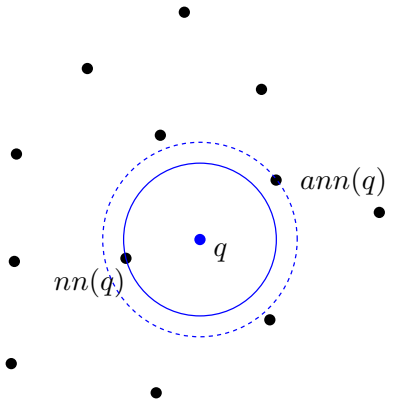

Figure 2.2: Approximate nearest neighbor.

Storage space	Query time	
$O(n)$	$O\left((\frac{c^d}{\epsilon^d})\log n\right)$	ANN [5]
$O(n)$	$O\left(c^d(\frac{1}{\epsilon^d}+\log n)\right)$	Compressed quadtrees [78]
$O\left(\frac{1}{\epsilon^d} n \log^2 n\right)$	$O(\log(\frac{n}{\epsilon}))$	Point location in balls [77, 82]

Table 2.2: Storage requirement and query time for a few approximate nearest neighbor algorithms.

applications can content themselves with an approximate answer to each nearest neighbor query providing a point that is not much farther from the query point than the exact nearest neighbor. Hence the idea to define approximate nearest neighbor queries.

Definition 2.2 (Approximate nearest neighbor) *Given a set of points P in e $\mathbb{R}^d$ and an approximation factor $\epsilon \in \mathbb{R}^+$, an ϵ-approximate nearest neighbor in P of a query point q is a point $ann(q)$ of P whose distance to q is at most $(1+\epsilon)d(q,P)$ where $d(q,P)$ is the distance from q to P, i.e. the distance $d(q,\mathrm{nn(q)})$ from q to its exact nearest neighbor in P. (See Figure 2.2).*

$$d(q, ann(q)) \leq (1+\epsilon)d(q,P)$$

Table 2.2 shows the storage space requirement and query time for a few approximate nearest neighbor algorithm.

2.3 Compressed quadtrees

A popular software for approximate nearest neighbors queries in high dimension, called ANN, is based on a data structure which is a variant of quadtrees. jd : quadtrees have not been defined The ANN data structure is a tree in which each node represent either a canonical box (i.e. a cell of a quadtree), or the set difference between two such canonical boxes. We have chosen to present here the related concept of *compressed quadtree*. The compressed quadtree is a general concept providing an elegant way to understand how ANN works.

We assume hereafter that the point set P has been rescaled to fit within the unit box $[0,1]^d$. We first describe what is the (standard) quadtree for point set P. Note that we use the term *quadtree* for point sets in any dimension although, obviously, this term was originally designed to apply in dimension 2. We call here *canonical box* any sub-box of $[0,1]^d$ obtained from $[0,1]^d$ by a sequence of *centered splits* that split boxes, through their center point, into 2^d equally sized sub-boxes. A quadtree is a tree in which each node corresponds to a canonical box and is either a leaf or linked to 2^d children nodes which correspond to the sub-boxes resulting from its centered split. The quadtree $\mathrm{QT}(P)$ for the point set P is obtained by recursively splitting canonical boxes that contain more than one point of P. The leaves of a quadtree form a space partition whose cells are the canonical boxes associated with leaves. The cells of the partition have a good aspect ratio (since they are hypercubes !) and the sizes of the cells adapts locally to the point set density. Cells are said to be adjacent when they share a piece of $d-1$-face. For the purpose of nearest or approximate nearest neighbor search, quadtrees have the major drawbacks that adjacent cells may have highly different sizes and therefore each cell has an unbounded number of adjacent cells. A node of the quadtree is said to be at *level* l when the edge length of the associated canonical box is 2^{-l}. Further splitting are sometimes performed on the quadtree $\mathrm{QT}(P)$ to get a regularized quadtree $\mathrm{RQT}(P)$ in which adjacent cells have levels that differ by at most one, or equivalently have edge lengths that differ by at most a factor 2. (See Figure 2.3). In a regularized quadtree, each cell has at most $d2^d$ adjacent cells. However the drawback is now that the tree has an unbounded number of empty nodes. In both cases, the tree may be quite unbalanced and its depth may be $\Omega(n)$. Compressed quadtrees are designed to remedy those drawbacks.

To describe the compressed quadtree $\mathrm{CQT}(P)$ of the point set P, we start

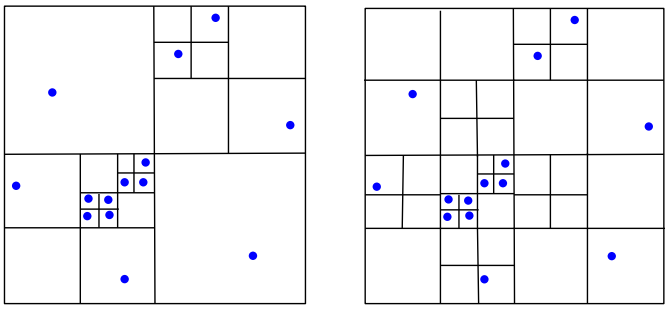

Figure 2.3: The quadtree of a set of points and the corresponding regularized quadtree.

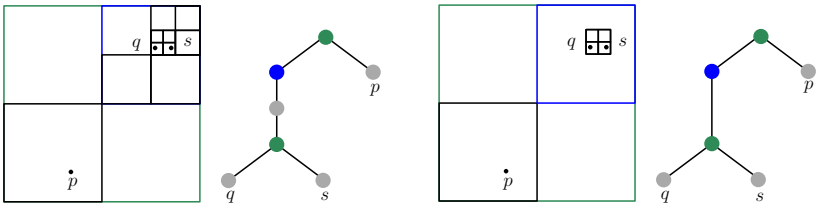

Figure 2.4: The compression of a sequence of nodes with a single child each.

from the quadtre $\mathrm{QT}(P)$ and first discard all the nodes that correspond to empty boxes, i.e. boxes containing no point of P. This yields a tree in which some nodes have a single child. There may even be sequences of nodes with a single child. In the compressed quadtree, each sequence of nodes with a single child is compressed into a single edge. Let us call *first* the node in the sequence with lowest level, i.e. corresponding to the biggest box, and *last* the node with the highest level, the compression involves the following : all nodes in the sequence except the first one are removed, all edges linking a node in the sequence to its single child are removed, the whole sequence is replaced by a single edge linking the first node to the child of the last one. See Figure 2.4.

As a consequence, the nodes of a compressed quadtree are either split nodes with more than one child, compressed nodes with a single child, or leaves with no child. See Figure 2.5. The space partition induced by the compressed quadtree is now composed by regions that are either associated to leaves compressed nodes or split nodes. The region associated to a leaf is just its canonical box while the region associated to a compressed or split

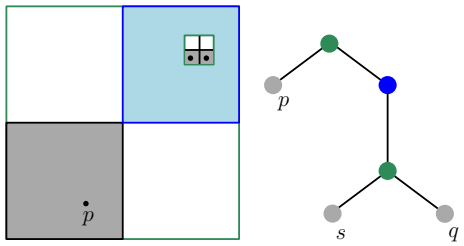

Figure 2.5: Compressed and uncompressed nodes: in blue, compressed nodes and associated regions, in grey leaves and associated regions, in green, split nodes

node is the set difference between the canonical box of the node and the canonical boxes its children.. Note that the regions of split or compressed nodes are empty of points of P. If the point set P has n points, the compressed quadtree $\mathrm{CQT}(P)$ has at most n leaves and $n-1$ split nodes. The single child of a compressed node is a split node. Charging each compressed node to its single child yields that the number of compressed nodes is at most $n-1$. The compressed quadtree $\mathrm{CQT}(P)$ has therefore at most $3n-2$ leaves.

The compressed quadtree is still not a balanced data structure and may have a $O(n)$ depth. To solve the balancing problem, we resort to a total order, called the Q-order, that can be defined on canonical boxes, hence on quadtree nodes, independentely of the actual tree structure. In a quadtree, all nodes are either leaves or have the same pattern of 2^d children and we may define a standard order for visiting the children of a node. Assuming such a standard ordering among the children of a quadtree node, the order in which the nodes of a quadtree are visited during a depth first traversal of the quadtree is defined and we define the Q-order of canonical boxes as follows. Let b and b' be two canonical boxes. We say that b preceeds b' according to the Q-order, and we write $b \prec b'$, iff b is visited before b' in the depth first traversal of any quadtree containing both b and b'. The smallest canonical box containing both b and b' correspond to a node $\mathrm{lca}(b, b')$ which is the least common ancestor of b and b' in any quadtree containing b and b'. We can say more precisely that $b \prec b'$ in the Q-order if either b is included in b' or, the child of $\mathrm{lca}(b, b')$ including b preceeds the child of $\mathrm{lca}(b, b')$ including b' in the standard ordering of the children of $\mathrm{lca}(b, b')$. Observe that the ordering of the visit of b and b' does not depend on the actual quadtree

containing b and b'. The Q-order may be naturally extended to compare the order of a point and a box or the order of two points. Given two canonical boxes b and b', their least common ancestor $\mathrm{lca}(b, b')$ and their Q-order can be computed in constant time if we assume that the operator $\mathrm{bit}_\Delta(\alpha, \beta)$ on two real numbers α and β in $[0, 1]$ is available and takes a constant time. The operator $\mathrm{bit}_\Delta(\alpha, \beta)$ provides the smallest index l where the coefficients of the power two expansions $\alpha = \sum \alpha_l 2^{-l}$ and $\beta = \sum_l \beta_l 2^{-l}$ are different.

The main idea consists in storing in a dictionary the nodes of the compressed quadtree $\mathrm{CQT}(P)$, ordered along the Q-order of their canonical boxes. The dictionary is implemented using any balanced data structure for ordered sets, like a red-black tree or a skip list. In the following, we call compressed quadtree dictionary, $\mathrm{CQD}(P)$, any balanced dictionary storing the nodes of the compressed quadtree $\mathrm{CQT}(P)$ according to their Q-order.

Theorem 2.3 *The compressed quadtree dictionary $CQD(P)$ for a set P of n points in $\mathbb{R}^d$ requires $O(dn)$ memory space. Assuming that Q-order comparisons are performed in constant time, $CQD(P)$ can be constructed in $O(dn \log n)$ time and dynamically maintained under points insertion and deletion in $O(d \log n)$ per operation.*

Proof The memory bound results from the number of nodes of a compressed quadtree. The main issue for the construction and dynamical maintenance of $\mathrm{CQD}(P)$ is the location of a query point q in the space partition induced by the compressed quadtree. Given a query point q, we want to find the node $u(q)$ of $\mathrm{CQT}(P)$ whose region contains q, whereas a search for q in dictionary $\mathrm{CQD}(P)$ provides the last node that preceeds q in Q-order. Algorithm 1, solves the problem in $O(\log n)$ time issuing at most two queries in the dictionary $\mathrm{CQD}(P)$. □

We proove below that the algorithm is correct. Different cases are shown in Figure 2.6 to give some intuition. Let v be the last node of $\mathrm{CQD}(P)$ preceeding q in Q-order. We claim that, if point q belongs to the canonical box of v, it belongs to the region of v (case a, b and c of Figure 2.6). The claim is obvious if v is a leaf (case a). Otherwise (case b and c in Figure 2.6), the claim follows from the fact that q belongs to canonical box of v and to none of the canonical boxes the children of v, because q preceeds the boxes of those nodes in Q-order.

If q does not belong to the canonical box of v, (case d, e of Figure 2.6) the

Algorithm 1 Location of point q in a compressed quadtree CQT(P) : find the node $u(q)$ of CQT(P) whose region contains Q

Require: The dictionary CQD(P) storing the nodes of CQT(P) along their Q-order
Ensure: Returns the node of CQT(P) whose region contains q
  Get v the last node of CQD(P) that preceedes q in Q-order
  if q is in the canonical box of v **then**
    return v.
  Compute $w = \text{lca}(q, v)$
  if w belongs to CQD(P) **then**
    return w
  return the last node of CQD(P) before w in Q order

algorithm computes the node w that is the least common ancestor of q and v in a standard quadtree. We claim that if node w belongs to CQT(P), it is the correct answer to the location query q. Indeed q belongs to the canonical box of w and we prove that it does not belong to the box of any children of w. Assume for contradiction that q belongs to the box of some child w' of w, then, in Q-order, w' would be before q and after v, which contradicts the definition of v.

A similar argument shows that in the case where $w = lca(v, q)$ does not belongs to CQT(P), q belongs to the region of the highest level node of CQT(P) whose box contains v and q, which is just the the last node of CQT(P) preceeding w in Q order.

2.4 Approximate nearest neighbor algorithm

Assume that the set of points P has been rescaled to fit in the unit box $[0, 1]^d$. To answer nearest neighbor queries in P, we maintain the compressed quadtree dictionary CQD(P) storing the node of CQT(P). Furthermore, the data stucture is slightly augmented by adding to each node u of CQT(P) a pointer to a point $p(u)$ of P lying in the canonical box of node u and called below the *representative* of node u. As before, a node u of CQT(P) is said to be at level l if its associated canonical box has size 2^{-l}.

To find the nearest neighbor of a query point q, the nodes in CQT(P) are visited in order of increasing level while keeping track of the point p_c,

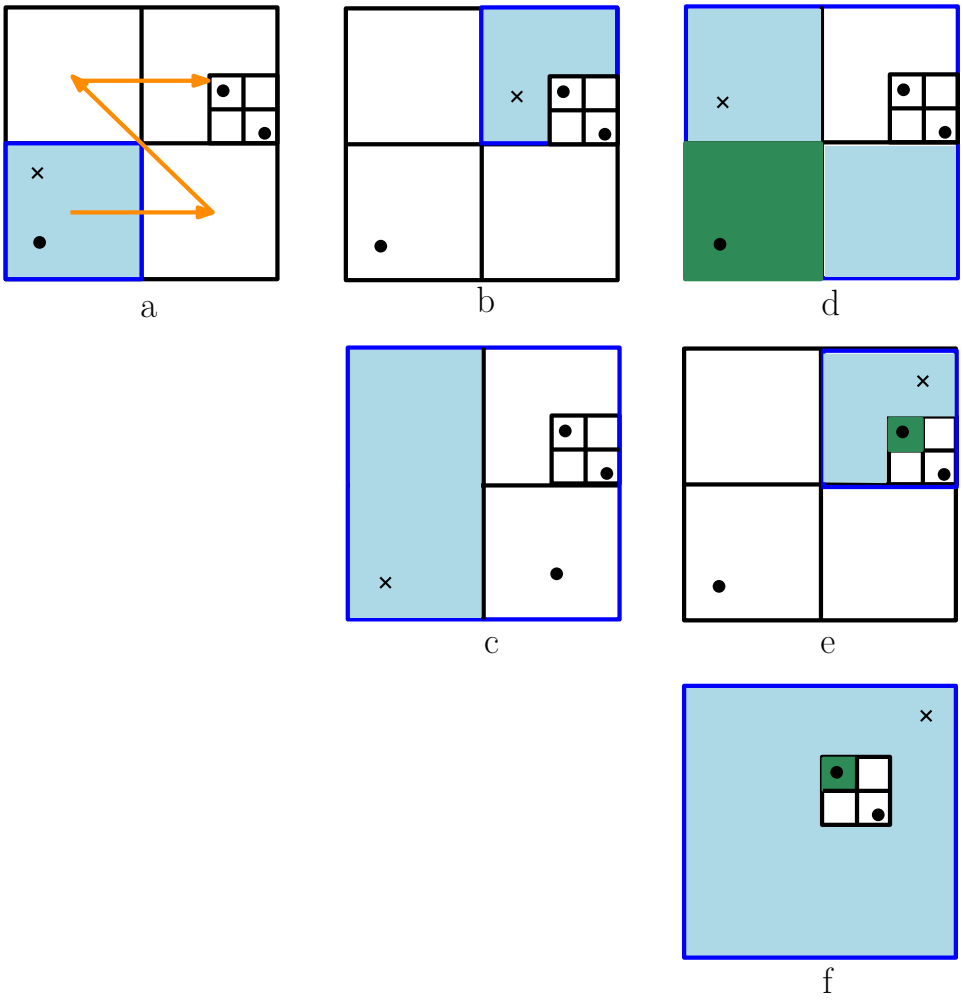

Figure 2.6: Location of point q in the partition induced by the compressed octree CQT(P). Different cases are shown in subfigure a,b,c,d,e,f. Point q is represented by a cross, points in P are shown as dots. In each case, the canonical box of the node u whose region contains q is stroked in blue, and the region of u is filled in light blue. The box of v, the last node preceeding q in Q order, is colored green when v is different from u . Orange arrows shows the standard order to visit the children of an octree node.

that is the closest point to q among the representatives of the so-far visited nodes. The efficiency of the method comes from the fact, that the algorithm does not visit all the nodes of CQT(P) but only those that are likely to improve significantly the current guess p_c. More precisely, assume that node w with level $l(w)$ and representative $p(w)$ is visited at a time where p_c is the current nearest neighbor approximation. The algorithm skips the subtree of CQT(P) rooted at w unless

$$d(q, p(w)) - \text{diam}(w) \leq \left(1 - \frac{\epsilon}{2}\right) d(q, p_c) \tag{2.1}$$

where $\text{diam}(w) = \sqrt{d}2^{-l}$ is the diameter of the canonical box associated to w.

Algorithm 2, given below, starts at the root u_0 of CQT(P) and computes, for each level i, a list A_i of nodes at level i whose children have to be visited. When the nodes in all the lists are have been visited, the algorithm ends and returns the current approximation.

Algorithm 2 Approximate nearest neighbor

Require: The dictionary CQD(P) and a query point q.
Ensure: The returned point is an ϵ-approximate nearest neighbor of q in P.

```
A_0 = {u_0}, A_i = ∅ for i > 0, p_c = p(u_0)
for i = 0, ..., i_m do {i_m is the max level of to be visited nodes}
  for each node u ∈ A_i do
    for each child w of u do
      if d(q, p(w)) ≤ d(q, p_c) then
        p_c = p(w)
      if d(q, p(w)) − diam(w) ≤ (1 − ε/2) d(q, p_c) then
        add w in A_{l(w)}.
return p_c
```

Theorem 2.4 *Algorithm 2 returns an ϵ-approximate nearest neighbor of q in P*

Proof Let v be the last visited node whose canonical box contains the exact nearest neighbor nn(q). We have

$$d(q, \text{nn(q)}) \geq \text{d(q, p(v))} - \text{d(p(v), nn(q))} \geq \text{d(q, p(v))} - \text{diam(v)} \geq (1 - \frac{\epsilon}{2})\text{d(q, p}_\text{c}),$$

where p_c is the current nearest neighbor approximation and the last equation arises from the fact that v does not fulfill Condition 2.1. Therefore,

$$d(q, p_c) \leq \frac{1}{(1-\frac{\epsilon}{2})} d(q, \mathrm{nn(q)}) \leq (1+\epsilon)\mathrm{d(q,P)} \text{ for } \epsilon < 1.$$

□

Lemma 2.5 *The maximum level i_m of compressed quadtree nodes inserted in the lists of Algorithm 2 is bounded as follows :*

$$i_m \leq \min\left(\log \phi(P), \log \frac{2\sqrt{d}}{\epsilon d(q,P)}\right),$$

where $\phi(P)$ is the spread *of P, that is the ratio from the longest to the shortest distance between two points of P.*

Proof The first component of the bound is just trivial since there is no node in CQT(P) with level higher than $\log \phi(P)$. The second component of the bound comes from Condition 2.1. Let $l(w)$ be the level of a listed node w and $diam(w) = 2^{-l(w)}$ the diameter of its box.

$$d(q, p(w)) - \mathrm{diam}(w) \leq \left(1-\frac{\epsilon}{2}\right))d(q, p_c) \leq \left(1-\frac{\epsilon}{2}\right) d(q, p(w))$$

Hence,

$$\sqrt{d}2^{-l(w)} = \mathrm{diam}(w) \geq \frac{\epsilon}{2} d(q, p(w)) \geq \frac{\epsilon}{2} d(q, P).$$

□

Lemma 2.6 *The total number of nodes of the compressed octree visited by Algorithm 2 is $O\left[c^d\left(\log \Phi(P) + \frac{1}{\epsilon^d}\right)\right]$, where c is a constant.*

Proof Algorithm 2 visits nodes that are children of nodes registered in the list A_i, therefore the total number of nodes visited by Algorithm 2 is at most $2^d \sum_i |A_i|$ where $|A_i|$ is the maximal size reached by list A_i. To bound the sum $\sum_i |A_i|$ we bound the number n_j of nodes inserted in the lists as the children of nodes in A_j. At the time nodes in A_j are considered by the

algorithm, the current approximation satisfies $d(q, p_c) \leq d(q, P) + \sqrt{d}2^{-j}$ and the children enlisted satisfy Condition 2.1 meaning that they are at distance at most r_j from q where

$$r_j \leq \left(1 - \frac{\epsilon}{2}\right) d(q, p_c) \leq d(q, P) + \sqrt{d}2^{-j}.$$

Each node enlisted as a child of a node A_j is included in a canonical box of size $2^{-(j+1}$ at distance at most r_j from q, and since at most one child is enlisted for each such box, we get

$$n_j \leq \left(\left\lceil \frac{2r_j}{2^{-j-1}} \right\rceil \right)^d \leq \left(1 + 4\sqrt{d} + 4d(q, P)2^j\right)^d.$$

Then, using that for any a and b in $\mathbb{R}$, $(a + b)^d \leq 2^d(a^d + b^d)$,

$$n_j = O\left[c^d \left(1 + 2^{jd} d(q, P)^d\right)\right]$$

where c is a constant. It follows that the total number of visited nodes N is

$$N \leq 2^d \sum_i |A_i| \leq 2^d \sum_j n_j = O\left[c^d \sum_j \left(1 + 2^{jd} d(q, P)^d\right)\right]$$

and using Lemma 2.5

$$N = O\left[c^d \left(i_m + d(q, P)^d 2^{i_m d}\right)\right] = O\left[c^d \left(\log \Phi(P) + \left(\frac{2\sqrt{d}}{\epsilon}\right)^d\right)\right].$$

□

Lemma 2.6 ensures that the cost of a nearest neighbor query is constant whenever one is entitled to consider that the point set P has a bounded spread. Note however that the constant depends exponentially on the dimension. Using the above algorithm as a building block it is possible to get a more sophisticated algorithm which computes approximate nearest neighbors in time in $O(c^d \log n)$ for any point set. This algorithms uses an additionnal data structure, called a ring tree, able to find in $O(\log n)$ time a very rough $(1 + 4n)$-nearest neighbor approximation of the query point q. (See Exercises 2.1 to 2.4). The algorithm runs as follows:

1. First, using the ring tree, find a point $p_0 \in P$ such that

$$d(q, P) \leq d(q, p_0) \leq (1 + 4n)d(q.P)$$

2. Second, let $r_0 = d(q, p_0)$ and $l_0 = \lceil \log r_0 \rceil$. Let B be the set of canonical boxes of size 2^{l_0} that are within distance r_0 from q. For each canonical box b in B, we find the highest level node v of the compressed octree CQD(P) such that the subset of points $P(v)$ stored in the subtree rooted at v correspond to the subset $b \cap P$. Let V be the set of nodes of CQD(P) thus found. Observe that node v is at level at least l_0 and that there is a constant number of canonical boxes in B and therefore a constant number of nodes in V.

3. For each node in V, the algorithm performs an approximate nearest neighbor query for q in $P(v)$ using Algorithm 2. In the end, the approximate nearest neighbor found closest to q is returned. In each node v of V, the subset of points $P(v)$ has diameter at most $\sqrt{d}2^{l_0} = \sqrt{d}r_0$ and is such that $d(q, P(v)) \geq d(q, P) \geq \frac{r_0}{(1+4n)}$. Therefore, owing to Lemma 2.6, the approximate nearest neighbor query time in $P(v)$ is: $O\left[c^d\left(\log \frac{\text{diam}(P(v))}{d(q,P(v))} + \frac{1}{\epsilon^d}\right)\right] = O\left[c^d\left(\log n + \frac{1}{\epsilon^d}\right)\right]$.

2.5 Probabilistic algorithms

As it is clear from Table 2.1 and 2.2, data structures for approximated nearest neighbor queries significantly improve the storage space-query time tradeoff compared to data structures for exact nearest neighbor queries. Indeed those data structures basically achieve $O(\log n)$ query time for $O(n)$ storage space, and the exponential behaviour with respect to the ambiant space dimension is now restricted to the constants in the big O notations. Obviously something more than only resorting to approximation is needed to completely get rid of the exponential behaviour with respect to the dimension. This is where probabilistic algorithms come into play. Probabilistic algorithms promise a valid output, in the present case a valid approximated nearest neighbor, but only with a certain probability. Table 2.3 shows the performances of a few probabilistic data structures for approximated nearest neighbor queries.

A popular method to obtain probabilistic data structures for approximate nearest neighbor queries is to rely on the dimension reduction methods and especially on the Johnson-Lindenstrauss theorem presented below.

Let f be a mapping $f : \mathbb{R}^d \longrightarrow \mathbb{R}^k$. The mapping f is said to be a λ-

Storage space	Query time	
$O\left(n^{O(\frac{1}{\epsilon^2}\log(1/\epsilon))}\right)$	$\text{poly}(d, \log n, \epsilon^{-1})$	Har Peled, Indyk, Motwani [79]
$n^{O(\epsilon^{-2})}$	$O\left(d\log d + \frac{\log^2 n}{\epsilon^3}\right)$	Ailon, Chazelle [2]

Table 2.3: Storage requirements and query time for a few probabilisic approximate nearest neighbor data structures.

embedding for the subset $X \subset \mathbb{R}^d$ iff:

$$\forall p, q \in X, \quad \frac{1}{\lambda}\|p-q\| \leq \|f(p)-f(q)\| \leq \lambda\|p-q\|.$$

Theorem 2.7 (Johnson-Lindenstrauss) *For any subset P of n points in $\mathbb{R}^d$, any $\epsilon \in \mathbb{R}$ such that $0 < \epsilon < 1$, and any k such that*

$$k \geq 4\frac{\log n}{\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3}}, \tag{2.2}$$

there is a $(1+\epsilon)$-embedding of P in $\mathbb{R}^k$.

Proof The proof is only sketched here. Let us consider the map $f = \sqrt{\frac{d}{k}}\omega$ where ω is the projection on a subspace $\mathbb{R}^k$ of $\mathbb{R}^d$. For a given pair (p,q) of points in P, and a random subspace $\mathbb{R}^k$ of $\mathbb{R}^d$, the probabiltiy that the map f embeds the pair (p,q) with a distortion $\max\{\frac{\|f(p)-f(q)\|}{\|p-q\|}, \frac{\|p-q\|}{\|f(p)-f(q)\|}\}$ larger than $1+\epsilon$ may be upper bounded by $\frac{2}{n^2}$ provided that k satisfies the bound in Equation 2.2. Therefore, for such a value of k satisfying 2.2, the map f has a non zero probabilty to be an $(1+\epsilon)-$embedding of P . $\square$

To understand the somewhat strange complexities of probabilistic algorithms for approximate nearest neighbor queries (see Table 2.3), note that $c^{d'}$ becomes $n^{O(\epsilon^{-2})}$ if we set $d' = O\left(\frac{\log n}{\epsilon^2}\right)$.

2.6 Exercises

Exercise 2.1 (The ring tree) *Let P be a set of n points in $\mathbb{R}^d$. A t-ring tree for P is a binary tree splitting the point of P as follows. At each node v*

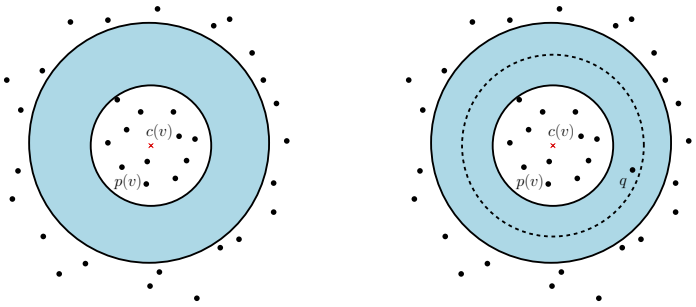

Figure 2.7: Left : A ring tree. Right: Searching an approxinate nearest neighbor using a ring tree.

of T, there is a ball $b(v) = b(c(v), r(v))$ with radius $r(v)$ and center $c(v)$, such that, if $P(v)$ is the set of points attached to node v and $P_{in}(v) = P(v) \cap b(v)$, the complementary set $P_{out}(v) = P(v) \setminus P_{in}(v)$ is included in the complementary $b^c(c(v), (1+t)r(v))$ of the ball with radius $(1+t)r(v)$ centered at $c(v)$. Subsets $P_{in}(v)$ and $P_{out}(v)$ are the subsets attached respectively to the children of node v.

Ann searching in a ring tree *Assume we are given a t-ring for P where each node v is additionally equipped with a representive point $p(v) \in P_{in}(v)$. The algorithm follows a path in the tree starting from the root to a leaf while maintaining a current guess p_c and its distance $d(q, p_c)$ to the query point. At each node v, the algorithm proceeds as follows:*

if $d(q, p(v)) \leq r_c$, ***then***
  update r_c, p_c.
if $d(q, c(v)) \leq \left(1 + \frac{t}{2}\right) r(v)$, ***then***
  search in $P_{in}(v)$
else
  search in $P_{out}(v)$

Show that the above algorithm yields a $(1 + \frac{4}{t})$-approximate nearest neighbor of the query point q in time proportionnal to the depth of T.

Exercise 2.2 *Given a set P of n points in $\mathbb{R}^d$, let $r_k(P)$ be the radius of the smallest ball containing k points of P. Prove that one can find in time $O(n \left(\frac{n}{k}\right)^d)$ a ball including k points of P with radius at most $2r_k(P)$ in time.*

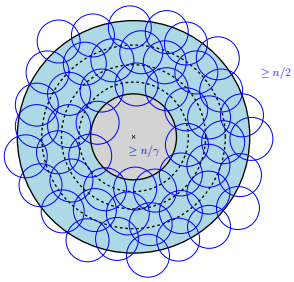

Figure 2.8: Finding a rough splitting ring

Hint. *For $i = 1, \ldots, d$, compute a set H_i of hyperplans normal to the i-axis and such that there are at most $k/2d$ points between two consecutive hyperplans. For each vertex v of the arrangement $\mathcal{A} = \mathcal{A}(\{H_i, i = 1, \ldots d\}$, consider the smallest ball centered on v and including k-points of P.*

Exercise 2.3 (A rough splitting ring) *Let P be a set of n points in $\mathbb{R}^d$ and t be a parameter. Show that, for a constant γ large enough, one can find in $O(n)$ time a ball $b(c, r)$ such that:*
$|b(c, r) \cap P| \geq \frac{n}{\gamma}$
$|b(c, r(1 + \frac{1}{t}) \setminus b(c, r) \cap P| \leq \frac{n}{2t}$
$|b^c(c, r(1 + \frac{1}{t}) \cap P| \geq \frac{n}{2}$.

Hint. *Using the result of exercise 2.2, find a ball $b(c, r_0)$ that contains n/γ points of P with radius $r_0 \leq 2r_{n/\gamma}(P)$.*
Observe that the ball $b(c, 8r_0)$ is covered by $m = O(16^d)$ balls of radius $r_0/2 \leq r_{n/\gamma}(P)$.
Choose γ large enough to have $m/\gamma \leq 1/2$ and observe that:
$|b(c, 8r_0) \cap P| \leq \frac{mn}{\gamma} \leq \frac{n}{2}$, $|P \cap b^c(c, 8r_0)| \geq \frac{n}{2}$
$|P \cap b(c, r_0)| \geq \frac{n}{\gamma}$.

Then, consider the radii:
$r_i = r_0(1 + 1/t)^i, i = 0, \ldots t$.
We have $r_t = r_0(1 + \frac{1}{t})^t \leq er_0 \leq 4r_0$ and t disjoint rings $R_i = b(p, r_i) \setminus b(p, r_{(i-1)})$ containing altogether at most $n/2$ points. One of those ring contains less than $\frac{n}{2t}$ points.

Exercise 2.4 *Given a set P of n points in $\mathbb{R}^d$, show that one can compute a $(1/n)$- ring tree of P with depth $O(\log n)$ in $O(n \log n)$ time.*

Hint. *Use recursively the result of exercise 2.3 setting* $t = n$.

2.7 Bibliographic notes

The ANN library for approximate nearest neighbor query in high dimension is based on a paper by Arya and Mount [5]. Compressed octrees are described in the book of Har Peled [78]. The reduction from nearest neighbor query to point location in balls was proposed by several authors [82, 77] and is also described in the book of Har Peled. The solution of Exercises 2.1 - 2.4 can also be found in [78]. The original proof of Johnson-Lindenstrauss theorem appeared in [83]. The elementary proof sketched in section 2.5 is given in a paper by Dasgupta and Gupta [52]. The probabilistic approximate nearest neighbor algorithms appearing in Table 2.3 are due to HarPeled, Indyk, Motwany [79] and to Ailon and Chazelle [2].

Chapter 3

Polytopes

3.1 Definitions

3.1.1 Convexity, convex hulls

The linear combination $\sum_{i=1}^{n} \lambda_i p_i$, where, for all index i, p_i is a point in $\mathbb{R}^d$ and λ_i a scalar in $\mathbb{R}$, is a *convex combination* iff

$$\lambda_i \geq 0, \quad \forall i = 1 \dots n, \quad \text{and} \quad \sum_{i}^{n} \lambda_i = 1.$$

A subset of $\mathbb{R}^d$ is convex iff it is stable by convex combinations. In particular, a convex set contains the segment joining any two points in the set. The intersection of convex set is convex and the whole space $\mathbb{R}^d$ is convex.

Let P be a set of points in $\mathbb{R}^d$. The *convex hull* $\text{conv}(P)$ of P is the set of all convex combinations of points in P. The convex hull of P is the smallest convex subset of $\mathbb{R}^d$ that contains P, where smallest refers to the inclusion relation. It can also be defined as the intersection of all convex subsets of $\mathbb{R}^d$ containing P.

define affine hull, affinely independent

3.1.2 Polytopes

The convex hull of a finite set of points in $\mathbb{R}^d$ is a *polytope*. Hence, a polytope is a closed bounded subset of $\mathbb{R}^d$. In particular, simplices (see 4) are polytopes.

The *dimension* of a polytope is the dimension of the affine subspace spanned by the polytope.

The convex hull of a set affinely independant point is a special polytope called *a simplex*. The convex hull of $k+1$ affinely independant points has dimension k and is also called a k-simplex. 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.

3.1.3 Polytope faces

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 the space $\mathbb{R}^d$ 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.

Let $P = \{p_1, \ldots, p_n\}$ be a finite set of points in $\mathbb{R}^d$ and $\mathcal{P} = \text{conv}(P)$ the convex hull of P. A hyperplane h of $\mathbb{R}^d$ is *a supporting hyperplane* of $\mathcal{P}$ 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 3.1 *A face of a polytope is a polytope.*

Proof Let $f = \mathcal{P} \cap h$ be a face of the polytope $\mathcal{P} = \text{conv}(P)$. We prove that f 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 as the intersection of two convex sets. Furthermore f contains $P \cap h$, so it contains the convex hull $\text{conv}(P \cap h)$. The reverse inclusion $f \subset \text{conv}(P \cap h)$ comes from the fact that any convex combinaison $\sum_{i=1}^{n} \lambda_i \, p_i$ of points in half-space h^+ belongs to h iff $\lambda_i = 0$ for all points p_i that are not in h. □

Therefore a polytope has a finite number of faces and each face of a polytope is itself a polytope. The boundary of $\mathcal{P}$ 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 that belong to P. The following lemma, whose proof is left as an exercise (Exercise 3.2) is a well known result from the theory of polytopes.

Lemma 3.2 *Any polytope is the convex hull of its vertices.*

simplices are defined in chapter 4 and also the description of their faces

The facial structure of a simplex can be easily described. Indeed, if σ is a simplex that is the convex hull of the set $S = \{p_1, \ldots, p_k\}$ 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 $\begin{pmatrix} k+1 \\ j+1 \end{pmatrix}$ 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 3.5 and 3.6).

Lemma 3.3 *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 3.3 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 $\mathcal{P}$.

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

do we need to define complexes here?

To describe the facial structure of polytopes, we intoduce here the notions of *simplicial complex* and *cell complex* which will be further studied in Chapter 4.

Definition 3.4 (Simplicial complex) *A simplicial complex is a set C of simplices which satisfies the two following properties::*

1. *Any face of any simplex of C is a simplex of C.*

2. *The intersection of any two simplices of C is either empty or a common face of both.*

Definition 3.5 (Cell complex) *A cell complex is a set C of polytopes which satisfies the two following properties*

1. *Any face of any polytope of C is a polytope of C.*
2. *The intersection of any two polytopes of C is either empty or a common face of both.*

Any face of a face of a polytope $\mathcal{P}$ is a face of $\mathcal{P}$ and the intersection of any two faces of $\mathcal{P}$ is either empty or a common face of both faces. The set of faces of polytopes $\mathcal{P}$ is therefore a cell complex which is called the *boundary complex* of $\mathcal{P}$ and denoted by $\mathrm{bd}(\mathcal{P})$. If the set of points P is in general position, the boundary complex of the polytope $\mathcal{P} = \mathrm{conv}(P)$ is a simplicial complex.

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 space of dimension $d - k$. A polytope defined as the intersection of n half-spaces bounded by hyperplanes in general position is called *simple*.

In the sequel, we will assume that point sets and hyperplanes sets are in general position. why ?

3.2 Duality

3.2.1 Point-hyperplane duality

The duality (or polarity) that is defined here, is related to the unit paraboloid of $\mathbb{R}^d$ and gives a special role to the last coordinate axis, that we call the *vertical axis*. We denote here by $x(x_1, \ldots, x_d)$ the point $x \in \mathbb{R}^d$ whose coordinates are $(x_1, \ldots, x_d)$. The unit paraboloid $\mathcal{Q}$ is defined as :

$$\mathcal{Q} = \{x(x_1, \ldots, x_d) \in \mathbb{R}^d : x_d = \textstyle\sum' x_i^2\},$$

where, here and in the following, we write $\sum'$ for $\sum_{i=1}^{d-1}$.

Duality is easier to describe if we consider the projective version $\mathcal{Q}_{proj}$ of the unit paraboloid :

$$\mathcal{Q}_{proj} = \{X(X_1, \ldots, X_{d+1}) \in \mathbb{P}^d : \sum{}' X_i^2 - X_d X_{d+1} = 0\}$$

where $\mathbb{P}^d$ is the projective space of dimension d, and $X(X_1, \ldots, X_{d+1})$ is the projective point associated to the euclidean point $x(x_1, \ldots, x_d)$ by the relations $x_i = X_i/X_{d+1}$ for $i = 1, \ldots d$. The equation of the projective unit paraboloid $\mathcal{Q}_{proj}$ is therefore a quadratic form that can be rewritten in matrix form

$$\mathcal{Q}_{proj} = \{X \in \mathbb{P}^d : X^t M X = 0\}$$

with

$$M = \begin{bmatrix} I_{d-1} & 0 & 0 \\ 0 & 0 & -1/2 \\ 0 & -1/2 & 0 \end{bmatrix}$$

where I_{d-1} is the identity matrix of dimension $(d-1)$. We consider below the bilinear form $Y^t M X$ associated to the quadratic form $X^t M X$ defining $\mathcal{Q}_{proj}$. Such a bilinear form defines a duality associating to a point Y of $\mathbb{P}^d$ the linear form $X \longrightarrow Y^t M X$, whose kernel is a hyperlan Y^* of $\mathbb{P}^d$:

$$Y^* = \{X \in \mathbb{P}^d : Y^t M X = 0\}.$$

In Euclidean space $\mathbb{R}^d$, duality associates to the point $y(y_1, \ldots, y_d)$ of $\mathbb{R}^d$, the non vertical hyperplane y^*:

$$y^* = \{(x \in \mathbb{R}^d : x_d - 2\sum{}' y_i x_i + y_d = 0\}.$$

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 $h^*(-h_1/2, \ldots - h_{d-1}/2, h_d)$.

Since $x^{**} = x$, 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:

$$x \in y^* \iff x_d - 2\sum{}' y_i x_i + y_d = 0. \iff y \in x^*,$$

Let h be a non vertical hyperplane, with normal form equation $h(x) = 0$. We say that point p is above hyperplane h or that h is below p if $h(p) > 0$. We say that point p is below hyperplane h or that h is above p if $h(p) < 0$. For a non vertical hyperplane h (like x^* or y^*), we denote by h^+ the half-space bounded by h that is above h and 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 vertical ordering, i.e. above-below relations, between points and hyperplane meaning :

$$\begin{aligned} x \in y^{*+} &\iff x_d - 2\sum{}' y_i x_i + y_d > 0 \iff y \in x^{*+} \\ x \in y^{*-} &\iff x_d - 2\sum{}' y_i x_i + y_d < 0 \iff y \in x^{*-}. \end{aligned}$$

3.2.2 Convex polyhedra and duality

In the sequel, we somehow extend the notion of polytopes and consider convex polyhedra. A *convex polyhedron* is the intersection of a finite number of half-spaces. It may be bounded or not. The notion of supporting hyperplanes and faces introduced above for 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. The notion of cell complexes can also be extended to include sets of convex polyhedra that satisfy the properties in Definition 3.5. The set of faces of a convex polyhedra may then be considered as a cell complex, the *boundary complex* of the convex polyhedra.

In the following, we call *upper half-space* an half-space h^+ that has a non vertical bounding hyperplane h and is above its bounding hyperplane..

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$. Let $H(P)$ be the set of hyperplanes supporting $\text{conv}(P)$. We distinguish the subset $H_+(P)$ of lower supporting hyperplanes where an hyperplane h of $H(P)$ is a lower supporting hyperplane if h is a non vertical hyperplane and $\text{conv}(P)$ is included in the in the upper half-space h^+ bounded by h. The lower part of the boundary of the convex hull, $\text{bd}^-(\text{conv}(P))$, is the subcomplex of the convex hull boundary, formed by faces of $\text{conv}(P)$ included in lower supporting hyperplanes:

$$\text{bd}^-(\text{conv}(P)) = \{\text{conv}(P) \cap h : h \in H_+(P)\}$$

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$, and $P^* = \{p_1^*, \ldots, p_n^*\}$ the set of dual hyperplanes. Duality yields a bijective map between the faces of the lower part of the boundary of the convex hull $\mathrm{bd}^-(\mathrm{conv}(P))$ and the faces of the convex polyhedra $\bigcap P^{*+}$, intersection of the upper half-spaces bounded by the hyperplanes in P^*. Indeed, let f be a $(k-1)$-face in $\mathrm{bd}^-(\mathrm{conv}(P))$. Face f is a $(k-1)$-face of $\mathrm{conv}(P)$, and, since general position is assumed, it includes k-vertices of $\mathrm{conv}(P)$ which are points of P. Let us assume, witout loss of generality, that $\{p_1, \ldots, p_k\}$ are included in f. We consider a lower supporting hyperplane h that intersects $\mathrm{conv}(P)$ along f. Hyperplane h includes $\{p_1, \ldots, p_k\}$ and, since it is a lower supporting hyperplane, other points of P are included in the upper half-space h^+. Let h^* be the dual point of h. Then, from the properties of the point-hyperplane duality introduced above, we get:

$$p_i \in h^+ \iff h^* \in p_i^{*+}, \quad \forall i = 1, \ldots, n \tag{3.1}$$

$$p_i \in h \iff h^* \in p_i^*, \quad \forall i = 1, \ldots, k \tag{3.2}$$

$$p_i \notin h \iff h^* \notin p_i^*, \quad \forall i = k+1, \ldots, n \tag{3.3}$$

Equations 3.1 , 3.2 and 3.3 show that h^* belongs to a $(d-k)$-face f^* of $\bigcap P^{*+}$ that is the intersection of $\bigcap P^{*+}$ with the k hyperplanes $\{p_i^*, i = 1, \ldots k\}$. Therefore, duality maps the k-face $f = \mathrm{conv}(\{p_1, \ldots, p_k\})$ of $\mathrm{bd}^-(\mathrm{conv}(P))$ to the $(d-k)$-face $f^* = \bigcap_{i=1,\ldots,k} p_i^* \bigcap P^{*+}$ in $\mathrm{bd}\,(\bigcap P^{*+})$. Since $f^{**} = f$, this map is bijective. Furthermore, it reverses inclusion meaning that if f and g are two faces in $\mathrm{bd}^-(\mathrm{conv}(P)$ such that $g \subset f$, the corresponding faces f^* and g^* in $\mathrm{bd} \bigcap P^{*+}$ are such that $f^* \subset g^*$. The two cell complexes $\mathrm{bd}^-(\mathrm{conv}(P)$ and $\mathrm{bd}\,(\bigcap P^{*+})$, related by a bijective map that reverses inclusion are said to be *dual*. useful? in ch 4?

3.3 Combinatorial bounds

3.3.1 Polytopes of dimension 3

Theorem 3.6 (Euler formula) *Let $\mathcal{P}$ be a polytope of dimension 3.. The number of vertices n_v, edges n_e and facets n_f of $\mathcal{P}$ are related by the following relation:*

$$n_v - n_e + n_f = 2.$$

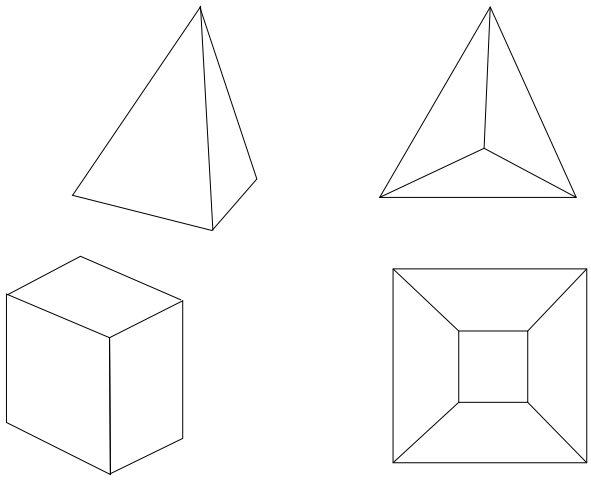

Figure 3.1: Schlegel diagrams of a tetrahedron and of a cube.

Proof The Schlegel diagram of $\mathcal{P}$ is a planar model of the complex $\mathrm{bd}(\mathcal{P})$ obtained when $\mathcal{P}$ is projected on a facet of $\mathcal{P}$ from a point which is outside $\mathcal{P}$ but arbitrary close to a point inside the facet. The Schlegel diagram is a connected map where the projection facet appears as a large polygon with all the remaining faces filling its interior.

We prove Euler's formula for any connected planar map where n_v, n_e and n_f are respectively the number of vertices, edges and 2-faces. We build the map from a single vertex by adding edges one-by-one. When an edge is added in the current map C, the vertices of the edge not yet in C are added at the same time. Sine the map is connected, it is possible to choose the order in which edges are inserted in such a way that the current map C is always connected. Then, at each step, the edge to be inserted to the current map C has one or two vertices already in C. In the first case, (case 1 in Figure 3.2), n_f does not change while n_v and n_e both increase by 1. In the other case, (case 2 in Figure 3.2), n_v does not change while n_e and n_f both increase by 1. In both cases, $n_v - n_e + n_f$ remains invariant. At the beginning we have, counting the unbounded face of the planar map, $n_v = n_f = 1$, and $n_e = 0$, and thus $n_v - n_e + n_f = 2$. This value 2 is maintained throughout the whole construction. The formula is thus true for any connected planar map and in particular for Schlegel diagrams. Any face of the Schlegel diagram except the unbounded facet is the projection of a polyhedron facet and any

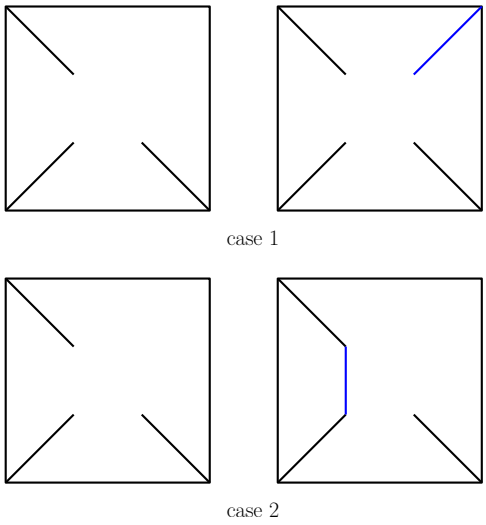

Figure 3.2: The induction step in the proof of Euler's formula.

polyhedron face correspond to a face in Schlegel diagram except for the facet on which the projection takes place. Therefore to get from the faces in the Schlegel diagram to the faces of the polyhedron, we need to withdraw 1 from n_f for the unbounded face and add 1 to n_f for the projection face. This leaves unchanged the number of faces and we get that the formula is true for any polytope of $\mathbb{R}^3$.

Each edge is incident to two facets and each facet has at least three edges. We thus have

$$2n_e \geq 3n_f,$$

with an equality in the case where all facets are triangles. Using this inequality in Euler's formula, we get

$$\begin{aligned} n_e &\leq 3n_v - 6 \\ n_f &\leq 2n_v - 4 \end{aligned}$$

with an equality when all facets are triangles. □

3.3.2 Beyond the third dimension

The following theorem considers the worst case complexity of a polytope in $\mathbb{R}^d$ which is known to have either n facets or n vertices. Though, for historical reasons, the theorem is called the upper bound theorem, it provides an optimal bound, i.e. a bound which is both an upper and a lower bound on such a worst case complexity.

Theorem 3.7 (Upper Bound Theorem) *The total number of faces of a polytope 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 3.10.

Let $\mathcal{P}$ be a polytope 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 then in general position. During this process, the number of faces of $\mathcal{P}$ can only increase. Hence, any upper bound that is valid for hyperplanes in general position holds for any set of hyperplanes. 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. Two faces f and g of a polytope are said to be *incident* if one of them is included in the other. should be said before Let p be one of the vertices of $\mathcal{P}$. Because general position is assumed, p, as any other vertex of $\mathcal{P}$, is incident to exactly d of the bounding hyperplanes and to d edges of $\mathcal{P}$. Thus there are at least $\lceil \frac{d}{2} \rceil$ edges incident to p 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 a simple polytope and the affine hull of any subset of k edges incident to a vertex of $\mathcal{P}$ contains a $(d-k)$-face of the polytope (Exercice 3.8). 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 3.7). We deduce that the number of k-faces is at most $\begin{pmatrix} n \\ d-k \end{pmatrix} = O(n^{d-k})$, 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 $O(n^{\lfloor \frac{d}{2} \rfloor})$.

At last, the number of k-faces incident to each vertex of the simplicial polytope $\mathcal{P}$ is $\begin{pmatrix} d \\ d-k \end{pmatrix}$ so that 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 number of faces of polytopes defined as the convex hull of n points. □

3.4 Convex hull algorithms

3.4.1 Hasse diagram and adjacency graph

Recall that two faces f and g of a polytope are said to be *incident* if one is included in the other.

A polytope is most commonly represented by its *Hasse diagram* including a node for each face and an edge between pairs of incident faces whose dimensions differ by 1.

Two facets of a d-polytope are said to be adjacent if they share a $(d-2)$-face. Any $(d-2)$-face of a d-polytope is incident to exactly two facets, as easily follows from the facts that the dual of a $(d-2)$-face is an edge, the dual of a facet is a vertex and an edge is incident to two vertices. Hence the Hasse diagram of a polytope encodes also its *adjacency graph* with one node for each facet and one edge for a pair of adjacent facets.

3.4.2 An incremental algorithm

The incremental algorithm for convex hull consists in maintaining the Hasse diagram of $\text{conv}(P)$ while the points of P are inserted one by one. We

assume for simplicity that P is in general position.

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. When point p_i is handled, the faces of $\text{conv}(P_{i-1})$ may be classified in the following way (see Figure 3.3:

- A facet f of $\text{conv}(P_{i-1})$ is *red* with respect to p_i if the hyperplane h_f supporting $\text{conv}(P_{i-1})$ along f *separates* p_i from P_{i-1} meaning that p_i belongs to the open half-space h_f^+ that does not intersect $\text{conv}(P_{i-1})$. Otherwise, as general position is assumed, p_i is included in the half-space h_f^- whose closure contains $\text{conv}(P_{i-1})$, and the facet f is said to be *blue*.
- A k-face, with $k < d$ is said to be *red* if all the incident facets are red, it is said to be *blue* if all the incident 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 $\text{conv}(P_{i-1})$ into $\text{conv}(P_i)$ when adding point p_i.

1. The set of facets of $\text{conv}(P_{i-1})$ that are red with respect to p_i, is connected by adjacency relations.
2. The set of faces of $\text{conv}(P_i)$ includes the blue and purple faces of $\text{conv}(P_{i-1})$ together with additionnal *new faces* which are the convex hulls $\{conv(g \cup p_i), \}$ where g is a purple face of $\text{conv}(P_{i-1})$.
3. The set of purple faces of $\text{conv}(P_{i-1})$ with their incidence relation is isomorphic to the set of faces of a $(d-1)$-polytope with at most $(i-1)$ vertices. (See Figure 3.4

Fact 1 is illustrated on Figure 3.3. Fact 2 and 3 are illustrated on Figure 3.4 showing that the faces of $\text{conv}(P_{i-1})$ incident to p_i are in bijection with the set of faces of a $(d-1)$-polytope obtained as the intersection of $\text{conv}(P_i)$ with any hyperplane h that separates $\text{conv}(P_{i_1})$ from p_i.

The incremental algorithm sorts points in P and introduces them in lexicographic order. This order ensures that, at step i where p_i is inserted, at least one of the facets of $\text{conv}(P_{i-1})$ incident to p_{i-1} is red with respect to p_i. Indeed fron the lexicographic ordering, p_i is separated from P_{i-1} by at

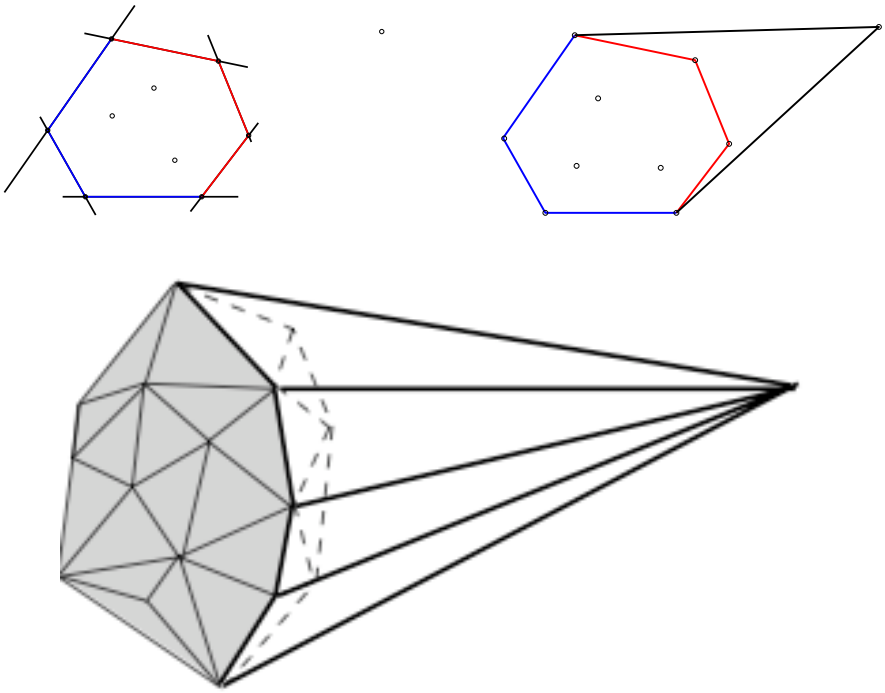

Figure 3.3: Incremental convex hull algorithm. Top: incremental algorithm in 2d, face coloring (left) and update of the convex hull (right) when adding a new point. Bottom: incremental algorithm in 3d.

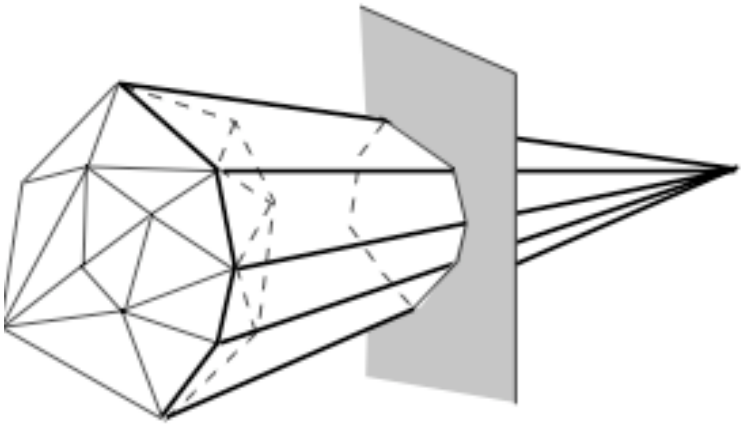

Figure 3.4: When adding point p_i, the set of purple facets of $\text{conv}(P_{i-1})$ is isomorphic to a $(d-1)$-polytope with at most $(i-1)$ vertices.

least one of the hyperplanes supporting $\text{conv}(P_{i-1})$ along a facet incident to p_{i-1}. The algorithm visits the facets of $\text{conv}(P_{i-1})$ incident to p_{i-1} until it finds a first red facet. Then the algorithm visits all red facets following the adjacency graph and identifies all red and purple faces following the Hasse diagram. The Hasse diagram is updated by removing red faces and creating a new faces $conv(g \cup p_i)$ for each purple face p_i. Incidence relation for new faces are deduced from incidence relation of purple faces in $\text{conv}(P_{i-1})$, details are omitted here.

Updating the Hasse diagram has a complexity proportionnal to the number of new and removed faces which are respectively the number of purple and red faces of $\text{conv}(P_{i-1})$. To find the first red facets, the algorithm visits facets incident to p_{i-1}. Therefore each facet will be visited only once for this purpose during the whole algorithm. Identifying red and purple faces is also clearly proprotionnal 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 facet g can be charged on the new face $\text{conv}(g \cup p_i)$. In summary, the cost of step i is proportionnal 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 proprotionnal to the total number of created faces. From above fact 3 and the upper bound theorem, the number of faces

created when inserting p_i is $O\left((i-1)^{\lfloor \frac{d}{2} \rfloor}\right)$ and the total cost of updating the Hasse diagram is:

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

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

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

3.4.3 Can we do better ?

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 $\mathbb{R}^d$.

Furthermore, it is known that sorting n elements requires at least $\log_2(n!) = \Omega(n \log n)$ comparisons. The following shows that sorting n number reduces in linear time to the computation of the convex hull of n points in $\mathbb{R}^2$.

Take n real numbers $x_1, \ldots, x_n$. We associate to each x_i the point $p_i = (x_i, x_i^2)$. The p_i lie on the parabola $y = x^2$. 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.

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. See Exercise 3.12 for an example in $\mathbb{R}^3$ where the incremental algorithm takes $\Omega(n^2)$ times.

3.4.4 Randomization helps

We present here a randomized version of the incremental algorithm whose expected complexity, matches the lower bound. Expectation here is related to the internal random choices performed by 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 feature of the randomized algorithm is to introduce the points in random order. Therefore the algorithm cannot rely on the lexicographic order of the input points to find a first red facet of the current hull and it has to resort to an additionnal data structure called the *conflict graph*.

In the following, points in $P = \{p_1, p_2 \dots p_n\}$ are assumed to be indexed according to their insertion order. We denote by P_i the subset formed by the i first inserted points. 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 $\text{conv}(P_i)$ in conflict with p_j. When a new point p_{i+1} is inserted, the conflict edge incident to p_{i+1} gives in constant time access to a first red facet of $\text{conv}(P_i)$ and the algorithm updates the Hasse diagram of the convex hull exactly as in the deterministic case. Of course, the algorithm needs now to update the conflict graph.

To update the conflict graph, it is required to find a new conflict edge for each not yet inserted point p_j ($j > i+1$) that is linked by its previous conflict edge to a red facet f_j of $\text{conv}(P_i)$. This is done by traversing the adjacency graph of $\text{conv}(P_i)$ visiting red facets of $\text{conv}(P_i)$ that are in conflict with p_j. Such facets are thus in conflict with two points, namely p_{i+1} and p_j. Let R denote the set of red facets of $\text{conv}(P_i)$ and F_j denote the set of facets in conflict with p_{i+1} and p_j. R is known to be connected by adjacency relations. The same is true for F_j if p_j lies inside $\text{conv}(P_{i+1})$. Otherwise, p_j is outside $\text{conv}(P_{i+1})$, and the new facets of $\text{conv}(P_{i+1})$ in conflict with p_j are necessaritly connected to F_j by purple $(d-2)$-faces meaning that there are at least one facet in F_j sharing a purple $(d-2)$-face with a new facet of $\text{conv}(P_{i+1})$ in conflict with p_j. See Figure 3.5 for an illustration in $\mathbb{R}^2$. Therefore, to update the conflict graph, the algorithms proceeds as follows, The current conflict graph yields the not yet inserted points whose associated conflicting facets in the conflict graph are red. Each such point is handled in turn. Let p_j, $j > i+1$, be such a point and let f_j be the red facet of $\text{conv}(P_i)$ in conflict with p_j. Starting from f_j, we walk on the adjacency graph Hasse diagram? of $\text{conv}(P_i)$, visiting red facets that are in conflict

with p_j. For each visited facet, we check its adjacent facets in turn. Let f be a visited facet and g be a facet of $\mathrm{conv}(P_i)$ adjacent to f. If one of the two following situations occur, p_j gets a new conflicting facet and the walk is stopped:

Case 1 : g conflicts with p_j but not with p_{i+1}. In this case, a new conflict edge is created between g and p_j.

Case 2 : g does not conflict with p_j nor with p_{i+1}, and the new facet f' containing the purple $(d-2)$-face $g \cap f$ conflicts with p_j. In this case, a new conflict edge is created between f' and p_j.

If the walk traverses all red facets in conflict with p_j without encountering one of the above two cases, p_j lies inside $\mathrm{conv}(P_{i+1})$ and is discarded.

We deduce from the above discussion that the procedure will find a new conflict for all points p_j, $j > i+1$, that lie outside $\mathrm{conv}(P_{i+1})$.

To analyze this algorithm, we assume that the points are inserted in random order. Let us choose as origin a point o in $\mathrm{conv}(P)$. For the purpose of randomized analysis, we call *configurations* the subsets of d independent points in $\mathbb{R}^d$. A configuration σ is said to be defined over a set of points P if the points in σ belong to P, it is said to be in conflict with a point p if the hyperplane which is the affine hull of σ separates o from p. We denote by $C(P)$ the set of configurations defined over P and by $C_j(P)$ (resp. $C_{\leq k}(P)$), the set of configurations defined and with j conflicts (resp. at most k conflicts) over P. Note that a configuration that is defined and without conflict over P, is just the vertex set of a facet of $\mathrm{conv}(P)$ so that there is a bijection between $C_0(P)$ and the facets of $\mathrm{conv}(P)$.

In the following, we consider random subsets S of P and we are interested in the number of configurations defined over S and P. The following theorem whose proof is sketched in Exercise 3.14 is a classical result.

Theorem 3.8 (Sampling theorem) *Assume that each configuration is defined by d points, and that P is a set of n points. The number $|C_{\leq k}(P)|$ of configurations defined and with at most k-conflicts over P is bounded as follows:*

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

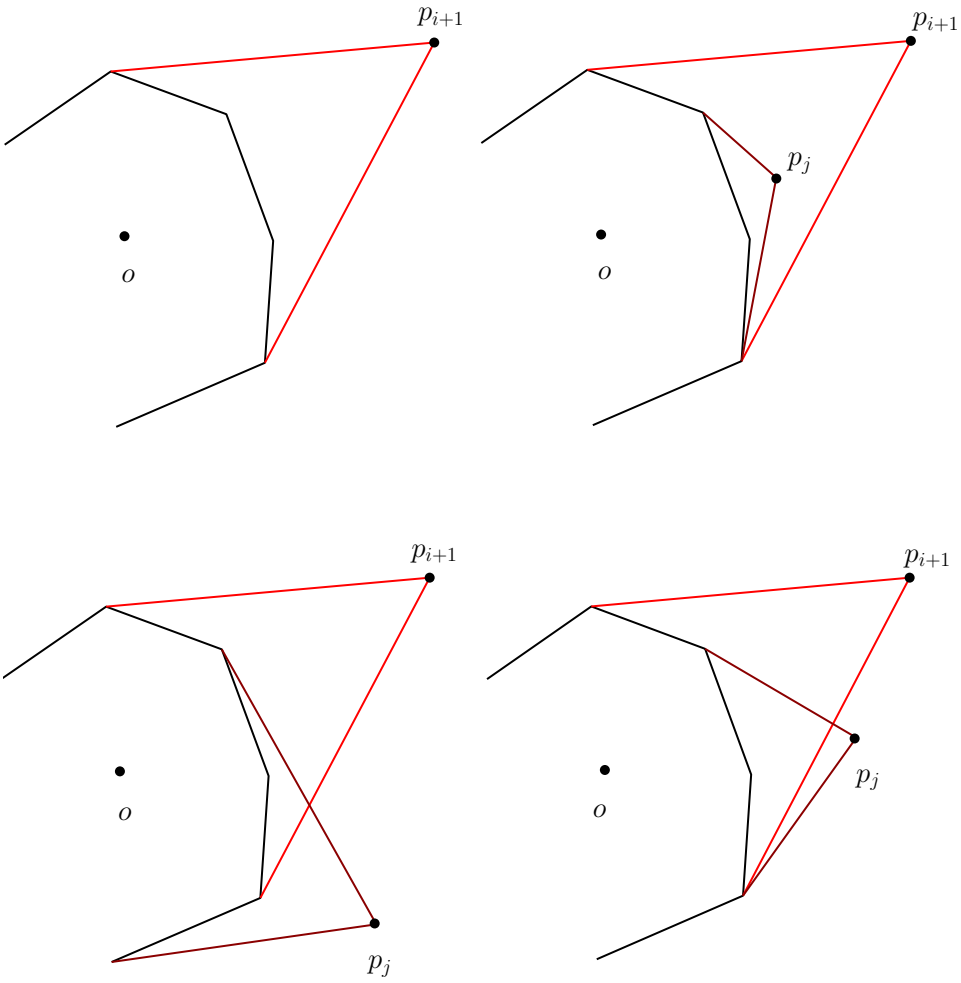

Figure 3.5: Incremental construction of the convex hull. Updating the conflict edge of p_j when inserting p_{i+1}.

where $\nu_o(\lfloor \frac{n}{k} \rfloor, P)$ is the expected number of configurations defined and without conflict over a random sample of P of size $\lfloor \frac{n}{k} \rfloor$. Given our definition of configurations and the upper bound theorem, we get that the number of configurations defined and without conflict over a set of i points is at most $O\left(i^{\lfloor \frac{n}{2} \rfloor}\right)$. Therefore

$$\nu_o(\lfloor \frac{n}{k} \rfloor, P) = O\left(\lfloor \frac{n}{k} \rfloor^{\lfloor \frac{d}{2} \rfloor}\right).$$

Plugging this 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 :

$$|C_{\leq k}(P)| = O\left(k^{\lceil \frac{d}{2} \rceil} n^{\lfloor \frac{d}{2} \rfloor}\right). \tag{3.5}$$

We analyze first the cost of updating the Hasse diagram of the convex hull and then the cost of maintaining the conflict graph.

As in the case of the incremental algorithm, the cost for 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 $n(i)$ of facets that are created at step i of the algorithm. Since the algorithm inserts the points of P in random order, P_i is a random subset of P of size i. Notice that a facet created at step i corresponds to a configuration of $C_0(P_i)$. Given P_i, a configuration σ of $C_0(P_i)$ corresponds to a facet created at step i iff one of the point in this configuration is the point p_i inserted at step i, which happens with probability $\frac{d}{i}$. Thus, we have

$$n(i) = \sum_{\sigma \in C(P)} \text{proba}(\sigma \in C_0(P_i)) \times \frac{d}{i} = \frac{d}{i}\nu_0(i, P) = O\left(i^{\lfloor \frac{d}{2} \rfloor - 1}\right).$$

MY : JD asked to add the following remark. Then I think that the lower bound for deterministic incr algo should be given at the end of deterministic algorithm. Is there an obvious or easy proof ? cyclic polytopes ? This is to be compared with the lower bound $\Omega\left(n^{\lfloor \frac{d+1}{2} \rfloor}\right)$ of any deterministic incremental algorithm.

By summing over all steps i and using linearity of expectation, we obtain that the expected total number of facets created by the algorithm and there-

fore the expected cost of updating the Hasse diagram is:

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

We now bound the cost of updating the conflict graph. As explained above, when inserting p_{i+1} at step $i+1$, the algorithm visits the facets of $\text{conv}(P_i)$ in conflict with p_{i+1} and p_j, to either find a new conflict for p_j with $j > i+1$ or discover that this point may be discarderd. This process has a cost proportional to the number of visited facets.

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 traversed 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 with size $i+2$ is given. 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

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

and, using Equation 3.5,

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

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 traversed to restore a conflict for p_j at step $i+1$ is also $O(i^{\lfloor \frac{d}{2} \rfloor - 2})$, and the expected total cost of updating the conflict graph is

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

The following theorem sums up the discussion about the expected complexity of the randomized incremental convex hull,

Theorem 3.9 *The randomized incremental algorithm computes the convex hull of n points in $\mathbb{R}^d$ in expected time $O\left(n \log n + n^{\lfloor \frac{d}{2} \rfloor}\right)$.*

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 of the convex hull, the randomized incremental construction of a convex hull is often qualified as optimal. Note the the expectation concern only internal choices made by the algorithm, namely the insertion order, and not the data. No hypothesis are made on the data.

3.5 Exercises

Exercise 3.1 Let $\mathcal{P}$ be a polytope that is the convex hull of the point set P, let h be a supporting hyperplane of $\mathcal{P}$ and $f = \mathcal{P} \cap h$ the corresponding face of $\mathcal{P}$.

Show that f is a polytope that is the convex hull of the subset P_h of points of P that are included in the hyperplane h.

Exercise 3.2 Show that a polytope $\mathcal{P}$ is the convex hull ot 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 3.3 Show that the intersection of any finite set of faces of a polytope is also a face of the polytope.

Exercise 3.4 Show that any face of a polytope $\mathcal{P}$ is the intersection of facets of $\mathcal{P}$.

Exercise 3.5 Let $\mathcal{P}$ be a 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 $\mathcal{P} = \cap_{h \in H} h^+$.

Exercise 3.6 Show that, if it is bounded, the intersection of a finite set of half-spaces is a polytope.

Exercise 3.7 Let H be a set of n hyperplanes $h_1, ..., h_n$ and $\mathcal{P}$ be the polyhedron defined as the intersection of the n half-spaces $h_1^+, ..., h_n^+$ 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 $1, ..., n$ and $F_l = \cap_{i \in I} h_i$. Show that, if it is non empty, the intersection and $\mathcal{P} \cap F_l$ is a face of $\mathcal{P}$ and that all 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, $\mathcal{P} \cap F_l$ is a face of dimension $d - k$ if $|I| = k$.

Exercise 3.8 Prove that if $\mathcal{P}$ is a simple polyhedron and p a vertex of $\mathcal{P}$, the affine hull of any subset of $k < d$ edges incident to p contains a face of P of dimension $d - k$. (Hint: use Exercise 3.3 and duality)

Exercise 3.9 Show that n hyperplanes are in general position iff their dual points are in general position.

Exercise 3.10 (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 $\begin{pmatrix} n \\ k \end{pmatrix}$ $(k-1)$-faces for any k such that $0 \leq k \leq d/2$.

Exercise 3.11 Extend the discussion of section 3.2 to the case of the intersection H of n general half-spaces (not necessarily all above their bounding hyperplane) assuming we know a point o in H. (Hint : To the point p, we associate the polar hyperplane p^* $p^* = \{\forall x \in p^*, (x - o) \cdot (p - o) = 1\}$. To a hyperplane h not passing through o, we associate the polar point h^* : $(h^* - o) \cdot (x - o) = 1, \forall x \in h$. If we denote by h^+ the half-space bounded by h not containing o, we can easily adapt the above discussion and get a bijection between the faces of the intersection of the half-spaces and the convex hull of the polar points.)

polar introduced?

Exercise 3.12 Show that any incremental algorithm that constructs the convex hull of n points of $\mathbb{R}^3$ takes $\Omega(n^2)$ time in the worst-case. (Hint : take half of the points on a circle in the plane $x = 0$ centered at O and the other points on the x-axis, $x > 0$. Insert the points on the circle before the points on the x-axis.)

Exercise 3.13 Show that we can remove a point from the convex hull of a set of n points in expected time $O(\log n)$.

Exercise 3.14 (Sampling theorem) Show that the number of configurations ....

3.6 Bibliographical notes

A modern introduction to the theory of polytopes can be found in Ziegler's book [110]. The original proof of the upper bound theorem has been established by McMullen in 1970. The simple asymptotic version given in Theorem 3.7 is due to Seidel [103]. MY: ajouter quelquechose sur l'algo incremental et sur l'algo randomise. Chazelle [44] 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 [92].

Chapter 4

Simplicial complexes

revoir la defintion du geometric simplex

Geometric shapes like curves, surfaces or more generally manifolds 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 contraints 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.

4.1 Geometric simplicial complexes

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

The set of convex combinations of $k+1$ affinely independent points $p_0, \ldots, p_k$

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

is called a k-dimensional simplex or a *k-simplex* for short. 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 a simplex σ whose vertex set is P are the simplices of lower dimensions obtained as convex combinations of subsets of P. The k-faces of a l-simplex σ, $l > k$, are obtained by taking the convex combinations of all the subsets of $k+1$ vertices of σ. For example, the faces of a triangle $[p_0, p_1, p_2]$ 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 4.1 (Simplicial complex) *A (finite) simplicial complex K is a finite union of simplices such that:*

1. *any face of any 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.*

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* $\mathrm{Sk}_j(K)$ of K is the subcomplex of K consisting of the simplices of dimension at most j; $\mathrm{Sk}_j(K)$.

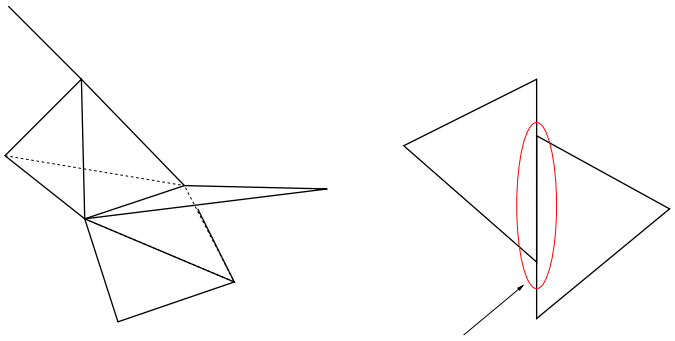

Figure 4.1: Left: an example of a simplicial complex. Right: a union of simplices which is not a simplicial complex

For a simplicial complex K in $\mathbb{R}^d$, its *underlying space* $|K|$ is the subset of $\mathbb{R}^d$ that 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.

4.2 Abstract simplicial complexes

In the definition above, the faces of K are geometric simplices of $\mathbb{R}^d$ and K is embedded in $\mathbb{R}^d$. However, since each face of K is determined by its vertices, a simplicial complex can be defined in a purely combinatorial way.

Definition 4.2 (Abstract simplicial complex) *Let $P = \{p_1, \cdots p_n\}$ be a finite set of elements. An abstract simplicial complex K with vertex set P is a set of subsets of P satisfying the two conditions :*

1. *The elements of P belong to K.*

2. *If* $\tau \in K$ *and* $\sigma \subseteq \tau$, *then* $\tau \in K$.

The elements of K are called the simplices or the faces of K. Here the simplices K are not considered as geometric simplices in $\mathbb{R}^d$ but just as sets of vertices. If $\sigma \in K$ has precisely $k+1$ elements, the dimension of σ is k and we say that σ is a k-simplex. The dimension of K is the maximal dimension of its simplices.

Any simplicial complex K naturally determines an abstract complex K_{abs}: the vertex set of K_{abs} is the set of vertices of K and the simplices of K_{abs} are the sets of vertices of the simplices of K. Conversely, any finite abstract simplicial complex K can be realized as a geometric simplicial complex K_g in an Euclidean space in the following way. Let $\{v_1, v_2, \cdots, v_N\}$ be the vertex set of K where N is the number of vertices of K, and let $\sigma \subset \mathbb{R}^N$ be the simplex spanned by $\{e_1, e_2, \cdots, e_N\}$ where, for any $i = 1, \cdots, N$, e_i is the vector whose coordinates are all 0 except the i^{th} one which is equal to 1. Then K_g is the subcomplex of σ defined by $[e_{i_0}, \cdots, e_{i_k}]$ is a k-simplex of K_g if and only if $[v_{i_0}, \cdots, v_{i_k}]$ is a simplex of K. Such a complex K_g is called a *geometric realization* of K. Note that the geometric realization of K is not uniquely defined. For example, any finite abstract simplicial complex can be realized as a simplicial complex in $\mathbb{R}^{2d+1}$ (Exercice 4.4). However, one can prove that if K_{g_1} and K_{g_2} are two geometric realizations of the same finite abstract complex K, then K_{g_1} and K_{g_2} are homeomorphic.

In this book, we will often encounter abstract simplicial complexes whose vertices are points in $\mathbb{R}^d$. Let K be such a complex and P the set of points associated to the vertices of K. Let ϕ be the map that maps the vertices of a simplex $\sigma \subset K$ to the convex hull of the associated points of P. If this mapping is injective, then $\phi(K)$ is a realization of K *embedded* in $\mathbb{R}^d$.

Examples. A *triangulated d-ball* is a simplicial complex K whose realization is homeomorphic to the unit d-ball of $\mathbb{R}^d$. A *triangulated $(d-1)$-sphere* is a simplicial complex whose realization is homeomorphic to the unit $(d-1)$-sphere of $\mathbb{R}^d$. Examples of triangulated $(d-1)$-spheres are given by the boundary complexes of simplicial d-polytopes.

4.3 Nerve

As noticed in previous the section, simplicial complexes can be seen at the same time as topological spaces and as purely combinatorial objects. Given an open cover of a topological space X, i.e. a collection $\mathcal{U} = (U_i)_{i\in I}$ of open subsets $U_i \subseteq X$, $i \in I$ where I is a set, we associate a simplicial complex $C(\mathcal{U})$ whose vertex set is $\mathcal{U}$ and such that

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

Such a simplicial complex is called *the nerve or Çech complex of the cover* $\mathcal{U}$. When all the open sets U_i 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 4.3 (Nerve Theorem) *Let $\mathcal{U} = (U_i)_{i\in I}$ be an open cover of a paracompact space X such that any finite intersection of the U_i's is either empty or contractible. Then X and $C(\mathcal{U})$ are homotopy equivalent.*

Similar to open covers, closed covers of a topological space X are defined as families of closed sets whose union is X. 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 4.4 (Nerve Theorem for closed covers) *Let $X \subset \mathbb{R}^d$ be a subset of $\mathbb{R}^d$ homeomorphic to a finite simplicial complex and let $\mathcal{F} = (F_i)_{i\in I}$ be a closed cover of X such that any finite intersection of the F_i's is either empty or contractible. 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 topology (homotopy type) of continuous topological space X by a simplical complex describing the combinatorics of a good cover. In particular, when X is a (finite) union of (closed or open) balls in $\mathbb{R}^d$ or, more generally, of convex sets, it is homotopy equivalent to the nerve of the cover made by this union.

4.4 Filtrations of simplicial complexes

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

Definition 4.5 *A filtration of a finite simplicial complex K is a sequence of sub-complexes defined by*

$$(i) \qquad \emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K$$
$$(ii) \quad K^{i+1} = K^i \cup \sigma^{i+1} \quad \textit{where } \sigma^{i+1} \textit{ is a simplex of } K$$

Equivalently, a filtration of K can be seen as an ordering of the simplices such that for any $i \geq 0$, the union 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.

Lemma 4.6 *$\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K$ with $K^{i+1} = K^i \cup \sigma^{i+1}$ for all $i = 0, \cdots m-1$ is a filtration of K if and only if for any $i = 0, \cdots m-1$, all the faces of σ^{i+1} are contained in K^i.*

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: $\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \cdots \subset K^{\alpha_m} = K$. For example, when a function is defined on the vertices of K, on 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 a real valued function defined on the vertices of K. For any simplex $\sigma = [v_0, \cdots v_k]$ one defines $f(\sigma)$ by

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

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.

4.5 Rips and Čech complexes

Filtrations are often built on top of finite sets of data points to reveal the underlying topological structure of data (see Chapter 13).

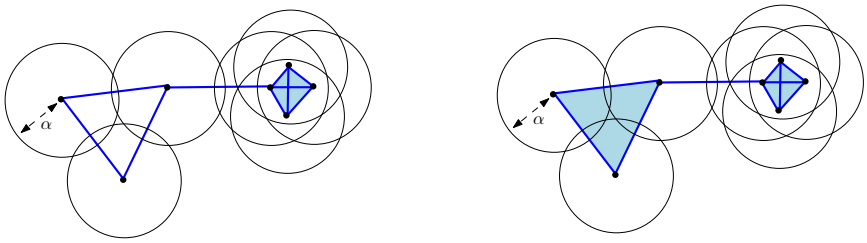

Figure 4.2: The Čech (left) and Vietoris-Rips (right) complexes built on top of a finite set of points in $\mathbb{R}^2$. Note that they both contains a 3-simplex and are thus not embedded in $\mathbb{R}^2$.

Given a finite set of points P in $\mathbb{R}^d$, the nerves $\check{\text{C}}\text{ech}(P,\alpha)$ of the unions of balls centered on P with radius α (see Chapter 4.3) also define a filtration as α goes from 0 to $+\infty$. This filtration of the simplex of dimension $|P|-1$ is called the **Čech complex** $\check{\text{C}}\text{ech}(P)$ built on top of P. The simplices of $\check{\text{C}}\text{ech}(P,\alpha)$ are characterized by the following condition:

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

The topology of the Čech is closely related to the topology of the alpha complex filtration introduced in Section 6 (they are homotopy equivalent). However, when d is larger than 3, their computation quickly becomes intractable in practice. Another classically used, closely related, and easier to compute filtration is the *Vietoris-Rips complex*, $\mathbb{R}\text{ips}(P)$, built on top of P and defined by the following condition

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

The Vietoris-Rips complex is the largest simplicial complex that has the same 1-skeleton as the Čech complex. Moreover, it satisfies the following interleaving property that plays a fundamental role in topological data analysis (see Chapter **??**).

Lemma 4.7 *Let P be a finite set of points in $\mathbb{R}^d$. for any $\alpha \geq 0$,*

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

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

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

Remark that the Čech and Vietoris-Rips complexes 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 given interleaving of Lemma 4.7 is not tight and can be slightly improved (see Exercise 13.2).

4.6 Representation of simplicial complexes

Hasse diagram. simplex tree ?

4.7 Manifold complexes

Definition 4.8 (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* $\text{link}(p, K)$*.*

Definition 4.9 (Pure complex) *A simplicial k-complex K is pure if every simplex in K is the face of a k-simplex.*

Definition 4.10 (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 4.11 (Manifold complex) *A simplicial complex K is a k-manifold complex if*

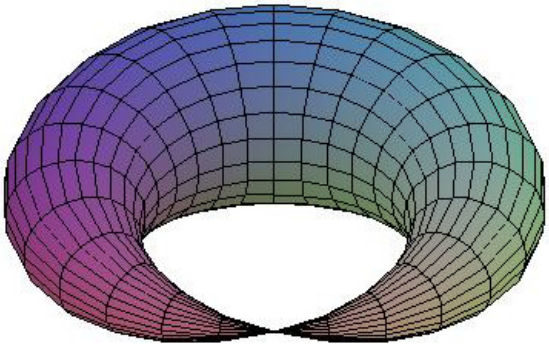

Figure 4.3: A pinched torus is a pseudomanifold complex.

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*

A pseudomanifold looks like a manifold at most of the points, but may contain singularities (see Figure **??**).

Definition 4.12 (Pseudomanifold complex) *A simplicial complex K is a k-pseudomanifold if*

1. *K is a pure k-complex*
2. *every $(k-1)$-simplex is the face of exactly two k-simplices*
3. *the adjacency graph of the k-simplices (i.e. the graph whose nodes are the k-simplices and whose edges are the $(k-1)$-simplices) is connected.*

We have the implications : polytope boundary $\Rightarrow$ triangulated sphere $\Rightarrow$ manifold complex $\Rightarrow$ pseudomanifold

Definition 4.13 (Triangulation of a point set) *A triangulation of a finite point set $P \in \mathbb{R}^d$ is a geometric simplicial complex K with vertices P such that the underlying space $|K|$ of K is the convex hull of P.*

Definition 4.14 (Triangulation of a manifold) *A simplicial complex K triangulates a k-manifold M if there is a homeomorphism between a realization of K and M.*

Lemma 4.15 (Pseudomanifold criterion) *Let K be a triangulation of P. Suppose $K \subseteq K'$, where K' is also a simplicial complex with the same vertex set. If K' is a pseudomanifold, then $K' = K$.*

Proof Let p be a point in $L = P \setminus \partial K$. Since $\mathrm{link}(p, K)$ is a simplicial $(d-1)$-sphere, it is a $(d-2)$ connected pseudomanifold. A $(d-2)$-connected pseudomanifold C cannot be properly contained in another $(d-2)$-connected pseudomanifold because a $(d-1)$-simplex that does not belong to C cannot share a $(d-2)$ face with a simplex of C: every $(d-2)$-simplex of C is already the face of two $(d-1)$-simplices of C.

Therefore, since $\mathrm{link}(p, K) \subseteq \mathrm{link}(p, K')$, we must have $\mathrm{link}(p, K) = \mathrm{link}(p, K')$. It follows that $\mathrm{star}(p, K') = \mathrm{star}(p, K)$ for every $p \in L$. If $b \in P$ is a boundary vertex, then q belongs to $\mathrm{link}(p, K)$ for some $p \in L$, and it follows that $\mathrm{star}(q, K') = \mathrm{star}(q, K)$. Therefore $K' = K$. $\square$

4.8 Exercises

Exercise 4.1 *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.*

Exercise 4.2 *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 $\gamma : [0, 1] \to C$ 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 4.3 *Give examples of simplicial complexes in $\mathbb{R}^3$ that are homeomorphic to a ball, a sphere, and a torus.*

Exercise 4.4 *Prove that any abstract simplicial complex K of dimension d can be realized as an isomorphic geometric simplicial complex in $\mathbb{R}^{2d+1}$.*

(Hint : map the vertices of K to points on the moment curve $C = \{(x, x^2, ..., x^{2d+1}) \in \mathbb{R}^{2d+1}, x \in \mathbb{R}\}$. 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 3.10.)

4.9 Bibliographical notes

Our presentation of simplicial complexes follows the one in Munkres [93]. The nerve theorem and its variants are classical results in algebraic topology. A proof is given in Hatcher [80], Section 4G.

Part II

Delaunay complexes

Chapter 5

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 $\mathbb{R}^d$ and is therefore a 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 triangulation of P is the projection onto $\mathbb{R}^d$ of the boundary complex of the convex hull of the lifted points. Using then the results of Chapter 3, we can bound the combinatorial complexity of Delaunay triangulations and Voronoi diagrams, and obtain optimal algorithms for their construction.

These results are worst-case optimal but, when the set of P is well distributed, better results can be obtained. This is shown in Section 5.5 where we introduce the notion of net of a bounded domain.

In Section 5.4.1, we introduce the notion of witness complex, a weak version of the Delaunay triangulation that can be defined and constructed in any finite metric spaces. In addition to enlarging the set of spaces where Delaunay-like complexes can be defined, witness complexes can be computed using only distance comparisons. This is a decisive advantage over Delaunay triangulations in high dimensions. We will see in Chapter **??** a perturbation scheme that can be used to make the witness complex of the perturbed point set identical to the non-relaxed one.

We lastly show that Voronoi diagrams and Delaunay triangulations can be used to search nearest neighbours in finite point sets, hence complementing the results of Chapter 2.

5.1 Voronoi diagrams

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

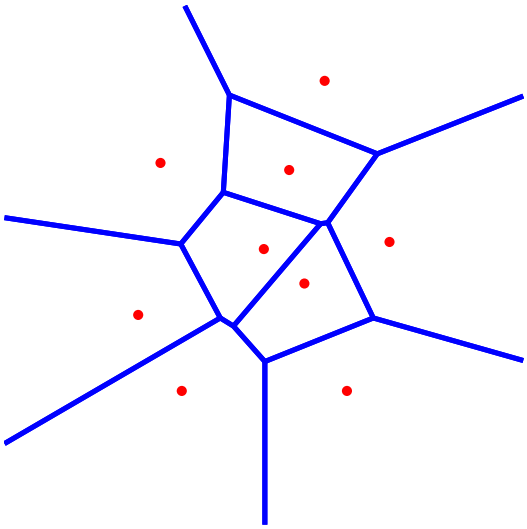

Figure 5.1: The Voronoi diagram of a set of 9 points.

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

Observe that $V(p_i)$ contains p_i and therefore is not empty. The Voronoi regions of P have disjoint interiors and, since any point of $\mathbb{R}^d$ belongs to at least one Voronoi region, they cover the entire space $\mathbb{R}^d$. The region $V(p_i)$ is the intersection of the $n-1$ half-spaces bounded by the bisector hyperplanes of p_i and each of the other points of P. $V(p_i)$ is therefore a convex polyhedron, polytope? possibly unbounded. The collection of the Voronoi regions and their faces constitute a *cell complex* called the Voronoi diagram of P.

We now establish a useful correspondence between Voronoi diagrams of $\mathbb{R}^d$ and a class of polyhedra of $\mathbb{R}^{d+1}$.

We associate to a point $x \in \mathbb{R}^d$ the so-called lifted point $\phi(x) = (x, x^2) \in \mathbb{R}^{d+1}$. Hence the lifting map is a bijection from $\mathbb{R}^d$ to the paraboloid of revolution $\mathcal{Q} = \{(x, x^2), x \in \mathbb{R}^d\}$. The hyperplane h_{p_i} tangent to $\mathcal{Q}$ at $\phi(p_i)$ is defined as

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

It follows that $(x - p_i)^2 = x^2 - 2p_i \cdot x + {p_i}^2$ is equal to the vertical distance

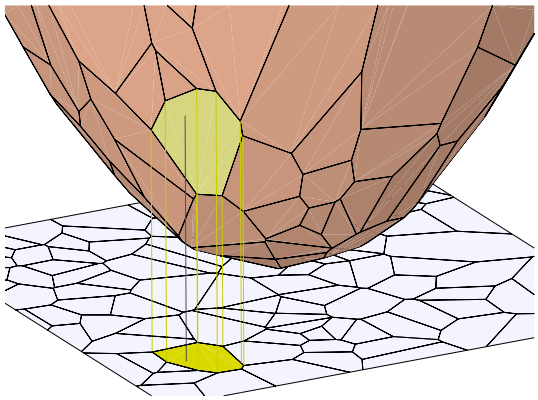

Figure 5.2: The polyhedron $\mathcal{V}(P)$, with one of its faces projected onto $\mathbb{R}^d$.

between $\phi(x)$ and h_{p_i} (see Fig. **??**).

notation : x^2

Let $h^+_{p_i}$ denote the half-space lying above h_{p_i} and define the polyhedron

$$\mathcal{V}(P) = h^+_{p_1} \cap \cdots \cap h^+_{p_n}.$$

(See Fig. 5.2). It follows from the above discussion that *the faces of the Voronoi diagram* $\mathrm{Vor}(P)$ *of* P *are the vertical projections of the faces of the convex polyhedron* $\mathcal{V}(P)$.

We then 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. In particular, the Voronoi diagram of n points of $\mathbb{R}^3$ may be quadratic (see Exercise 5.6).

5.2 Delaunay complexes

in ch. 4 : A *triangulation* of a finite set of points P of $\mathbb{R}^d$ is a simplicial complex embedded in $\mathbb{R}^d$ that covers the convex hull of P.

Let P be a finite set of points in $\mathbb{R}^d$ and $\mathrm{Vor}(P)$ its Voronoi diagram. The nerve of the collection of all Voronoi faces of $\mathrm{Vor}(P)$ is an abstract simplicial

complex called the *Delaunay complex* of P. Specifically, let f be a face of dimension k of $\mathrm{Vor}(P)$ and $\mathrm{int}f$ be the set of points of f that do not belong to any proper subface of f. All the points of $\mathrm{int}f$ have the same subset P_f of closest points in P, and f is the intersection of the Voronoi cells of the points in P_f. Accordingly, there is a simplex f^* in the Delaunay complex.

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 triangulation.

Lemma 5.1 (The empty ball property) *Any subset $\sigma \subset P$ is a simplex of the Delaunay complex of P iff it has a circumscribing 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 $m - 1 > d$.

However, the Delaunay complex can be embedded in $\mathbb{R}^d$ when the points of P are in *general position*. This is Delaunay's fundamental result.

Definition 5.2 (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.*

canonical devient natural. Faire une definition If P is in general position wrt spheres, the canonical embedding that maps a vertex of the Delaunay complex to its corresponding point of P and a simplex to the convex hull of its vertices embeds the Delaunay complex in $\mathbb{R}^d$. Moreover, the union of its simplices covers the convex hull of P.

Theorem 5.3 (Delaunay triangulation) *If a finite set of points $P \in \mathbb{R}^d$ is in general position wrt spheres, the canonical embedding in $\mathbb{R}^d$ of its Delaunay complex is a triangulation of P.*

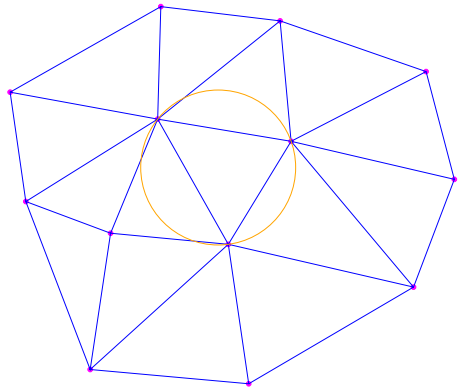

Figure 5.3: The empty ball property.

Proof We introduce a mapping between hyperspheres of $\mathbb{R}^d$ and points in $\mathbb{R}^{d+1}$.

We associate to the hypersphere σ of $\mathbb{R}^d$ of equation $\sigma(x) = x^2 - 2c \cdot x + s = 0$ the point $\phi(\sigma) = (c, s)$ of $\mathbb{R}^{d+1}$. Moreover, we associate to this point $\phi(\sigma)$ the so-called *polar hyperplane* h_σ of $\mathbb{R}^{d+1}$ of equation $x_{d+1} = 2c \cdot x - s$. Observe that if σ is reduced to a point c, h_σ is identical to the hyperplane h_c that has been introduced in Section 5.1. It should also be observed that the intersection of h_σ with the paraboloid $\mathcal{Q}$ projects vertically onto σ (see Figure 5.5). We deduce the remarkable following property : $x \in \sigma$ if and only if $\phi(x) = (x, x^2) \in h_\sigma$ and σ encloses x if and only if $\phi(x)$ is below h_σ.

Consider now a set $P = \{p_1, \ldots, p_n\}$ of n points in general position wrt spheres. It follows from the discussion above that any simplex with an empty circumscribing ball is mapped by ϕ to a face of the lower convex hull of $\phi(P)$. The lower convex hull is the subcomplex of the convex hull whose faces are supported by hyperplanes lying below $\phi(P)$.

Write $\mathcal{D}(P)$ and $\mathcal{D}'(P)$ respectively for the convex hull and the lower convex hull of the points $\phi(p_i)$, $i = 1, \ldots, n$. Under the general position assumption, $\mathcal{D}'(P)$ is a simplicial complex embedded in $\mathbb{R}^{d+1}$ and, from the discussion and Lemma 5.1, its projection onto $\mathbb{R}^d$ coincides with the Delaunay complex of P. Because the projection is 1-1, this projected cell complex is properly embedded in $\mathbb{R}^d$ and, since the projection preserves convexity, it covers the convex hull of P. This concludes the proof of Delaunay's theorem. □

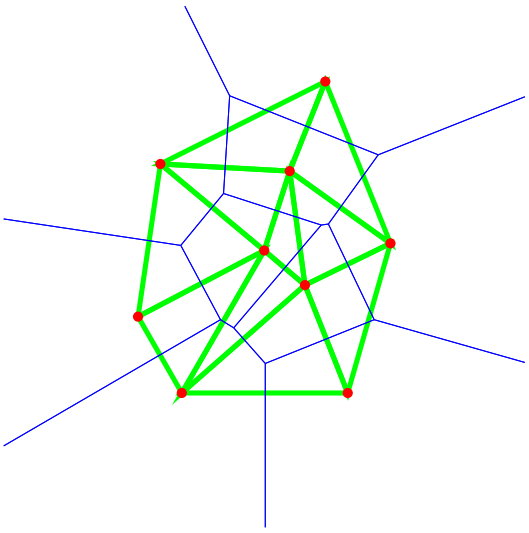

Figure 5.4: The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines).

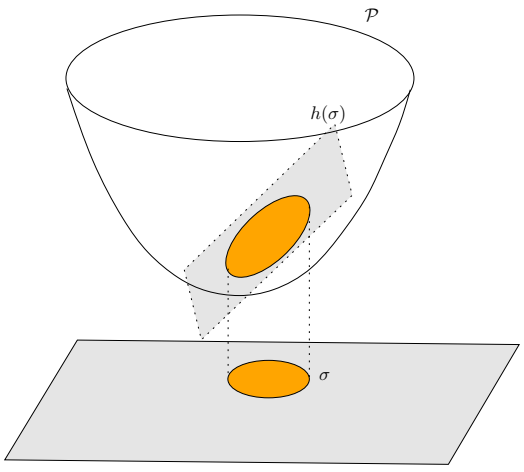

Figure 5.5: The polar hyperplane of a hypersphere.

We can use the polarity between polytopes introduced in Section 3.2. Let $\mathcal{V}(P)$ denote, as in Section 5.1, the polyhedron defined as the intersection of the n halfspaces above the n polar hyperplanes h_{p_i}, and let $\mathcal{V}'(P)$ denote the boundary complex of $\mathcal{V}(P)$. By the results of Section 3.2, $\mathcal{V}'(P)$ and $\mathcal{D}'(P)$ are dual polytopes. Hence, we have established a dual correspondence between $\mathrm{Vor}(P)$ and $\mathrm{Del}(P)$. We sum up our construction in the following commutative diagram that commutes when the point set P is in general position with respect to spheres.

$$\begin{array}{ccc} \mathcal{V}'(P) = \mathrm{bd}\mathcal{V}(P) & \stackrel{\text{duality}}{\longrightarrow} & \mathcal{D}'(P) = \mathrm{lowerhull}(\phi(P)) \\ \uparrow & & \downarrow \\ \text{Voronoi Diagram } \mathrm{Vor}(P) & \stackrel{\text{nerve}}{\longrightarrow} & \text{Delaunay Triangulation } \mathrm{Del}(P) \end{array}$$

If the points of P are not in general position wrt spheres, some simplices may be of dimension greater than the dimension d of the embedding space and the Delaunay complex cannot be embedded in $\mathbb{R}^d$. We define the *shadow* of a simplex σ of the Delaunay complex as the convex hull of of its vertices in the canonical embedding. If the points are in general position wrt spheres, it follows from Theorem 5.3 that the shadow of each simplex of the Delaunay complex is a simplex of $\mathbb{R}^d$. If the points are not in general position wrt spheres, the shadow of some simplices of the Delaunay complex may not be simplices but more general polytopes. Nevertheless, it is easy to see that the above diagram still holds. The only difference is that the complexes $\mathcal{D}'(P)$ and $\mathrm{Del}(P)$ are no longer simplicial. An embedded triangulation can then be obtained by triangulating the shadows that are not simplices (see Exercise 5.10). Any such triangulation is called *a* Delaunay triangulation. Since there are several ways of triangulating the faces of a polytope, the Delaunay triangulation of P is no longer unique.

It follows from the duality established in the previous section that the combinatorial complexity of the Delaunay triangulation of n points is the same as the combinatorial complexity of its dual Voronoi diagram. Moreover, the Delaunay triangulation of n points of $\mathbb{R}^d$ can be deduced from the dual Voronoi diagram or vice versa in time proportional to its size. We also deduce from what precedes that computing the Delaunay triangulation 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 known results on convex hulls (see Theorems 3.7 and 3.9).

Theorem 5.4 *The combinatorial complexity of the Voronoi diagram of n points of $\mathbb{R}^d$ and of their Delaunay triangulation 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 particular, the Voronoi diagram of n points of $\mathbb{R}^3$ may be quadratic (see Exercise 5.6). These bounds are worst-case bounds. Under some assumptions on the point distribution, better bounds can be obtained (see Exercise 5.8).

5.3 Nets

In this section, consider a bounded domain $\Omega \subset \mathbb{R}^d$ and a finite point set P that samples Ω. We introduce a notion of quality of a sample which is captured through the concept of a net and show that nets have Delaunay triangulations of linear size that can be computed efficiently in any fixed dimension.

Let P be a finite set of points in. 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-sample* for Ω.

We further call $\eta = \min_{p,q \in P} \|p - q\|$ the *separation* of P and $\bar{\eta} = \eta/\varepsilon$ the *sparsity* of P.

A finite point set P of Ω whose sparsity is lower bounded by a positive constant is called a *net*, or more precisely an *$(\varepsilon, \bar{\eta})$-net* if one wants to specify the sampling radius and the sparsity of P. 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.

It is easy to see that $\bar{\eta} \leq 2$ and also that $(\varepsilon, 1)$-nets exists (see Exercise 5.12).

We are interested in computing the Delaunay complex of a finite vertex set $P \subset \mathbb{R}^d$. If P is an ε-sample for Ω, then every simplex in $\mathrm{Del}(P)$ that has a Delaunay centre in Ω must have a circumradius not greater than than ε. The subcomplex of all such simplices is called the *restriction* of $\mathrm{Del}(P)$ to Ω and denoted $\mathrm{Del}_{|\Omega}(P)$. Observe that all Delaunay simplices with a vertex greater than 2ε from the boundary of Ω belong to $\mathrm{Del}_{|\Omega}(P)$ (Exercise 5.13).

In the sequel, we will avoid boundary considerations and restrict our attention to the restricted Delaunay triangulation $\mathrm{Del}_{|\Omega}(P)$.

Theorem 5.5 *Let Ω be a bounded domain of $\mathbb{R}^d$ and P a net for Ω. The restriction of the Delaunay triangulation of P to Ω has linear size $O(n)$ where $n = |P|$. Moreover, $\mathrm{Del}_{|\Omega}(P)$ can be computed in time $O(n \log n)$. The constants in the O depends exponentially on d.*

Proof Assume without loss of generality that P is an $(\varepsilon, \bar{\eta})$-net. Let p be a point of P and let σ be a simplex of $\mathrm{Del}_{|\Omega}(P)$ in the star of p. The diameter of σ is at most 2ε. Indeed, since its circumcenter c is in Ω and P is an ε-sample of Ω, we have

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

for any $q \in \sigma$. Moreover, all the open balls $B(x, \frac{\eta}{2})$, $x \in P$, are disjoint by definition of the sparsity. Hence, the number of neighbours of p in $\mathrm{Del}_{|\Omega}(P)$ is at most

$$N = \frac{\mathrm{vol}\left(B(\varepsilon + \frac{\bar{\eta}\varepsilon}{2})\right)}{\mathrm{vol}\left(B(\frac{\bar{\eta}\varepsilon}{2})\right)} = \left(1 + \frac{2}{\bar{\eta}}\right)^d,$$

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

A simple algorithm can be easily deduced. We first compute for all points p in P, the subset $N(p)$ of points that lie at distance at most 2ε from p for a total cost of $O(n \log n)$ (See Section **??**). Then, we compute for each point p the star of p in $\mathrm{Del}_{|\Omega}(N(p))$, which can be done in $O(n)$ time since $|N(p)| = O(1)$. It then remains to glue all the stars. □

5.4 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 a finite set of points and the distances between any two points. 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 is

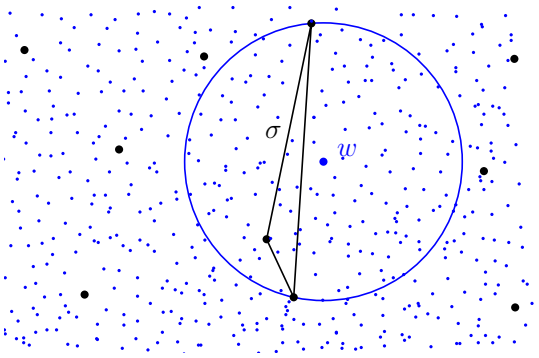

Figure 5.6: A simplex and one of its witnesses.

when the points come from some Euclidean space but their actual positions are not known.

The witness complex is defined from two sets of points L and W. The first one, called the set of *landmarks*, is finite. The other one, called the set of *witnesses*, serves an an approximation of the ambient space. The witness complex $\mathrm{Wit}(L, W)$ can be seen as a weak notion of Delaunay triangulation $\mathrm{Del}(L)$ which is easy to compute, even in high dimensions, since it only involves comparisons of distances between input points.

We will take $\mathbb{T}^d$ as our working space. CHANGE? We will provide a sufficient condition under which $\mathrm{Wit}(L, W) = \mathrm{Del}(L)$. In Exercise **??**, we will show that the input point set L can be perturbed so that the conditions are fulfilled, hence resulting in a witness complex of a perturbed point set L' that is identical to its exact restricted Delaunay triangulation $\mathrm{Del}(L')$.

We always assume that L contains a non-degenerate d-simplex, and we demand that $\lambda < 1/4$ so as to ensure that the Delaunay complex will be a simplicial complex.

Definition 5.6 (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 5.7 (Witness complex) *The witness complex $\mathrm{Wit}(L, W)$ is*

the complex consisting of all simplexes σ such that for any simplex $\tau \subseteq \sigma$, τ has a witness in W.

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 $\mathrm{Wit}(L, W)$ 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, on or outside the sphere circumscribing a d-simplex, whose degree depends on d (see Exercise 5.9).

5.4.1 Identity of witness and Delaunay complexes

The witness complex can be seen as a weak Delaunay complex. The results below make this connection more precise.

Lemma 5.8 *If $W' \subseteq W$, then* $\mathrm{Wit}(L, W') \subseteq \mathrm{Wit}(L, W)$.

Lemma 5.9 $\mathrm{Del}(L) \subseteq \mathrm{Wit}(L, \mathbb{T}^d)$.

Proof Any simplex σ of $\mathrm{Del}(L)$ has an empty circumscribing ball whose center c is a witness of σ, and c is also a witness for all the faces of σ. □

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 $\mathbb{T}^d$.

Theorem 5.10 (Weak characterization) $\mathrm{Wit}(L, \mathbb{T}^d) = \mathrm{Del}(L)$.

Proof We have already proved that $\mathrm{Del}(L) \subseteq \mathrm{Wit}(L, \mathbb{T}^d)$. We prove now the converse inclusion. Let $\tau = [p_0, ..., p_k]$ be a k-simplex of $\mathrm{Wit}(L)$ witnessed by a ball B_τ, i.e. $B_\tau \cap L = \tau$. Write S_τ for the sphere bounding B_τ. We prove that $\tau \in \mathrm{Del}(L)$ by a double induction on k and $l := |S_\tau \cap \tau|$. The claim holds for $k = 0$ since a witness for a vertex is also a Delaunay center,

i.e. the center of an emptyball passing through the vertex. We also note that $\tau \in \mathrm{Del}(L)$ when $l = k+1$ since B_τ then circumscribes τ. We can also assume that $|S_\tau \cap \tau| \geq 1$ since we can reduce the radius of B_τ so that it still witnesses τ and S_τ contains at least a vertex of τ. Hence the claim holds for $k = 0$ and $l = 1$. Assume that the claim holds for all simplices of $\mathrm{Wit}(L, W)$ of dimension up to $k-1$ and for all $l \leq k$, and refer to the figure.

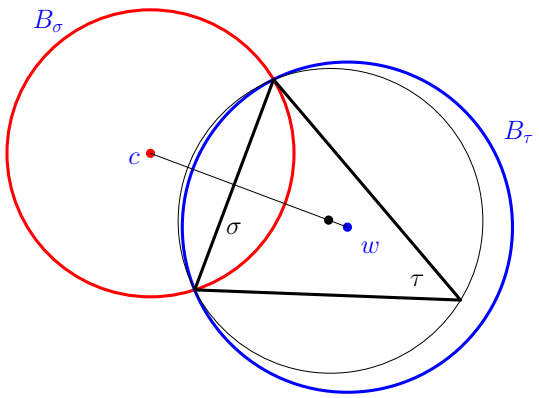

Figure 5.7: Proof of Lemma 5.10.

We will show that one can find a new witness ball of τ, we still call B_τ, such that the number of points of $S_\tau \cap \tau$ gets increased. Write $\sigma = S_\tau \cap \tau$ and w for the center of B_τ. By the induction hypothesis, σ is a Delaunay simplex and therefore there exists a Delaunay ball B_σ circumscribing σ. Write c for its center. Consider the set of balls F centered on the line segment $s = [wc]$ and circumscribing σ. Any ball in F is included in $B_\tau \cup B_\sigma$ and its bounding sphere 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 on s such the associated ball of F witnesses τ and contains $l+1$ points of τ on its boundary. Then, by carrying on the induction, we obtain a witness ball that contains all the vertices of τ and thus is a Delaunay ball of τ. □

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 5.8, this is mandatory for the theorem to hold.

We deduce from Lemma 5.8 and Theorem 5.10 the following corollary

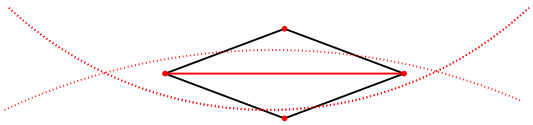

Figure 5.8: Two triangles that have a witness but not their common edge, even if $W = \mathbb{R}^d$. These two triangles are not Delaunay triangles.

Corollary 5.11 $\mathrm{Wit}(L, W) \subseteq \mathrm{Del}(L)$.

If the points L are in general position with respect to spheres, we know that $\mathrm{Del}(L)$ is embedded in $\mathbb{R}^d$ by Delaunay's theorem **??**. It therefore follows from Corollary 5.11 that the same is true for $\mathrm{Wit}(L, W)$. In particular, the dimension of $\mathrm{Wit}(L, W)$ is at most d. We also deduce the following lemma.

Lemma 5.12 (Locality) *If L is a λ-sample of $\mathbb{T}^d$, the circumradius R_σ of any simplex σ in $\mathrm{Wit}(L, W)$ is at most λ.*

Proof By Corollary 5.11, σ is a simplex of $\mathrm{Del}(L)$ and, since L is a λ-sample of $\mathbb{T}^d$, the circumradius R_σ is at most λ. □

When W is not the whole space $\mathbb{T}^d$ but a finite set of points, Theorem 5.10 no longer holds. However, the following lemma provides an easy-to-check condition under which the witness complex and the Delaunay complex are identical. The lemma follows from Corollary 5.11 and Lemmas 4.15 and **??**. Observe that the lemma does not assume L to be in general position. In the sequel, we say that the link of a vertex in a complex is *good* if it is a pseudomanifold **??**.

Lemma 5.13 (Identity from good links) *If L is in general position with respect to spheres and the vertices of $\mathrm{Wit}(L, W)$ have good links, then $\mathrm{Wit}(L, W) = \mathrm{Del}(L)$.*

5.4.2 Computing witness complexes

We first show how to compute $\mathrm{Wit}(L, W)$ where L and W are finite point sets.

Algorithm 3 Construction of witness complexes

Input: L, W, $\|p-w\|$ for all $p \in L, w \in W$
$\mathrm{Wit}(L,W) := \emptyset$
$W' := W$
for each $w \in W'$ **do**
  compute the list $N(w) = (p_0(w), ..., p_{|L|-1}(w))$ of the points of L sorted by distance from w
for $i = 0, ..., |L|-1$ **do**
  for each $w \in W'$ **do**
    if the i-simplex $\sigma(w) = (p_0(w), p_2(w), ..., p_i(w)) \notin \mathrm{Wit}(L,W)$ **then**
      if the $(i-1)$-faces of $\sigma(w)$ are in $\mathrm{Wit}(L,W)$ **then**
        add to $\mathrm{Wit}(L,W)$
      else
        $W' := W' \setminus \{w\}$
Output: $\mathrm{Wit}(L,W)$

If the points lie in $\mathbb{R}^d$ and are in general position with respect to spheres, we know that $\mathrm{Wit}(L,W)$ is embedded in $\mathbb{R}^d$ and therefore of dimension d. This will be sufficient in the next section. We can theferore only consider the first $d+1$ neighbors of each witness. The complexity of the algorithm is then $O(|W| \times |L| \log |L|) + (|\mathrm{Wit}(L,W)| + |W|) \log |L|$. The first term is for the first **for** loop. The second is for the second **for** loop and comes from the fact that we either construct a new simplex or remove a point from the set of active witnesses W', and at each step we can decide if a simplex has already been constructed in time $O(\log |\mathrm{Wit}(L,W)|) = O(\log |L|)$.

We can improve the complexity of the first **for** loop of the algorithm if L is a $(\lambda, \bar{\mu})$-net for a convex bounded domain Ω and W is an ε-sample of Ω. As observed above (Lemma 5.12), the diameter of any simplex in $\mathrm{Wit}(L,W)$ is at most 2λ. Hence the elements of $N(w)$, for any $w \in W$, have only to be searched in the ball $B(w, 2\lambda)$. Since the sparsity ratio of L is $\bar{\mu}$, we have

$$|L \cap B(w, 2\lambda)| \leq \frac{\mathrm{vol}(B(\lambda(2+\frac{\bar{\mu}}{2})))}{\mathrm{vol}(B(\lambda \frac{\bar{\mu}}{2}))} = \left(1 + \frac{4}{\bar{\mu}}\right)^d.$$

Note that the algorithm only requires to compare (squared) distances between points.

update witness complexes

5.5 Nearest neighbor

Nearest neighbor queries are defined in Chapter 2. If we assume that the distance under consideration is the Euclidean distance, the Voronoi diagram $\mathrm{Vor}(P)$ appears as the natural geometric structure to answer nearest neighbor queries concerning the point set P. Indeed, given a query point q, finding the nearest neighbor for q in P just amounts to localize q in $\mathrm{Vor}(P)$ which means finding the cell in $\mathrm{Vor}(P)$ that contains q. The rest of this section shows that, despite the exponential combinatorial bound on the size of Voronoi diagrams and Delaunay triangulations, those concepts lead to a data structure, the 1-skeleton Delaunay hierarchy, that provides in most practical situations a logarithmic nearest neighbor query time $O(c^d \log n)$ with a linear memory footprint $O(c^d n)$, where c is a constant independant of the ambiant space dimension.

5.5.1 Walking in the Voronoi diagram

Since the nearest neighbor of q in P is the site of the cell of $\mathrm{Vor}(P)$ containing q, we first handle the problem of point location in a Voronoi diagram. An easy way to perform such a point location consist in walking in the Voronoi diagram. Assume we want to locate a query point p in the Voronoi diagram $\mathrm{Vor}(P)$. As described in Algorithm 4 below (see also Figure **??**, the walk starts from any cell $V(p)$ of the Voronoi diagram $\mathrm{Vor}(P)$. At the current step, if the walk has currently reached the cell $V(p)$ of p, the algorithm scans the points adjacent to p in $\mathrm{Vor}(P)$, i.e. the points in P whose Voronoi cells share a $(d-1)$-face with $V(p)$. If one of them, say $n(p)$, is closer to q than p, the walk steps into the cell $V(n(p))$. Otherwise, the walk is ended and point p is returned.

Algorithm 4 Find the nearest neighbor of q in set P by walking in the Voronoi diagram $\mathrm{Vor}(P)$.

Start from any cell $V(p)$ of $\mathrm{Vor}(P)$.
while there is a neighbor $n(p)$ of p in $\mathrm{Vor}(P)$, closer to q than p **do**
  $p = n(p)$
return p.

To prove that the returned point is actually the closest to q in P, we consider the subset P' of P including p and the points adjacent to p in $\mathrm{Vor}(P)$. The

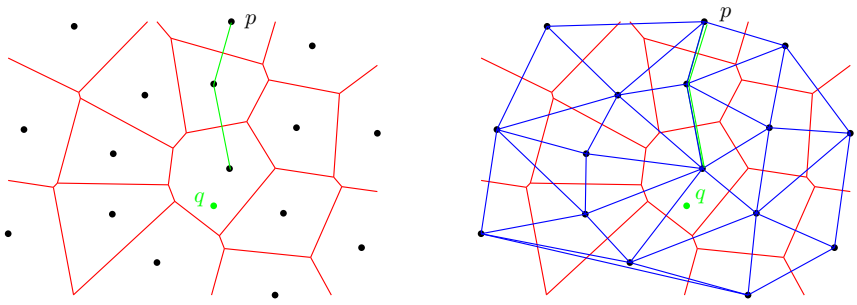

Figure 5.9: Right: walking the Voronoi diagram $\mathrm{Vor}(P)$ to find the nearest neighbor of q in set P. Left: the walk follows the edges of the Delaunay triangulation $\mathrm{Del}(P)$.

algorithm ensures that the query point q belongs to the Voronoi cell $V(p, P')$ of p in the diagram $\mathrm{Vor}(P')$. Furthermore, since the shape of the Voronoi cell $V(p, P)$ is completely determined by the subset P' of P, the Voronoi cells $V(p, P)$ and $V(p, P')$ coincide, which implies that the query point q belongs to $V(p, P)$.

5.5.2 The Delaunay hierarchy

Walking in the Voronoi diagram is a nice and straightforward strategy to locate a query point. However it is rather inefficient since, in the worst case, the walk might have to visit all the Voronoi cells or at least a large fraction of them. Fortunately, this drawback can be avoided by using a date structure called the Delaunay hierarchy.

First, let us notice, that walking in the Voronoi diagram implies only steps between adjacent Voronoi cells, i.e. Voronoi cells sharing a $(d-1)$-face. A pair of adjacent Voronoi cells corresponds to a pair of vertices in the dual Delaunay triangulation that are the endpoints of an edge of the triangulation. Therefore, the walk in the Voronoi diagram is easily implemented using the dual Delaunay triangulation. Furthermore, the walk only needs the 1-skeleton, i.e. the edges of the Delaunay triangulation.

Let P be a set of points. A Delaunay hierarchy is a sequence of Delaunay triangulations computed for a nested sequence of point sets:

$$P_0 = P \supset P_1 \ldots \supset P_m.$$

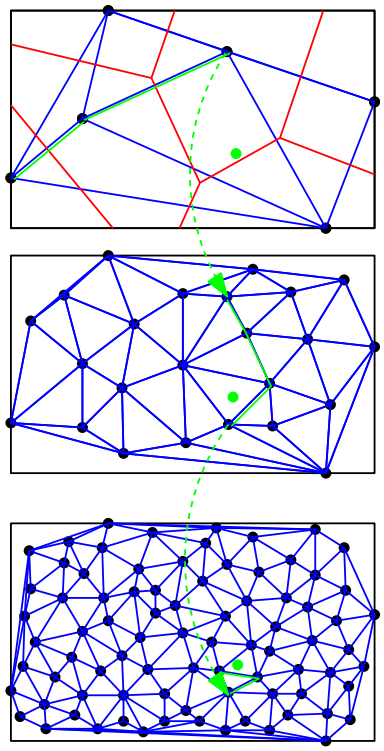

Figure 5.10: The Delaunay hierarchy

See Figure **??**. The first set P_0 is just the whole set P. Each point $p \in P_l$ is included in the next level subset P_{l+1} with probability $\beta = 1/\alpha$. The Delaunay hierarchy includes the Delaunay triangulation of each subset P_l. Furthermore, each vertex in the Delaunay triangulation $\mathrm{Del}(P_{l+1})$ at level $l+1$ has a pointer to the vertex of the previous level triangulation $\mathrm{Del}(P_l)$ corresponding to the same point.

To locate the query point q in the Voronoi diagram $\mathrm{Vor}(P)$, we proceed by decreasing levels of the hierarchy. The query point q is first located in the Voronoi diagram $\mathrm{Vor}(P_m)$ with the highest level by walking in the Voronoi diagram (Algorithm 4). We explain now how to proceed from level $l+1$ to level l. At level $l+1$, we have located q and therefore know the point $n_{l+1}(q) \in P_{l+1}$ closest to q. To locate q in $\mathrm{Vor}(P_l)$ we first use the pointer from $n_{l+1}(q) \in P_{l+1}$ to $n_{l+1}(q) \in P_l$ and, still using Algorithm 4, we walk in level l from $n_{l+1}(q)$ to $n_l(q)$, the point closest to q in P_l.

The Delaunay hierarchy is a randomized data structure since every subset P_{l+1} is a random subset of the previous level set P_l. The efficiency of this data structure comes from the following lemma.

Lemma 5.14 *When locating a query point in a Voronoi diagram, using the walking Voronoi strategy at each level of a Delaunay hierarchy, the expected number of steps performed at each level is constant.*

Proof Let m_l be the number of steps performed at level l of the hierarchy for the location of the query point q, and let $\exp(m_l)$ be the expectation of m_l. Since at each step the distance to the query point q decreases, m_l is at most k, if $n_{l+1}(q)$ is the k-th nearest vertex to q in P_l. Such a situation occurs with propability $\beta(1-\beta)^{k-1}$. Indeed, $n_{l+1}(q)$ is the k-th nearest vertex to q in P_l iff $n_{l+1}(q)$ has been inserted at level $l+1$ while all the $k-1$ vertices of P_l closest to q have not. Hence,

$$\begin{aligned} \exp(m_l) &\leq \sum_{k=1}^{n_l} k(1-\beta)^{k-1}\beta \\ &\leq \beta\big[-\frac{\partial}{\partial\beta}\sum_k (1-\beta)^k\big] = \frac{1}{\beta} \end{aligned}$$

□

In other words, if $1/\alpha$ points of each level are selected and appear in the next level, the walk performs on average at most α steps at each level. Since the number of levels is $O(\log n)$, the total number of steps performed by the algorithm is also $O(\log n)$. However, one should not conclude that the location of a query point in $\mathrm{Vor}(P)$ is performed in $O(\log n)$ time. Indeed, at each step, the algorithm may have to look at all the cells adjacent to the current cell in the Voronoi diagram, which might $n-1$ in number. However, as shown in the next section, this worst-case situation does not reflect cases of practical importance.

Another drawback of the walking strategy is the space required to store the data structure. Notice however that one may only store the 1-skeleton of each triangulation in the Delaunay hierarchy, leading to $O(n^2)$ storage, where the constant in the O is an absolute constant that, in particular, does not depend on the dimension of the ambient space.

5.5.3 Nearest neighbor queries in practical cases

In many practical cases, the hierarchy of 1-skeletons of Delaunay triangulations behaves much better than in the worst-case. This happens when the number of Voronoi cells that are incident to a given cell is bounded by a constant or, equivalently, when the degree od each node in the corresponding Delaunay triangulation is bounded. Such a situation arises for any point set whose *spread* is bounded. The spread of a point set is defined as the ratio of the diameter of the point set and the shortest distance between any two points. As shown in Exercise 5.16, a more general situation arises also when P is an (ϵ, δ)-sample of some bounded domain D and the ratio ϵ/δ is bounded. P is said to be an (ϵ, δ)-sample of D if there is a Lifschitz field $\mathrm{lfs}(x)$ defined on D such that any point x of D is at distance at most $\epsilon\mathrm{lfs}(x)$ from P and if any point p of P is at distance at least $\delta\mathrm{lfs}(p)$ from $P \setminus \{p\}$.

5.6 Exercises

Exercise 5.1 Show that the combinatorial complexity of the upper envelope of n univariate functions whose graphs intersect pairwise in at most two points is $O(n)$. Show that the envelope can be computed in optimal time $\Theta(n \log n)$.

Exercise 5.2 Let $\mathcal{Q}$ be the paraboloid of $\mathbb{R}^{d+1}$ of equation $x_{d+1} = x^2$. Show that the hyperplane h_{p_i} is tangent to $\mathcal{Q}$ at the point (p_i, p_i^2).

Exercise 5.3 To a hypersphere σ of center c and radius r of equation

$$\sigma(x) = (x - c)^2 - r^2 = x^2 - 2c \cdot x + s = 0,$$

where $s = \sigma(0) = c^2 - r^2$, we associate the point of $\mathbb{R}^{d+1}$ $\phi(\sigma) = (c, s)$. Show that the image by ϕ of a point, considered as a hypersphere of radius 0, is a point of the paraboloid $\mathcal{Q}$.

Show that the image by ϕ of the hyperspheres that pass through a given p of $\mathbb{R}^d$ is the *hyperplane* h_p of $\mathbb{R}^{d+1}$

$$x_{d+1} = 2p \cdot x - p^2.$$

Exercise 5.4 What are the preimages by ϕ of the points of $\mathbb{R}^{d+1}$ that lie

1. above $\mathcal{Q}$?
2. on the boundary of $\mathcal{V}(P)$? in the interior of $\mathcal{V}(P)$?
3. on a line ?

Exercise 5.5 Consider the diagram obtained by projecting the faces of $h^-_{p_1} \cap \cdots \cap h^-_{p_n}$ vertically. Characterize the points that belong to a face of this diagram.

Dualy, project vertically the faces of the *upper* convex hull of the $\phi(p_i)$. Show that we obtain a triangulation of the vertices of $\mathrm{conv}(P)$ such that each ball circumscribing a simplex contains all the points of P.

Exercise 5.6 Show that if we take points on two non coplanar lines of $\mathbb{R}^3$, say n_1+1 on one of the lines and n_2+1 on the other, their Delaunay triangulation has $n_1 n_2$ tetrahedra (or, equivalently, that their Voronoi diagram has $n_1 n_2$ vertices).

Exercise 5.7 Show, using Euler's formula, that *any* triangulation of a set P of n points in the plane has $2n-2-n_e$ triangles if n_e is the number of points of P that are vertices of $\mathrm{conv}(P)$. Show, using also Euler's formula, that, for any triangulation $T(P)$ of a set P of points of $\mathbb{R}^3$, we have $t < e-n$ where e and t are respectively the number of edges and tetrahedra of $T(P)$. Prove that any set of n points of $\mathbb{R}^3$ with no 3 points on a same line admits a triangulation of size $O(n)$.

Exercise 5.8 Let D be a bounded domain of $\mathbb{R}^d$ and let P be an ε-net of D, i.e. a set of points of D such that any point of D is at distance less than ε from a point of P and such that no two points of P are at distance less than $\kappa\,\varepsilon$, where κ is some positive constant. Denote by p a point of D lying at distance at least ε from the boundary of D. Show that the circumradius of any Delaunay simplex incident to p is less than ε. Deduce that all the Delaunay simplices incident to p lie in the ball $B(p, 2\varepsilon)$. Show next, using a packing argument, that the number of Delaunay edges incident to p is a constant that depends exponentially on d. Bound the total complexity of the star of p.

Exercise 5.9 Let S be a hypersphere of $\mathbb{R}^d$ passing through $d+1$ points $p_0, \ldots, p_d$. Show that a point p_{d+1} of $\mathbb{R}^d$ lies on S, in the interior of the ball B_S bounded by S or outside B_S, depending whether the determinant of the $(d+2) \times (d+2)$ matrix

$$\texttt{in\_sphere}(p_0, \ldots, p_{d+1}) = \begin{vmatrix} 1 & \cdots & 1 \\ p_0 & \cdots & p_{d+1} \\ p_0^2 & \cdots & p_{d+1}^2 \end{vmatrix}$$

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

Exercise 5.10 Describe an algorithm to triangulate a bounded polytope (Hint: proceed by faces of increasing dimensions).

Exercise 5.11 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 $\mathrm{MST}(P) \subset \mathrm{Del}(P)$.

Exercise 5.12 (ε-nets) Let D a bounded domain of $\mathbb{R}^d$. Show that the output set P of the following greedy algorithm an $(\varepsilon, 1)$-net for D. Pick as a first point any point of D. Then, as long as it is possible, add to P a point of $D \setminus B$, where B is the union of the balls of radius $\varepsilon/2$ centered at the points of the current set P.

Exercise 5.13 (Restricted Delaunay triangulation) Let P and D be as in Exercise 5.12. Show that if a simplex $\sigma \in \mathrm{Del}(P)$ has a vertex at distance greater that 2ε from the boundary of D, it belongs to $\mathrm{Del}_{|D}(P)$, the restriction of $\mathrm{Del}(P)$ to D.

Exercise 5.14 (Delaunay complex contains a triangulation) For any finite set of points $P \in \mathbb{R}^d$, there is a subcomplex $K \subseteq \mathrm{Del}(P)$ that is a triangulation of P.

Exercise 5.15 (Natural coordinates) Let E be a finite set of points $p_1, ..., p_n$ of $\mathbb{R}^d$. As usual, $\mathrm{Vor}(E)$ denotes the Voronoi diagram of E and $V(p_i)$ the cell of p_i in $\mathrm{Vor}(E)$. Given a point $x \in \mathrm{conv}(E)$, we write

$E^+ = E \cup \{x\}$, $V^+(x)$ for the Voronoi cell of x in $\mathrm{Vor}(E^+)$, $V^+(x, p_i) = V^+(x) \cap V^+(p_i)$ and $W(x, p_i) = V^+(x) \cap V(p_i) = V(p_i) \setminus V^+(p_i)$. Now we define $v_i(x) = \mathrm{vol}(V^+(x, p_i))$, $\bar{v}_i(x) = v_i(x)/\|x - p_i\|$ and $\bar{v}(x) = \sum_{i=1}^n \bar{v}_i(x)$. In addition, we define $w_i(x) = \mathrm{vol}(W(x, p_i))$ and $w(x) = \sum_{i=1}^n w_i(x)$.

We call Laplace coordinates the n functions $\lambda_1, \ldots, \lambda_n$ defined by $\lambda_i(x) = \bar{v}_i(x)/\bar{v}(x)$ for $x \notin E$, and $\lambda_i(p_j) = \delta_{ij}$ otherwise, where δ_{ij} is the Kronecker delta. We call Sibson's coordinates the n functions $\varsigma_i(x) = w_i(x)/w(x)$, $i = 1, \ldots, n$.

Show that the set of λ_i is a partition of unity. Same question for the set of ς_i. (Hint : for the Laplace coordinates, apply exercise 6.3 to $V^+(x)$. For the Sibson's coordinates, apply exercise 6.3 to the polytope $\mathcal{D}(E^+) \setminus \mathcal{D}(E)$ of $\mathbb{R}^{d+1}$, where $\mathcal{D}(E)$ is defined in Section 5.2).

Exercise 5.16 ((ϵ, δ)-samplings) (ϵ, δ)-samples are defined in Section 5.5.3. Show that if a set of points P is an (ϵ, δ)-sample of a bounded domain $D \subset \mathbb{R}^d$, then any Vornoi cell in $\mathrm{Vor}(P)$ has $O(c^d)$ incident Voronoi cells, where c depends on the ratio ε/δ and d.

Exercise 5.17 (Identity of the witness and the Delaunay complexes from protection) *Let W be an ε-sample for $\mathbb{T}^d$ L a $(\lambda, \bar{\mu})$-net, and let $p \in L$. If all the d-simplices in $\mathrm{star}^2(p, \mathrm{Del}(L))$ are δ-protected with $\delta \geq \frac{8d\varepsilon}{\bar{\mu}}$, then $\mathrm{star}(p; \mathrm{Wit}(L, W)) = \mathrm{star}(p; \mathrm{Del}(L))$.*

Proof By Corollary 5.11, we have $\mathrm{star}(p; \mathrm{Wit}(L, W)) \subseteq \mathrm{star}(p; \mathrm{Del}(L))$. We now prove the other inclusion. Let σ be a simplex in $\mathrm{star}(p; \mathrm{Del}(L))$ and set $\delta' = \frac{\bar{\mu}\delta}{4d} \geq 2\varepsilon$. Let $\tau \subseteq \sigma$ be a face of σ. Then, by Lemma **??**, τ is δ'-protected at a point z_τ such that

1. $\|z_\tau - p_i\| = \|z_\tau - p_j\| = r \quad \forall p_i, p_j \in \tau$
2. $\|z_\tau - p_l\| > r + \delta' \quad \forall p_l \in L \setminus \tau$

For any $x \in B(z_\tau, \delta'/2)$, any $p_i \in \tau$ and any $p_l \in L \setminus \tau$, we have

$$|x - p_i| \leq |z_\tau - p_i| + |z_\tau - x| \leq r + \frac{\delta'}{2}$$

and

$$|x - p_l| \geq |z_\tau - p_l| - |x - z_\tau| > r + \delta' - \frac{\delta'}{2} = r + \frac{\delta'}{2}$$

Hence, x is a witness of τ. Since $\varepsilon \leq \delta'/2$, there must be a point $w \in W$ in $B(z_\tau, \delta'/2)$ which witnesses τ. Since this is true for all faces $\tau \subseteq \sigma$, we have $\sigma \in \mathrm{star}(p; \mathrm{Wit}(L, W))$. $\square$

5.7 Bibliographical notes

To know more about the space of spheres, one may read the books by Pedoe [98] and Berger [10]. An entire book is devoted to Voronoi diagrams [96]. One may also look at the survey by Aurenhammer and Klein [8] and the part of the textbook by Boissonnat and Yvinec [23, 24] devoted to Voronoi diagrams. The result of Exercise 5.8 holds (in expectation) for random sets of points, as shown by Dwyer [56].

Witness complexes and relaxed Delaunay triangulations have been introduced in the seminal work of de Silva [54] who first proved Theorem 5.10. The proof presented in Section 5.4.1 follows the proof of Attali et al. [6]. The identity of witness and Delaunay complexes and the perturbation algorithm is taken from recent unpublished work by Boissonnat, Dyer and Ghosh.

The Delaunay hierarchy has been invented by O. Devillers [].

Natural coordinates (Exercise 5.15) have been introduced by Sibson [106, 105].

The CGAL library [**?**] is a good example of a rigorous and very efficient implementation of the exact paradigm in Computational Geometry. CGAL has, in particular, fully reliable and very efficient implementations of algorithms to construct Delaunay triangulations in low dimensions.

Chapter 6

Weighted Delaunay

A construction similar to what we did for the Euclidean Voronoi diagrams of points and their dual Delaunay triangulations can be done for the so-called *weighted Voronoi diagrams*. Here we take as our finite set of objects a set of weighted points (which can be considered as (hyper)spheres) and consider as distance function of a point x to a sphere σ the power of x to σ, i.e. the squared distance from x to the center of the sphere minus the squared radius of the sphere. 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. As the unweighted Voronoi diagrams, weighted diagrams are maximization diagrams of affine functions. We call them affine diagrams and we will show that any affine Voronoi diagram of $\mathbb{R}^d$ is in fact identical to the weighted Voronoi diagram of some set of spheres of $\mathbb{R}^d$, therefore showing that the class of affine diagrams is identical to the class of weighted Voronoi diagrams. Lastly, we will introduce some important examples of affine diagrams.

6.1 Weighted Delaunay complexes

6.1.1 Weighted points and weighted distance

A weighted point $\hat{p} = (p, w)$ is an element of $R^d \times R$ where the point $p \in \mathbb{R}^d$ is called the center of $\hat{p}$ and $w \in \mathbb{R}$ is called its weight. An unweighted point is identified with a point of weight 0. When w is non negative, $\hat{p}$ can be considered as a sphere centered at p of squared radius w, hence of equation $(x-p)^2 - w = 0$. Since considering negative weights will cause no problem in our developments, we adopt the hopefully less confusing name of weighted point.

We define the weighted distance or simply 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 coincides with the squared Euclidean distance when $w_1 = w_2 = 0$. Two weighted points are said to be *orthogonal* if their distance is zero. They are said to be further (resp., closer) than orthogonal if their distance is positive (resp. negative). Observe that, when the weights are positive, the notion of orthogonality between weighted points is the same as the notion of orthogonality between spheres. Specifically, two

positively weighted points are orthogonal iff the corresponding spheres are orthogonal.

We also say that a weighted point $\hat{p}$ is orthogonal to a finite set $\hat{P}$ of weighted points if $\hat{p}$ is orthogonal to all the weighted points of $\hat{P}$. Accordingly, we say that $\hat{p}$ is further or closer than orthogonal to $\hat{P}$ if it is further or closer than orthogonal to any ball $\hat{p}_i \in \hat{P}$.

6.1.2 Weighted Voronoi diagrams

Let $\hat{P} = \{\hat{p}_1, \ldots, \hat{p}_n\}$ be a set of weighted points of $\mathbb{R}^d \times \mathbb{R}$. To each $\hat{p}_i$, we associate the region $L(\hat{p}_i)$ consisting of the points x of $\mathbb{R}^d$ whose weighted distance $D(x, \hat{p}_i)$ is not larger than the weighted distance to the other weighted points of $\hat{P}$:

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

The set of points of $\mathbb{R}^d$ that are at equal distance to two weighted points $\hat{p}_i$ and $\hat{p}_j$ is a hyperplane, noted π_{ij}, called *the radical hyperplane* of $\hat{p}_i$ and $\hat{p}_j$. Hyperplane π_{ij} is orthogonal to the line joining the centers of $\hat{p}_i$ and $\hat{p}_j$. We denote by π^i_{ij} the half-space bounded by π_{ij} consisting of the points whose power to $\hat{p}_i$ is smaller than their power to $\hat{p}_j$. The region $L(\hat{p}_i)$ is the intersection of all half-spaces π^i_{ij}, $j \neq i$. If this intersection is not empty, it is a convex polyhedron, possibly not bounded. We call *weighted Voronoi regions* the non empty regions $L(\hat{p}_i)$.

We define the *weighted Voronoi diagram* of $\hat{P}$, noted $\mathrm{Vor}(\hat{P})$, as the cell complex whose cells are the weighted Voronoi regions and their faces. When all weighted points have the same weight, their weighted Voronoi diagram 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 functions $D(x, hatp_i), \ldots, D(x, \hat{p}_n)$. Observing that for any x

$$\arg\min_i D(x, \hat{p}_i) = \arg\max_i (2p_i \cdot x - w_i),$$

we obtain that the weighted Voronoi diagram of $\hat{P}$ is the maximization diagram of the set of affine functions

$$d_i(x) = 2c_i \cdot x - s_i$$

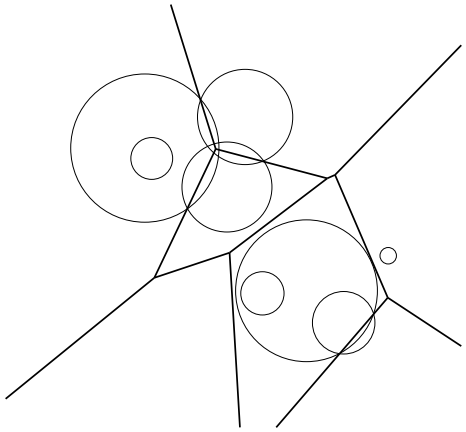

Figure 6.1: A weighted Voronoi diagram.

whose graphs are the *polar hyperplanes* $h_{\hat{p}_i} \subset \mathbb{R}^{d+1}$ of $\hat{p}_i$ (see section 5.2). Let $h^+_{\hat{p}_i}$ denote the half-space lying above $h_{\hat{p}_i}$. The maximization diagram of the d_i is obtained by projecting vertically the faces of the convex polyhedron

$$\mathcal{L}(\mathcal{S}) = h^+_{\hat{p}_1} \cap \cdots \cap h^+_{\hat{p}_n}.$$

Hence, the faces of the weighted Voronoi diagram $\mathrm{Vor}(\hat{P})$ are the vertical projections of the faces of the convex polyhedron $\mathcal{L}(\hat{P})$.

Weighted Voronoi diagrams are very similar to Voronoi diagrams: the main difference is that some hyperplane may not contribute a face. In other words, some weighted point $\hat{p}_i$ may have an empty Voronoi region (see the small circle in the upper left corner of figure 6.1).

6.1.3 Weighted Delaunay triangulations

We proceed now as in Section 5.2. To a weighted point $\hat{p} = (p, w)$ of $\mathbb{R}^d$, we associate the point $\phi(\hat{p}) = (p, p^2 - w)$ of $\mathbb{R}^{d+1}$. Consider the convex polyhedron $\mathcal{R}(\hat{P}) = \mathrm{conv}(\phi(\hat{P}))$. By the results of Section 3.2, $\mathrm{lowerhull}(\mathcal{R}(\hat{P}))$ is dual to $\mathrm{bd}(\mathcal{L}(\hat{P}))$. The vertical projection of $\mathrm{lowerhull}(\mathcal{R}(\hat{P}))$ is a cell complex which, in general, is a simplicial complex. We call such a complex the *weighted Delaunay triangulation* of $\hat{P}$ and denote it by $\mathrm{Del}(\hat{P})$. We have the following diagram :

$$\begin{array}{ccc} \mathcal{L}'(\hat{P}) = \mathrm{bd}\left(\mathcal{L}(\hat{P})\right) & \longleftrightarrow & \mathcal{R}'(\hat{P}) = \mathrm{lowerhull}(\phi(\hat{P})) \\ \updownarrow & & \updownarrow \\ \text{Weighted Voronoi diagram } \mathrm{Vor}(\hat{P}) & \longleftrightarrow & \text{Weighted Delaunay triangulation } \mathrm{Del}(\hat{P}) \end{array}$$

The following lemma is the analogue of Lemma 5.1 for unweighted Delaunay triangulations.

Lemma 6.1 *Let $\hat{P}$ be a finite set of weigted points in $\mathbb{R}^d \times \mathbb{R}$. A simplex τ whose vertices correspond to a subset $\hat{P}_\tau$ of $\hat{P}$, is a simplex of the weighted Delaunay triangulation* $\mathrm{Del}(\hat{P})$ *iff there is a weighted point orthogonal to $\hat{P}_\tau$ and further than orthogonal to $\hat{P} \setminus \hat{P}_\tau$.*

A triangulation T whose vertices correspond to a subset $\hat{P}' \subset \hat{P}$ is a weighted Delaunay triangulation of $\hat{P}$ iff for any d-simplex τ of T there is a weigthed point orthogonal to the subset $\hat{P}_\tau$ of $\hat{P}$ corresponding to the vertices of τ and further than orthogonal to $\hat{P} \setminus \hat{P}_\tau$.

6.1.4 Complexity of weighted diagrams

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 6.1.3 and of results on convex hulls (see Theorems 3.7 and 3.9).

Theorem 6.2 *The combinatorial complexity of the weighted Voronoi diagram of n weigthed 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 $\Theta\left(n \log n + n^{\lceil \frac{d}{2} \rceil}\right)$.*

6.2 Affine diagrams

We have seen that the Voronoi diagram of n weighted points is the maximization diagram of n affine functions. The converse is also true.

Lemma 6.3 *The maximization diagram of n affine functions defined over $\mathbb{R}^d$ is the weighted Voronoi diagram of n weighted points of $\mathbb{R}^d$.*

Proof Let $h_i(x) = 2c_i \cdot x - s_i$ be n affine fonctions defined over $\mathbb{R}^d$. Write h_i for the graph of $h_i(x)$, i.e. the hyperplane of $\mathbb{R}^{d+1}$ defined by $x_{d+1} = 2c_i \cdot x - s_i$, and let $P = h_1^+ \cap \cdots \cap h_n^+$. We associate to h_i the weighted point $\hat{p}_i$ of $\mathbb{R}^d$ whose center is p_i and whose weight is $p_i^2 - w_i$. h_i is the polar hyperplane not def of σ_i and, as shown in section 6.1, the faces of the weighted Voronoi diagram of the σ_i are the vertical projections of the faces of P. □

Let us call for short *affine diagram* the minimization diagram of a finite set of affine functions. The above lemma and the converse statement just show that the class of weighted Voronoi diagrams is identical to the class of affine diagrams. Trivial consequences of this correspondence are the following facts

Fact 6.4 *The image of an affine diagram by an affine transformation is an affine diagram.*

Fact 6.5 *The intersection of an affine diagram of $\mathbb{R}^d$ by an affine subspace $A \subset \mathbb{R}^d$ is an affine diagram of A. (see exercise 6.4).*

6.2.1 Diagrams for quadratic distance

We equip $\mathbb{R}^d$ with a non Euclidean metric defined by

$$\|x - y\|_Q = (x - y)\ Q\ (x - y)^t,$$

where Q is a symmetric matrix, i.e. $Q = Q^t$. Given n points $p_1, \ldots, p_n$ of $\mathbb{R}^d$, the Voronoi diagram of those points for the quadratic distance is the cell complex consisting of the n regions

$$V_Q(p_i) = \{x \in \mathbb{R}^d : (x - p_i)\ Q\ (x - p_i)^t \leq (x - p_j)\ Q\ (x - p_j)^t, \forall 1 \leq j \leq n\}.$$

When Q is taken to be the inverse of the covariance matrix, the quadratic distance is known as the *Mahalanobis distance*, extensively used in computer vision.

includegraphics[width=7cm]Fig/diagvorordre2

Figure 6.2: The 2-order Voronoi diagram of a set of points (in bold line) and the corresponding 1-order Voronoi diagram (in thin line).

Lemma 6.6 *The Voronoi diagram of a set of n points $P = \{p_1, \ldots, p_n\}$ for the distance $\|.\|_Q$ is the weighted Voronoi diagram of n weighted points $\hat{q}_1, \ldots, \hat{q}_n$ of $\mathbb{R}^d$. The center of $\hat{q}_i$ is the point $q_i = p_i\, Q$ and its weight is $w_i = p_i \cdot p_i - p_i\, Q\, p_i^t$.*

Proof

$$\begin{aligned} \|x - p_i\|_Q < \|x - p_j\|_Q &\iff -2\, p_i\, Q\, x^t + p_i\, Q\, p_i^t < -2\, p_j\, Q\, x^t + p_j\, Q\, p_j^t \\ &\iff D(x, \hat{q}_i) < D(x, \hat{q}_j). \end{aligned}$$

□

6.2.2 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 region 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 regions are not empty. These regions are polytopes and form a cell complex whose domain is $\mathbb{R}^d$ called the *k-order diagram of P* (see Figure 6.2). For $k = 1$, we obtain the usual Voronoi diagram.

Theorem 6.7 *The k-order diagram of P is the weighted Voronoi diagram of a set of n weighted points of $\mathbb{R}^d$.*

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

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

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 $w_i = c_i^2 - s_i$ with $s_i = \frac{1}{k}\sum_{p\in K_i^k} p^2$. Hence, $x \in V_k(K_i)$ iff x lies in the weighted Voronoi region of $\hat{p}_i$. $\square$

6.2.3 Bregman diagrams

Let $\mathcal{X}$ 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 $\mathcal{X}$. For any two points $p = (p_1, \ldots, p_d)$ and $q = (q_1, \ldots, q_d)$ of $\mathcal{X}$, the Bregman divergence $D_F(\cdot||\cdot) : \mathcal{X} \mapsto \mathbb{R}$ of p to q associated to F is defined as

$$D_F(p||q) = F(p) - F(q) - \langle \nabla F(q), p - q\rangle, \tag{6.1}$$

where $\nabla F = [\frac{\partial F}{\partial x_1} \ \ldots \ \frac{\partial F}{\partial x_d}]^T$ denotes the gradient operator, and $\langle p, q\rangle$ the inner (or dot) product: $\sum_{i=1}^d p_i q_i$.

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 between $\hat{p} = (p, F(p))$ and the hyperplane H_q tangent to the graph $\mathcal{F}$ of F at point $\hat{q}$: $D_F(p||q) = F(p) - H_q(p)$.

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 generator function F implies that, for any p and q in $\mathcal{X}$, $D_F(p||q) \geq 0$, 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 for $F(x) = \|x\|^2$) and the quadratic distance function $F(x) = x^T Q x$ where Q is a symmetric positive definite matrix (see section 6.2.1).

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 probabil-

ity 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

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

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

Bregman diagrams

Let $P = \{p_1, ..., p_n\}$ be a finite point set of $\mathcal{X} \subset \mathbb{R}^d$. To each point p_i, 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, $i = 1, \ldots, n$, is called the Bregman Voronoi diagram of P, which we denote by $\mathrm{Vor}_F(P)$. 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

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

It is easy to see that the minimization diagram of the n functions $D_F(x||p_i)$ $i = 1, \ldots, n$, is the maximization of the n affine functions $h_i(x) = \langle x - p_i, p_i' \rangle$, $i = 1, \ldots, n$. Hence, Bregman diagrams are *affine* diagrams. More precisely, we have

Theorem 6.8 *The Bregman Voronoi diagram of n sites is identical to the restriction to $\mathcal{X}$ of the weighted Voronoi diagram of the n Euclidean spheres of equations*

$$\langle x - p_i', x - p_i' \rangle = \langle p_i', p_i' \rangle + 2(F(p_i) - \langle p_i, p_i' \rangle), \ i = 1, \ldots, n.$$

where $p_i' = \nabla F(p_i)$.

Proof We have

$$D_F(x||p_i) \leq D_F(x||p_j) \iff -F(p_i) - \langle x - p_i, p_i' \rangle \leq -F(p_j) - \langle x - p_j, p_j' \rangle.$$

Multiplying twice the last inequality, and adding $\langle x, x \rangle$ to both sides yields

$$\begin{aligned} & \langle x, x \rangle - 2\langle x, p_i' \rangle - 2F(p_i) + 2\langle p_i, p_i' \rangle \leq \langle x, x \rangle - 2\langle x, p_j' \rangle - 2F(p_j) + 2\langle p_j, p_j' \rangle \\ \iff \quad & \langle x - p_i', x - p_i' \rangle - r_i^2 \leq \langle x - p_j', x - p_j' \rangle - r_j^2, \end{aligned}$$

where $r_i^2 = \langle p_i', p_i' \rangle + 2(F(p_i) - \langle p_i, p_i' \rangle)$ and $r_j^2 = \langle p_j', p_j' \rangle + 2(F(p_j) - \langle p_j, p_j' \rangle)$. The last inequality means that the distance of x to the weighted point (p_i', r_i) is no more than its distance to the weighted point (p_j', r_j). □

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 triangulations

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

For the same reasons as for Delaunay triangulations (see section 5.2), the vertical projection of $\mathcal{T}$ onto $\mathcal{X}$ is a triangulation $\mathrm{Del}_F(E)$ of E embedded in $\mathcal{X} \subseteq \mathbb{R}^d$. We call $\mathrm{Del}_F(E)$ the *Bregman triangulation* of E, noted also $\mathrm{Breg}(E)$ when we don't want to insist on a specific F . When $F(x) = ||x||^2$, $\mathrm{Del}_F(E)$ is the Delaunay triangulation of E.

We now show that the empty sphere property of Delaunay triangulations (Lemma 5.1) naturally extends to Bregman triangulations. We define the *Bregman ball* centered at c and of radius r as

$$B_F(c, r) = \{x \in \mathcal{X} \mid D_F(x||c) \leq r\}.$$

The boundary of a Bregman ball is called a *Bregman sphere*. It is easy to see that any Bregman sphere σ is obtained as the vertical projection of the

intersection of $\mathcal{F}$ with a hyperplane noted h_σ (see Fig. 5.5 for the case where $F = \|x\|^2$).

A Bregman sphere is said to be *empty* if it does not enclose any point of E. Let $s = \text{conv}(p_0, \ldots, p_d)$ be a d-simplex of $\text{Breg}(E)$. The affine hull of the lifted points $\hat{p}_0, \ldots, \hat{p}_d$ is a hyperplane h of $\mathbb{R}^{d+1}$ whose intersection with $\mathcal{F}$ projects vertically on the (unique) Bregman sphere $\sigma(s)$ that circumscribes s. Since, by construction, $\text{conv}(\hat{p}_0, \ldots, \hat{p}_d) \in h$ is a facet of the lower hull of $\hat{E}$, $\sigma(s)$ must be empty.

Lastly, we show that Bregman triangulations and diagrams are dual structures. Let $\text{Breg}(E)$ be the Bregman diagram of E, $\text{Vor}(E') = \text{Breg}(E)$ the associated weighted Voronoi diagram and let $T(E')$ denote the weighted Delaunay triangulation that is dual to $\text{Vor}(E')$. The image by $\nabla^{-1}F$ of $T(E')$ is an embedded curved triangulation $T'(E)$ of E. We claim that $T'(E)$ and $\text{Del}_F(E)$ are isomorphic. Indeed, both triangulations are embedded in $\mathbb{R}^d$ and their d-simplices are in bijection with the vertices of the Bregman diagram of E.

6.3 Exercises

Exercise 6.1 Let $\sigma = (c, r)$ be a hypersphere of $\mathbb{R}^d$ and, as in exercise 5.2.2, $\phi(\sigma)$ be the point $(c, s = c^2 - r^2)$ of $\mathbb{R}^{d+1}$. We can define a distance between two spheres $\sigma_1 = (c_1, r_1)$ and $\sigma_2 = (c_2, r_2)$ as

$$d(\sigma_1, \sigma_2) = \sqrt{(c_1 - c_2)^2 - r_1^2 - r_2^2}.$$

Show that if $d(\sigma_1, \sigma_2) = 0$, σ_1 and σ_2 are orthogonal. In addition, show that the image by ϕ of the spheres that are orthogonal to a given sphere σ is the hyperplane h_σ.

Exercise 6.2 Show that the only numerical operation that is required to check if a triangulation is the weighted Delaunay triangulation of a set of hyperspheres σ_i is the evaluation of the sign of the determinant of the $(d + 2) \times (d + 2)$ matrix

$$\texttt{power\_test}(\sigma_0, \ldots, \sigma_{d+1}) = \begin{vmatrix} 1 & \cdots & 1 \\ c_0 & \cdots & c_{d+1} \\ c_0^2 - r_0^2 & \cdots & c_{d+1}^2 - r_{d+1}^2 \end{vmatrix}$$

where c_i and r_i are respectively the center and the radius of σ_i.

Exercise 6.3 Prove the following conservation law for flows entering a weighted Voronoi region $L(\sigma_i)$ normally to the facets of the region: if f_{ij}, $j \in J$, are the facets of $L(\sigma_i)$, we have $\sum_{j \in J} \mathrm{vol}(f_{ij}) \, \frac{c_j - c_i}{\|c_j - c_i\|} = 0$. This property makes Voronoi and weighted Voronoi diagrams useful when applying finite volume methods in fluid dynamics. (Hint : show that if P is a convex polyhedron and f_j, $j \in J$, are its facets, we have $\sum_{j \in J} \mathrm{vol}(f_j) \, n_j = 0$, where n_j is the unit normal vector to f_j oriented towards the outside of P).

Exercise 6.4 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 hyperspheres of $\mathbb{R}^d$ is the weighted Voronoi diagram of n hyperspheres (i.e., $(k-1)$-dimensional spheres) of H. Show how to modify the predicates of an algorithm that computes a weighted Voronoi diagram in $\mathbb{R}^k$ so that the algorithm will compute L_H.

Exercise 6.5 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.

Exercise 6.6 (**) 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.

6.4 Bibliographical notes

Weighted Voronoi diagrams appear in the literature under various names, weighted Voronoi diagrams, power diagrams.

To know more about weighted Voronoi diagrams, one may look at the survey paper by Aurenhammer [7]. A recent survey on affine and curved Voronoi diagrams can be found in [22]. The solution to exercise 6.6.1 is due to Clarkson and Shor [49] (see also [23, 24]). Bregman Voronoi diagrams have been studied by Boissonnat, Nielsen and Nock [20].

Alpha shapes were introduced by Edelsbrunner and Mücke [64, 58] and widely used in the early algorithms of shapes reconstruction from sets of data points measured on the surface of an object [65] . Alpha shapes are also famous for their ability to represent union of balls through the nerve theorems [59]. They are widely used in study of macro molecular structure and docking, see e.g [62, 87, 61]

Part III

Reconstruction of smooth submanifolds

Triangulating an object $\mathbb{M}$ consists in computing a simplicial complex K with the same topology type as $\mathbb{M}$. This is a demanding quest and, in this part, we will make the strong assumption that $\mathbb{M}$ 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 four chapters. The first one provides conditions under which one can triangulate a submanifold $\mathbb{M}$ of $\mathbb{R}^d$, i.e. construct a simplicial complex embedded in $\mathbb{R}^d$ which is homeomorphic to $\mathbb{M}$. Prior to stating and proving the main theorem, we introduce the central notions of ε-net on a manifold and of thick simplices.

In Chapter 8, 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 $\mathbb{M} \subset \mathbb{R}^d$. While the alpha complex of P usually does not triangulate $\mathbb{M}$, we will show that it has the same homotopy type as $\mathbb{M}$ under appropriate α and sampling conditions.

Avoiding flat simplices is crucial both in theory and in practice. Unfortunately, Delaunay simplices are not guaranteed to be thick even if the point set is a nice net of $\mathbb{R}^d$. We need techniques, to be presented in Chapter 9, to remove non-thick simplices in Delaunay complexes.

The ultimate goal of Part III of the book is to solve the manifold reconstruction problem, which is to be done in Chapter 10. Given a finite set of points P on an unknown manifold $\mathbb{M}$, the goal is to compute a simplicial complex $\hat{\mathbb{M}}$ that triangulates $\mathbb{M}$. This problem is of primary importance when $\mathbb{M}$ is a surface of $\mathbb{R}^3$ (it is then known as the surface reconstruction problem). It also founds 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 $\mathbb{M}$ is much smaller than the dimension d of the ambient space, we will see in Chapter 10 that we can walk around the curse of dimensionality and

present an efficient algorithm that reconstructs a submanifold $\mathbb{M}$ 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

In this section as well as in the rest of the chapter, $\mathbb{M}$ denotes a k-dimensional submanifold of $\mathbb{R}^d$ that is compact, closed and sufficiently smooth (in a sense to be precisely defined below). We write T_x for the tangent space at point $x \in \mathbb{M}$.

7.1 Reach and ε-nets on manifolds

7.1.1 Projection map, medial axis and reach

We define the *medial axis* of $\mathbb{M}$, noted $\mathrm{ax}(\mathbb{M})$, as the closure of the set of points $x \in \mathbb{R}^d$ that have more than one closest point on $\mathbb{M}$ (see Figure 7.1). Equivalently, the medial axis can be defined as the closure of the locus of centers of open balls that are tangent to $\mathbb{M}$ in at least two points and do not intersect $\mathbb{M}$. Such a ball will be called a *medial ball*. Another equivalent definition is to define the medial axis of $\mathbb{M}$ as the closure of the locus of the centers of the balls that do not intersect $\mathbb{M}$ and are maximal for the inclusion (i.e. are not contained in a bigger ball that does not intersect $\mathbb{M}$).

We define the *projection onto* $\mathbb{M}$ as the mapping

$$\Pi : \mathbb{R}^d \setminus \mathrm{ax}(\mathbb{M}) \to \mathbb{M}$$

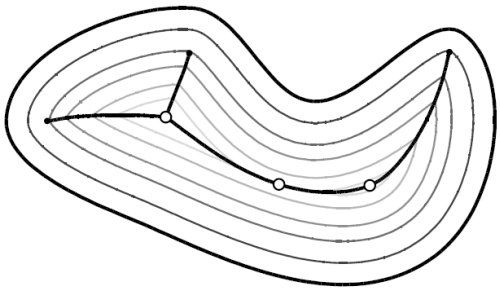

Figure 7.1: The medial axis of a closed curve. We only show the component of the medial axis that is contained in the domain bounded by the curve. Various offsets of the curve are also shown in thin line.

that maps a point x to its unique closest point on $\mathbb{M}$.

The distance of a point $x \in \mathbb{M}$ to $\text{ax}(\mathbb{M})$ is called the *local feature size* of x and denoted by $\text{lfs}(x)$. It is easy to show that lfs is 1-Lipschitz, i.e.

$$\forall x, y \in \mathbb{M}, \quad |\text{lfs}(x) - \text{lfs}(y)| \leq \|x - y\|.$$

We call *reach* of $\mathbb{M}$ and write $\text{rch}(\mathbb{M})$ for the minimum of lfs over $\mathbb{M}$, which exists since $\mathbb{M}$ is compact.

The following result is an easy consequence of the definitions above. We state it as a lemma for further references.

Lemma 7.1 *Let $\mathbb{M}$ be a differentiable submanifold and $x \in \mathbb{M}$. Any open ball that is tangent to $\mathbb{M}$ at x whose radius is less than $\text{lfs}(x)$ does not intersect $\mathbb{M}$.*

It follows from the lemma that the local feature at x size cannot be greater than the smallest radius of curvature at x. Moreover, if the radius of curvature is defined everywhere on $\mathbb{M}$ and does not vanish, lfs is strictly positive.

In this chapter, we will restrict our attention to the class of submanifolds that are smooth in the sense that their reach is strictly positive. This class of submanifolds include all submanifolds of class C^2 and also some submanifolds whose curvature may be discontinuous along subsets of dimension 0. An example of such a submanifold is the *r-offset* of a cube, i.e. the set of points at distance at most r from the cube.

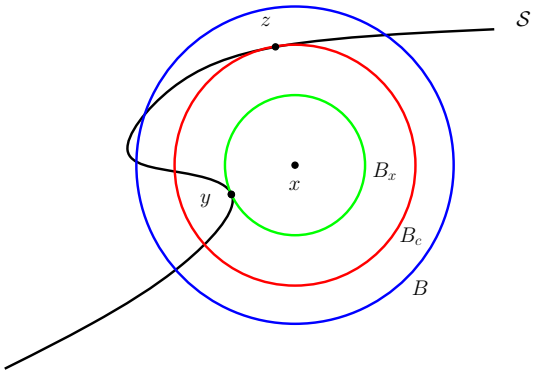

Figure 7.2: For the proof of Lemma 7.2.

We now state a topological lemma that may give an indication of the importance of the concept of reach in the context of submanifold approximation.

Lemma 7.2 *Let B be a closed ball that intersects $\mathbb{M}$. If $B \cap \mathbb{M}$ is not a topological ball, then B contains a point of the medial axis of $\mathbb{M}$.*

notation fig

Proof Write r for the radius of B and x for its center. The result is trivial when x belongs to the medial axis of $\mathbb{M}$. Therefore assume that $x \notin \mathrm{ax}(\mathbb{M})$.

We denote by $B(x, r_x)$ the largest closed ball centered at x whose interior does not intersect $\mathbb{M}$ (if $x \in \mathbb{M}$, $B(x, r_x)$ is reduced to point x). $B(x, r_x)$ is tangent to $\mathbb{M}$ in a unique point y since x does not belong to $\mathrm{ax}(\mathbb{M})$ (see Figure 7.2). Hence $B(x, r_x) \cap \mathbb{M} = \{y\}$ is a topological ball. Since $B \cap \mathbb{M}$ is not a topological ball, there exists a point $z \neq y$ of $\mathbb{M}$ at distance $r_c > r_x$ from x such that the ball $B(x, r_c)$ is tangent to $\mathbb{M}$ at z. Consider the set $\mathcal{B}$ of closed balls that are tangent to $\mathbb{M}$ at z and are centered on the line segment $[zx]$. Note that $B(x, r_c)$ is the ball of $\mathcal{B}$ centered at x. Since the interior of $B(x, r_c)$ contains y and therefore intersects $\mathbb{M}$, there must exist a ball $B_z \in \mathcal{B}$ maximal for the inclusion whose interior does not intersect $\mathbb{M}$. B_z is a medial ball and we have $B_z \subset B(x, r_c) \subset B$. Since the center of B_z belongs to $\mathrm{ax}(\mathbb{M})$, the lemma is proved. □

If $x \in \mathbb{M}$, $B(x, r)$ cannot intersect the medial axis of $\mathbb{M}$ for any $r < \mathrm{rch}(\mathbb{M})$. Lemma 7.2 thus implies

Lemma 7.3 *For any x of $\mathbb{M}$, and any $r < \mathrm{rch}(\mathbb{M})$, the intersection of $\mathbb{M}$ with the ball $B(x,r)$ centered at x of radius r is a topological ball.*

7.1.2 ε-nets on a submanifold

We extend the definition of ε-nets introduced in Section 5.3 to the case of point samples on submanifolds.

Definition 7.4 ($(\varepsilon,\bar{\eta})$-net) *Let $\mathbb{M}$ be a submanifold of $\mathbb{R}^d$ of positive reach. We say that a finite point set $P \subset \mathbb{M}$ is an*

ε-dense sample of $\mathbb{M}$ *if any point x of $\mathbb{M}$ is at distance at most $\varepsilon\,\mathrm{rch}(\mathbb{M})$ from a point of P (the distance is the Euclidean distance in $\mathbb{R}^d$).*

η-separated sample of $\mathbb{M}$ *if, for any two points p,q of P, $\|p-q\| \geq \eta\,\mathrm{rch}(\mathbb{M})$.*

$(\varepsilon,\bar{\eta})$-net of $\mathbb{M}$ *if it is ε-dense and $\bar{\eta}\varepsilon$-separated. We call ε the sampling radius and $\bar{\eta}$ the separation ratio of P.*

This definition does not allow the sampling radius to vary over the submanifold. Hence, we are confined with uniform samples, which may be quite restrictive in practice. A way to allow non-uniform samples is to replace, in the definition above, the global quantity $\mathrm{rch}(\mathbb{M})$ by the local feature size lfs. We will say that a finite set of points $P \in \mathbb{M}$ is a *non-uniform*

ε-dense sample of $\mathbb{M}$ if any point x of $\mathbb{M}$ is at distance at most $\varepsilon\,\mathrm{lfs}(x)$ from a point of P.

η-separated sample of $\mathbb{M}$ if, for any two points p,q of P, $\|p-q\| \geq \eta\,\min(\mathrm{lfs}(p),\mathrm{lfs}(q))$.

$(\varepsilon,\bar{\eta})$-net of $\mathbb{M}$ if it is ε-dense and $\bar{\eta}\varepsilon$-separated.

The results of this chapter can be extended to such 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 let the extension to non-uniform ε-nets as an exercise (Exercise 7.3). See also the bibliographic notes.

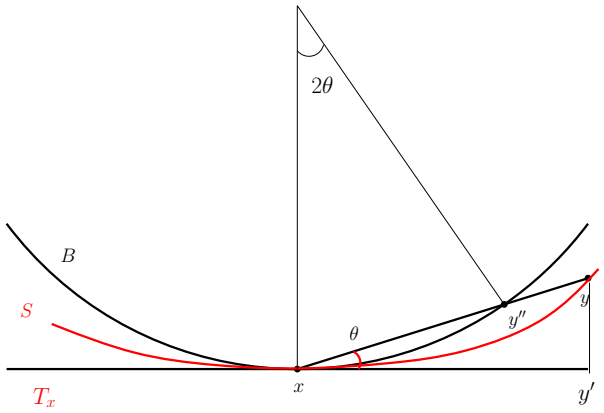

Figure 7.3: For the proof of Lemma 7.5. $\theta = \angle(xy, T_x)$.

7.2 Some elementary differential geometry

Lemma 7.5 *Let x and y be two points of $\mathbb{M}$. We have*

1. $\sin\angle(xy, T_x) \leq \frac{\|x-y\|}{2\,\mathrm{lfs}(x)}$;
2. *the distance from y to T_x is at most* $\frac{\|x-y\|^2}{2\,\mathrm{lfs}(x)}$.

Proof 1. Let y' be the orthogonal projection of y onto T_x and let H be the plane (xyy'). Let in addition B be the disk of H of radius $\mathrm{lfs}(x)$ that is tangent to $\mathbb{M}$ at x and whose center is on the same side of $T_x \cap H$ as y (Refer to Figure 7.3). By definition of the local feature size, the open medial ball tangent to $\mathbb{M}$ at x does not intersect $\mathbb{M}$. Hence y does not belong to the interior of B (Lemma 7.1). Let y'' be the intersection other than x between the line segment $[xy]$ and the boundary of B. We have

$$\|x - y\| \geq \|x - y''\| = 2\,\mathrm{lfs}(x)\,\sin\angle(xy, T_x)$$

which proves the first statement.

2. We have $\|y - y'\| = \|x - y\|\,\sin\angle(xy, T_x) \leq \frac{\|x-y\|^2}{2\,\mathrm{lfs}(x)}$.

□

make θ more visible

We now show that the projection $\Pi : \mathbb{R}^d \setminus \mathrm{ax}(\mathbb{M}) \to \mathbb{M}$ that maps a point to its closest point on $\mathbb{M}$ is continuous in a neighborhood of $\mathbb{M}$.

Lemma 7.6 (Continuity of Π) *Let x and y be two points of $\mathbb{R}^d \setminus \mathrm{ax}(\mathbb{M})$. Write $x' = \Pi(x)$ and $y' = \Pi(y)$ for their (unique) projections on $\mathbb{M}$. If $\|x - x'\| \leq \alpha\, \mathrm{rch}(\mathbb{M})$ and $\|y - y'\| \leq \alpha\, \mathrm{rch}(\mathbb{M})$ with $\alpha < 1$, then $\|x' - y'\| \leq \frac{1}{1-\alpha}\, \|x - y\|$.*

Add a fig. Notations : keep p, q for important points (in lemmas), x,y,z for intermediate points (in proofs) (everywhere)

Proof Using Lemma 7.5 (1), we have

$$\begin{aligned}
(x - x') \cdot (x' - y') &\geq -\|x - x'\|\, \|x' - y'\|\, |\cos\left(x - x', x' - y'\right)| \\
&= -\|x - x'\|\, \|x' - y'\|\, |\sin\left(x' - y', T_{x'}\right)| \\
&\geq -\alpha\, \mathrm{rch}(\mathbb{M})\, \|x' - y'\|\, \frac{\|x' - y'\|}{2\mathrm{lfs}(x')} \\
&\geq -\frac{\alpha}{2}\, \|x' - y'\|^2
\end{aligned}$$

Similarly

$$(y - y') \cdot (y' - x') \geq -\frac{\alpha}{2} \|\, x' - y'\|^2.$$

Therefore,

$$\begin{aligned}
\|x - y\| \times \|x' - y'\| &\geq (x - y) \cdot (x' - y') \\
&= \left((x' - y') + (x - x') + (y' - y)\right) \cdot (x' - y') \\
&\geq (1 - \alpha)\, \|(x' - y')\|^2
\end{aligned}$$

from which the lemma follows. □

The following lemma bounds the angle between two tangent spaces. Angles between affine spaces are defined in Section 7.3.

Lemma 7.7 (Angle between tangent spaces) *Let $p, q \in \mathbb{M}$ and $\|p - q\| = t\, \mathrm{rch}(M)$ with $t \leq 1/2$. Then*

$$\cos \angle(T_p, T_q) \geq 1 - 2t^2 \quad \textit{and therefore} \quad \sin \angle(T_p, T_q) \leq 2t. \tag{7.1}$$

Proof See [95][Proposition 6.2] and the improvement in [9][Lemma 3.4]. □

Notation $B(p,r)$ and $\bar{B}$

def embedding in ch. 1 : a homeomorphism onto its image.

Lemma 7.8 (Π_p is a local embedding) *Let Π_p be the orthogonal projection onto the tangent space T_p of $\mathbb{M}$ at p. Let $B_M(p,r) = B(p,r) \cap \mathbb{M}$ and assume that $r < \frac{\mathrm{rch}(\mathbb{M})}{3}$. Then we have*

(i) *Π_p restricted to $\overline{B}_M(p,r)$ is an embedding.*

(ii) *For any $r' \leq 0.98\, r$, we have $\overline{B}(p,r') \cap T_p \subseteq \Pi_p(\overline{B}_M(p,r))$.*

Proof We first observe that, according to Lemma 7.3, $B_M(p,r)$ is a k-dimensional topological ball where k is the dimension of $\mathbb{M}$.

(i) For all x and $y \in \overline{B}_{\mathbb{M}}(p,r)$, we have from Lemma 7.7 and Lemma 7.5

$$\sin\angle(x\,T_p) \leq \sin\angle(xy,\,T_x) + \sin\angle(T_x,\,T_p) \leq \frac{3r}{\mathrm{rch}(\mathbb{M})} < 1. \tag{7.2}$$

This implies that Π_p is 1-1. Indeed, otherwise, there would exist x and $y \neq x$ in $\overline{B}(p,r) \cap \mathbb{M}$ such that $\Pi_p(x) = \Pi_p(y)$ then $\angle(xy, T_p) = \pi/2$, contradicting Eq. 7.2. The fact that Π_p is a homeomorphism between $\overline{B}_{\mathbb{M}}(p,r)$ and its image $\Pi_p\left(\overline{B}_{\mathbb{M}}(p,r)\right)$ follows the continuity of Π_p and of its inverse (which is ensured since $\overline{B}_{\mathbb{M}}(p,r)$ is compact and its image is Hausdorff (see the discussion after Definition 1.2).

(ii) Since $\overline{B}_{\mathbb{M}}(p,r)$ is a closed topological ball, and Π_p restricted to $\overline{B}_{\mathbb{M}}(p,r)$ is an embedding, it follows that $\Pi_p(\overline{B}_{\mathbb{M}}(p,r))$ is also a topological ball containing p in its interior. Moreover, any point $y \in \partial\Pi_p(\overline{B}_{\mathbb{M}}(p,r))$ is the image of a point $\Pi_p^{-1}(y) \in \partial\overline{B}_{\mathbb{M}}(p,r)$. Using Pythagoras theorem and Lemma 7.5, we find

$$\|p-y\|^2 = \|p-\Pi_p^{-1}(y)\|^2 - \|y-\Pi_p^{-1}(y)\|^2 = r^2\left(1 - \left(\frac{r}{2\mathrm{rch}(\mathbb{M})}\right)^2\right) > (0.98\,r)^2$$

This proves the second part of the lemma. □

Lemma 7.9 *Let $r' < \mathrm{rch}(\mathbb{M})/4$. The restriction of Π to $B(p,r') \cap T_p$ is a non-singular embedding (i.e. a local diffeomorphism). Moreover, for any $x \in B(p,r') \cap T_p$, we have*

$$\|x - \Pi(x)\| \leq (1 - \sqrt{1-t^2})\,\mathrm{rch}(\mathbb{M}) \quad \text{where } t = \frac{\|x-p\|}{\mathrm{rch}(\mathbb{M})}.$$

Proof Let $r = r'/0.98$ and note that $r < \mathrm{rch}(\mathbb{M})/3$. We write $B_M(p,r) = B(p,r) \cap \mathbb{M}$ and $B_T(p,r') = B(p,r') \cap T_p$. Let $x' = \Pi_p^{-1}(x) \cap B_M(p,r)$ and let $x'' = \Pi(x)$ be the point of $\mathbb{M}$ closest to x. Since $\|x - x''\| \leq \|x - x'\| < \mathrm{rch}(\mathbb{M})$, x'' is uniquely defined. Moreover, the restriction of Π to $B_T(p,r')$ is continuous by Lemma 7.6. We prove that it is 1-1 by contradiction. Assume that two distinct points $x, y \in B_T(p,r')$ have the same projection on $\mathbb{M}$, say x''. The line segment $[xy]$ must belong to the normal space at x''. But since $\|p - x''\| \leq 2\|p - x\| \leq 2r'$, we deduce from Lemma 7.7 that $\cos(\angle(T''_x, T_p)) \geq 1 - 8(r'/\mathrm{rch}(\mathbb{M}))^2 > 0$. We obtain a contradiction with the fact that $[xy] \subset T_p$.

Let us show that the restriction of Π to $B_T(p,r')$ is non-singular. Being singular means that $[xx'']$ is contained in T''_x. But this is impossible since $\|p - x''\| < 2r'$ and thus $\cos(\angle(T''_x, T_p)) > 0$ by Lemma 7.7. This proves the first part of the lemma.

Let us prove now the second part. Since, from Lemma 7.8(ii), $\overline{B}(p,r') \cap T_p \subseteq \Pi_p\left(\overline{B}_M(p,r)\right)$, it follows that there exists $x' \in \overline{B}_M(p,r)$ such that $\Pi_p(x') = x$. Writing $t = \frac{\|p-x\|}{\mathrm{rch}(\mathbb{M})}$ and $t' = \frac{\|p-x'\|}{\mathrm{rch}(\mathbb{M})}$ and using Lemma 7.5, we have

$$\begin{aligned} \|x - x'\| &\leq \frac{\|p-x'\|^2}{2\mathrm{rch}(\mathbb{M})} = \frac{t'^2}{2}\,\mathrm{rch}(\mathbb{M}) \\ \|p-x'\|^2 = t'^2\,\mathrm{rch}(\mathbb{M})^2 &= \|p-x\|^2 + \|x-x'\|^2 \leq t^2\mathrm{rch}(\mathbb{M})^2 + \frac{t'^4}{4}\mathrm{rch}(\mathbb{M})^2. \end{aligned}$$

The above inequality together with $t' < \frac{1}{3}$ leads to $\frac{t'^2}{2} \leq 1 - \sqrt{1-t^2}$ and this, together with Lemma 7.5, implies

$$\mathrm{dist}(x,\mathbb{M}) \leq \|x - x'\| \leq \frac{t'^2}{2}\mathrm{rch}(\mathbb{M}) \leq (1 - \sqrt{1-t^2})\,\mathrm{rch}(\mathbb{M}).$$

□

Lemma 7.10 *Let σ be a geometric k-simplex with its vertices on $\mathbb{M}$ and assume that its diameter $\Delta(\sigma)$ is at most $\mathrm{rch}(\mathbb{M})/3$. Then, for any point x in σ, we have* $\|\Pi(x)-x\| \leq \frac{\Delta^2(\sigma)}{2\mathrm{rch}(\mathbb{M})} + \left(1-\sqrt{1-\left(\frac{\Delta(\sigma)}{\mathrm{rch}(\mathbb{M})}\right)^2}\right)\mathrm{rch}(\mathbb{M}) \approx \frac{\Delta^2(\sigma)}{\mathrm{rch}(\mathbb{M})}$.

Terminology : make sure in ch 0 that $x \in \sigma$ makes sense if σ geometric and x not a vertex and that is done consistently in the book

Proof Let p be a vertex of σ, $x' = \Pi_p(x)$, the orthogonal projection of x on T_p, and $y = \Pi(x')$. Denote by $\lambda_q(x)$, $q \in \sigma$, the barycentric coordinates of x. Using Lemma 7.5 and the lineartity of Π_p, we have

$$\|x-x'\| = \|\lambda_q(x)\,(q-q')\| \leq \left(\sum_{q\in\sigma}\lambda_q(x)\right)\frac{\Delta^2(\sigma)}{2\mathrm{rch}(\mathbb{M})} = \frac{\Delta^2(\sigma)}{2\mathrm{rch}(\mathbb{M})}.$$

Then we get from Lemma 7.9,

$$\|y-x'\| \leq \left(1-\sqrt{1-\left(\frac{\Delta(\sigma)}{\mathrm{rch}(\mathbb{M})}\right)^2}\right)\mathrm{rch}(\mathbb{M}).$$

Since $d(x,\mathbb{M}) \leq \|x-y\| \leq \|x-x'\| + \|y-x'\|$, the lemma follows. □

7.3 Thick simplices

For a given set of points $P \in \mathbb{R}^2$, $\mathrm{Del}(P)$ maximizes, over all possible triangulations of P, the smallest angle of the triangles (Exercise 9.4) and it can be easily shown that, if P is a net, then all angles are lower bounded by some positive constant. However, the property does not hold for higher dimension Delaunay complexes and one cannot bound the dihedral angles of higher dimensional simplices as shown in Fig. 7.4.

For any vertex p of a simplex σ, the *face oppposite* p is the face determined by the other vertices of σ, and is denoted by σ_p. The *altitude* of p in σ is the distance $D(p,\sigma) = d(p,\mathrm{aff}(\sigma_p))$, and 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*

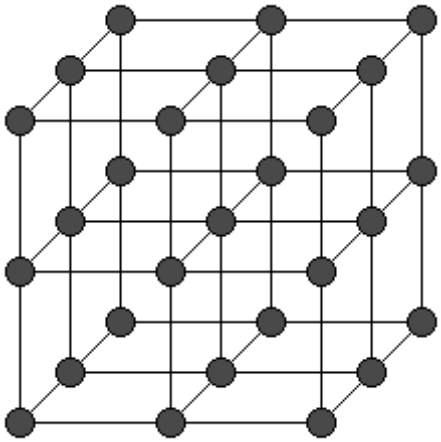

Figure 7.4: The four vertices of a squared face f of a uniform grid are cocircular and can be circumscribed by a sphere centered in the middle of the face. The sphere does not enclose any other vertex of the grid. If we slightly perturb the vertices of f form a tetrahedron of small volume whose circumscribing 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.

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.

7.3.1 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. If A is a $d \times j$ matrix, $j \leq d$, we denote its 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 7.11 *If $\eta > 0$ is an upper bound on the norms of the columns of A, then $s_1(A) = \|A\| \leq \sqrt{j}\eta$.*

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 an $d \times j$ matrix of rank $j \leq d$, then the *pseudo-inverse* $A^\dagger = (A^\mathsf{T}A)^{-1}A^\mathsf{T}$ 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:

Lemma 7.12 *Let A be a $d \times j$ matrix of rank $j \leq d$ and let $A^\dagger$ be its* pseudo inverse $= (A^\mathsf{T}A)^{-1}A^\mathsf{T}$. We have

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

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

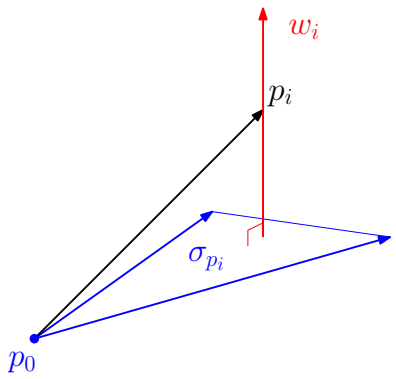

Figure 7.5: Choosing p_0 as the origin, the edges emanating from p_0 in $\sigma = [p_0, \ldots, p_j]$ form a basis for $\mathrm{aff}(\sigma)$. The proof of Lemma 7.13 demonstrates that the dual basis $\{w_i\}$ consists of vectors that are orthogonal to the facets, and with magnitude equal to the inverse of the corresponding altitude.

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 7.13 (Thickness and singular value) *Let $\sigma = \mathrm{conv}(p_0, \ldots, p_j)$ 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*

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

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

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

Indeed, by the definition of $\mathsf{P}^\dagger$, 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 $w_i^\mathsf{T}(p_i - p_0) = 1 = \|w_i\| u_i^\mathsf{T}(p_i - p_0)$. By the definition of the altitude of a vertex, we have $u_i^\mathsf{T}(p_i - p_0) = D(p_i, \sigma)$. 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 7.11, yields

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

The stated bound on $s_j(\mathsf{P})$ follows from Lemma 7.12. □

The proof of Lemma 7.13 shows that the pseudoinverse of P has a natural geometric interpretation in terms of the altitudes of σ, and thus the altitudes provide a convenient lower bound on $s_j(\mathsf{P})$. By Lemma 7.11, $s_1(\mathsf{P}) \leq \sqrt{j}\Delta(\sigma)$, and thus $\Theta(\sigma) \leq \frac{s_j(\mathsf{P})}{s_1(\mathsf{P})}$. In other words, $\Theta(\sigma)^{-1}$ provides a convenient upper bound on the *condition number* of 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.

7.3.2 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.

If U and V are vector subspaces of $\mathbb{R}^m$, with $\dim U \leq \dim V$, the *angle* between them is defined by

$$\sin\angle(U,V) = \sup_{\substack{u\in U\\ \|u\|=1}} \|u - \pi_V u\|, \tag{7.3}$$

where π_V is the orthogonal projection onto V. The angle between affine subspaces K and H is defined as the angle between the corresponding parallel vector subspaces.

Lemma 7.14 (Whitney's angle bound) *Suppose σ is a j-simplex, $j < d$, whose vertices all lie within a distance h from a k-dimensional affine space $H \subset \mathbb{R}^d$ with $k \geq j$. Then*

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

Proof Suppose $\sigma = \text{conv}(p_0, \ldots, p_j)$. Choose p_0 as the origin of $\mathbb{R}^d$ and let $\pi : \mathbb{R}^d \to H$ be the orthogonal projection onto H. Let u be any unit vector in $\text{aff}(\sigma)$. Since the vectors $v_i = (p_i - p_0)$, $i \in \{1, \ldots, j\}$ form a basis for $\text{aff}(\sigma)$, we may write $u = \mathsf{P}a$, 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 $X = \mathsf{P} - \pi\mathsf{P}$, we get

$$\|u - \pi u\| = \|Xa\| \leq \|X\| \, \|a\|.$$

Since $d(p_i, H) \leq h$ for all $0 \leq i \leq j$, $\|v_i - \pi v_i\| \leq 2h$. It follows then from Lemma 7.11 that

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

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

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

and the result follows from Lemma 7.13. □

Whitney's angle bound is especially useful when H is a tangent space to a smooth manifold $\mathbb{M}$. Let σ be a thick simplex whose vertices are close (relatively to the diameter of the simplex) to a tangent space T_p, $p \in \mathbb{M}$. Then the 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 Fig. 7.6. 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 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.

7.4 Triangulation of manifolds

Theorem 7.15 *Let $\mathbb{M}$ be a k-submanifold of $\mathbb{R}^d$ without boundary and of positive reach* $\text{rch}(\mathbb{M})$. *Let in addition $\hat{\mathbb{M}}$ be a k-manifold complex without*

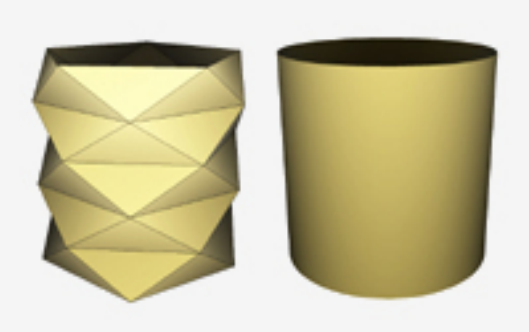

Figure 7.6: The Schwarz lantern.

boundary embedded in $\mathbb{R}^d$ *with vertex set* $P \subset \mathbb{M}$. *Assume that* P *is an* ε*-net of* $\mathbb{M}$ *and that all the simplices of* $\hat{\mathbb{M}}$ *have a diameter at most* $2\varepsilon\,\mathrm{rch}(\mathbb{M})$ *for some* $\varepsilon < \frac{1}{9}$. *Assume in addition that there exists on each connected component* $\mathbb{M}_i$ *of* $\mathbb{M}$ *a point* $y_i \in \mathbb{M}_i$ *such that* $\Pi_{|\hat{\mathbb{M}}}^{-1}(y_i)$ *contains exactly one point in* $\hat{\mathbb{M}}$. *Then the following facts hold.*

1. $\Pi_{|\hat{\mathbb{M}}} : \hat{\mathbb{M}} \to \mathbb{M}$ *is a homeomorphism.*

2. *The Hausdorff distance between* $\hat{\mathbb{M}}$ *and* $\mathbb{M}$ *is* $O(\varepsilon^2 \mathrm{rch}(\mathbb{M}))$ *where the constant in the big-O is close to* 4.

3. *If* σ *is a* k*-simplex of* $\hat{\mathbb{M}}$ *and* p *one of its vertices, we have*

$$\sin \angle(\mathrm{aff}(\sigma), T_p) \leq \frac{\Delta(\sigma)}{\Theta(\sigma)\,\mathrm{rch}(\mathbb{M})} \leq \frac{2\varepsilon}{\Theta(\sigma)}$$

where $\Delta(\sigma)$, $L(\sigma)$ *and* $\Theta(\sigma)$ *are respectively the longest edge, the shortest edge and the thickness of* σ.

Proof Observe that $\Pi_{|\hat{\mathbb{M}}} : \hat{\mathbb{M}} \to \mathbb{M}$ is well defined since $\hat{\mathbb{M}}$ does not intersect the medial axis of $\mathbb{M}$ since $\varepsilon < 1$. We will prove that $\Pi_{|\hat{\mathbb{M}}}$ is a homeomorphism, which implies that $\hat{\mathbb{M}}$ is a triangulation of $\mathbb{M}$. This is the first part of the theorem. The second part of the theorem follows from Lemma 7.10 and the third part from Lemmas 7.14 and 7.5. We now prove the first part of the theorem.

The proof is organized as follows. We first observe that $\Pi_{|\hat{\mathbb{M}}}$ is continuous (by Lemma 7.6). We then prove that $\Pi_{|\hat{\mathbb{M}}}$ embeds any star of $\hat{\mathbb{M}}$ (Lemma 7.16). We then prove that $\Pi_{|\hat{\mathbb{M}}}$ is surjective (Lemma 7.17) and injective (Lemma 7.18). Since $\hat{\mathbb{M}}$ is compact and $\mathbb{M}$ is Hausdorff, it follows that $\Pi_{|\hat{\mathbb{M}}}$ is a homeomorphism, and that $\hat{\mathbb{M}}$ and $\mathbb{M}$ are homeomorphic.

Lemma 7.16 (Π embeds the stars) *Under the same conditions as in Theorem 7.15, Π embeds the star of any $p \in P$ in $\hat{\mathbb{M}}$.*

Proof TO DO

□

The previous lemma states that the restriction of Ψ to the star of any point $p \in P$ in $\hat{\mathbb{M}}$ is injective and therefore a local homeomorphism onto its image. We now prove that Ψ is surjective.

Lemma 7.17 (Surjectivity of $\Pi_{|\hat{\mathbb{M}}}$) *Under the same conditions as in Lemma 7.16, $\Pi_{|\hat{\mathbb{M}}}$ is surjective on $\mathbb{M}$.*

Proof By the invariance of domain theorem, $\Pi(\hat{\mathbb{M}})$ has no boundary since $\hat{\mathbb{M}}$ has no boundary by hypothesis. To prove that it is equal to $\mathbb{M}$, it remains to show that $\hat{\mathbb{M}}$ has vertices on all the connected components of $\hat{\mathbb{M}}$.

Since the vertex set of $\hat{\mathbb{M}}$ is an ε-sample of $\mathbb{M}$, every connected component C of $\hat{\mathbb{M}}$ contains at least one vertex that is incident to a facet of $\hat{\mathbb{M}}$ (see Exercise 7.7). This ensures that $\Pi(\hat{\mathbb{M}})$ contains that vertex. Hence, $\Pi(\hat{\mathbb{M}})$ intersects all the connected components of $\mathbb{M}$. □

Lemma 7.18 (Injectivity of Ψ) *Under the same conditions as in Theorem 7.15, Π is injective on $\hat{\mathbb{M}}$.*

Proof Let x be a point of $\Pi(\hat{\mathbb{M}})$. $\Pi^{-1}_{|\hat{\mathbb{M}}}(x)$ is non empty and finite since Π is 1-1 on each simplex (Lemma 7.16) and there are only a finite number of simplexes in $\hat{\mathbb{M}}$. For each point $y \in \Pi^{-1}(x) \cap \hat{\mathbb{M}}$, we choose an open neighborhood U_y, homeomorphic to a disk and small enough so that U_y is contained only in simplices that contain y. U_y is therefore entirely contained

in the star of some point q, star$(\hat{q})$. Since, as already noticed, $\Pi_{|\hat{\mathbb{M}}}$ is a local homeomorphism that maps the star of any vertex q in $\hat{\mathbb{M}}$ homeomorphically onto its image, each U_y is mapped homeomorphically onto its image. Hence the preimage under $\Pi_{|\hat{\mathbb{M}}}$ of any sufficiently small open neighborhood U_x of a point x of $\mathbb{M}$ is a union of disjoint open sets each of which is contained in a star of $\hat{\mathbb{M}}$ and mapped homeomorphically onto $U(x)$ by Π. In topological terms, $(\hat{\mathbb{M}}, \Pi)$ is a covering space of $\Pi(\hat{\mathbb{M}})$.

A standard result in topology then asserts that Ψ covers all the points of a connected component of its image $\Psi(\hat{\mathbb{M}})$ the same number of times. This number is 1 by hypothesis.

can we remove this hypothesis ?

□

This ends the proof of Theorem 7.15. □

7.5 Exercises

Exercise 7.1 (Sparse samples) Let P be an ε-sample of a submanifold $\mathbb{M}$ 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 ε, insert p in P'. Show that P' is an $(\varepsilon, 1)$-net of P and a $(2\varepsilon, \frac{1}{2})$-net of $\mathbb{M}$.

Exercise 7.2 (Bound on $\bar{\mu}$) Let P be an $(\varepsilon, \bar{\mu})$-net of a submanifold $\mathbb{M}$ of positive reach. Show that $\bar{\eta} \leq 2$ (Hint : consider a maximal empty d-ball point centered on $\mathbb{M}$)

Exercise 7.3 (Non uniform samples) Using the fact that the local feature size lfs is a 1-Lipschitz function, extend the results of this chapter to sufficiently dense non uniform ε-nets as defined in Section 7.1.2.

Exercise 7.4 (Computing rch$(\mathbb{M})$) Let $\mathbb{M}$ be a smooth submanifold and write ax$(\mathbb{M})$ for the medial axis of $\mathbb{M}$. Show that rch$(\mathbb{M})$ is the radius of a ball $B(m)$ centered at a point $m \in$ ax$(\mathbb{M})$ and tangent to $\mathbb{M}$ in at least two points (called contact points). Show that m is a critical point of the distance function to $\mathbb{M}$. According to Exercise 11.6, m is contained in the

convex hull of the contact points $B(m) \cap \mathbb{M}$. In particular, if $B(m)$ has only one contact point, $B(m)$ is osculating $\mathbb{M}$ at the contact point. If $B(m)$ has two distinct contact points, the two contact points are the endpoints of a diameter of $B(m)$. Show how to compute $\mathrm{rch}(\mathbb{M})$ when $\mathbb{M}$ 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 7.5 Let x be a point of $B_M = B(p, r) \cap \mathbb{M}$ where $r = \rho\,\mathrm{rch}(\mathbb{M})$, $\rho < 1$. Show that there exists a point $y \in T_p$ such that $\Pi(y) = x$, where Π denotes the projection on $\mathbb{M}$, and $\|x - y\| \leq \frac{r}{1-\rho}$.

Exercise 7.6 (Size of nets) Show that the size of an $(\varepsilon, \bar{\eta})$-sample of a k-submanifold with a bounded sampling ratio $\bar{\eta}$ depends exponentially on k.

Exercise 7.7 (Distance between components) Let P be an ε-sample of a submanifold $\mathbb{M}$ of positive reach. Prove that each connected component of $\mathbb{M}$ contains at least one point of P and therefore at least one vertex of $\mathrm{Del}_{T\mathbb{M}}(P)$. (Hint : show that the distance between any two connected components of $\mathbb{M}$ is at least $2\mathrm{rch}(\mathbb{M})$).

Exercise 7.8 (Ambient isotopy) *Show that Ψ, the restriction to $\hat{\mathbb{M}}$ of the orthogonal projection onto $\mathbb{M}$, induces an ambient isotopy*

$$\Phi^* : \mathbb{R}^d \times [0, 1] \longrightarrow \mathbb{R}^d$$

such that the map $\Psi^*(\cdot, 0)$ restricted to $\hat{\mathbb{M}}$ is the identity map on $\hat{\mathbb{M}}$ and $\Phi^*(\hat{\mathbb{M}}, 1) = \mathbb{M}$. The isotopy does not move the points by more than $O(\varepsilon^2\,\mathrm{rch}(\mathbb{M}))$.

Proof 1. Let

$$\Phi : \hat{\mathbb{M}} \times [0, 1] \longrightarrow \mathbb{R}^3, \quad (x, t) \mapsto x + t\,(\Pi(x) - x)$$

Note that $\Phi(\cdot, 0)$ is an identity map on $\hat{\mathbb{M}}$ and $\Phi(\cdot, 1) = \Psi$. The map Φ is an isotopy because the maps

$$\Phi_t : \hat{\mathbb{M}} \longrightarrow \mathbb{R}^3, \quad x \mapsto \Phi(x, t)$$

are homeomorphisms between $\hat{\mathbb{M}}$ and $\Phi_t(\hat{\mathbb{M}})$.

Isotopy Φ can be extended to an ambient isotopy $\Phi^* : \mathbb{R}^3 \times [0, 1] \longrightarrow \mathbb{R}^3$ such that $\Phi^*(\cdot, 0) \mid_{\hat{\mathbb{M}}} = \Phi(\cdot, 0)$ and $\Phi^*(\cdot, 1) \mid_{\hat{\mathbb{M}}} = \Phi(\cdot, 1)$ (see, e.g. [81]).

2. Since any facet $f \in \hat{\mathbb{M}}$ is contained in a $\mathbb{M}$-centered Delaunay ball of radius at most ε, the isotopy does not move the points by more than ε. □

7.6 Bibliographical notes

The notion of reach has been introduced by Federer [67] who proved Lemmas 7.5 and Exercise 7.6. The related notion of local feature size was introduced by Amenta and Bern in their seminal paper on surface reconstruction [4]. See also[21, 75].

The notion of thick triangulations goes back to the early work on differential topology by Cairns [26], Whitehead [108], Whitney [109], Munkres [93]. Thick triangulations also play a central role in the work of Cheeger et al. [46] and Fu [70] 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 [18]. The proof of Lemma 7.14 is due to Whitney [109]. More recently, thick triangulations have been found important in mesh generation where numerical simulations require meshes to be thick [66].

Variants of Theorem 7.15 have been proved by Cheng and al. [], Boissonnat and Ghosh [19] and Dyer et al. [57].

Chapter 8

Alpha complexes and homological reconstruction

Est ce que le titre ne devrait as plutot etre "Alpha complexes and homotopy reconstruction.

Alpha shapes, hereafter often called α-shapes for short, are families of shapes defined from sets of points. They closely related to subcomplexes of the Delaunay complex called the α-complexes. Weighted alpha shapes, or weighted α-shapes, related to subcomplexes of the weighted Delaunay complex are defined similarly from sets of weighted points Those families of shapes are parameterized by a scalar usually called α, hence their name.

α-shapes and weighted α-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. Owing to their capacity of representing union of balls, cavities and pockets formed by union of balls, α-shapes play an important role in the description of proteins and macro molecules, docking studies and drug design. As we show in this chapter, α-shapes have the essential property to capture the homotopy type of union of balls and this property stands in any dimension. Such a property becomes even more interessing for geometric inference in high dimension when related to a theorem by Niyogi et al [95] stating that, under some conditions on regularity and sampling, manifolds have the same homotopy types than union of balls centered on the sample. Furthermore, since α-shapes and α-complexes are parametrized

by the scalar α, they allow to define filtrations. Filtrations, defined in chapter **??**, are special sequences of nested complexes that may be used to infer the homology groups of sampled manifolds, as will be seen in chapter 13.

8.1 Definitions

Alpha complexes. Let P be a set of points in $\mathbb{R}^d$. From Lemma 5.1 we know that simplices in the Delaunay complex $\mathrm{Del}(P)$ are characterized by the *empty ball property* meaning that a simplex with vertices in P belong to $\mathrm{Del}(P)$ 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 out the simplices of $\mathrm{Del}(P)$ admitting small empty circumscribing balls. More precisely, for any $\alpha \in \mathbb{R}$, we consider the subset $\mathcal{A}(P, \alpha)$ of simplices in $\mathrm{Del}(P)$ admitting a circumscribing ball with square radius at most α. Because a ball circumscribing a simplex circumscribes any face of this simplex, $\mathcal{A}(P, \alpha)$ is a subcomplex of $\mathrm{Del}(P)$. It is called the α-complex.

The α-complex evolves when α increases, from the empty set for $\alpha < 0$, to the set of vertices of $\mathrm{Del}(P)$ when $\alpha = 0$, and finally to the whole Delaunay triangulation $\mathrm{Del}(P)$ when α is big enough.

The α-shape of the set P for the parameter value α is simply defined as the underlying space of the α-complex $\mathcal{A}(P, \alpha)$.

Under general position assumption, the maximal dimension of the complex $\mathcal{A}(P, \alpha)$ is the dimension d of the embedding space. However the α-complex may have a lower dimension. It may also not be a pure complex, having some simplices which are not faces of simplices of maximal dimension. Some applications, requiring a pure complex consider the *regularized version* of the α-complex defined as the subcomplex $\mathcal{A}_r(\mathcal{P}, \alpha)$ of $\mathcal{A}(\mathcal{P}, \alpha)$ formed by the simplices of maximal dimensions and their faces.

Weigthed alpha complex. The definition of α-complexes and α-shapes extend to the weighted case. Let $\hat{P}$ be a set of weighted points. From Lemma 6.1, a simplex τ with vertex set $\hat{P}_\tau \subset \hat{P}$ belongs to the weighted Delaunay triangulation $\mathrm{Del}(\hat{P})$ iff there is a weighted point orthogonal to $\hat{P}_\tau$ and further than orthogonal to $\hat{P} \setminus \hat{P}_\tau$. For any value of $\alpha \in \mathbb{R}$, we consider the subset $\mathcal{A}(\hat{P}, \alpha)$ of simplices in $\mathrm{Del}(\hat{P})$ for which there is a weighted

point with weight at most α, orthogonal to $\hat{P}_\tau$ and further than orthogonal to $\hat{P} \setminus \hat{P}_\tau$. The simplices in $\mathcal{A}(\hat{P}, \alpha)$ form a subcomplex of $\mathrm{Del}(\hat{P})$ which is called the weighted α-complex. The weigthed α-shape of the set $\hat{P}$ for the parameter value α is still defined as the underlying space of the weighted α-complex $\mathcal{A}(\hat{P}, \alpha)$.

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 weigthed point set are equal. Therefore in the following we call simply α-complexes and α-shapes, complexes and shapes obtained from weighted or non weighted point sets.

The alpha filtration. Filtrations have been defined in Section 4.4 as subsequences of nested complexes such that each complex is obtained from the previous one by the addition of a single simplex. Let P be a set of weighted points, and $\mathrm{Del}(P)$ the Delaunay complex of P. If we order the simplices of $\mathrm{Del}(P)$ according to the value α at which they enter the α-complex $\mathcal{A}(P, \alpha)$, breaking ties according to increasing dimensions, we get a filtration of $\mathrm{Del}(P)$, called the α-filtration. The sequence of complexes in the α-filtration of P run over all the subcomplexes of $\mathrm{Del}(P)$ that can be obtained as an α-complex of P for a certain value of the parameter α. The α-filtration of a set of weighted points $\hat{P}$ is obtained in the same way by ordering the simplices of the regular complex $\mathrm{Del}(\hat{P})$ according to the value α at which they enter the α-complex $\mathcal{A}(\hat{P}, \alpha)$

8.2 Computing alpha complexes and filtrations

The unweighted case. Consider first, for simplicity, the case of unweighted points. Let P be a set of points in $\mathbb{R}^d$ and $\mathrm{Del}(P)$ 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 $\mathrm{Del}(P)$ 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 now a simplex of dimension $k < d$. The centers of the

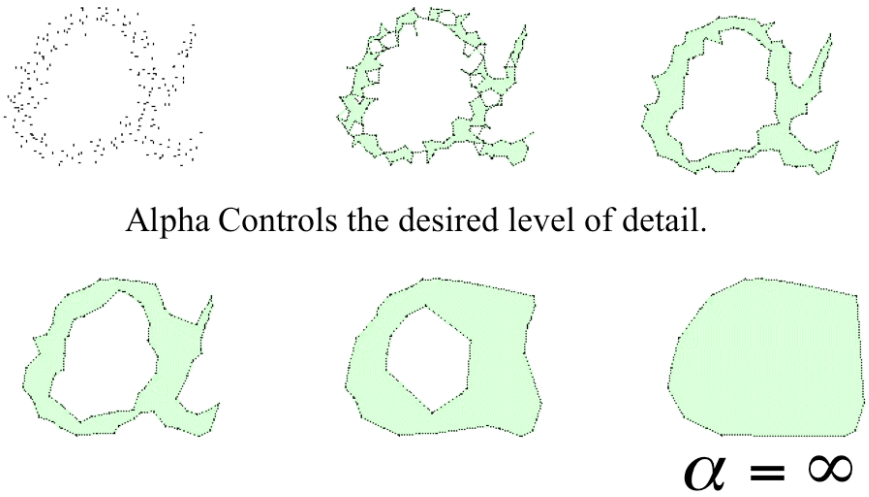

Figure 8.1: α-shapes in 2D. Courtesy of

d-balls circumscribing τ form a $(d-k)$-flat $h(\tau)$ in $\mathbb{R}^d$. We consider the circumscribing ball of τ with smallest radius. This ball denoted by $b(c(\tau), r(\tau))$ is often called the *smallest circumball* , the center $c(\tau)$ is called the center of τ and the radius $r(\tau)$), the smallest circumradius of τ. The center $c(\tau)$ is the point where $h(\tau)$ intersects the k-flat aff(τ) spanned by τ. A k-simplex τ of Del(P) is said to be a *Gabriel* simplex if its smallest circumball, $b(c(\tau), r(\tau))$, includes no point of P in its interior. Equivalently the k-simplex τ of Del(P) is a Gabriel simplex iff $c(\tau)$ belongs to the face $V(\tau)$ of the Voronoi diagram Vor(P) that is the dual face of τ.

Therefore, if τ is a *Gabriel* simplex, it belongs to the α-complex $\mathcal{A}(P, \alpha)$ iff $\alpha \geq r(\tau)^2$. Otherwise we claim that τ enters in the α-complex for the same value of α as one of its cofaces. To prove this, we consider the function $w(x)$ giving fo each point x in $h(\tau)$ its squared distance to the vertices in τ. Finding the critical value $\alpha(\tau)$ at which τ enters the α-complex $\mathcal{A}(P, \alpha)$ amounts to minimizing $w(x)$ with the condition that $x \in V(\tau)$. Since $w(x)$ is a convex fonction and $V(\tau)$ is a convex polytope 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 τ whose dual Voronoi face contains $c(\tau)$ in its interior, we have $\alpha(\tau) = \alpha(\sigma)$.

For any $\tau \in$ Del(P) with dimension $k < d$, we denote by $U(\tau)$ the set of

cofaces of τ in Del(P) with dimension $k+1$. Algorithm 5 computes for each simplex τ in Del(P), the critical value $\alpha(\tau)$ at which τ enters the α- complex $\mathcal{A}(P, \alpha)$.

Algorithm 5 Computing α-complexes and filtration

Input: the set of points P in $\mathbb{R}^d$
Compute the Delaunay complex Del(P)
for each d-simplex $\tau \in$ Del(P) **do**
  set $\alpha(\tau) = r(\tau)^2$ (the squared circumradius of τ)
for $k = d-1, ..., 0$ **do**
  for each d-simplex $\tau \in$ Del(P) **do**
    if the 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 Del(P) has been computed

The weighted case. The above considerations and algorithm to compute the critical α-values of simplexes extend almost verbatim to the case of weighted points, provided that we replace the circumballs of simplexes by balls orthogonal to the weighted points associated to the vertices of the simplex. Thus, if $\hat{P}$ is a set of weighted points in $\mathbb{R}^d$, a simplex τ of Del($\hat{P}$) is a Gabriel simplex iff the ball with smallest radius orthogonal to the vertices of τ (more precisely to the weighted points associated to the vertices of τ) is further than orthogonal to any other weighted point in $\hat{P}$. Likewise, the function $w(x)$ used to justify the algorithm 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.

8.3 Application to union of balls

Lemma 8.1 *Let B be a 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 8.2 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 $\mathrm{Del}(B)$, $\mathrm{Vor}(B)$ and $\mathcal{A}(B, \alpha)$ respectively the Delaunay complex, Voronoi diagram and α-shape of the balls in B. Let $V(b)$ the cell of b in $\mathrm{Vor}(B)$. We show here that $\{b \cap V(b), b \in B\}$ forms a good closed cover of $U(B)$ as defined in Section 4.3, i.e. a closed cover such that the intersection of any finite subsets of the cover is either empty or contractible.

Est ce qu'on choisit un open cover or a closed cover here? Closed seems more easy, as voronoi cells are defined as closed, however what do we do about the required assumption on the union. From Section 4.3, the union has to be homeomorphic to a finite simplicial complex ? MY

First, we show 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 a 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 $b(p) \cap V(b(p))$, hence to $\bigcup_{b \in B} b \cap V(b)$.

Let us now show that $\{b \cap V(b), b \in B\}$ forms a good closed cover of $U(B)$. Indeed each set in $\{b \cap V(b), b \in B\}$ is convex as it is the intersection of two convex sets. It is therefore contractible and the same is true for any finite intersection of subsets in $\{b \cap V(b), b \in B\}$.

It then follows from the Nerve theorem (Theorem 4.4) 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 that there is a point x in $\bigcap_{b \in B'} b \cap V(b)$ or a point x having an equal negative weighted distance $w(x)$ to balls in B' and a greater weighted distance to any other ball of B. This means that the weighted point $\hat{x} = (x, w(x))$ is orthogonal to any ball in B' and further than

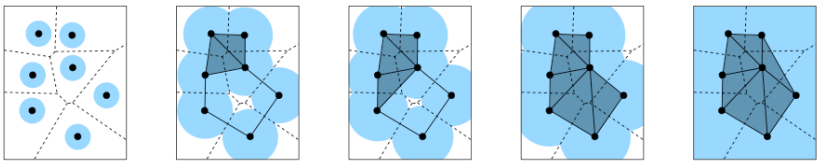

Figure 8.2: Homotopy between union of balls and α-shapes. Courtesy of ?

orthogonal to any ball in $B \setminus B'$ and therefore that B' belongs to $\mathrm{Vor}(B)$ 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 the Čech complex $\check{\mathrm{C}}\mathrm{ech}(P, r)$ defined (in Section 4.5) as the nerve of the cover of $U(B(P, r))$ by balls in $B(P, r)$. Therefore, the Čech complex $\check{\mathrm{C}}\mathrm{ech}(P, r)$ 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 $\check{\mathrm{C}}\mathrm{ech}(P, r)$ is often much bigger than the α-complex $\mathcal{A}(B(P, r), 0)$. In particular, the dimension of the Čech complex $\check{\mathrm{C}}\mathrm{ech}(P, r)$ may be larger than d, in fact as large as the number of balls in $B(P, r))$. The Čech complex $\check{\mathrm{C}}\mathrm{ech}(P, r)$ is likely not to be naturally embedded in $\mathbb{R}^d$. Differently, the α-complex $\mathcal{A}(B(P, r), 0)$, being a subcomplex of $\mathrm{Del}(B)$, is, under general position assumption, naturally embedded in $\mathbb{R}^d$.

8.4 Homology reconstruction

Unions of balls play a central role in manifold reconstruction. Indeed, while the previous section 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 $\mathbb{M}$ with positive reach can be obtained from union of balls centered at the points of a sample of $\mathbb{M}$ provided that the sample is dense enough with respect to the reach.

Theorem 8.2 *Let $\mathbb{M}$ be a manifold with positive reach* $\mathrm{rch}(\mathbb{M})$*, and $P \subset$*

$\mathbb{M}$ be a point sample of $\mathbb{M}$ with sampling radius $\frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M})$, meaning that any point x of $\mathbb{M}$ is at distance less than $\frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M})$ from the closest sample point. If $\varepsilon < \sqrt{\frac{3}{5}}$, the union of balls with radius $\varepsilon\,\mathrm{rch}(\mathbb{M})$ centered on P, is homotopy equivalent to $\mathbb{M}$.

Proof In the following, we denote by $b(p,r)$ the ball with radius r centered on p, by $B(P,\varepsilon\mathrm{rch}(\mathbb{M})))$ the set of balls with radius r centered on points of P and we write for short U for the union of balls with radius $\varepsilon\mathrm{rch}(\mathbb{M})$ centered on P:

$$U = U(B(P,\varepsilon\mathrm{rch}(\mathbb{M}))) = \bigcup_{p\in P} b(p,\varepsilon\mathrm{rch}(\mathbb{M})).$$

Obviously, $\mathbb{M}$ is included in U. To prove the homotopy equivalence we prove below that U deformation retracts to $\mathbb{M}$. 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 $\mathbb{M}$. F is continuous from Lemma 7.6. For all $x \in U$, $F(x,0) = x$, and $F(x,1) = \Pi(x)$ is in $\mathbb{M}$, and for all $x \in \mathbb{M}$ and $t \in [0,1]$, $F(x,t) = x$. Therefore F is a deformation retracts from U to $\mathbb{M}$, 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 $y \in \mathbb{M}$ is

$$\Pi_U^{-1}(y) = N_y \cap U \cap b(y,\mathrm{rch}(\mathbb{M})), \tag{8.1}$$

where N_y is the normal subspace (orthogonal to the tangent space) of $\mathbb{M}$ at y. The ball with radius $\mathrm{rch}(\mathbb{M})$ centered at y, $b(y,\mathrm{rch}(\mathbb{M}))$, 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 $\mathbb{M}$ is curved. Therefore,

$$\Pi_U^{-1}(y) = \bigcup_{p\in P} b(p,\varepsilon\mathrm{rch}(\mathbb{M})) \cap N_y \cap b(y,\mathrm{rch}(\mathbb{M})).$$

We also consider the subset $st(y)$ defined as

$$st(y) = \bigcup_{p\in P\cap b(y,\varepsilon\mathrm{rch}(\mathbb{M}))} b(p,\varepsilon\mathrm{rch}(\mathbb{M})) \cap N_y \cap b(y,\mathrm{rch}(\mathbb{M})).$$

Obviously, $st(y) \subset \Pi_U^{-1}(y)$. Then Lemma 8.3 below proves that $st(y)$ is star shaped with respect to y and Lemma 8.4 proves that $st(y) = \Pi_U^{-1}(y)$. It follows that for any $y \in \mathbb{M}$, $\Pi_U^{-1}(y)$ is star shaped with respect to y and that for any $(x,t) \in U \times [0,1]$, $F(x,t)$ belongs to U. □

Lemma 8.3 *The subset $st(y)$ is star shaped with respect to y.*

Proof Let z be an arbitrary point in $st(y)$. Then $z \in b(p, \varepsilon \text{rch}(\mathbb{M})) \cap N_y \cap b(y, \text{rch}(\mathbb{M}))$ for some $p \in P \cap b(y, \varepsilon \text{rch}(\mathbb{M}))$. Since $p \in b(y, \varepsilon \text{rch}(\mathbb{M}))$, $y \in b(p, \varepsilon \text{rch}(\mathbb{M}))$. Since z and y are both in $b(p, \varepsilon \text{rch}(\mathbb{M}))$, the segment zy is entirely contained in $b(p, \varepsilon \text{rch}(\mathbb{M}))$. At the same time, zy is entirely contained in N_y and in $b(y, \text{rch}(\mathbb{M}))$ and therefore in $st(y)$. □

Lemma 8.4 *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 $b(p, \varepsilon \text{rch}(\mathbb{M})) \cap N_y \cap b(y, \text{rch}(\mathbb{M}))$ where p is a point of P such that $p \notin b(y, \varepsilon \text{rch}(\mathbb{M}))$. Lemma 8.5 shows that the distance from z to y is at most $\varepsilon^2 \text{rch}(\mathbb{M})$ and Lemma 8.6 shows that if P is $(\frac{\varepsilon}{2}\text{rch}(\mathbb{M}))$-dense in $\mathbb{M}$ with $\varepsilon < \frac{3}{5}$, there is some point $q \in P \cap b(y, \varepsilon \text{rch}(\mathbb{M}))$ such that $z \in b(q, \varepsilon \text{rch}(\mathbb{M}))$ which achieves the proof. □

Lemma 8.5 *Let z be a point in $b(p, \varepsilon \text{rch}(\mathbb{M})) \cap N_y \cap b_{\text{rch}(\mathbb{M})}(y)$ where p is a point of P such that $p \notin b(y, \varepsilon \text{rch}(\mathbb{M}))$. The distance $d(y,z)$ from z to y is at most $\varepsilon^2 \text{rch}(\mathbb{M})$.*

Proof Let us consider the plane H through y, z and p. See Figure 8.3. The tangent and orthogonal subspaces at y intersect H according to two orthogonal lines through y. Point z belongs to $H \cap N_y$. and point p lies anywhere outside of the two balls with radius $\text{rch}(\mathbb{M})$ tangent to T_y at y. Since the distance $d(y,z)$ is constrained by the fact that $d(p,z) \leq \varepsilon \text{rch}(\mathbb{M})$ it is obvious that the maximum of the distance $d(y,z)$ occurs in the configuration where points p and z lie on the same side of $T_y \cap H$, as shown in Figure 8.3. We note r' the radius of the circle tangent to $\mathbb{M}$ in y and going through p, and θ the angle $\angle(T_y, yp)$ between T_y and yp. We have

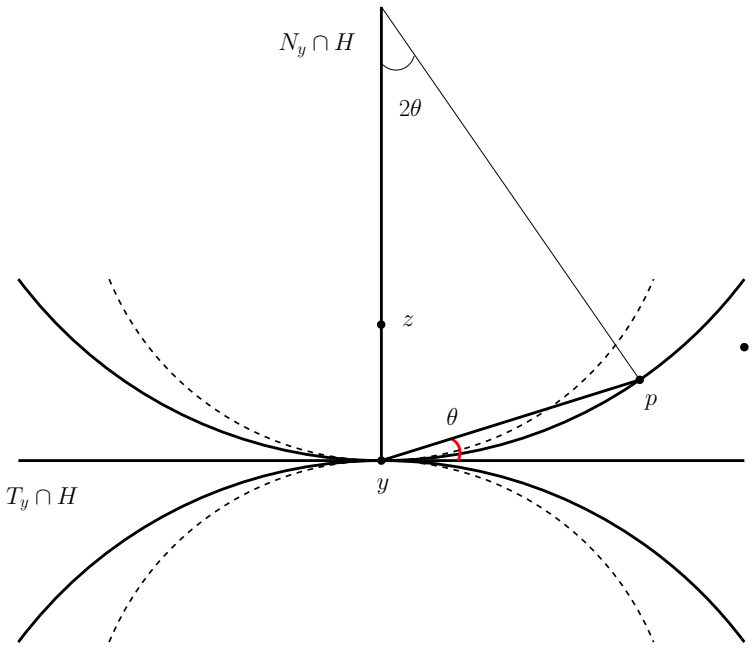

Figure 8.3: For the proof of Lemma 8.4.

$r' \geq \text{rch}(\mathbb{M})$ and $d(y,p) = 2r' \sin\theta$. Furthermore we have, see Figuree 8.3 :

$$\begin{aligned} d(y,z) &= d(y,p)\sin\theta + \sqrt{d(z,p)^2 - d(y,p)^2\cos^2\theta} \\ &= 2r'\sin^2\theta + \sqrt{d(z,p)^2 - r'^2\sin^2 2\theta} \\ &\leq 2r'\sin^2\theta + \sqrt{\varepsilon^2\text{rch}(\mathbb{M})^2 - r'^2\sin^2 2\theta}. \end{aligned} \tag{8.2}$$

Therefore $d(y,z)$ is upper bounded by the function $f(r',\theta)$ defined by the right member of Equation 8.2:

$$d(y,z) \leq f(r',\theta) \stackrel{\text{def}}{=} 2r'\sin^2\theta + \sqrt{\varepsilon^2\text{rch}(\mathbb{M})^2 - r'^2\sin^2 2\theta}. \tag{8.3}$$

We have:

$$\begin{aligned} \frac{df}{d\theta} &= 2r'\sin 2\theta - \frac{2r'^2\sin 2\theta\cos 2\theta}{\sqrt{\varepsilon^2\text{rch}(\mathbb{M})^2 - r'^2\sin^2 2\theta}} \\ &= 2r'\sin 2\theta\left(1 - \frac{r'\cos 2\theta}{\sqrt{\varepsilon^2\text{rch}(\mathbb{M})^2 - r'^2\sin^2 2\theta}}\right). \end{aligned}$$

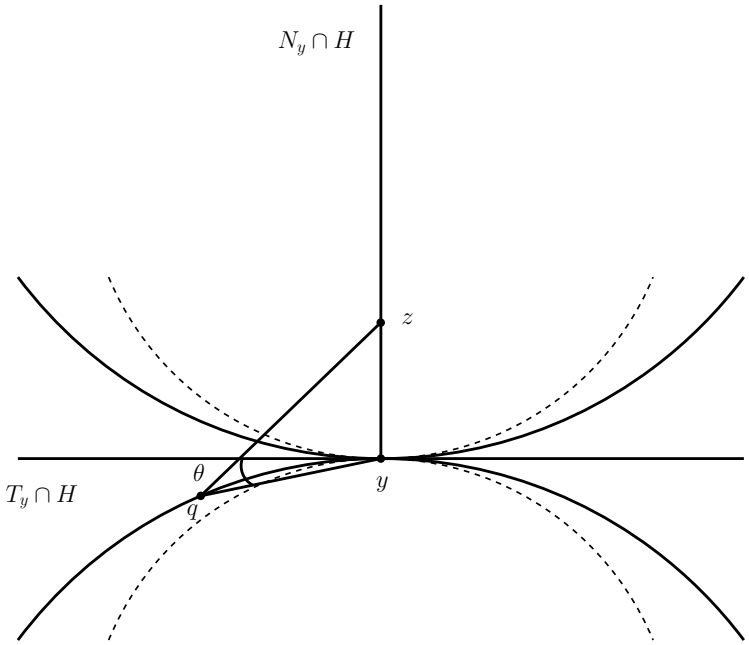

Figure 8.4: For the proof of Lemma 8.5.

Since $\varepsilon < \sqrt{\frac{3}{5}} < 1$ and $\mathrm{rch}(\mathbb{M}) \leq r'$, the function f is monotonically decreasing with respect to θ and reaches its maximum when θ is minimum, i. e. when $d(y,p) = 2r' \sin\theta = \varepsilon \mathrm{rch}(\mathbb{M})$. We have then

$$f(r',\theta) = \varepsilon \mathrm{rch}(\mathbb{M}) \sin\theta + \sqrt{\varepsilon^2 \mathrm{rch}(\mathbb{M})^2 - \varepsilon^2 \mathrm{rch}(\mathbb{M})^2 \cos^2\theta} \quad (8.4)$$

$$= 2\varepsilon \mathrm{rch}(\mathbb{M}) \sin\theta = \frac{\varepsilon^2 \mathrm{rch}(\mathbb{M})^2}{r'} \leq \varepsilon^2 \mathrm{rch}(\mathbb{M}), \quad (8.5)$$

which, owing to Equation 8.3 achieves the proof. □

Lemma 8.6 *Let z be a point in $b(p, \varepsilon \mathrm{rch}(\mathbb{M})) \cap N_y \cap b_{\mathrm{rch}(\mathbb{M})}(y)$ where p is a point of P such that $p \notin b(y, \varepsilon \mathrm{rch}(\mathbb{M}))$. If P is $(\frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M}))$-dense in $\mathbb{M}$ with $\varepsilon < \sqrt{\frac{3}{5}}\mathrm{rch}(\mathbb{M})$, 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 $(\frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M}))$-dense in $\mathbb{M}$, there is a point $q \in P$ at distance at most $\frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M})$ from y. We prove that $z \in b(q, \varepsilon)$. Figure 8.4 drawn in the plane through y, z and q shows the worst situation for the

distance $d(z,q)$. Denoting now by θ the angle $\angle(T_y, yq)$ between T_y and yq, we have

$$\begin{aligned} d(z,q)^2 &\leq d(y,q)^2\cos^2\theta + (d(y,q)\sin\theta + d(y,z))^2 \\ &\leq d(y,q)^2 + 2d(y,q)d(y,z)\sin\theta + d(y,z)^2. \end{aligned}$$

Since $d(y,q) \leq \frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M})$ and, from Lemma 8.5, $d(y,z) \leq \varepsilon^2\mathrm{rch}(\mathbb{M})$, we get

$$d(z,q)^2 \leq \frac{\varepsilon^2}{4}\mathrm{rch}(\mathbb{M})^2 + \varepsilon^3\mathrm{rch}(\mathbb{M})^2\sin\theta + \varepsilon^4\mathrm{rch}(\mathbb{M})^2.$$

We have $d(y,q) = 2r'\sin\theta \leq \frac{\varepsilon}{2}\mathrm{rch}(\mathbb{M})$ where r' is now the radius of the circle tangent to $\mathbb{M}$ in y and going through q. Since point q lies outside the two balls with radii $\mathrm{rch}(\mathbb{M})$ tangent to T_y at y, r' is greater than $\mathrm{rch}(\mathbb{M})$ and therefore: $\sin\theta \leq \frac{\varepsilon}{4}$. Thus,

$$\begin{aligned} d(z,q)^2 &\leq \frac{\varepsilon^2}{4}\mathrm{rch}(\mathbb{M})^2 + \frac{\varepsilon^4}{4}\mathrm{rch}(\mathbb{M})^2 + \varepsilon^4\mathrm{rch}(\mathbb{M})^2 \\ &\leq \varepsilon^2\mathrm{rch}(\mathbb{M})^2\left(\frac{1}{4} + \frac{5}{4}\varepsilon^2\right), \end{aligned}$$

which is not gretar than $\varepsilon^2\mathrm{rch}(\mathbb{M})^2$ if $\varepsilon^2 \leq \frac{3}{5}$. □

8.5 Exercises

Exercise 8.1 (Computing regularized α-complexes.) *Let $\hat{P}$ be a set of weighted points in $\mathbb{R}^d$. Modify slightly Algorithm 5 so that it computes for each simplex τ of the Delaunay triangulation* $\mathrm{Del}(\hat{P})$ *the critical value of paramater α at which the simplex τ becomes part of the regularized α-complex $\mathcal{A}_r(\hat{P},\alpha)$*

Exercise 8.2 (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* $\mathrm{Del}(\hat{P})$ *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* $\mathrm{Del}(\hat{P})$ *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 5 so that it computes for all simplex τ of $\mathrm{Del}(\hat{P})$*, the at most three critical values of paramater α where the status of the simplex change from external to singular and then to boundary and interior.*

8.6 Bibliographical notes

α-shapes were introduced by Edelsbrunner and Mücke [64, 58] and widely used in the early algorithms of shapes reconstruction from sets of data points measured on the surface of an object [65] . α-shapes are also famous for their ability to represent union of balls through the nerve theorems [59]. They are widely used in study of macro molecular structure and docking, see e.g [62, 87, 61] Theorem 8.2 about reconstructing the homotopy type of a manifold from a point sampling is due to Niyogi, Smale and Weinberger [95].

Chapter 9

Making Delaunay complexes thick

Although Delaunay triangulations have many beautiful properties, their simplices, in dimension greater than 2, may be arbitrarily flat 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 7.3) has been recognized since the early days of differential topology. They play a central role in many works on the triangulation of manifolds (see Part III) and appear to be crucial in numerical simulations to ensure the convergence of numerical methods solving partial differential equations.

We describe two ways of perturbing Delaunay complexes so that their simplices gain a guaranteed thickness. In Section 9.1, we perturb the metric and weight the points. Accordingly, we replace the Delaunay complex by its weighted version. In Section 9.2, instead of perturbing the metric, we perturb the position of the vertices of the complex. We introduce the related notion of *protection* that measures how far a point set is from being degenerate with respect to spheres. Protection implies thickness and we will show that Delaunay triangulations can be protected and therefore thickened by slightly perturbing their vertices.

To keep the exposition simpler, we will most of the time work in the standard flat torus $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ instead of $\mathbb{R}^d$ as we did in the other chapters. As a consequence, $\mathrm{Del}(P)$ where P is a finite set of points in $\mathbb{T}^d$ has no boundary. Boundary issues obscurate 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 far away from the boundary of the domain (see the exercise section).

9.1 Thickness from weighting

The notion of thickness introduced in Section 7.3 is an important measure of the shape of a simplex. A simplicial complex is thick when all its simplices are thick. It is well known that Delaunay triangulations are not necessarily thick complexes 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 affecting (relatively small) weights to the points of the net and considering the weighted Delaunay triangulation of the resulting set of weighted points.

As in the previous sections, 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 affects to each point $p \in P$ a weight $w(p) \in \mathbb{R}$. We note $\hat{P}$ the resulting set of weighted points, that is $\hat{P} = \{(p, w(p) : p \in P\}$, and $\mathrm{Del}(\hat{P})$ the corresponding weighted Delaunay triangulation.

We restrict the weighting scheme to positive 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 neighbour in $P \setminus p$. We also wish to ensure that each point of P appears as a vertex of $\mathrm{Del}(\hat{P})$. This is in particular achieved, if the spheres with equations $\{(x-p)^2 - w(p) : p \in P\}$ that represent the weighted points of $\hat{P}$, are bounding disjoint balls. Such disjoint balls are granted if, for each $p \in P$, $\sqrt{w(p)} < L(p)/2$. Therefore, we constraint the weighting scheme to 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, we can find a weighting scheme with relative amplitude smaller than $\tilde{w_0} < \frac{1}{4}$ and such that the weighted Delaunay triangulation $\mathrm{Del}(\hat{P})$ has no j-simplex with thickness less than Θ_0^j, for $j = 1, \ldots, d$.

In the following, we use the same notation σ for a 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 $\hat{P}$: $\hat{\sigma} = \{(p, w(p) : p \in \sigma\}$, that we call hereafter a *weighted (abstract) simplex*. If σ is a simplex with dimension j, the weighted points orthogonal to the vertices of $\hat{\sigma}$ are centered on an affine subspace of dimension $d-j$ orthogonal to the affine subpace aff(σ). 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 the vertices of $\hat{\sigma}$. Note that $(c(\hat{\sigma}), R^2(\hat{\sigma}))$ is the weighted point with minimal weight among the weighted point orthogonal to the vertices of $\hat{\sigma}$. The point $c(\hat{\sigma})$ is called the *(weighted) center* of the weigthed simplex $\hat{\sigma}$ and the $R(\hat{\sigma})$) is called *the smallest orthogonal radius* or more simply *the radius* of the weigthed simplex $\hat{\sigma}$. We begin with an easy lemma about the radii of weighted simplexes in Del$(\hat{P})$.

Lemma 9.1 (Radii of weighted simplexes in Del$(\hat{P})$**)** *If $\hat{P}$ is a set of weighted points obtained by assigning positive weights to the points of an $(\varepsilon, \bar{\eta})$-net P, the radius of any weighted simplex of* Del$(\hat{P})$ *is at most ϵ.*

Proof Let $\hat{\sigma}$ be a weighted simplex of Del$(\hat{P})$. The weighted simplex $\hat{\sigma}$ is included in some maximal weighted simplex $\hat{\tau}$ of Del$(\hat{P})$ that correspond to a simplex τ with dimension d. The weighted radius $R(\hat{\sigma})$ is at most the weighted radius $R(\hat{\tau})$ and we prove that $R(\hat{\tau})$ is at most ϵ. Indeed otherwise, 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 of P whose weighted distance to $(c(\hat{\tau}), R^2(\hat{\tau}))$ is negative, contradicting the fact that $\hat{\sigma}$ belongs to Del$(\hat{P})$. □

9.1.1 Θ_0-thickness and flakes

We need here to somehow classify non-thick simplices. A simplex that is not thick has a relatively small altitude, here we will focus 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 9.2 (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 $j < 2$.

Ensuring that all simplices are Θ_0-thick is the same as ensuring that there are no flakes. Indeed, if σ is not Θ_0-thick, then it is either a flake or it has a proper j-face $\sigma_j \subseteq \sigma$ that is not Θ_0-thick. By considering such a face with minimal dimension, we arrive at the following observation:

Lemma 9.3 *A simplex is not Θ_0-thick if and only if it has a face that is a Θ_0-flake.*

We obtain an upper bound on the altitudes of a Θ_0-flake through a consideration of dihedral angles. In particular, we observe the following general relationship between simplex altitudes:

Lemma 9.4 *If σ is a j-simplex with $j \geq 2$, then for any two vertices $p, q \in \sigma$, the dihedral angle between $\sigma_p = \sigma \setminus \{p\}$ and $\sigma_q = \sigma \setminus \{q\}$ defines an equality between ratios of altitudes:*

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

Proof Let $\sigma_{pq} = \sigma_p \cap \sigma_q$, and let p' and p'' be the projections of p into $\mathrm{aff}(\sigma_q)$ and $\mathrm{aff}(\sigma_{pq})$ respectively. We have $\|p - p'\| = \|p - p''\| \sin \angle(\mathrm{aff}(\sigma_p), \mathrm{aff}(\sigma_q))$. Since $\|p - p'\| = D(p, \sigma)$ and $\|p - p''\| = D(p, \sigma_q)$, the first equality is proved. A symmetric argument is carried out with q to obtain the result. □

We arrive at the following important observation about flake simplices:

Lemma 9.5 (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*

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

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

Proof Recalling Lemma 9.4 we have

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

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

$$D(q,\sigma) = k\Theta(\sigma)\Delta(\sigma) < k\Theta_0^k\Delta(\sigma),$$
$$D(q,\sigma_p) \geq (k-1)\Theta(\sigma_p)\Delta(\sigma_p) \geq (k-1)\Theta_0^{k-1}L(\sigma),$$
$$D(p,\sigma_q) \leq \Delta(\sigma_q) \leq \Delta(\sigma),$$

and the bound is obtained. □

9.1.2 The weight range of a flake with small radius

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

Lemma 9.6 *Let σ be a Θ_0 flake and assume a positive weighting scheme on the vertices of σ. If the radius $R(\hat{\sigma})$ of the weighted simplex $\hat{\sigma}$ is smaller than ε, the weight $w(p)$ of any vertex p of σ belongs to an interval $I(\sigma,p)$ whose measure $|I(\sigma,p)|$ is such that:*

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

In particular, if the vertices of σ belongs to an $(\varepsilon,\bar{\eta})$-net, the measure of $I(\sigma,p)$ is such that:

$$|I(\sigma,p)| \leq |I| \stackrel{\text{def}}{=} \frac{16\epsilon^2}{\bar{\eta}}\Theta_0.$$

Proof Let $c(\hat{\sigma})$ and $R(\hat{\sigma})$ be respectively the center and radius of the weighted simplex $\hat{\sigma}$. Likewise, we use $c(\hat{\sigma}_p)$ and $R(\hat{\sigma}_p)$ for respectively the center and the radius of $\hat{\sigma}_p$, where $\hat{\sigma}_p$ is just $\hat{\sigma} \setminus \hat{p}$ as usual. See Figure 9.1. We have :

$$R^2(\hat{\sigma}) = d^2(c(\hat{\sigma}), c(\hat{\sigma}_p)) + R^2(\hat{\sigma}_p), \tag{9.1}$$

Let us now introduce the subspace $N(\hat{\sigma}_p)$ that is the affine subspace of points with equal weighted distances to the vertices of σ_p and the distance $d(p, N(\hat{\sigma}_p))$ from p to $N(\hat{\sigma}_p)$. We have:

$$R^2(\hat{\sigma}) + w(p) = d^2(p, c(\hat{\sigma})) = d^2(p, N(\hat{\sigma}_p)) + \left(D(p,\sigma) - H(p,\hat{\sigma})\right)^2, \tag{9.2}$$

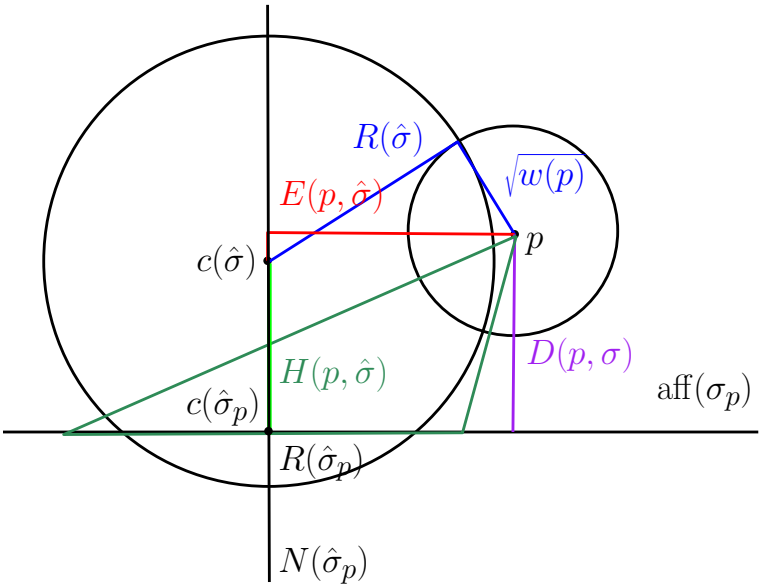

Figure 9.1: For the proof of Lemma 9.6. MY: notations on the figure to be fixed

where $D(p,\sigma)$ is as usual 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 aff($sigma_p$) and $H(p,\hat{\sigma}) = -d(c(\hat{\sigma}), c(\hat{\sigma}_p))$ otherwise.

Using Equations 9.1 and 9.2 together, we get:

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

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

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

Hence,

$$|w(p) - F(p,\hat{\sigma})| \leq 2D(p,\sigma)d(c(\hat{\sigma}), c(\hat{\sigma}_p)). \tag{9.3}$$

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

From Lemma 9.5, $D(p,\sigma) \leq \Theta_0 \frac{\Delta(\sigma)^2}{L(\sigma)}$. From Equation 9.1, we get that $d(c(\hat{\sigma}), c(\hat{\sigma}_p)$ is at most $R(\hat{\sigma})$ which is at most ε by hypothesis. Therefore

$$|w(p) - F(p,\hat{\sigma})| \quad \leq \quad 2\frac{\Delta(\sigma)^2}{L(\sigma)}\Theta_0\varepsilon, \tag{9.4}$$

which achieves to prove the lemma.

In the case where σ belongs to an $(\varepsilon,\bar{\eta})$-net, $\Delta(\sigma) \leq 2\varepsilon$ and $L(\sigma) \geq \bar{\eta}\varepsilon$.

□

9.1.3 The number of flakes incident to p

Assume that the weights of all points in P are assigned except the weight of point p. Our goal here is to show that it is possible to choose for p a weight less than $\tilde{w_0}L^2(p)$ and such that $\mathrm{Del}(\hat{P})$ includes no Θ_0-flake incident to p.

Let $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$ denote the set of all simplices $\sigma \in P$ such that $\hat{\sigma}$ belongs to $\mathrm{Del}(\hat{P})$ for some weighting scheme with relative amplitude smaller than $\tilde{w_0}$. We wish to bound the number of simplices of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$ incident to p, which gives a crude bound on the number of flakes of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$ that are incident to p.

We first bound the number of points in P linked to p by an edge of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$.

Lemma 9.7 *Assume that P is an $(\varepsilon,\bar{\eta})$-net and that weighting schemes are restricted to relative amplitude smaller than $\tilde{w_0} < 1/4$. Then, the number N_1 of edges incident to p in $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$ satisfies:*

$$N_1 \leq \left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^d \tag{9.5}$$

Proof Let $\tau = pq$ be an edge of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w_0})$. Then, for some weighting scheme, the weighted simplexe $\hat{\tau}$ belongs to $\mathrm{Del}(\hat{P})$ and therefore is included in at least a d-simplex $\hat{\sigma}$ of $\mathrm{Del}(\hat{P})$. From lemma 9.1, $\hat{\sigma}$ has a radius $R(\hat{\sigma})$

that is at most ε by. Let $c(\hat{\sigma})$ be the center of $\hat{\sigma}$. We have:

$$\begin{aligned} d(p,q) &\leq d(p,c(\hat{\sigma})) + d(c(\hat{\sigma}),q) \\ &\leq \sqrt{R^2(\hat{\sigma}) + w(p)} + \sqrt{R^2(\hat{\sigma}) + w(q)} \\ &\leq \sqrt{\varepsilon^2 + \tilde{w}_0 L^2(p)} + \sqrt{\varepsilon^2 + \tilde{w}_0 L^2(q)} \\ &\leq 2\sqrt{\varepsilon^2 + 4\tilde{w}_0\varepsilon^2} \\ &\leq 2\sqrt{2}\varepsilon. \end{aligned}$$

Therefore, if pq is an edge of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w}_0)$ point q belongs to a ball with radius $R = 2\sqrt{2}\varepsilon$ centered in p. As P is $\varepsilon\bar{\eta}$-separated, each point of P may be surrounded by disjoint balls with radii $r = \frac{1}{2}\varepsilon\bar{\eta}$. When pq is an edge of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w}_0)$ the ball with radius r around q is included in a ball with radius $R + r$ centered in p and, by a volume argument, the number of N_1 of edges incident to p in $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w}_0)$ satisfies:

$$N_1 \leq \left(\frac{R+r}{r}\right)^d = \left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^d.$$

□

Lemma 9.8 *If P is an $(\varepsilon,\bar{\eta})$-net and if the weighting schemes are restricted to relative amplitude smaller than $\tilde{w}_0 < 1/4$, the number N of simplices incident to p in $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w}_0)$ satisfies:*

$$N \leq 2\left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^{d^2} \tag{9.6}$$

Proof N_1 is the number of edges of $\mathrm{U}_{\mathrm{Del}}(P,\tilde{w}_0)$ incident to p. The number if j-simplexes incident to p is at most N_1^j and the number N of simplices incident to p is less than $\Sigma_{i=1}^{i=d} N_1^j$.

$$\begin{aligned} N &\leq \Sigma_{i=1}^{i=d} N_1^j = N_1 \frac{N_1^d - 1}{N_1 - 1} \leq 2N_1^d \\ &\leq 2\left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^{d^2}. \end{aligned}$$

□

Lemma 9.9 *Assume that P is an $(\varepsilon, \bar{\eta})$-net and that weighting schemes are restricted to relative amplitude smaller than $\tilde{w_0} < 1/4$. If the weight of a point p is reassigned randomly using a uniform probability in $[0, \tilde{w_0} L(p)^2]$, the probability ϖ to get a Θ_0-flake incident to p in $\mathrm{Del}(\hat{P})$ satisfies:*

$$\varpi \leq \left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^{d^2} \frac{16}{\tilde{w_0}\bar{\eta}^3}\Theta_0.$$

Proof Since each flake of $\mathrm{U}_{\mathrm{Del}}(P, \tilde{w_0})$ incident to p occurs in $\mathrm{Del}(\hat{P})$ only when the weight of p belongs to an interval of measure less than $|I|$ and since there at most N flakes in $\mathrm{U}_{\mathrm{Del}}(P, \tilde{w_0})$ incident to p, the probability ϖ to get in $\mathrm{Del}(\hat{P})$ a flake incident to p is less than $\frac{N|I|}{\tilde{w_0} L(p)^2}$. The result then follows from Lemma 9.6, Lemma 9.8 and the fact that $L(p) \geq \bar{\eta}\varepsilon$. □

Observe that the probability ϖ can be made arbitrary small if Θ_0 is small enough.

9.1.4 Algorithm

Based on the results of the previous subsections, Algorithm 6 takes as input an $(\varepsilon, \bar{\eta})$-net P, a constant $\tilde{w_0} < 1/4$ and a constant Θ_0 such that the probability ϖ given in Lemma 9.9 is less than 1. It outputs a weighting scheme on P whose relative amplitude is smaller than $\tilde{w_0}$ and such that the weighted Delaunay triangulation $\mathrm{Del}(\hat{P})$ is Θ_0-thick.

Algorithm 6 Thick weighted Delaunay triangulation

Input: P, $\tilde{w_0}$, Θ_0
Initialize all weights to 0 and compute $\mathrm{Del}(\hat{P}) = \mathrm{Del}(P)$
while there is a Θ_0-flake σ in $\mathrm{Del}(\hat{P})$ **do**
  choose a vertex p of σ
  while there is a Θ_0-flake incident to p **do**
    reassign the weight of p uniformly at random in $[0, \tilde{w_0} L^2(p)]$
    update $\mathrm{Del}(\hat{P})$
Output: A weighting scheme on P and the corresponding weighted Delaunay triangulation which is granted to be Θ_0-thick.

Theorem 9.10 *If P is an $(\varepsilon, \bar{\eta})$-net P, $\tilde{w_0}$ a constant less than $1/4$ and Θ_0 a constant such that:*

$$\left(1 + \frac{4\sqrt{2}}{\bar{\eta}}\right)^{d^2} \frac{16}{\tilde{w_0}\bar{\eta}^3}\Theta_0 < 1, \tag{9.7}$$

Algorithm 6 outputs a weighting scheme $\hat{P}$ on P whose relative amplitude is smaller than $\tilde{w_0}$ and such that the weighted Delaunay triangulation $\mathrm{Del}(\hat{P})$ *is Θ_0-thick. Its expected complexity is linear with respect to the size of P.*

Proof We first show by induction that, at each execution of the outer while loop triggered by the existence of a flake σ, all the vertices p of σ have current weight 0. The base case is trivial since all weights are initialized to 0. Let us assume that the claim is true up to the beginning of the *ith* execution of the loop, and assume that the *ith* execution of the loop is trigerred by the flake σ in the weighted Delaunay triangulation of the current weighting scheme $\hat{P}_i$. Lemma 9.9 and Equation 9.7 ensure that the inner while loop terminates after a constant number of trials for the weight of the vertex p of σ with a a new weighting scheme $\hat{P}_{i+1}$ and no flake incident to p in $\mathrm{Del}(\hat{P}_{i+1})$. Since the weight of p has increased from 0 to some value $w(p)$ in $[0, \tilde{w_0}L^2(p)]$ between $\hat{P}_i$ and $\hat{P}_{i+1}$, all the simplices in $\mathrm{Del}(\hat{P}_{i+1}) \setminus \mathrm{Del}(\hat{P}_i)$ are incident to p, and therefore are not flakes. This proves that $\mathrm{Del}(\hat{P}_{i+1})$ does not include flakes that were not in $\mathrm{Del}(\hat{P}_i)$. As a consequence, each flake in $\mathrm{Del}(\hat{P}_{i+1})$ already occurs in $\mathrm{Del}(P)$ and cannot have vertices with a non zero weight.

The outer loop is therefore executed at most once for each point in P. At each execution of the inner loop, the number of trials for the weight $w(p)$ of the considered vertex is constant, which proves the termination of the algorithm. Furthermore, since P is an $(\varepsilon, \bar{\eta})$-net and since all the weightings have a bounded relative amplitude, the reassignement of a single $w(p)$ entails a bounded number of changes in the weighted Delaunay triangulation $\mathrm{Del}(\hat{P})$ and updating the triangulation can be performed in constant time, which proves that the complexity of the algorithm is linear with respect to the number of points in P. □

9.2 Thickness from perturbation

As in the previous sections, we consider an $(\varepsilon, \bar{\eta})$-net for $\mathbb{T}^d$. The goal of this section is to show that one can compute a perturbed point set P' close to P so that the simplices have a bounded thickness.

In fact, the algorithm will protect $\mathrm{Del}(P')$, which, by Theorem 9.13, is sufficient to ensure that the simplices have a bounded thickness. The analysis of the perturbation algorithm that constructs P' relies on the Moser-Tardos constructive proof of the general Lovász local lemma (see the bibliographical notes).

9.2.1 Thickness from protection

In this section, we introduce the notion of protection of a simplex which provides a parametrized notion of genericity for Delaunay triangulations.

Definition 9.11 (δ-protection) *We say that a simplex $\sigma \subset P$ is δ-protected at $x \in \mathbb{R}^d$ if*

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

If we don't want to specify the point x at which σ is δ-protected, we simply say that σ is δ-protected. We say that $K \subseteq \mathrm{Del}(P)$ is δ-protected when each Delaunay d-simplex of K is δ-protected at its circumcenter.

This definition is to be compared to the notion of general position of a point set with respect to spheres that has been introduced in Section 5.2. For a point set P, being in general position with respect to spheres is the same as $\mathrm{Del}(P)$ being 0-protected.

Our goal here is to demonstrate that if the d-simplices are δ-protected, their vertices are well separated and they have a bounded thickness. As already mentionned, to focus on the central properties and avoid boundary issues, we now assume that the working space is the flat torus $\mathbb{T}^d$.

We let as an exercice to prove the following lemma (Exercise 9.5).

Lemma 9.12 (Separation from protection) *If $P \in \mathbb{T}^d$ and $\mathrm{Del}(P)$ is δ-protected, its edges have length greater than δ.*

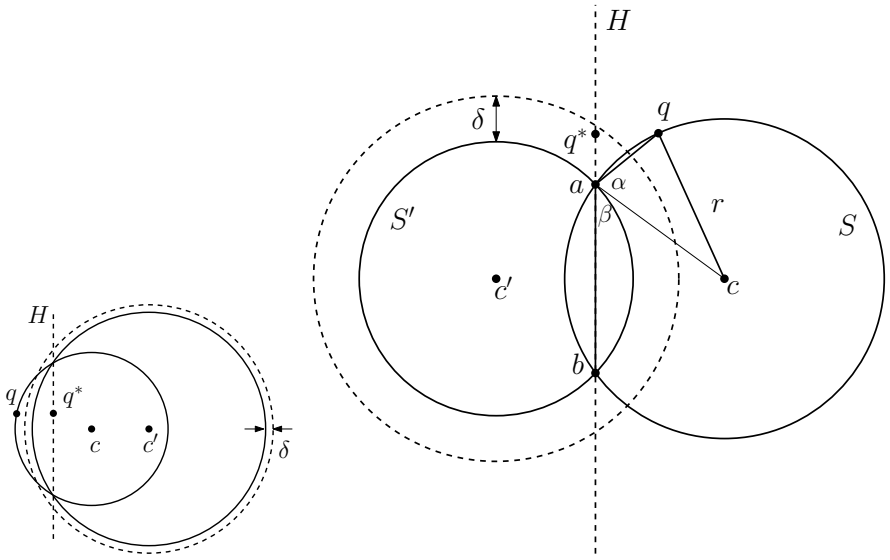

(a) c and c' are on the same side of H

(b) H separates c and c'

Figure 9.2: Diagram for the proof of Theorem 9.13.

Theorem 9.13 (Thickness from protection) *Let P be an $(\varepsilon, \bar{\eta})$-net of $\mathbb{T}^d$. If $\mathrm{Del}(P)$ is $\bar{\delta}\varepsilon$-protected, then the thickness of any simplex σ (of any dimension) satisfies*

$$\Theta(\sigma) \geq \frac{\bar{\delta}\,(\bar{\eta} + \bar{\delta})}{8d}.$$

Proof We will write $\delta = \bar{\delta}\varepsilon$ and $\eta = \bar{\eta}\varepsilon$. Observe that, since P is an ε-sample, $\bar{\delta} \leq 2$. Since we are working in $\mathbb{T}^d$, $\mathrm{Del}(P)$ has no boundary and we can write σ as $\sigma = q * \tau$ where τ is a facet of σ and $q \in P \setminus \sigma$. Moreover, there exists another simplex $\sigma' \in \mathrm{Del}(P)$ such that $\tau \subset \sigma'$ and $q \notin \sigma'$.

Suppose that $B = B(c, r)$ (resp. $B' = B(c', r')$) is a Delaunay ball for σ (resp. σ'). Since σ is in $\mathrm{Del}(P)$ and P is an ε-sample of $\mathbb{T}^d$, we have $r < \varepsilon$.

Let $B'' = B(c', r' + \delta)$ be the δ-protected Delaunay ball for σ'. Let H be the $(d-1)$-flat $\mathrm{aff}(\tau)$ and observe that H contains $\partial B \cap \partial B'$. Our geometry will be performed in the plane Q defined by c, c', and q. This plane is orthogonal to H, and it follows that the distance $d(q, H)$ is realized by a segment in the

plane Q: the projection, q^*, of q onto H lies in Q, and $d(q, H) = \|q - q^*\|$. We now prove that

$$\|q - q^*\| \geq \frac{\delta}{4\varepsilon}(\eta + \delta) \tag{9.8}$$

If c and c' lie on the same side of H (Fig. 9.2(a)), then $\partial B'$ separates q from q^*, and $\|q - q^*\| > d(q, \partial B') > \delta$, since B' is δ-protected. Equation 9.8 then follows since η and δ are each no larger than 2ϵ. Thus assume that H separates c from c', as shown in Fig. 9.2(b). Let $S' = Q \cap \partial B'$, and $S = Q \cap \partial B$, and let a and b be the points of intersection $S' \cap S$. Thus $H \cap Q$ is the line through a and b.

We will bound $\|q - q^*\|$ by finding an upper bound on the angle $\gamma = \angle qab$. Without loss of generality, we may assume that $\gamma \geq \angle qba$, and we will assume that $\gamma \geq \pi/2$ since otherwise $q^* \in B'$ and thus $\|q - q^*\| > \delta$ and Equation 9.8 is again trivially satisfied.

Since $\|q - a\| > \delta$, we have $\|q - q^*\| = \|a - q\| \sin\gamma > \delta \sin\gamma$. Also note that $\|a - b\| \geq 2R(\tau) \geq L(\tau) \geq \eta$, where $R(\tau)$ denotes the radius of the smallest ball circumscribing τ and $L(\tau)$ denotes the length of a shortest edge of τ. Let $\alpha = \angle qac$. Then $\cos\alpha = \frac{\|a-q\|}{2r} \geq \frac{\delta}{2\varepsilon}$, which means that $\alpha \leq \arccos\frac{\delta}{2\varepsilon} \leq \frac{\pi}{2}$. Similarly, with $\beta = \angle cab$, we have $\beta \leq \arccos\frac{\eta}{2\varepsilon} \leq \frac{\pi}{2}$. Thus $\frac{\pi}{2} \leq \gamma = \alpha + \beta \leq \gamma'$, where

$$\gamma' = \arccos\frac{\delta}{2\varepsilon} + \arccos\frac{\eta}{2\varepsilon}.$$

Since $\sin\gamma \geq \sin\gamma'$, when $\frac{\pi}{2} \leq \gamma \leq \gamma' \leq \pi$, we have

$$\begin{aligned}
\|q - q^*\| &> \delta\sin\gamma \geq \delta\sin\gamma' \\
&= \delta\sin\left(\arccos\frac{\delta}{2\varepsilon} + \arccos\frac{\eta}{2\varepsilon}\right) \\
&\geq \delta\left(\frac{\eta}{2\varepsilon}\sin\left(\arccos\frac{\delta}{2\varepsilon}\right) + \frac{\delta}{2\varepsilon}\sin\left(\arccos\frac{\eta}{2\varepsilon}\right)\right) \\
&\geq \frac{\delta}{4\varepsilon}(\eta + \delta),
\end{aligned}$$

where the last inequality follows from $\eta \leq 2\varepsilon$ and $\delta \leq 2\varepsilon$, and

$$\sin\left(\arccos\frac{\delta}{2\varepsilon}\right) = \sqrt{1 - \frac{\delta^2}{4\varepsilon^2}} \geq 1 - \frac{\delta^2}{8\varepsilon^2} \geq \frac{1}{2}.$$

This ends the proof of Equation 9.8.

Since $\mathrm{aff}(\tau) \subset H$, the bound also applies to $D(q,\sigma) \geq d(q,H)$. Since the arguments above hold for any simplex σ and for any facet τ of σ, we deduce a bound on the thickness of all the simplices. Since the altitudes of the simplices are at least $\frac{\delta}{4\varepsilon}(\eta+\delta)$, and since $\Delta(\sigma) \leq 2\varepsilon$, we conclude that $\Theta(\sigma) \geq \frac{\bar{\delta}\,(\bar{\eta}+\bar{\delta})}{8d}$ for any σ.$\square$

From Lemma 9.12, we know that the lengths of the edges of the simplices are at least $\bar{\delta}$. So, δ-protection is sufficient to get a bound on the thickness of the simplices: $\Theta(\sigma) \geq \frac{\bar{\delta}^2}{4d}$.

9.2.2 Lovász local lemma

The celebrated Lovász local lemma 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 bounded probability and there is limited dependence among them. Here is the lemma in a simple form that will be sufficient for our purposes.

Lemma 9.14 (Lovász local lemma) *Let $\mathcal{A} = \{A_1, \ldots, A_N\}$ 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 A_j, and that $\mathsf{Pr}[A_i] \leq \varpi$ for all $1 \leq i \leq N$. If*

$$\varpi \leq \frac{1}{e(\Gamma+1)} \tag{9.9}$$

where e denotes the base of the natural logarithm, then

$$\mathsf{Pr}\left[\bigwedge_{i=1}^{N} \neg A_i\right] > 0$$

Assume that the events depend on a finite set of mutually independent variables in a probability space. Moser and Tardos gave a constructive proof of Lovász lemma leading to a simple and natural algorithm that checks whether some event $A \in \mathcal{A}$ is violated and randomly picks new values for the random variables on which A depends. We call this a resampling of the event 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 $\mathcal{A}$. Specifically, the randomized algorithm described above

resamples an event $A \in \mathcal{A}$ at most an expected $1/\Gamma$ times before it finds such an evaluation. Thus the expected total number of resampling steps is at most N/Γ.

9.2.3 Protecting the simplices via perturbation

The input is an $(\varepsilon, \bar{\eta})$-net P for $\mathbb{T}^d$. The algorithm will perturb the points of P, possibly several times, and output the Delaunay triangulation of a perturbed point set P'. By perturbing a point $p \in P$, we mean to randomly and independently pick a point p' from the ball of radius $\rho = \bar{\rho}\varepsilon$ centered at p. $B(p, \rho)$ will be called the *picking region* for p.

Let us refer now to the terminology of the Lovász local lemma. We denote by P' a set of random points, one in each picking region $B(p, \rho), p \in P$. We therefore have $n = |P|$ random variables and we call an assignement of positions of the points of P' inside their picking regions a *resampling*.

We now define the events. An event happens when there exists a so called *bad configuration* $\phi' = (\sigma', p')$ of $d+2$ points of P' such that the circumradius of the d-simplex σ' is at most $\varepsilon + \rho$ and p' lies in the *δ-protection zone* $Z_\delta(\sigma')$ of σ'. The *protection zone* $Z_\delta(\sigma')$ of a d-simplex σ' is defined as the spherical shell

$$B(c_{\sigma'}, R_{\sigma'} + \delta) \setminus B(c_{\sigma'}, R_{\sigma'}).$$

bad events could be defined for Delaunay simplices only and the analysis could be improved with a bound on the size of a star.

The algorithm is given in Algorithm 7. The algorithm depends on two parameters, the radius ρ of the picking region, and the protection bound δ. These two parameters have to satisfy Equation 9.10 to be given below. The algorithm outputs the perturbed point set P' and its Delaunay triangulation which is δ-protected and therefore has a bounded thickness by Theorem 9.13.

Algorithm 7 Thick Delaunay triangulation

Input: P, ρ, δ
while there exists a bad configuration $\phi' = (\sigma', p')$ **do**
  perturb the points of ϕ'
  update Del(P')
Output: P' and Del(P') which is δ-protected

9.2.4 Correctness of the algorithm

In this section, we will write $\rho = \bar{\rho}\varepsilon$, $\bar{\delta} = \delta/\rho$ and $\eta = \bar{\eta}\varepsilon$, and assume $\bar{\rho} \leq \bar{\eta}/4$. We will use the easy fact that $\bar{\eta} \leq 2$ and recall that we must have $\bar{\delta} \leq \bar{\eta}$ (Lemma 9.12). We will also use the two following simple geometric lemmas.

Lemma 9.15 *If P is an $(\varepsilon, \bar{\eta})$-net of $\mathbb{T}^d$ and $\bar{\rho} < \bar{\eta}$, then P' is an $(\varepsilon', \bar{\eta}')$-net, where*

$$\varepsilon' = \varepsilon(1 + \bar{\rho}) \text{ and } \bar{\eta}' = \frac{\bar{\eta} - 2\bar{\rho}}{1 + \bar{\rho}} \geq \frac{\bar{\eta}}{3}$$

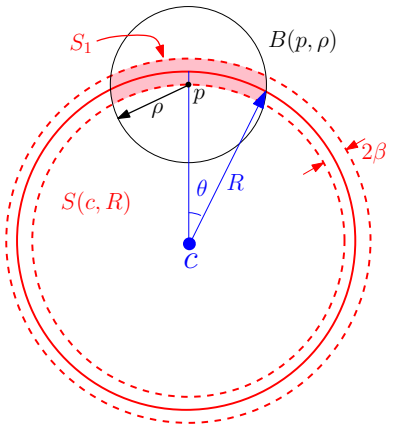

Figure 9.3: For Lemma 9.16.

Lemma 9.16 *Let $S(c, R)$ be a hypersphere of $\mathbb{R}^d$ of radius R centered at c and T_δ the spherical shell $T_\delta = B(c, R + \delta) \setminus B(c, R)$. Let, in addition, B_ρ denote any d-ball of radius $\rho < R$ (see Figure 9.3). We have*

$$\mathrm{vol}_d(T_\delta \cap B_\rho) \leq U_{d-1} \left(\frac{\pi}{2}\rho\right)^{d-1} \delta,$$

where vol_d denotes the d-dimensional volume and U_{d-1} the volume of an Euclidean $(d-1)$-ball with unit radius.

jd: change β to δ in the fig

Proof Consider a hypersphere S, concentric with S and with radius $\tilde{R}$, $R \leq \tilde{R} \leq R + \delta$. The intersection of $B(p, \rho)$ with S will be a geodesic ball $D \subset S$. Since $\rho < \tilde{R}$, the geodesic radius of D, say $r = \tilde{R}\theta$, is subtended by an angle θ that is less than $\pi/2$, and $\frac{2}{\pi}\theta \leq \sin\theta \leq \rho/\tilde{R}$. It follows that $r \leq \frac{\pi}{2}\rho$, a bound that is independent of R or $\tilde{R}$. Since the volume of a geodesic ball in a hypersphere is smaller than a Euclidean $(d-1)$-dimensional ball of the same radius [29, Theorem III.4.2], we have

$$\mathrm{vol}(D) \leq U_{d-1} \left(\frac{\pi}{2}\rho\right)^{d-1} \delta,$$

and the stated bound follows. □

Lemma 9.17 *Assume that $\bar{\rho} \leq \bar{\eta}/4$. The number of points of P' that fall in a ball B of radius $\bar{r}\varepsilon$ is at most $I(\bar{r}) = \left(1 + \frac{6\bar{r}}{\bar{\eta}}\right)^d$.*

Proof Since P' is η'-separated, we can use a volume argument and get using Lemma 9.15,

$$I(r) = \left(\frac{\bar{r}\varepsilon + \frac{\eta'}{2}}{\frac{\eta'}{2}}\right)^d \leq \left(1 + \frac{6\bar{r}}{\bar{\eta}}\right)^d$$

□

Lemma 9.18 *Assume that $\bar{\rho} \leq \bar{\eta}/4$. An event is independent of all but at most Γ other events where*

$$\Gamma = \frac{1}{(d+2)!} \left(\frac{65}{\bar{\eta}}\right)^{d(d+2)}.$$

Proof Let $\phi' = (\sigma', p')$ be a bad configuration. By definition, the distance from $c_{\sigma'}$ to any other point in ϕ' is at most

$$R_{\sigma'} + \delta \leq R = \varepsilon + \rho + \delta = \varepsilon\,(1 + \bar{\rho} + \bar{\delta}),$$

using Lemma 9.15. Hence two configurations (σ', p') and (τ', q') such that $\|c_{\sigma'} - c_{\tau'}\| > 2R$ cannot have a vertex in common and therefore the corresponding events are independent.

It follows that the number of events that may not be independent from an event (σ', p') is at most the number of $(d+1)$-simplices with vertices in $B(c_{\sigma'}, 3R)$. The number of points of P' in $B(c_{\sigma'}, 3R)$ is bounded by $I = I(3R/\varepsilon)$ (Lemma 9.17). We deduce that the number of events that may not be independent from an event (σ', p') is at most

$$\begin{pmatrix} I \\ d+2 \end{pmatrix} \leq \frac{I^{d+2}}{(d+2)!}.$$

Using $\bar{\rho} \leq 1/4$ and $\bar{\delta} \leq 2$, $\bar{R} \leq 7/2$. The announced bound follows. □

Lemma 9.19 *The probability that a point $p' \in P'$ falls into the δ-protection zone of a d-simplex σ is at most $\varpi = \frac{2^{d+1}}{\pi} \frac{\delta}{\rho}$.*

Proof The probability we want to bound is also the probability $\varpi(p', \sigma')$ that (p', σ') is a bad configuration.

The probability $\varpi(p', \sigma')$ is the ratio between the volume of the intersection of the spherical shell $Z_\delta(\sigma')$ with the picking region of radius ρ, and the volume of the picking region. Hence, using again Lemma 9.16, we obtain

$$\varpi(p', \sigma') \leq \frac{U_{d-1}}{U_d} \frac{2}{\pi} \frac{\rho^{d-1}\,\delta}{\rho^d},$$

where $\frac{U_{d-1}}{U_d} \sim \sqrt{\frac{2}{e^2 \pi\, d}}$ (see Exercise 9.9). □

We are now ready to apply the Lovász local lemma. We have bounded Γ in Lemma 9.18 and ϖ in Lemma 9.19. It remains to ensure that $\varpi \leq 1/(e(\Gamma+1))$. In view of the preceding inequality, this will be garanteed by putting an upper bound on δ/ρ. We also need $\bar{\rho} \leq \bar{\eta}/4$. We sum up the constraints on ρ and δ in the following equation

$$\frac{2^{d+1} e}{\pi} (\Gamma+1)\, \delta \leq \rho \leq \frac{\eta}{4} \quad \text{where} \quad \Gamma = \frac{1}{(d+2)!} \left(\frac{65}{\bar{\eta}}\right)^{d(d+2)} \tag{9.10}$$

Note that η and $\bar{\eta}$ can be measured from the input point set P and that Equation 9.10 can be satisfied for a small enough protection parameter δ. Under the condition above, we can apply Moser-Tardos analysis and show

that the algorithm terminates. The set of points P' output by the algorithm is at Hausdorff distance at most ρ from P. Moreover, the d-simplices of $\mathrm{Del}(P')$ are δ-protected and thus have a positive thickness by Theorem 9.13.

The expected number of resamplings executed by the algorithm is at most N/Γ, where N denotes the total number of events. Write $\Sigma(p')$ for the number of d-simplices that can possibly make a bad configuration with $p' \in P'$ for some perturbed set P', and let $R = \varepsilon+\rho+\delta = \bar{R}\varepsilon$. Using Lemma 9.17, we can bound the total number E of events as follows :

$$E \leq \sum_{p' \in P'} \Sigma(p') \leq |P'| \times \frac{I(2\bar{R})^{d+1}}{(d+1)!} \leq \frac{|P'|}{(d+1)!} \times \left(\frac{44}{\bar{\eta}}\right)^{d(d+1)}.$$

The expected number of resampling is therefore at most

$$\frac{E}{\Gamma} \leq \ (d+2)\,\frac{44^{d(d+2)}}{65^{d(d+1)}}\,\bar{\eta}^d\,|P|.$$

Each resampling consists in perturbing a constant number of points and the algorithm must update the Delaunay triangulation after each resampling which takes also constant time. It follows that the expected complexity of the algorithm is linear in the number of points. We sum up the results in the following theorem.

Theorem 9.20 *Let P an $(\varepsilon, \bar{\eta})$-net P of $\mathbb{T}^d$, and ρ and δ two constants satisfying Eq. 9.10. Then Algorithm 7 perturbs the points of P by at most ρ so that the perturbed set P' is δ-protected. As a consequence, the thickness of all simplices is at least $\frac{\bar{\delta}+\bar{\eta}}{8d}$ (Theorem 9.13). The expected complexity of the algorithm is linear with respect to the size of P, the constant of proportionality depending on d and $\bar{\eta}$.*

Remark. The above analysis is very crude and the guarantees on δ and the thickness of the simplices are pretty weak. We let as an open question to improve theses bounds and their dependence on d.

9.3 Exercises

Exercise 9.1 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 9.2 Let τ be a j-simplex. Show that $L(\tau) \geq D(p, \tau) \geq j! \, \Theta(\tau) \, \Delta(\tau)$.

Exercise 9.3 Show that simplices with a bounded thickness have bounded dihedral angles.

Exercise 9.4 (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 $V(T) = (\alpha_1, ..., \alpha_{3t})$ where the $\alpha_i \in [0, \pi]$ are the angles of the t triangles of T, sorted by increasing values. Show that $\mathrm{Del}(P)$ is, among all triangulations of P, the one that maximizes $V(T)$ for the lexicographic order. In particular, $\mathrm{Del}(P)$ maximizes the smallest angle.

Exercise 9.5 (Separation from thickness) Prove Lemma 9.12. (Hint : if pq is an edge of $\mathrm{Del}(P)$, there exists a d-simplex $\sigma \in \mathrm{Del}(P)$ containing p but not q.)

Exercise 9.6 (Bounded domain) Consider a bounded domain $\Omega \subset \mathbb{R}^d$ and an $(\varepsilon, \bar{\eta})$-net P for Ω. Recall that $\bar{\eta} \leq 2$. Show that the simplices of the Delaunay triangulation of P restricted to Ω, $\mathrm{Del}_{|\Omega}(P)$, have a circumradius no greater than ε. To avoid considering boundary issues, we will further restrict our attention to the simplices of $\mathrm{Del}_{|\Omega}(P)$ that have no face on the boundary of $\mathrm{Del}_{|\Omega}(P)$. We call those simplices the *safe* simplices.

Write $\Omega^{-2\varepsilon}$ for the subset of Ω defined as $\Omega^{-2\varepsilon} = \{x \in \Omega : d(x, \partial\Omega) > 2\varepsilon\}$. Show that any Delaunay simplex σ of $\mathrm{Del}_{|\Omega^{-2\varepsilon}}(P)$ is safe.

Extend Theorem 9.13 and the other results of Section 7.3 to safe simplices.

Exercise 9.7 (Extension of Theorem 5.10) Show that Theorem 5.10 extends when Ω is any bounded convex domain $\Omega \subset \mathbb{R}^d$, $L \subset \Omega$ and $W = \Omega$.

Exercise 9.8 (Witness complex) Show that the input point set L can be perturbed so that the conditions are fulfilled, hence resulting in a witness complex of a perturbed point set L' that is identical to its exact restricted Delaunay triangulation $\mathrm{Del}(L')$. As for the algorithm of Section 7.3, the notion of protection is central and the analysis of the perturbation algorithm relies on Exercise 5.17 and the Moser-Tardos constructive proof of the general Lovász local lemma (Section 9.2.2).

v

Exercise 9.9 (Bound on two unit-ball volume ratio) Let U_d denote the volume of a d-dimensional unit ball in $\mathbb{R}^d$. When d goes to infinity, we have:

$$\frac{U_d}{U_{d-1}} \sim \sqrt{\frac{e^2\pi}{2}}\sqrt{d}$$

Proof The volume U_d can be expressed with Γ-functions as $U_d = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}$ [97]. Now, using Stirling's formula to approximate Γ, we have for big d:

$$\frac{U_d}{U_{d-1}} = \sqrt{\pi}\frac{\Gamma(\frac{d}{2}+1)}{\Gamma(\frac{d-1}{2}+1)} \sim \sqrt{\pi}\frac{\sqrt{2\pi\left(\frac{d}{2}+1\right)}\left(\frac{\frac{d}{2}+1}{e}\right)^{\frac{d}{2}+1}}{\sqrt{2\pi\left(\frac{d-1}{2}+1\right)}\left(\frac{\frac{d-1}{2}+1}{e}\right)^{\frac{d-1}{2}+1}} = \sqrt{e\pi}\sqrt{\frac{d}{2}+1}\left(\frac{d+2}{d+1}\right)^{\frac{d-1}{2}+1}$$

Note that $\left(\frac{d+2}{d+1}\right)^{\frac{d-1}{2}+1} = \left(1+\frac{1}{d+1}\right)^{\frac{d+1}{2}} \xrightarrow[d\to\infty]{} \sqrt{e}$, therefore $\frac{U_d}{U_{d-1}} \sim \sqrt{\frac{e^2\pi}{2}}\sqrt{d}$.

□

9.4 Bibliographical notes

The notion of protection has been introduced by Boissonnat, Dyer and Ghosh to study the stability of Delaunay triangulations and the existence and construction of Delaunay triangulations on Riemannian manifolds [18,

16]. Earlier related results can be found in [sliver exudation]. Li, Edelsbrunner

The Lovász local lemma, proved initially by Lovász and Erdös [3], is a 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 [91].

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 [66] : a sliver is a flake with an upper bound on the ratio of its cicumradius to the length of its shortest edge. Sliver removal in higher dimensions has been discussed in [85, 47].

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 prefered to the weighting mechanism in the context of mesh generation [86, 13, 17].

Chapter 10

Reconstruction of submanifolds

Given a finite set of points P on an unknown manifold $\mathbb{M}$, we want to reconstruct a simplicial complex $\hat{\mathbb{M}}$ that triangulates $\mathbb{M}$.

Major difficulties arise when working in higher dimensions. First, as seen in Chapter 7, Delaunay facets of dense and sparse point sets sampling manifolds of dimension greater than 2 are not guaranteed to be thick (see Section 7.3) and therefore not guaranteed to approximate the tangent bundle of $\mathbb{M}$ (see Lemma 7.14). This is an issue : to be able to reconstruct manifolds of dimension greater than 2, we need to explicitly take care of non-thick simplices using techniques similar to what has been done in Chapter 9 in the Euclidean case.

The second major difficulty comes from the so-called curse of dimensionality. We have seen (Theorem 5.4, Exercise 5.8) 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 $\mathrm{Del}(P)$ or any other subdivision of $\mathbb{R}^d$. Instead, we will introduce a subcomplex of $\mathrm{Del}(P)$, called the tangential Delaunay complex, whose complexity depends on the intrinsic dimension of $\mathbb{M}$ and not on the dimension of the ambient dimension d. This complex is defined locally and the various local triangulations are glued together so as to constitute an manifold complex that is embedded in $\mathbb{R}^d$ and triangulates $\mathbb{M}$.

10.1 Tangential complex

We introduce in this section a data structure, named the tangential complex. Let $\mathbb{M}$ be a k-submanifold $\mathbb{M}$ of $\mathbb{R}^d$. The only knowledge we have on $\mathbb{M}$ is a finite set of points $P \in \mathbb{M}$ and the tangent spaces at each point of P. The tangential complex $\mathrm{Del}_{T\mathbb{M}}(P)$ is a k-dimensional subcomplex of the d-dimensional Delaunay complex $\mathrm{Del}(P)$. An important property is that $\mathrm{Del}_{T\mathbb{M}}(P)$ 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 10.2 that $\mathrm{Del}_{T\mathbb{M}}(P)$ can be used to reconstruct a triangulation of $\mathbb{M}$ naturally embedded in $\mathbb{R}^d$.

10.1.1 Definition

Let P be a finite set of $n > k+1$ points on $\mathbb{M}$. Let $\mathrm{Del}(P)$ 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. We will further assume that the points of P are in general position so that $\mathrm{Del}(P)$ is a triangulation of P naturally embedded in $\mathbb{R}^d$ (see Section 5.2).

For $p \in P$, we denote by $\mathrm{Del}_p(P)$ the Delaunay triangulation of P restricted to the tangent space T_p, i.e. the collection of all the simplices with vertices in P that admit an empty circumscribing ball centered on T_p. Equivalently, the simplices of $\mathrm{Del}_p(P)$ are the simplices of $\mathrm{Del}(P)$ whose Voronoi dual face intersect T_p. In addition to be in general position, we will assume in the rest of the chapter that P satisfies the following *transversality condition* : T_p contains no point that is equidistant from more than $k+1$ points of P. This condition implies that if T_p intersects a $(d-k)$-face f of $\mathrm{Vor}(P)$, the intersection consists of a single point c. The unique point c is the center of an empty ball circumscribing the k-simplex of $\mathrm{Del}(P)$ dual to f. It is easy to see that applying an infinitesimal perturbation to P is sufficient to ensure the transversality condition.

Assuming as usual that $\mathrm{Del}(P)$ is a triangulation, the assumption above implies that $\mathrm{Del}_p(P)$ is a k-dimensional triangulation. We write $\mathrm{star}(p)$ for the star of p in $\mathrm{Del}_p(P)$, i.e. the subcomplex of $\mathrm{Del}_p(P)$ consisting of the simplices of $\mathrm{Del}_p(P)$ that are incident to p together with their faces (see Figure 10.1). For a k-simplex in $\mathrm{star}(p)$, we write $B_p(\tau)$ for the ball centered on T_p that circumscribes τ, $c_p(\tau)$ for its center and $R_p(\tau)$ for its radius. Observe that $R_p(\tau) \geq R(\tau)$, where $R(\tau)$ is as usual the radius of the smallest ball circumscribing τ.

Definition 10.1 (Tangential Delaunay complex) *We call tangential Delaunay complex the simplicial complex* $\mathrm{Del}_{T\mathbb{M}}(P) = \{\tau, \tau \in \mathrm{star}(p), p \in P\}$.

Plainly, $\mathrm{Del}_{T\mathbb{M}}(P)$ is a subcomplex of $\mathrm{Del}(P)$ and is therefore a simplicial complex embedded in $\mathbb{R}^d$ if the points of P are in general position (see Section 5.2). The following lemma is crucial since it shows that computing the tangential Delaunay complex reduces to computing n weighted Delaunay triangulations in the k-dimensional flats T_p, $p \in P$.

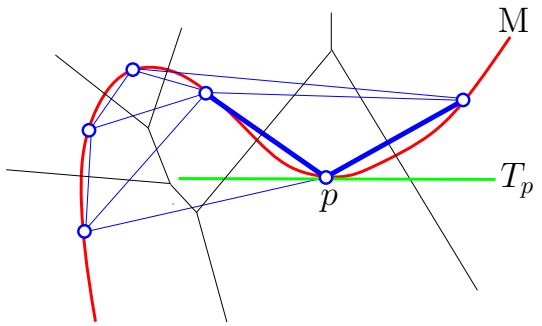

Figure 10.1: $\mathbb{M}$ is in red. The sample P is the set of blue circles. The green line is the tangent space T_p at p. The Voronoi diagram of the sample is in black. The Delaunay triangulation $\mathrm{Del}(P)$ is drawn in blue with, in bold, $\mathrm{star}(p)$.

We define a map $\psi_p : P \to T_p \times \mathbb{R}$ 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 $\mathrm{Vor}(P)$ intersects T_p along the weighted k-dimensional Voronoi diagram $\mathrm{Vor}(\psi_p(P))$ (see Exercise 6.4). Accordingly $\mathrm{Del}_p(P)$, the restriction of the d-dimensional Delaunay complex $\mathrm{Del}(P)$ to T_p, is isomorphic to the weighted Delaunay complex $\mathrm{Del}(\psi_p(P))$. Note that, under the transversality assumption, $\mathrm{Del}(\psi_p(P))$ is of dimension k and indeed a triangulation of T_p. Moreover, the simplices of $\mathrm{Del}(\psi_p(P))$ are obtained by projecting onto T_p the simplices of $\mathrm{Del}_p(P)$. Conversely, the simplices of $\mathrm{Del}_p(P)$ can be deduced from the simplices of $\mathrm{Del}(\psi_p(P))$ by a piecewise linear map we call the *lifting map*. Specifically, the lifting map lifts each weighted point (p', p'') associated to a vertex of $\mathrm{Del}(\psi_p(P))$ to the unique point $p \in P$ such that $(p', p'') = \psi_p(p)$. The lift of a simplex σ_p of $\mathrm{Del}_p(P)$ is then the geometric simplex σ whose vertices are the lifts of the vertices of σ_p. $\mathrm{Del}_p(P)$ is then the image of $\mathrm{Del}(\psi_p(P))$ by the *lifting map*. We summarize our discussion in the following lemma :

Lemma 10.2 *If the points of P are in general position and satisfy the transversality condition, $\mathrm{Del}_p(P)$ is the lift of $\mathrm{Del}(\psi_p(P))$, the k-dimensional weighted Delaunay triangulation of $\psi_p(P)$ in T_p.*

We deduce from the lemma an efficient algorithm to compute $\mathrm{star}(p)$: project P onto T_p, compute the star of p in $\mathrm{Del}(\psi_p(P))$, and then lift this star to $\mathrm{star}(p)$. 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 $\mathbb{M}$, we can even restrict our attention to the subset of P inside the ball of radius 2ε centered at p.

10.1.2 Inconsistent simplices

In general, the tangential Delaunay complex is *not* a manifold complex (see Definition 4.11). This is due to the presence of so-called *inconsistent simplices*.

Definition 10.3 (Inconsistent simplex) *A simplex $\tau \in \mathrm{Del}_{T\mathbb{M}}(P)$ is called inconsistent if τ does not belong to the stars of* all *its vertices. Let τ be an inconsistent simplex and let p_i and p_j be two vertices of σ so that σ is in* $\mathrm{star}(p_i)$ *but not in* $\mathrm{star}(p_j)$*, We say that (p_i, p_j) witnesses the inconsistent simplex σ.*

Refer to Figure 10.2. Let τ be a k-simplex with two vertices p_i and p_j such that τ belongs to the star of p_i but not to the star of p_j. Equivalently, the Voronoi $(d-k)$-dimensional face $\mathrm{Vor}(\tau)$ dual to τ intersects T_{p_i} (at a point $c_{p_i}(\tau)$) but does not intersect T_{p_j}. Observe that $c_{p_i}(\tau)$ is the center of an empty d-dimensional ball $B_{p_i}(\tau)$ circumscribing τ. Let $c_{p_j}(\tau)$ denote the intersection of $\mathrm{aff}(\mathrm{Vor}(\tau))$ with T_{p_j}. Differently from $B_{p_i}(\tau)$, the d-dimensional ball $B_{p_j}(\tau)$ centered at $c_{p_j}(\tau)$ that circumscribes τ contains a subset $P_j(\tau)$ of points of P in its interior. Therefore, the line segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ intersects the interior of some Voronoi cells (among which are the cells of the points of $P_j(\tau)$). We denote by p_l the point of $P \setminus \tau$ whose Voronoi cell is hit first by the segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$, when oriented from $c_{p_i}(\tau)$ to $c_{p_j}(\tau)$.

We now formally define an *inconsistent supersimplex*. See Figure 10.2.

Definition 10.4 (Inconsistent supersimplex) *Let τ be an inconsistent simplex witnessed by (p_i, p_j), and let p_l be the point of P such that its Voronoi cell* $\mathrm{Vor}(p_l)$ *is the first cell of* $\mathrm{Vor}(P)$ *whose interior is intersected by the line segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$, where $c_{p_i}(\tau) = T_{p_i} \cap \mathrm{Vor}(\tau)$ and $c_{p_j}(\tau) = T_{p_j} \cap$*

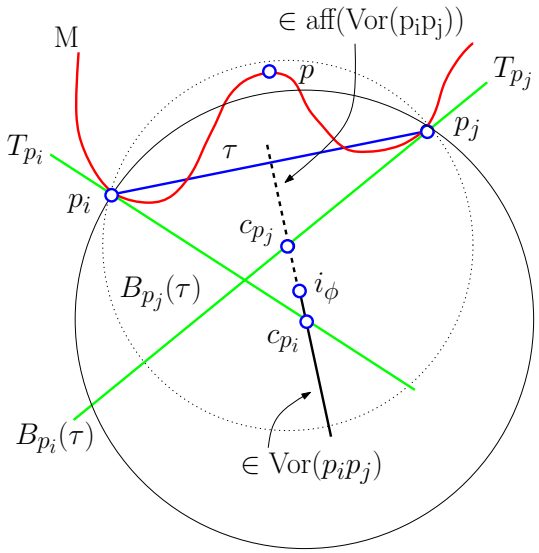

Figure 10.2: The figure shows an example of an inconsistent simplex $\tau = [p_i, p_j]$ that belongs to star(p_i) and does not belong to star(p_j). The solid black line segment belongs to Vor(τ) and the dashed black line segment belongs to aff(Vor(τ)) but not to Vor(τ).

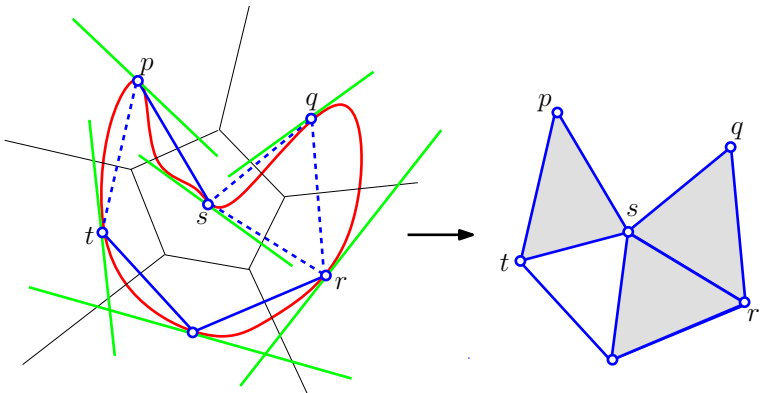

(a) $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ (b) $C(\mathcal{P})$

Figure 10.3: In Figure (a), $\mathbb{M}$ is in red, the sample P is the set of blue points, the green lines denote the tangent space at the point, the Voronoi diagram of the sample is in black, and $\mathrm{Del}_{T\mathbb{M}}(P)$ is drawn in blue. The dashed lines represent the inconsistent edges. In Figure (b), the blue lines denote $\mathrm{Del}_{T\mathbb{M}}(P)$ and the grey triangles denote the inconsistent supersimplices.

$\mathrm{aff}(\mathrm{Vor}(\tau))$, *and* $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ *is oriented from* $c_{p_i}(\tau)$ *to* $c_{p_j}(\tau)$. *We say that* $\phi = \tau \star p_l$ *is the inconsistent supersimplex of* $\mathrm{Del}_{T\mathbb{M}}(P)$ *associated to* τ.

Write $i(\phi)$ for the first point of $\mathrm{Vor}(p_l)$ hit by the oriented segment $[c_{p_i}(\tau)$ to $c_{p_j}(\tau)]$. Observe that $i(\phi)$ is the point on $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ that belongs to $\mathrm{Vor}(\phi)$. Hence, an inconsistent supersimplex is a $(k+1)$-simplex of $\mathrm{Del}(P)$. Since we assumed that the points satisfy the transversality condition, an inconsistent supersimplex cannot belong to the tangential complex (which does not contain faces of dimension greater than k). Observe also that some of the subfaces of an inconsistent supersimplex may not belong to the tangential complex.

Since inconsistent supersimplices are $(k+1)$-simplices, we will use the same notations for inconsistent supersimplices as for simplices, e.g. $R(\phi)$ for the circumradius of ϕ or $\Theta(\phi)$ for its thickness. We write $\mathrm{Inc}(p)$ for the set of inconsistent supersimplices incident to p and $\mathrm{Inc}(P) = \cup_{p\in P}\mathrm{Inc}(p)$. We also define the *completed tangential complex* $\mathrm{Del}^{+I}_{T\mathbb{M}}(P) = \mathrm{Del}_{T\mathbb{M}}(P) \cup \mathrm{Inc}(P)$. See Figure 10.3.

Calculating $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$ is easy once we know $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$. Indeed, it suffices to detect the inconsistent k-simplices that do not appear in the stars of all their vertices. Let τ be such a simplex, p_i and p_j two of its vertices as in the definition of an inconsistent supersimplex. To compute the associated inconsistent supersimplex, we need to compute p_l. This can be done by computing the restriction of $\mathrm{Vor}(P)$ to the line $L_{ij} = (c_{p_i}(\tau), c_{p_j}(\tau))$. Which in turn can be done by projecting the points of P onto L_{ij} and computing a 1-dimensional weighted Voronoi diagram (see Exercise 6.4).

10.1.3 Geometric properties of inconsistent supersimplices

We give some simple geometric lemmas that, in particular, bound the thickness of the inconsistent supersimplices. We will use angles between affine spaces as defined in Section 7.3.1.

Lemma 10.5 *Let P be an ε-dense sample of a submanifold $\mathbb{M}$ with $\alpha_0 < 0.32\,\mathrm{rch}(\mathbb{M})$ where $\alpha_0 = \frac{1}{2}\left(1+\varepsilon-\sqrt{1-2\varepsilon-\varepsilon^2}\right) \approx \varepsilon$. Then*

(i) $\mathrm{Vor}(p) \cap T_p \subseteq B(p, \alpha_0\,\mathrm{rch}(\mathbb{M}))$.

(ii) for any k-simplex $\tau \in \mathrm{star}(p)$, $R_p(\tau) \leq \alpha_0\,\mathrm{rch}(\mathbb{M})$ (recall that $R_p(\tau)$ is the radius of the ball centered on T_p that circumscribes τ)

(iii) for any $\tau \in \mathrm{Del}_{T\mathbb{M}}(P)$, $\Delta(\tau) \leq 2\alpha_0\,\mathrm{rch}(\mathbb{M})$.

Proof We prove (i). The other statements easily follow. Let $x \in \mathrm{Vor}(p) \cap T_p$ and write $\|p-x\| = \alpha\,\mathrm{rch}(\mathbb{M})$. Let x' be the point of $\mathbb{M}$ closest to x. Let us first assume that $\alpha < 0.32\,\mathrm{rch}(\mathbb{M})$. We then deduce from Lemma 7.9

$$\|x-x'\| \leq \left(1-\sqrt{1-\alpha^2}\right)\,\mathrm{rch}(\mathbb{M}).$$

Since P is an ε-sample, there exists a point $q \in P$, s.t. $\|x'-q\| \leq \varepsilon\,\mathrm{rch}(\mathbb{M})$. Together with $x \in \mathrm{Vor}(p)$, this implies

$$\|x-p\| \leq \|x-q\| \leq \|x-x'\| + \|x'-q\| \leq \left(\left(1-\sqrt{1-\alpha^2}\right)+\varepsilon\right)\,\mathrm{rch}(\mathbb{M}) \tag{10.1}$$

We must thus have

$$\alpha \leq 1-\sqrt{1-\alpha^2}+\varepsilon,$$

which is true for $\alpha \leq \alpha_0$. Hence $\mathrm{Vor}(p)$ cannot contain points whose distances to p lie in the range $(\alpha_0 \mathrm{rch}(\mathbb{M}), 0.32\,\mathrm{rch}(\mathbb{M}))$ (which is not empty according to the hypothesis of the lemma). Since $\mathrm{Vor}(p)$ is connected and contains p, $\mathrm{Vor}(p)$ is entirely contained in $B(p, \alpha_0 \mathrm{rch}(\mathbb{M}))$. □

The following lemmas bound the size and shape of inconsistent supersimplices with respect to the geometric properties of their facets. are small and thick, then it is small (Lemma 10.6) and cannot be thick (Lemma 10.7).

Lemma 10.6 *Let ϕ be an inconsistent supersimplex witnessed by p_i, p_j and p_l. Let $\tau = \phi \setminus \{p_l\}$ and write $\theta = \max_{p\in\tau}(\angle(\mathrm{aff}(\tau), T_p)$. We have*

$$\sin\theta \leq \frac{\Delta(\tau)}{\Theta(\tau)\,\mathrm{rch}(\mathbb{M})} \quad \text{and} \quad R(\phi) \leq \frac{R(\tau)}{\cos\theta}.$$

Proof From the definition of inconsistent supersimplices, τ belongs to $\mathrm{Del}_{p_i}(P)$ but not to $\mathrm{Del}_{p_j}(P)$. Accordingly, the vertices of ϕ lie in the closure of $B_{ij} = B_{p_j} \setminus B_{p_i}$. Observe also that c_{p_i}, c_{p_j} and $i(\phi)$ lie in the $(d-k)$-flat that contains $c(\tau)$ and is perpendicular to $\mathrm{aff}(\tau)$. Hence the orthogonal projection of these four points onto $\mathrm{aff}(\tau)$ is $c(\tau)$.

We now bound $\theta = \angle(\mathrm{aff}(\tau), T_p)$ for any $p \in \tau$. By Lemma 7.14 and Lemma 7.5 (2), we have

$$\sin\theta \leq \frac{2\frac{\Delta^2(\tau)}{2\mathrm{rch}(\mathbb{M})}}{\Theta(\tau)\,\Delta(\tau)} = \frac{\Delta(\tau)}{\Theta(\tau)\,\mathrm{rch}(M)}.$$

Since the orthogonal projection of c_{p_i} onto $\mathrm{aff}(\tau)$ is $c(\tau)$, $\omega = \angle(p_i - c_{p_i}, p_i - c(\tau)) = \min_{u\in\mathrm{aff}(\tau)} \angle(p_i - c_{p_i}, u) \leq \theta$. We thus have

$$R_{p_i}(\tau) = \|p_i - c_{p_i}\| = \frac{R(\tau)}{\cos\omega} \leq \frac{R(\tau)}{\cos\theta}.$$

and we get the same bound if p_i is replaced by p_j in the above inequality. Since $i(\phi) \in [c_{p_i} c_{p_j}]$, we also have $R(\phi) \leq \|i(\phi) - p_i\| \leq R(\tau)/\cos\theta$. □

We deduce from the lemma a bound on the thickness of inconsistent supersimplices. This is a crucial property to be used later to remove inconsistencies.

Lemma 10.7 *With the same notations as in Lemma 10.6, we can bound the thickness of an inconsistent supersimplex:*

$$\Theta(\phi) \leq \frac{\Delta(\phi)}{2(k+1)\mathrm{rch}(\mathbb{M})} \left(1 + \frac{2}{\Theta(\tau)}\right)$$

Proof We bound the altitude $D(p_l, \phi)$ of ϕ. Let $q \in \tau = \phi \setminus \{p_l\}$. We deduce from Lemmas 7.5 and 10.6

$$\begin{aligned} D(p_l, \phi) &= \|p_l - q\| \sin\angle(p_l - q, \mathrm{aff}(\tau)) \\ &\leq \Delta(\phi) \left(\sin\angle(p_l - q, T_q) + \sin\angle(T_q, \mathrm{aff}(\tau))\right) \\ &\leq \Delta(\phi) \left(\frac{\Delta(\phi)}{2\mathrm{rch}(\mathbb{M})} + \frac{\Delta(\tau)}{\Theta(\tau)\,\mathrm{rch}(\mathbb{M})}\right) \\ &\leq \frac{\Delta^2(\phi)}{2\mathrm{rch}(\mathbb{M})} \left(1 + \frac{2}{\Theta(\tau)}\right) \end{aligned}$$

The bound on the thickness then follows from the definition of thickness (see Section 7.3). □

Corollary 10.8 *Let P be an ε-sample of $\mathbb{M}$ and ϕ an inconsistent supersimplex whose facets are Θ_0-thick where Θ_0 satisfies*

$$\Theta_0^{k+1} > \frac{\alpha_0}{k+1} \left(\cos\arcsin\frac{2\alpha_0}{\Theta_0^k}\right)^{-1} \left(1 + \frac{2}{\Theta_0^k}\right) \tag{10.2}$$

and α_0 is defined in Lemma 10.5. Then ϕ is a Θ_0-flake.

Proof By definition, ϕ is a flake if $\Theta(\tau) \geq \Theta_0^k$ and $\Theta(\phi) < \Theta_0^{k+1}$. From Lemma 10.7 we have

$$\Theta(\phi) \leq \frac{\Delta(\phi)}{2(k+1)\mathrm{rch}(\mathbb{M})} \left(1 + \frac{2}{\Theta(\tau)}\right) \leq \frac{\Delta(\phi)}{2(k+1)\mathrm{rch}(\mathbb{M})} \left(1 + \frac{2}{\Theta_0^k}\right).$$

We now bound $\Delta(\phi)$

$$\begin{aligned} \Delta(\phi) &\leq 2R(\phi) \\ &\leq 2R(\tau) \left(\cos\arcsin\frac{2R(\tau)}{\Theta(\tau)\,\mathrm{rch}(\mathbb{M})}\right)^{-1} \qquad \text{(Lemma 10.6)} \\ &\leq 2\alpha_0\,\mathrm{rch}(\mathbb{M}) \left(\cos\arcsin\frac{2\alpha_0}{\Theta_0^k}\right)^{-1} \qquad (R(\tau) \leq R_p(\tau) \text{ and Lemma 10.5}) \end{aligned}$$

Hence, $\Theta(\phi)$ is not greater than the right hand side of Eq. 10.2. It follows that if Eq. 10.2 holds, $\Theta(\phi) < \Theta_0^{k+1}$ and ϕ is a Θ_0-flake. The corollary is proved. □

Note that the condition in the corollary is satisfied when the thickness of the supersimplex satisfies $\Theta(\phi) > \Theta_0^{k+1} = \Omega(\varepsilon^{1/k})$.

10.2 Submanifold reconstruction

Let $\mathbb{M}$ denote a submanifold of $\mathbb{R}^d$ that is compact, closed and differentiable and whose reach is positive. The only knowledge we have about $\mathbb{M}$ is its dimension k together with a finite point sample $P \subset \mathbb{M}$ and the tangent spaces at those points. From that knowledge, we want to construct a triangulation of $\mathbb{M}$. We assume that P is an $(\varepsilon, \bar{\eta})$-net of $\mathbb{M}$. 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.

To achieve our goal, we will first show how to remove inconsistent supersimplices from the completed tangential Delaunay complex $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$. We know from Lemma 10.7 that inconsistent supersimplices are Delaunay simplices of dimension $k+1$ that cannot be thick if their faces are small and thick. Hence to remove inconsistencies it suffices to remove non-thick simplices of dimension up to $k+1$ from $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$. This will be done by adapting Algorithm 6 of Section 9.1. The idea is to weight the points and to resort to their weighted Delaunay triangulation. We can easily extend the definition of tangential Delaunay complexes to the weighted case and, as we will see, weights can be chosen so that the weighted tangential Delaunay complex has no inconsistencies.

Remark: In Section 9.1, we assumed that the points were living in the flat torus $\mathbb{T}^d$ rather than in $\mathbb{R}^d$. This was mostly for convenience to avoid boundary issues. Since we are working here with a submanifold that has no boundary, the results of Section 9.1 can still be used in the context of this chapter.

10.2.1 Weight assignment

We recall for convenience the definition of a *weighting scheme* introduced in Section 9.1. Given a point set $P = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^d$, a weighting scheme on P is a function w that assigns to each point $p_i \in P$ a non-negative real weight $w(p_i)$: $w(P) = (w(p_1), ..., w(p_n))$. We write $\hat{p}_i = (p_i, w(p_i))$ and $\hat{P} = \{\hat{p}_1, \ldots, \hat{p}_n\}$.

We recall the definition of the *relative amplitude* of w as $\tilde{w} = \max_{p \in P} \frac{w(p)}{L^2(p)}$ where $L(p) = \min_{q \in P \setminus \{p\}} ||p - q|| \geq \bar{\eta}\varepsilon \mathrm{rch}(\mathbb{M})$ since P is an $(\varepsilon, \bar{\eta})$-net.

In the sequel, we assume that all weights are non negative and that $\tilde{w} \leq \tilde{w}_0 < 1/4$, for some constant $\tilde{w}_0$. Hence, for any point $p \in P$, $w(p) \in [0, \tilde{w}_0 L_p^2]$. The condition on $\tilde{w}$ implies in particular that all the weighted points, considered as balls, are disjoint.

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 10.1 the Delaunay triangulation of P by the weighted Delaunay triangulation of $\hat{P}$ (see Section 6.1.3). The role played above by Delaunay balls will be played by weighted Delaunay balls, i.e. balls orthogonal to weighted Delaunay simplices. Lemma 10.2 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, $\psi_p(\hat{p}) = (p_i, w_i - \|p_i - p_i'\|^2)$, where p_i' is the orthogonal projection of p_i on the tangent space T_p.

We need to extend the lemmas of Section 10.1 to the weighted case and, in particular, Corollary 10.8 that shows that the inconsistent supersimplices are flakes. Since the geometric quantities involved in the lemmas depend continuously on w and, since $\tilde{w}$ is bounded by a constant, the results of the previous section won't be deeply affected. Only the constants will be modified and will now depend on $\tilde{w}_0$.

Lemma 10.5 can be extended to the weighted case. Let $\hat{\tau}$ be a k-simplex of $\mathrm{star}(\hat{p})$ in $\mathrm{Del}_p(\hat{P})$. An easy adaptation of the calculations of Lemma 10.5 (i) shows that $\mathrm{Vor}(\hat{p}) \cap T_p \subset B(p, \hat{\alpha}_0 \,\mathrm{rch}(\mathbb{M}))$ where $\hat{\alpha}_0 \approx \varepsilon \sqrt{1 + \tilde{w}_0 \bar{\eta}^2}$ (see Exercise 10.3). The diameter of any k-simplex in $\mathrm{star}(\hat{p})$ is at most $2\hat{\alpha}_0 \,\mathrm{rch}(\mathbb{M})$. The other lemmas of Section 10.1.3 remain unchanged except for the value of α_0 that has to replaced by $\hat{\alpha}_0$.

10.2.2 Inconsistencies removal

Recall that inconsistent supersimplices are Delaunay simplices of dimension $k+1$ and that the completed tangential Delaunay complex $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$ is a k-complex identical to $\mathrm{Del}_{T\mathbb{M}}(P)$ if there is no inconsistencies, and is a $(k+1)$-complex that contains $\mathrm{Del}_{T\mathbb{M}}(P)$ otherwise.

Let Θ_0 be a constant satisfying the condition in Corollary 10.8. We know from the corollary that inconsistent supersimplices whose facets are Θ_0-thick are Θ_0-flakes. Hence, in order to have no inconsistencies, it suffices to ensure that $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$ contains no Θ_0-flakes (of dimension up to $k+1$).

To remove flakes, we apply a variant of Algorithm 6 that weights the points of P, leading to a set $\hat{P}$ whose weighted tangential Delaunay complex $\mathrm{Del}_{T\mathbb{M}}(\hat{P})$ has no inconsistencies, which eventually happens since the algorithm can remove any Θ_0-flake. More precisely, the algorithm removes flakes from the completed complex $\mathrm{Del}^{+I}_{T\mathbb{M}}(P)$ by weighting the points of P until there is no inconsistencies in $\mathrm{Del}_{T\mathbb{M}}(\hat{P})$. See Algorithm 8.

Algorithm 8 Inconsistencies removal

Input: P, $\{T_p, p \in P\}$, $\tilde{w}_0$, Θ_0
Initialize all weights to 0 and compute $\mathrm{Del}^{+I}_{T\mathbb{M}}(\hat{P}) = \mathrm{Del}^{+I}_{T\mathbb{M}}(P)$
while there is an inconsistent supersimplex ϕ **do**
  choose a vertex ϕ of σ
  reassign the weight of p uniformly at random in $[0, \tilde{w}_0 L^2(p)]$
  update $\mathrm{Del}^{+I}_{T\mathbb{M}}(\hat{P})$
Output: A weighting scheme on P and the corresponding weighted tangential Delaunay complex $\mathrm{Del}^{+I}_{T\mathbb{M}}(\hat{P})$ which is granted to be Θ_0-thick and to have no inconsistencies.

If Θ_0 satisfies the conditions in Theorem 9.10 and Corollary 10.8, then Algorithm 8 outputs, a weighting scheme $\hat{P}$ on P whose relative amplitude is smaller than $\tilde{w}_0$ and such that the weighted tangential Delaunay complex $\mathrm{Del}_{T\mathbb{M}}(\hat{P})$ output by the algorithm has no inconsistencies. Its expected complexity is linear with respect to the size of P.

The following theorem is a direct consequence of Theorem 9.10 and of the above discussion. Θ_0 has to be chosen so as to satisfy Eq. 9.7 (Theorem 9.10) and also Eq. 10.2 (Corollary 10.8). Observe that Eq. 9.7 provides an upper bound on Θ_0 that does not depend on the sampling density ε while Eq. 10.2 require Θ_0 to be large enough with respect to ε. Hence the algorithm is only

valid for sufficiently dense samples.

Theorem 10.9 (Inconsistencies removal) *Let $\mathbb{M}$ a submanifold of positive reach $\mathbb{M}$ and P be an $(\varepsilon, \bar{\eta})$-net of $\mathbb{M}$ for a sufficiently small ε. Let in addition $\tilde{w}_0$ and Θ_0 be two constants. If $\tilde{w}_0 < 1/4$ and Θ_0 satisfies Eq. 9.7 and 10.2, then Algorithm 8 assigns weights to the points of P leading to a set of weighted points $\hat{P}$ so that the output (weighted) tangential complex $\hat{\mathbb{M}} = \mathrm{Del}_{T\mathbb{M}}(\hat{P})$ is free of inconsistencies. The expected time complexity of the algorithm is $O(|P|)$.*

10.2.3 Guarantees on the reconstruction

In addition to removing all inconsistencies in the tangential Delaunay complex, the simplicial complex $\hat{\mathbb{M}}$ output by the algorithm is a good approximation of $\mathbb{M}$ as stated in the following theorem.

Theorem 10.10 (Guarantees) *Under the same hypotheses as in Theorem 10.9, the Delaunay tangential complex $\hat{\mathbb{M}}$ output by the algorithm satisfies the following properties:*

1. *All the simplices in $\hat{\mathbb{M}}$ are Θ_0-thick.*
2. *$\hat{\mathbb{M}}$ is a piecewise linear k-submanifold without boundary;*
3. *$\hat{\mathbb{M}}$ is homeomorphic to $\mathbb{M}$;*
4. *The Hausdorff distance between $\hat{\mathbb{M}}$ and $\mathbb{M}$ is $O(\varepsilon^2 \mathrm{rch}(\mathbb{M}))$;*
5. *If τ is a k-simplex of $\hat{\mathbb{M}}$ and p is a vertex of τ, we have*

$$\sin \angle(\mathrm{aff}(\tau), T_p) = O(\varepsilon).$$

The constant in the big-O depend on k, w_0, $\bar{\eta}$ and Θ_0.

Proof Property 1 directly follows from the algorithm. Properties 4 and 5 follows from the geometric lemmas of Section 10.1.3. Property 2 will be proved next.

Proving Property 2 reduces to proving that the link of any vertex of $\hat{\mathbb{M}}$ is a topological $(k-1)$-sphere. We first observe that, since $\hat{\mathbb{M}}$ contains no

inconsistencies, the star of any vertex p in $\hat{\mathbb{M}}$ is identical to $\text{star}(\hat{p})$, the star of p in $\text{Del}_p(\hat{P})$. Hence, to prove Property 2, 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 10.11 *$\hat{\mathbb{M}}$ is a simplicial manifold complex.*

Proof It is sufficient to prove that, for any $p \in P$, the link of p in $\hat{\mathbb{M}}$ is a topological $(k-1)$-sphere. By Lemma 10.2, $\text{star}(\hat{p})$ 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 10.5, 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 $\text{star}(\hat{p})$. □

The proof of Property 3 then follows from Theorem 7.15. This ends the proof of the theorem. □

10.3 Exercises

Exercise 10.1 Given is an $(\varepsilon, \bar{\eta})$-net P of a differentiable (unknown) submanifold $\mathbb{M} \in \mathbb{R}^d$. Propose a method to approximate the tangent space T_p of $\mathbb{M}$ at $p \in P$.

Exercise 10.2 Let $\tau = (p_0, ..., p_k)$ be a geometric simplex and let $\hat{\tau} = ((p_0, w_0), ..., (p_k, w_k))$ be the associated weighted simplex. Write $R(\tau)$ for the radius of the minimal circumscribing ball of τ and $R(\hat{\tau})$ for the radius of the minimal ball orthogonal to $\hat{\tau}$. Show that, if $0 \leq \tilde{w} \leq \tilde{w}_0 < \frac{1}{2}$, we have

$$\sqrt{1 - 4\tilde{w}_0}\, R(\tau) \leq R(\hat{\tau}) \leq R(\tau).$$

Proof Let $c(\tau)$ denote the circumcenter of τ and $c(\hat{\tau})$ the orthocenter of $\hat{\tau}$. Let q be a vertex of τ closest to $c(\hat{\tau})$. Since the weights are non negative,

$$R(\hat{\tau})^2 = \|c(\hat{\tau}) - q\|^2 - w(q) \leq \|c(\tau) - q\|^2 = R(\tau)^2.$$

Let p be a vertex of τ furthest from $c(\hat{\tau})$. We have $\|p - c(\hat{\tau})\| \geq \|p - c(\tau)\|$ and thus

$$\begin{aligned} R(\hat{\tau})^2 &= \|c(\hat{\tau}) - p\|^2 - w(p) \\ &\geq R(\tau)^2 - \tilde{w}_0 L(\tau)^2 \\ &\geq (1 - 4\tilde{w}_0)\, R(\tau)^2. \end{aligned}$$

□

Exercise 10.3 Extend Lemma 10.5 to the case of weighted points with bounded weights $w_p \in [0, \tilde{w}_0]$, $p \in P$.

Proof Let $x \in V(\hat{p}) \cap T_p$, $\|p - x\| = \hat{\alpha}\, \mathrm{rch}(\mathbb{M})$, $\hat{\alpha} < 0.32$. Denote by x' the point of $\mathbb{M}$ closest to x.

We have

$$\hat{\alpha}^2\, \mathrm{rch}^2(\mathbb{M}) - w_p = \|p - x\|^2 - w_p \leq \left(\varepsilon + (1 - \sqrt{1 - \hat{\alpha}^2})\right)^2\, \mathrm{rch}^2(\mathbb{M}) - w_q.$$

Up to higher order terms in ε, this inequality holds for

$$\hat{\alpha}^2 \leq \varepsilon^2 + \frac{w_p}{\mathrm{rch}^2(\mathbb{M})} \leq \varepsilon^2\, (1 + \tilde{w}_0 \bar{\eta}^2).$$

□

10.4 Bibliographical notes

The tangential complex has been independently defined by Freedman [68] and by Boissonnat and Flottoto [11]. Boissonnat and Ghosh later showed how to remove inconsistencies in the tangential complex by star stitching [12] and proved Theorems 10.9 and 7.15. The tangential complex can be seen as a light variant of the cocone algorithm of Cheng, Dey and Ramos [47].

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 [104] and by Boissonnat, Wormser and Yvinec to generate anisotropic meshes [13]. The central

question behind this approach is the stability of Delaunay triangulations and the existence and construction of Delaunay triangulations on manifolds [15, 14, 17].

A first proof of Theorem 7.15 has been given in [12]. A more general result can be found in the recent work of Dyer, Vegter and Wintraecken [57].

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 [75]. Estimating the tangent space can be done using principal component analysis (PCA) [84].

Part IV

Distance-based inference

Chapter 11

Stability of distance functions

In the sequel, all the considered shapes and their approximations are represented by compact subsets [1] of an Euclidean space $\mathbb{R}^d$. We use indifferently the words shape and compact set. In this section, we address the general problem consisting in recovering the topology of a shape from an approximation of it. A typical example of such a problem is illustrated on figure 11.1.

General inference problem: given an approximation K' (typically a point cloud) of an unknown shape K, can we infer the topology of K from K'?

This question is far too general to be directly rigorously addressed. Different points need to be precised. First, it is necessary to formalize the notion of approximation. This is done by defining a notion of distance between shapes; in this section we use the so-called *Hausdorff distance*. Second, it is also necessary to explain what is meant by "infering the topology". In the best case, one may be able to reconstruct a shape that is homeomorphic or even isotopic to the underlying shape K. Sometimes one can only reconstruct a topological space that is homotopy equivalent to K or one can only recover some topological invariants of K. It is also important to be aware that there is no hope to provide a general relevant answer to the inference problem without making any assumptions on the "regularity" of the underlying shape and on the quality of the approximation. Determining relevant assumptions leads to sampling conditions that are sufficient to recover the topology of K from K'.

Note that the classical reconstruction problems for smooth curves and surfaces sampled by point cloud data sets that have been widely studied in the last decades are particular cases of the general inference problem. The goal of this section is to introduce a general mathematical framework allowing to provide positive answers to the above mentionned problem in a wide range of settings encountered in practical applications. In particular, our framework allows to work with a wide class of non smooth shapes containing sharp edges and singularities. Our approach is based on the notion of distance function to a compact set.

[1] recall that a subset K of $\mathbb{R}^d$ is compact if and only if it is closed and bounded.

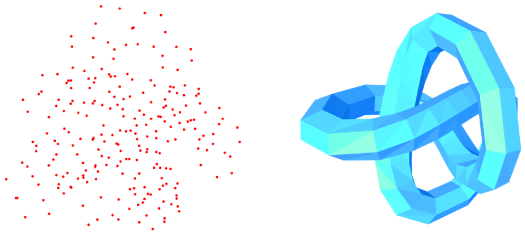

Figure 11.1: A point cloud data set (left) sampled around the surface of a knotted torus (right). How can we recover the topologie of the underlying torus from the point cloud?

11.1 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

$$d_K(x) = \inf_{y \in K} d(x, y) \quad \text{for all} \quad x \in \mathbb{R}^d$$

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

$$K^r = d_K^{-1}([0, r]) = \{x \in \mathbb{R}^d : d_K(x) \leq r\}$$

To quantify the notion of approximation for shapes it is necessary to introduce a distance measuring the closeness between two different shapes. Intuitively, two shapes K and K' could be considered as close to each other if any point of K is close to K' and if any point of K' is close to K. This means that K is contained in a small offset of K' and K' is contained in a small offset of K or equivalently, the restriction of $d_{K'}$ to K is bounded by some small constant and the restriction of d_K to K' is bounded by some

small constant. This intuition is quantified by the Hausdorff distance introduced in Definition 1.7.

Proposition 11.1 *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:*

- $d_H(K, K') = \max\left(\sup_{y \in K'}(\inf_{x \in K} \|x - y\|), \sup_{x \in K}(\inf_{y \in K'} \|x - y\|)\right)$
- $d_H(K, K')$ *is the smallest number r such that $K \subset K'^r$ and $K' \subset K^r$.*
- $d_H(K, K') = \max\left(\sup_{x \in K} d_{K'}(x), \sup_{x \in K'} d_K(x)\right)$.
- $d_H(K, K') = \|d_K - d_{K'}\| := \sup_{x \in \mathbb{R}^d} |d_K(x) - d_{K'}(x)|$.

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

In general, the topological and geometric features of a shape cannot be directly extracted from the features of an approximation. 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': for example, if $d_H(K, K')$ is small enough, the number of connected components of K is obviously different from the number of connected components (i.e. the cardinality) of K'. Worse, the occurence 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. More generally, point clouds in themselves do not carry any non trivial topological or geometric structure. 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 the sequel consists in considering the distance functions to compact sets and to compare the topology of their sublevel sets (i.e. the offsets) for 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 11.2. The goal of this section is to turn this intuition into a formal formal framework with formal statements. This requires to procceed in two steps. First, one needs to understand how the topology of the offsets K^r of a given compact

set K evolves with the parameter r. The answer to this question is given by the theory of critical points for distance functions that has been developped in Riemannian Geometry [45, 76]. 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.

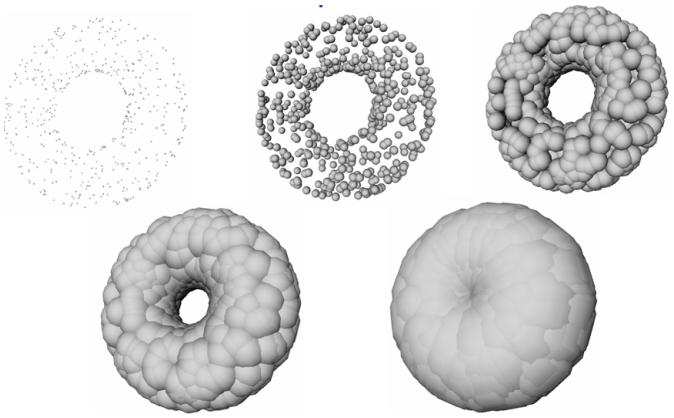

Figure 11.2: Various offsets of a point cloud data set sampled around a torus in $\mathbb{R}^3$.

11.2 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 a (empty) square in the plane, d_K is not differentiable along the diagonals of K. Nevertheless, it is possible to define a generalized gradiend vector field $\nabla_K : \mathbb{R}^d \to \mathbb{R}^d$ 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:

$$\Gamma_K(x) = \{y \in K : d(x, y) = d_K(x)\}$$

This is a non empty compact subset of K.

Let $\sigma_K(x)$ be the smallest closed ball enclosing $\Gamma_K(x)$ and let $\theta_K(x)$ be its center and $F_K(x)$ its radius (see figure 11.3). For $x \in \mathbb{R}^d \setminus K$, the generalized

gradient $\nabla_K(x)$ is defined by

$$\nabla_K(x) = \frac{x - \theta_K(x)}{R_K(x)}$$

and for $x \in K$, $\nabla_K(x) = 0$.

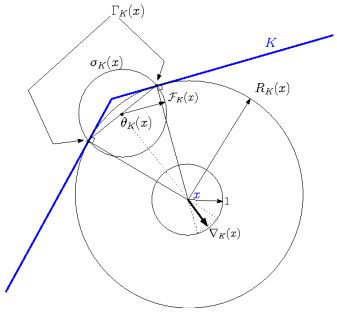

Figure 11.3: A 2-dimensional example with 2 closest points.

The norm of the gradient is given by

$$\|\nabla_K(x)\|^2 = 1 - \frac{F_K(x)^2}{R_K(x)^2} \tag{11.1}$$

Equivalently, the norm of $\nabla_K(x)$ is the cosine of the half angle of the smallest ("circular") cone with apex x that contains $\Gamma_K(x)$. Intuitively, the direction of $\nabla_K(x)$ 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 11.4).

Definition 11.2 *The set of points $x \in \mathbb{R}^d$ such that $\Gamma_K(x)$ contains more than one point is called the medial axis of K. It is denoted by* $\mathrm{ax}(K)$.

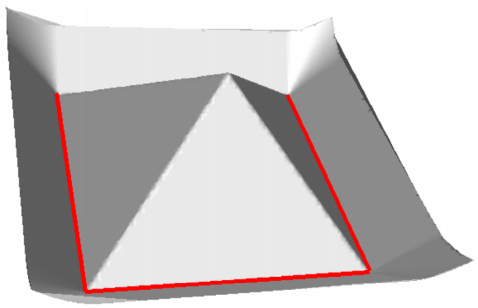

Figure 11.4: The graph of the distance to a square (in red) in the plane. Along the diagonal of the square, the direction of the gradient is given by the diagonals and its norm is the slope of the "edges" of the graph above the diagonals.

Note that since $F_K(x) = 0$ if and only if $\Gamma_K(x)$ contains only one point it follows from equation (11.1) that $\mathrm{ax}(K) = \{x \in \mathbb{R}^d : |\nabla_K(x)| < 1\}$.

The map $x \in \mathbb{R}^d \to \nabla_K(x) \in \mathbb{R}^d$ is in general not continuous. In other words, ∇_K is a discontinuous vector field. Nevertheless it is possible to show (see [88], or [100] for a more general statement) that $x \to |\nabla_K(x)|$ is a lower semi-continuous function [2] and ∇_K is integrable, that is the differential equation

$$\frac{dX}{dt} = \nabla_K(X(t))$$

has some solutions. More precisely, the Euler scheme associated to this equation converges uniformely, as the integration step tends to 0, to a continuous flow $\mathfrak{C} : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^d$ such that $t \to \mathfrak{C}(t, x)$ is the trajectory of ∇_K starting from x at $t = 0$. Moreover this continuous trajectory can be parametrized by arc length $s \to \mathfrak{C}(t(s), s)$ and one has

$$d_K\left(\mathfrak{C}(t(l), x)\right) = d_K(x) + \int_0^l \|\nabla_K(\mathfrak{C}(t(s), x)\| ds \tag{11.2}$$

The above equation implies that d_K is increasing along the trajectories of ∇_K. It can also be shown [88] that F_K is also non decreasing along the trajectories of ∇_K.

[2] a function $f : \mathbb{R}^d \to \mathbb{R}$ is lower semi-continuous if for any $a \in \mathbb{R}$, $f^{-1}((-\infty, a])$ is closed subset of $\mathbb{R}^d$

The gradient ∇_K allows to define the notion of critical point for d_K in the same way as for differentiable functions.

Definition 11.3 *A point x is a critical point of d_K if $\nabla_K(x) = 0$. 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.

11.3 Topology of the offsets

Using the notion of critical point defined in the previous section, it appears that distance functions to compact sets share some properties similar to 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 that are proven in [76].

Theorem 11.4 *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 $(n-1)$-dimensional topological submanifold of $\mathbb{R}^d$.*

Theorem 11.5 (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 "annulus" $A(r_1, r_2) = \{x \in \mathbb{R}^d : r_1 \leq d_K(x) \leq r_2\}$ 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 n-dimensional topological manifold with boundary. In particular, when one considers "small" offsets K^r, $r > 0$, their topology 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* first introduced in [38, 39, 40].

Definition 11.6 *Let $K \subset \mathbb{R}^d$ be a compact set. The weak feature size* $\mathrm{wfs}(K)$ *of K is the infimum of the positive critical values of d_K. If d_K does not have critical values,* $\mathrm{wfs}(K) = +\infty$.

It follows from the Isotopy Lemma 11.5 that if $0 \leq \alpha \leq \beta < \mathrm{wfs}(K)$, then K^α and K^β are isotopic. In a more intuitive way, the knowledge of K up to a precision of α, or at scale α, gives the same information for any choice of $0 < \alpha < \mathrm{wfs}(K)$. Moreover, the following result allows to compare the topology of the offsets of two close compact sets with positive weak feature sizes.

Theorem 11.7 ([39, 40]) *Let $K, K' \subset \mathbb{R}^d$ and $\varepsilon > 0$ be such that $d_H(K, K') < \varepsilon$,* $\mathrm{wfs}(K) > 2\varepsilon$ *and* $\mathrm{wfs}(K') > 2\varepsilon$. *Then*
(i) $\mathbb{R}^d \setminus K$ and $\mathbb{R}^d \setminus K'$ have the same homotopy type,
(ii) for any $0 < \alpha \leq 2\varepsilon$, K^α and K'^α are homotopy equivalent.

The proof of this result, ommitted in these notes, uses the gradient vector fields of d_K and $d_{K'}$ to build a map H that fullfills the hypothesis of proposition 1.6. It complete proof can be found in [39] (prop. 3.3).

The theorem 11.7 shows that the compact sets with positive weak feature size provides 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 inequations) all have positive weak feature size (see [71] p. 1045 and [39], proposition 3.6).
Nevertheless, the previous theorem suffer from an important weakness that prevent it to be really useful in practice. Indeed, the assumtion made on the wfs involving both $\mathrm{wfs}(K)$ and $\mathrm{wfs}(K')$ is hardly satisfied in practical situations. For example if K is some "continuous" shape (e.g. a surface) approximated by a finite point cloud K', $\mathrm{wfs}(K')$ 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 11.5. As a consequence, even if the weak feature size of K is large, it may happen that whatever the quality of the approximation by K' the hypothesis of the theorem 11.7 are never satisfied. This phenomenum can also be interpreted as a lack of continuity of the map $K \to \mathrm{wfs}(K)$ or as an instability property of the critical points of distance functions.

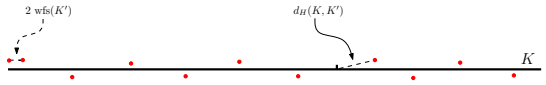

Figure 11.5: A segment K approximated by a point cloud K' (in red). The weak feature size of K' is obviously smaller than two times the Hausdorff distance between K and K'.

11.4 Stability of the critical points

Since the topology of the offsets of a compact can only change at critical points 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 11.6.

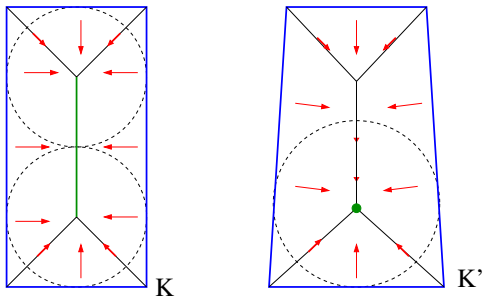

Figure 11.6: When K is a rectangle, there is a segment of critical points of d_K along one of the axes of symmetry of K (the bold segment). This segment collapses to one point as soon as one stretches the bottom side of K to obtain K'. Nevertheless, along the previously critical segment, the norm of the gradient of $d_{K'}$ remains small.

To overcome this unstability problem we introduced a "parametrized" notion of critical point.

Definition 11.8 *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 $|\nabla_K(x)| \leq \mu$.*

Note that a 0-critical points are exactly the above defined critical points of d_K. Unlike the 0-critical points alone, the family of μ-critical points satisfies the following fundamental stability property.

Theorem 11.9 (critical point stability theorem [33]) *Let K and K' be two compact subsets of $\mathbb{R}^n$ 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.*

Roughly, this theorem states that in a neighborhood of size $O(\sqrt{\varepsilon})$ of a given μ-critical point of K, there is a $(\mu + O(\sqrt{\varepsilon})$-critical point for any ε-approximation K' of K. In particular, any ε-approximation K' of K has at least one $O(\sqrt{\varepsilon})$-critical point in a $O(\sqrt{\varepsilon})$-neighborhood of each critical point of K.

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 11.10 *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:*

$$d_K(y)^2 \leq d_K(x)^2 + 2\mu d_K(x)||x - y|| + ||x - y||^2$$

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 $c = \theta_K(x)$ be the center of the minimal enclosing ball of Γ, and $\alpha = \cos^{-1}(\mu)$ (see figure 11.7).

For any $x' \in \mathbb{R}^d$ we have

$$\begin{aligned} d_K(y)^2 \leq \|y - x'\|^2 &= < (y - x) + (x - x'), (y - x) + (x - x') > \\ &= \|y - x\|^2 + \|x - x'\|^2 + 2 < (y - x), (x - x') > \\ &= d_K(x)^2 + 2d_K(x)\|x - y\| \cos(y - x, x - x') + \|x - y\|^2 \end{aligned}$$

To prove the lemma it is thus sufficient to prove
(*) there exists a point $x' \in \Gamma$ such that the angle between $(y - x)$ and $(x' - x)$ is not greater than $\pi - \alpha$.
We distinguish between two cases.

Case 1: $\mu \neq 0$
Assume that (*) is not satisfied. Then for any $x' \in \Gamma$ the angle between $x' - x$ and $x - y$ 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 $x - y$ and apex angle α'. This cone intersects S along

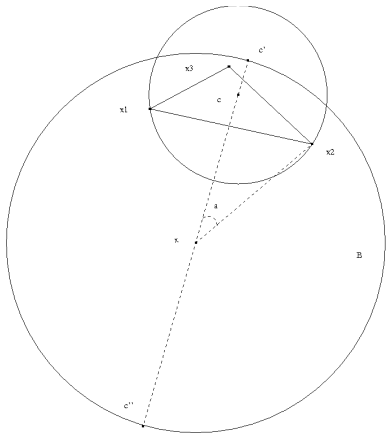

Figure 11.7: Proof of lemma 11.10

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 larger 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 11.11 *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:*

$$\mu' \leq \frac{\rho}{2d_K(x)} + 2\frac{\varepsilon}{\rho} + \mu$$

Proof Let us consider an integral line C of the vector field $\nabla_{K'}$ parameterized by arc length and starting at x. If C reaches a critical point of

K' before length ρ, the lemma holds. Assume this is not the case. Letting $y = C(\rho)$, we have:

$$d_{K'}(y) - d_{K'}(x) = \int_0^\rho \|\nabla_{K'}(C(s))\| ds$$

Therefore, there must exist a point p on the curve C between $s = 0$ and $s = \rho$ such that:

$$\|\nabla_{K'}(p)\| \leq \frac{d_{K'}(y) - d_{K'}(x)}{\rho} \tag{11.3}$$

The curve C being parametrized by arc length, note that $||p - x|| \leq \rho$. Now Lemma 11.10 applied to x, y, and K reads:

$$d_K(y) \leq \sqrt{d_K(x)^2 + 2\mu d_K(x)||x - y|| + ||x - y||^2}$$

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:

$$\begin{aligned} d_{K'}(y) - d_{K'}(x) &\leq \sqrt{d_K(x)^2 + 2\mu d_K(x)||x - y|| + ||x - y||^2} \\ &\quad -d_K(x) + 2\varepsilon \\ &\leq d_K(x)[\sqrt{1 + \frac{2\mu||x - y||}{d_K(x)} + \frac{||x - y||^2}{d_K(x)^2}} - 1] \\ &\quad + 2\varepsilon \\ &\leq \mu||x - y|| + \frac{||x - y||^2}{2d_K(x)} + 2\varepsilon \end{aligned}$$

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 (11.3) shows that p satisfies the desired requirements. □

Proof [of Theorem 11.9] 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. □

Remark 11.12 *Note that since $d_{K'}$ is increasing along the trajectories of $\nabla_{K'}$ (see equation (11.2)), the μ'-critical point p for $d_{K'}$ of lemma 11.11 can be chosen such that $d_{K'}(p) \geq d_{K'}(x)$.*

11.5 The critical function of a compact set

Despite its simplicity the critical point stability theorem 11.9 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 11.13 *The critical function of a compact set $K \subset \mathbb{R}^d$, $\chi_K : (0, +\infty) \to \mathbb{R}_+$ is defined by*

$$\chi_K(r) = \inf_{x \in d_K^{-1}(r)} \|\nabla_K(x)\|$$

An example of critical function is given on figure 11.8. Note that from the Isotopy Lemma 11.5 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 informations 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 11.9, we easily get the following stability result for the critical function.

Theorem 11.14 (critical function stability theorem [32]) *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:*

$$\inf\{\chi_{K'}(u) \,|\, u \in I(r, \varepsilon)\} \leq \chi_K(r) + 2\sqrt{\frac{\varepsilon}{r}}$$

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 d cannot be smaller than $\alpha - 2\sqrt{\frac{\varepsilon}{d}}$. In particular, if $\alpha > 2\sqrt{\frac{\varepsilon}{d}}$ then d cannot be a critical value of d_K. Since the topology of the offsets of K can only

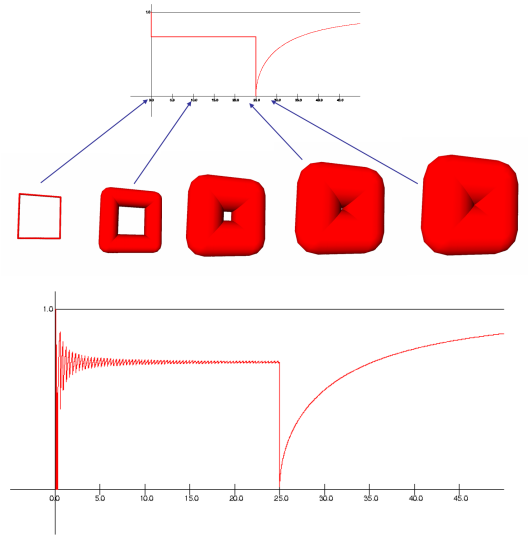

Figure 11.8: The critical function of a square embedded in $\mathbb{R}^3$ with edge length equal to 50 (top) and the critical function of a point cloud sampling this square (bottom).

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 11.8 and 11.9 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 Voronoï diagram of K.

Proof Let $r \geq 0$ and let $x \in d_K^{-1}(r)$ be such that $\|\nabla_K(x)\| = \chi_K(r)$ [3]. The critical point stability theorem 11.9 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 the lemma 11.10 to x, p and K we get

$$\begin{aligned} d_K(p) &\leq \sqrt{r^2 + 4\chi_K(r)d\sqrt{\varepsilon r} + 4\varepsilon r} \\ &\leq r\sqrt{1 + 4\chi_K(r)\sqrt{\varepsilon/r} + 4\varepsilon/r} \\ &\leq r + 2\chi_K(r)\sqrt{\varepsilon r} + 2\varepsilon \end{aligned}$$

[3] Note that 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 $\|\nabla_K\|$ and the compactness of $d_K^{-1}(r)$.

Now, according to Remark 11.12, 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. □

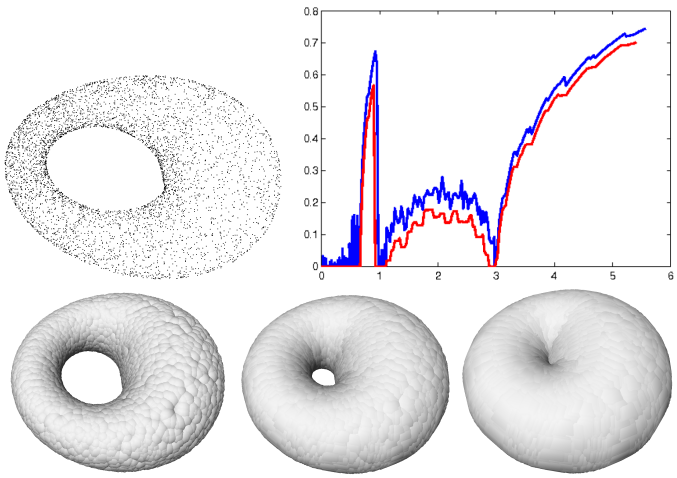

Figure 11.9: A 4000 points set (left) sampled around a torus shape in $\mathbb{R}^3$ (which is not a torus of revolution) and its critical function (the upper curve). The lowest curve represents the lower bound for the critical function of any shape K at distance less than some fixed threshold (here 0.001, the diameter of the torus being 10) from the point cloud. We distinguish 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).

11.6 Sampling conditions and μ-reach

The theorem 11.7 allows to give a first topological stability result for the offsets of compact sets when their Hausdorff distance is a fraction of their weak feature sizes. To get stronger results, in this section we introduce the notion of μ-reach that can be seen as a parametrized version of the weak feature size.

Definition 11.15 *For $0 < \mu \leq 1$, the μ-reach $r_\mu(K)$ of a compact set*

$K \subset \mathbb{R}^d$ is defined by

$$r_\mu(K) = \inf\{r > 0 : \chi_K(r) < \mu\}$$

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* and has been introduced by H. Federer [67] in Geometric Measure Theory. The function $\mu \to r_\mu(K)$ is non increasing and we have

$$\lim_{\mu \to O^+} r_\mu(K) \leq \text{wfs}(K)$$

Note that the above inequality can be strict (exercise 11.10).

It follows from the the critical point stability theorem 11.9 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 which proof is left to the reader (see also [33]).

Theorem 11.16 (critical values separation theorem) *Let K and K' be two compact subsets of $\mathbb{R}^d$, ε be the Hausdorff distance between K and K', and μ be 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*

$$]\frac{4\varepsilon}{(\mu - \mu')^2}\,,\, r_\mu(K') - 3\sqrt{\varepsilon r_\mu(K')}[$$

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 buid examples showing that the bounds of the interval in the above theorem are tight (see [33]).

The notion of μ-reach allows to introduce the following sampling condition.

Definition 11.17 *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 $K' \subset \mathbb{R}^d$ if*

$$d_H(K, K') \leq \kappa r_\mu(K')$$

The (κ, μ)-approximations generalize some sampling conditions introduced in the setting of surface reconstruction in $\mathbb{R}^3$. For example, when K' is a smooth compact surface in $\mathbb{R}^3$ and $K \subset \mathbb{R}^3$ is a point cloud, the notion of $(\kappa, 1)$-approximation is closely related to the r-samples introduced in [89].

11.7 Offset reconstruction

With the notion of (κ, μ)-approximations 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 11.18 (Isotopic Reconstruction Theorem) *Let $K' \subset \mathbb{R}^d$ be a compact set such that $r_\mu(K') > 0$ for some $\mu > 0$. Let K be a (κ, μ)-approximation of K' where*

$$\kappa < \min\left(\frac{\sqrt{5}}{2} - 1, \frac{\mu^2}{16 + 2\mu^2}\right)$$

and let r, r' be such that

$$0 < r' < \mathrm{wfs}(K') \quad \text{and} \quad \frac{4\kappa r_\mu}{\mu^2} \leq r < r_\mu(K') - 3\kappa r_\mu$$

Then the offset K^r and level set hypersurface $d_K^{-1}(r)$ are isotopic to $K'^{r'}$ and $d_{K'}^{-1}(r')$ respectively.

Indeed the proof of the isotopy between the offsets is beyond the scope of these notes (but it can be found in [42]). We prove here the following weaker version.

Theorem 11.19 *Let $K \subset \mathbb{R}^d$ be a (κ, μ)-approximation of 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

$$0 < \eta < \mathrm{wfs}(K') \quad \text{and} \quad \frac{4d_H(K, K')}{\mu^2} \leq \alpha < r_\mu(K') - 3d_H(K, K')$$

Proof The critical values separation theorem 11.16 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 11.5 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 (see Exercise 11.9)

$$\chi_{K^\alpha}(r) = \chi_K(r + \alpha)$$

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 the theorem 11.7, the conclusion of 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 $\kappa < \frac{\mu^2}{5\mu^2+12}$. □

11.8 Bibliographical notes

The distance functions framework introduced in this chapter does not restrict to topologically correct inference or shape reconstruction has been extended to further geometric inference problems. For example it has been used to prove stability results for normals [42] and curvatures [35] 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 (CAGD) are theoretically well-founded [34].

Distance functions have been widely studied and used in Riemannian Geometry [45, 76] and non-smooth analysis [48]. The above introduced notion of critical point coincide with the notion of critical point for distance function used in Riemannian Geometry and non-smooth analysis where the notion of Clarke gradient is closely related to the above defined gradient.
Distance functions are also particular cases of so-called semi-concave functions. Many of the results presented in this section can be deduced from general results for semi-concave functions [100]. 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 above defined gradient flow $\mathfrak{C}$ have been previously and independantly considered in the litterature [60, 74, 28].

11.9 Exercises

Exercise 11.1 *Prove that d_K is 1-Lipschitz.*

Exercise 11.2 *Prove that the three above definitions of Hausdorff distance are equivalent and that d_H is a distance [4] on the space of compact subsets of $\mathbb{R}^d$.*

Exercise 11.3 *Prove that $\sigma_K(x)$ the smallest enclosing ball containing $\Gamma_K(x)$ exists and is unique.*

Exercise 11.4 *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 11.5 *Show that for any $x \in \mathbb{R}^d$, $\theta_K(x)$ is the closest point to x on the convex hull of $\Gamma_K(x)$.*

Exercise 11.6 *Prove that for any $x \in \mathbb{R}^d$ one has the following equivalence: (x is a critical point of d_K) $\Leftrightarrow$ (x is in the convex hull of $\Gamma_K(x)$)*

Exercise 11.7 *Let K be a finite set of points in $\mathbb{R}^2$ (one may assume that the points are in general position). Prove that a point $x \in \mathbb{R}^2$ is a critical point of d_K isf and only if it satisfies one of the following conditions:*
- $x \in K$,
- x is an intersection point between a Voronoï edge and its dual Delaunay edge,
- x is a Voronoï vertex contained in its dual Delaunay triangle.
How does this result generelize for finite point clouds in higher dimensions.

[4] Recall that d_H is a distance if and only if it has the three following properties: (i) $d_H(K, K') = 0$ if and only if $K = K'$; (ii) $d_H(K, K') = d_H(K', K)$; (iii) $d_H(K, K'') \leq d_H(K, K') + d_H(K', K'')$.

Exercise 11.8 *Let $K = \{p_1, \cdots p_n\} \subset \mathbb{R}^d$ be a finite point set. Prove that* $\mathrm{wfs}(K) = \frac{1}{2} \min_{i \neq j} \|p_i - p_j\|$.

Exercise 11.9 *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 11.10 *Give an example of a compact set (e.g. a compact subset of $\mathbb{R}^2$) K such that $\lim_{\mu \to O^+} r_\mu(K) \neq \mathrm{wfs}(K)$.*

Chapter 12

Distance to probability measures

12.1 Back to distance-based inference: the problem of outliers

As we have seen in Chapter 11, the use of distance functions provides an interesting framework for the robust estimation of the topological and geometric properties of an approximated shape in $\mathbb{R}^d$. The general problem of geometric inference can be stated in the following general but not precise way.

Problem (topological and geometric inference): Given an *approximation* C of a *geometric object* K is it possible and how can we reliably estimate the topological and geometric properties of K from C?

Obviously, this problem needs to be stated in a more precise framework. In particular, it is necessary to define the class of geometric objects that are considered and the notion of distance that is used between these objects. For example in Chapter 11, the considered objects are a class of compact subsets of $\mathbb{R}^d$ (compact sets with positive wfs or positive μ-reach) and the Hausdorff distance is used to quantify the notion of approximation. 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 12.1). As a consequence, the distance-based approach of Chapter 11 fails for data corrupted by noise and outliers as illustrated on figure 12.2. To overcome this issue we adapt the distance-based framework for geometric inference to the general framework of data carrying "noise" and "outliers".

The three main ingredients for distance-based inference It is important to notice that all the inference results of Chapter 11 follow from only three fundamental properties of distance functions:

- Stability of the map $K \to d_K$: 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*: $x \to \|x\|^2 - d_K^2(x)$ is convex.

The first property is an obvious necessary 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 11.2) and the isotopy lemma of Section 11.3. These results still hold for general proper semiconcave functions [100]), motivating the following definition of functions that are of particular interest for geometric inference.

Definition 12.1 *A non-negative function* $\phi : \mathbb{R}^d \to \mathbb{R}^+$ *is* distance-like *if*
i) ϕ *is* 1-*Lipschitz,*
ii) ϕ^2 *is* 1-*semiconcave,*
iii) ϕ *is* proper *i.e., for any compact set* $K \subset \mathbb{R}$, $\phi^{-1}(K)$ *is compact.*

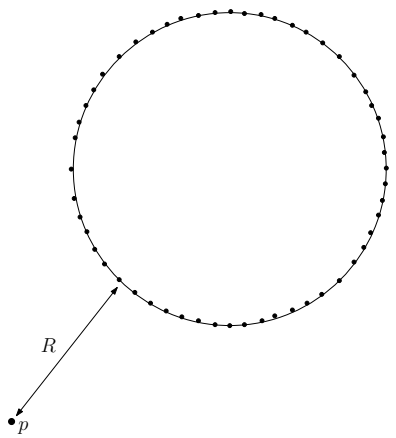

Figure 12.1: Adding just one point at distance R to a point cloud sampling a circle changes the Hausdorff distance between the shape and the sample by R.

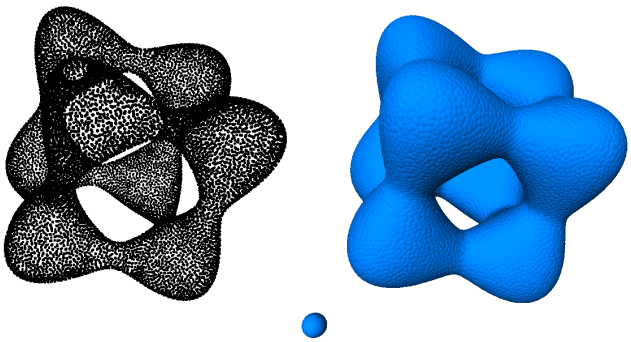

Figure 12.2: To a set of points densely sampled on the surface S of a tangle cube a single "outlier" located away from S has been added (left). When we consider the offsets, the new added point creates a connected component that makes the estimation of the topology of S from the offsets of the data imposible (right): e.g. the estimated number of connected components (two) is clearly wrong.

12.2 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. (Borel) measures, instead of purely geometric compact sets. Considering (probability) measures as the new class of studied objects allows us to quantify the notion of proximity using the so-called Wasserstein distances that are much better adaptated to cope with noise and outliers.

12.2.1 Replacing compact sets by measures

Definition 12.2 *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 numerable family of disjoint Borel subsets of* $\mathbb{R}^d$, $\mu\left(\cup_{i\in\mathbb{N}} B_i\right) = \sum_i \mu(B_i)$. *A* probability measure *is a measure whose total mass* $\mu(\mathbb{R}^d)$ *is equal to* 1.

The *support* of a measure μ is the smallest closed set K on which the mass of μ is concentrated, i.e. $\mu(\mathbb{R}^d \setminus K) = 0$. 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. Intuitively this is a unit of mass concentrated on x. Given a set of N points C, the *counting* or *uniform measure* μ_C, associated to C, is defined by $\mu_C(B) = \frac{1}{N}|B \cap C|$. 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 on M. As M is compact, $\mathrm{vol}_M(M)$ is finite and we define a probability measure μ_M supported on M by $\mu_M(B) = vol_M(B \cap M)/\mathrm{vol}_M(M)$, 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 $M_1, \cdots, M_k$ then we can define probability measures on M just by considering weighted sums of the measures μ_{M_i}.

12.2.2 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 distances W_p, for some technical reasons that become clear in the following we only consider the distance W_2.

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 $A, B \subseteq \mathbb{R}^d$ $\pi(A \times \mathbb{R}^d) = \mu(A)$ 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

$$\mathcal{C}(\pi) = \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 \mathrm{d}\pi(x, y) \right)^{1/2}$$

As an example, when μ and ν are two probability measures with finite supports:

$$\mu = \sum_{j=1}^{m} c_j \delta_{x_j} \quad \text{and} \quad \nu = \sum_{i=1}^{n} d_i \delta_{y_i}$$

with $\sum_{j=1}^{m} c_j = 1$ and $\sum_{i=1}^{n} d_i = 1$, a transport plan between μ and ν is then a $n \times m$ matrix $\Pi = (\pi_{i,j})$ with non negative entries such that

$$\sum_{i=1}^{n} \pi_{i,j} = c_j \quad \text{and} \quad \sum_{j=1}^{m} \pi_{i,j} = d_i.$$

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 12.3 *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 $W_2(\mu, \nu)$.*

Even for measures with finite support, the computation of the Wasserstein distance is practically intractable. However, it provides an interesting notion to quantify the resilience to a reasonable amount of outliers. To illustrate this, consider a set $C = \{x_1, x_2, \cdots, x_N\}$ 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 $i = 1, \cdots, n$. 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

$$W_2(\mu, \nu) \leq (\frac{n}{N})^{\frac{1}{2}} (R + \text{diam}(C))$$

while the Hausdorff distance between C and C' is at least R. To prove this inequality, let 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 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 ($n << N$) 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 cloud data 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 independant of any convergence property of empirical measures.

12.3 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 perturbation of the considered probability measure.

12.3.1 Definition

The distance function to a compact set K 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 12.4 *Let μ be a probability measure on $\mathbb{R}^d$ and $0 \leq m < 1$ a given a parameter. We denote $\delta_{\mu,m}$ 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, i.e. the map $\mu \to \delta_{\mu,m}$ is not continuous as illustrated in 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 $\varepsilon > 0$ one has

$\delta_{\mu_\varepsilon,1/2}(t) = |1-t|$ for $t < 0$ while if $\varepsilon = 0$, one obtains $\delta_{\mu_0,1/2}(t) = |t|$ 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 12.5 *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 $d_{\mu,m_0} : \mathbb{R}^d \to \mathbb{R}_+$ defined by :*

$$d^2_{\mu,m_0} : \mathbb{R}^n \to \mathbb{R}^+,\ x \mapsto \frac{1}{m_0}\int_0^{m_0} \delta_{\mu,m}(x)^2 dm$$

12.3.2 Distance function to empirical measures

An interesting property of the above defined functions is that they have a very 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 kth nearest neighbor in C, where k is the smallest integer larger than $m\,|C|$. Hence the function $m \mapsto \delta_{\mu_C,m}(x)$ is constant equal to the distance from x to its k-th nearest neighbor in C on each interval $(\frac{k}{N}, \frac{k+1}{N}]$. Integrating the square of this piecewise constant functions gives the following expression for d^2_{μ,m_0}, where $m_0 = k_0/\,|C|$:

$$\begin{aligned} d^2_{\mu,m_0}(x) &= \frac{1}{m_0}\int_0^{m_0} \delta_{\mu,m}(x)^2 = \frac{1}{m_0}\sum_{k=1}^{k_0}\frac{1}{N}\delta_{\mu,k/N}(x)^2 \\ &= \frac{1}{k_0}\sum_{p\in \mathrm{NN}_C^{k_0}(x)} \|p-x\|^2 \end{aligned}$$

where $\mathrm{NN}_C^{k_0}(x)$ denote the k_0 nearest neighbors of x in C. As a consequence the pointwise evaluation of $d^2_{\mu_C,k_0/n}(x)$ reduces to a k-nearest neighbor query in C.

********** Put a figure here to illustrate! **********

12.3.3 Equivalent formulation

In this paragraph, we prove that the distance function to a measure d_{μ,m_0} 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 12.3). This equivalent formulation will be used to deduce that $\mu \to \mathrm{d}_{\mu,m_0}$ is Lipschitz and $x \to \mathrm{d}^2_{\mu,m_0}$ is semiconcave.

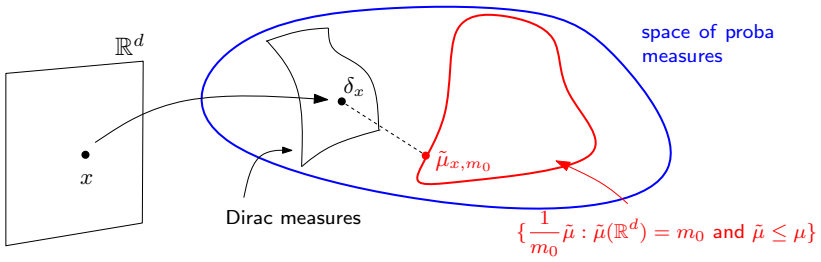

Figure 12.3: The distance function to a measure as a usual distance function in an infinite dimensional space.

Definition 12.6 *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* $\mathrm{Sub}(\mu)$*, while the set of submeasures of μ with a prescribed total mass $m_0 > 0$ is denoted by* $\mathrm{Sub}_{m_0}(\mu)$*.*

Proposition 12.7 *For any probability measure μ on $\mathbb{R}^d$, the distance function to μ at x is defined by the following equality:*

$$\mathrm{d}_{\mu,m_0}(x) = \min\left\{m_0^{-1/2}\,\mathrm{W}_2\left(m_0\delta_x, \nu\right); \nu \in \mathrm{Sub}_{m_0}(\mu)\right\} \tag{12.1}$$

Moreover, for any measure μ_{x,m_0} that realizes the above minimum one has:

$$\mathrm{d}_{\mu,m_0}(x) = \left(\frac{1}{m_0^{1/2}} \int_{\mathbb{R}^d} \|x-h\|^2 \, \mathrm{d}\mu_{x,m_0}(h) \right)^{1/2}$$

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 $\mathrm{Sub}_{m_0}(\mu)$. It is not empty but it might not be reduced to a single element. Indeed, it coincide with the set of submeasures μ_{x,m_0} of total mass m_0 whose support is contained in the closed ball $\bar{\mathbb{B}}(x, \delta_{\mu,m}(x))$, and whose restriction to the *open* ball $\mathbb{B}(x, \delta_{\mu,m}(x))$ coincides with μ.

12.3.4 Stability of the distance function to a measure

The characterization of d_{μ,m_0} given in Proposition 12.7 provides a rather easy way to prove the stability of $\mu \mapsto \mathrm{d}_{\mu,m_0}$.

Theorem 12.8 ((Distance function stability)) *If μ and μ' are two probability measures on $\mathbb{R}^d$ and $m_0 > 0$, then $\left\| \mathrm{d}_{\mu,m_0} - \mathrm{d}_{\mu',m_0} \right\|_\infty \leq \frac{1}{\sqrt{m_0}} \mathrm{W}_2(\mu,\mu')$.*

The proof of theorem 12.8 follows from the following proposition.

Proposition 12.9 *Let μ and μ' be two probability measures on $\mathbb{R}^d$. Then,*

$$d_H(\mathrm{Sub}_{m_0}(\mu), \mathrm{Sub}_{m_0}(\mu')) \leq \mathrm{W}_2(\mu,\mu')$$

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 $\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x-y\|^2 \pi(x,y) \mathrm{d}x\mathrm{d}y = \varepsilon^2$. Given a

submeasure ν of μ, one can find a submeasure π' of π that transports ν to a submeasure ν' of μ' [1]. Then,

$$\mathrm{W}_2(\nu,\nu')^2 \leq \int_{\mathbb{R}^d\times\mathbb{R}^d} \|x-y\|^2 \,\pi'(x,y)\mathrm{d}x\mathrm{d}y \leq \varepsilon^2$$

This shows that $\mathrm{dist}(\nu, \mathrm{Sub}_{m_0}(\mu')) \leq \varepsilon$ for every submeasure $\nu \in \mathrm{Sub}_{m_0}(\mu)$. The same hold by exchanging the roles of μ and μ', thus proving the bound on the Hausdorff distance. □

Proof of Theorem 12.8 The following sequence of equalities and inequalities, that follows from Propositions 12.7 and 12.9, proves the theorem:.

$$\begin{aligned}
\mathrm{d}_{\mu,m_0}(x) &= \frac{1}{\sqrt{m_0}}\mathrm{dist}_{\mathrm{W}_2}(m_0\delta_x, \mathrm{Sub}_{m_0}(\mu)) \\
&\leq \frac{1}{\sqrt{m_0}}(d_H(\mathrm{Sub}_{m_0}(\mu), \mathrm{Sub}_{m_0}(\mu')) + \mathrm{dist}_{\mathrm{W}_2}(m_0\delta_x, \mathrm{Sub}_{m_0}(\mu'))) \\
&\leq \frac{1}{\sqrt{m_0}}\,\mathrm{W}_2(\mu,\mu') + \mathrm{d}_{\mu',m_0}(x)
\end{aligned}$$

□

12.3.5 The distance to a measure is distance-like.

The subdifferential of a function $f : \Omega \subseteq \mathbb{R}^d \to \mathbb{R}$ 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, $f(x+h) \geq f(x) + \langle h|v\rangle$. This gives a characterization of convexity: a function $f : \mathbb{R}^d \to \mathbb{R}$ 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.

Proposition 12.10 *The function $v_{\mu,m_0} : x \in \mathbb{R}^d \mapsto \|x\|^2 - \mathrm{d}^2_{\mu,m_0}$ is convex, and its subdifferential at a point $x \in \mathbb{R}^d$ is given by*

$$\partial_x v_{\mu,m_0} = \left\{ 2x - \frac{2}{m_0}\int_{h\in\mathbb{R}^d} (x-h)\,\mathrm{d}\mu_{x,m_0}(h)\,;\, \mu_{x,m_0} \in \mathcal{R}_{\mu,m_0}(x) \right\}$$

[1] This 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.

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 12.7 we have the following sequence of equalities and inequalities:

$$\begin{aligned}
\mathrm{d}^2_{\mu,m_0}(y) &= \frac{1}{m_0}\int_{h\in\mathbb{R}^d} \|y-h\|^2 \,\mathrm{d}\mu_{y,m_0}(h) \\
&\leq \frac{1}{m_0}\int_{h\in\mathbb{R}^d} \|y-h\|^2 \,\mathrm{d}\mu_{x,m_0}(h) \\
&\leq \frac{1}{m_0}\int_{h\in\mathbb{R}^d} \|x-h\|^2 + 2\langle x-h|y-x\rangle + \|y-x\|^2 \,\mathrm{d}\mu_{x,m_0}(h) \\
&\leq \mathrm{d}^2_{\mu,m_0}(x) + \|y-x\|^2 + \langle v|y-x\rangle
\end{aligned}$$

where v is the vector defined by

$$v = \frac{2}{m_0}\int_{h\in\mathbb{R}^d} [x-h]\,\mathrm{d}\mu_{x,m_0}(h).$$

The inequality can be rewritten as:

$$(\|y\|^2 - \mathrm{d}^2_{\mu,m_0}(y)) - (\|x\|^2 - \mathrm{d}^2_{\mu,m_0}(x)) \geq \langle 2x - v|y-x\rangle$$

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. □

Corollary 12.11 *The function d^2_{μ,m_0} is 1-semiconcave. Moreover,*

(i) d^2_{μ,m_0} is differentiable almost everywhere in $\mathbb{R}^d$, with gradient defined by

$$\nabla_x \mathrm{d}^2_{\mu,m_0} = \frac{2}{m_0}\int_{h\in\mathbb{R}^d} [x-h]\,\mathrm{d}\mu_{x,m_0}(h)$$

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:

$$\nabla_x \mathrm{d}_{\mu,m_0} = \frac{\nabla_x \mathrm{d}^2_{\mu,m_0}}{2\mathrm{d}_{\mu,m_0}} = \frac{1}{\sqrt{m_0}} \frac{\int_{h\in\mathbb{R}^d}[x-h]\,\mathrm{d}\mu_{x,m_0}(h)}{(\int_{h\in\mathbb{R}^d}\|x-h\|^2\,\mathrm{d}\mu_{x,m_0}(h))^{1/2}}$$

Using the Cauchy-Schwartz inequality we find the bound $\|\nabla_x \mathrm{d}_{\mu,m_0}\| \leq 1$ which proves the statement. □

12.4 Applications to geometric inference

Reconstruction from point clouds with outliers was the main motivation for introducing the distance function to a measure. In this section, we show that the Reconstruction Theorem 11.18 of Chapter 11 can be generalized to compare the sub-level sets of two close distance-like functions. It is also possible to adapt most of the topological and geometric inference results of Chapter 11 in a similar way.

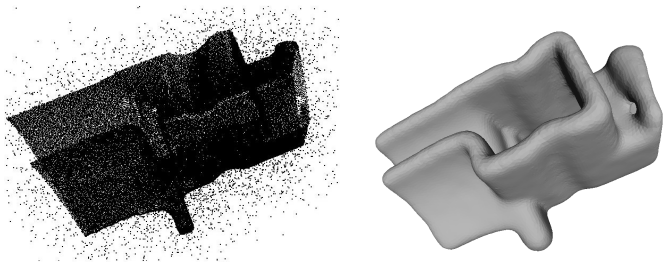

Figure 12.4: On the left, a point cloud sampled on a mechanical part to which 10% of outliers (uniformly sampled in a box enclosing the model) have been added. On the right, the reconstruction of an isosurface of the distance function d_{μ_C,m_0} to the uniform probability measure on this point cloud.

12.4.1 Extending the sampling theory for compact sets

Let $\phi : \mathbb{R}^d \to \mathbb{R}$ be a *distance-like* function. The 1-semiconcavity of ϕ^2 allows to define a notion of gradient vector field $\nabla_x\phi$ for ϕ, defined everywhere and satisfying $\|\nabla_x\phi\| \leq 1$. Although not continuous, the vector field $\nabla\phi$ is sufficiently regular to be integrated in a continuous locally Lipschitz flow

$\Phi^t : \mathbb{R}^d \to \mathbb{R}^d$, $t \geq 0$. 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 $t > 0$, $\frac{\mathrm{d}\gamma}{\mathrm{d}\varepsilon}\Big|_{t^-} = \nabla_{\gamma(t)}\phi$. Moreover, for any integral curve $\gamma : [a, b] \to \mathbb{R}^d$ parametrized by arc-length, one has:

$$\phi(\gamma(b)) = \phi(\gamma(a)) + \int_a^b \left\|\nabla_{\gamma(t)}\phi\right\| \mathrm{d}t.$$

Definition 12.12 *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 $\|\nabla_x\phi\|$ 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 minimum $r' > 0$ such that ϕ doesn't have any critical value between r and $r + r'$. We denote it by* $\mathrm{wfs}_\phi(r)$. *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. Obviously, the α-reach is always a lower bound for the weak-feature size, with $r = 0$.*

The Isotopy Lemma 11.5 extends to distance-like functions and as in Theorem 11.7 the offsets of two uniformly close distance-like functions with large weak feature size have the same homotopy type.

Proposition 12.13 (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 $r \in [r_1, r_2]$.*

Proposition 12.14 *Let ϕ and ψ be two distance-like functions, such that $\|\phi - \psi\|_\infty \leq \varepsilon$. Suppose moreover that $\mathrm{wfs}_\phi(r) > 2\varepsilon$ and $\mathrm{wfs}_\psi(r) > 2\varepsilon$. 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 11.9 also holds for distance-like functions.

Proposition 12.15 *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 11.9 □

Corollary 12.16 *Let ϕ and ψ be two ε-close distance-like functions, and suppose that $\mathrm{reach}_\alpha(\phi) \geq R$ for some $\alpha > 0$. Then, ψ has no critical value in the interval $\left]4\varepsilon/\alpha^2, R - 3\varepsilon\right[$.*

Proof The proof is almost verbatim the same as the proof of Theorem 11.16. □

Theorem 12.17 (Reconstruction theorem) *Let ϕ, ψ be two ε-close distance-like functions, with $\mathrm{reach}_\alpha(\phi) \geq R$ 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 isotopy lemma, 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:

$$\mathrm{wfs}_\phi(r) \geq R - 4\varepsilon/\alpha^2 \text{ and } \mathrm{wfs}_\psi(r) \geq R - 3\varepsilon - 4\varepsilon/\alpha^2$$

By Proposition 12.14, 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 $\frac{1}{2}\mathrm{wfs}_\phi(r)$ and $\frac{1}{2}\mathrm{wfs}_\psi(r)$. This is true, provided that $2\varepsilon \leq R - \varepsilon(3 + 4/\alpha^2)$. The theorem follows. □

Remark that in the above definition 12.12 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 12.17 would still hold under the weaker condition that the (r, α)-reach of ϕ is positive.

12.4.2 Distance to a measure vs. distance to its support

In this paragraph, 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 $x \in \mathbb{R}^d$, $d_S(x) \leq \mathrm{d}_{\mu,m_0}(x)$. 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 $\mathrm{d}_{\mu,m_0} - \mathrm{d}_S$ 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 $\mathbb{B}(p,r)$ centered at a point p of the support decreases with r. Let us define:

(i) We say that a non-decreasing positive function $f : \mathbb{R}^+ \to \mathbb{R}^+$ is a *uniform lower bound on the growth of* μ if for every point p in the support of μ and every $\varepsilon > 0$, $\mu(\mathbb{B}(p,\varepsilon)) \geq f(\varepsilon)$;

(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 12.18 *Let μ be a probability measure and f be a uniform lower bound on the growth of μ. Then $\|\mathrm{d}_{\mu,m_0} - \mathrm{d}_S\|_\infty < \varepsilon$ 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 $\|x-p\| = d(p,S)$. By assumption, $\mu(\mathbb{B}(x, \mathrm{d}_S(x)+\varepsilon)) \geq \mu(\mathbb{B}(p,\varepsilon)) \geq m_0$. Hence, $\delta_{\mu,m_0}(x) \leq \mathrm{d}_S(x) + \varepsilon$. The function $m \mapsto \delta_{\mu,m}(x)$ being non-decreasing, we get: $m_0 \mathrm{d}_S^2(x) \leq \int_0^{m_0} \delta_{\mu,m}^2(x)\mathrm{d}m \leq m_0(\mathrm{d}_S(x)+\varepsilon)^2$. Taking the square root of this expression proves the lemma. □

Corollary 12.19 *(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:*

$$\|d_{\mu,m_0} - d_S\|_\infty = \sup_{x\in\mathbb{R}^d} |d_{\mu,m_0}(x) - d_S(x)| \overset{m_0\to 0}{\to} 0$$

(ii) If the measure μ has dimension at most $k > 0$, then

$$\|\mathrm{d}_{\mu,m_0} - \mathrm{d}_S\|_\infty \leq C(\mu)^{-1/k} m_0^{1/k}$$

Proof (i) If S is compact, there exists a sequence $x_1, x_2, \cdots$ 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, $\eta(\varepsilon) = \min_{i=1\cdots n} \mu(\mathbb{B}(x_i, \varepsilon/2))$ is positive. Now, for any point $x \in S$, there is a x_i such that $\|x - x_i\| \leq \varepsilon/2$. Hence, $\mathbb{B}(x_i, \varepsilon/2) \subseteq \mathbb{B}(x, \varepsilon)$, which means that $\mu(\mathbb{B}(x, \varepsilon)) \geq \eta(\varepsilon)$. (ii) Follows straightforwardly from the Lemma. □

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 12.20 *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*

$$\|\mathrm{d}_S - \mathrm{d}_{\mu,m_0}\|_\infty \leq C(S)^{-1/k} m_0^{1/k}$$

for m_0 small enough and $C(S) = (2/\pi)^k \beta_k / \mathcal{H}^k(S)$ 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 [73, 3.101]).

12.4.3 Shape reconstruction from noisy data

The previous results lead to shape reconstruction theorems from 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 $d_K : \mathbb{R}^d \to \mathbb{R}_+$ 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

$$\begin{aligned} \left\|d_K - d_{\mu',m_0}\right\|_\infty &\leq \|d_K - d_{\mu,m_0}\|_\infty + \|d_{\mu,m_0} - d_{\mu',m_0}\|_\infty && (12.2) \\ &\leq C(\mu)^{-1/k} m_0^{1/k} + \frac{1}{\sqrt{m_0}} \mathrm{W}_2(\mu, \mu') && (12.3) \end{aligned}$$

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 12.8 and the corollary 12.19. 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 m_0 make the distance functions to measures more stable. Eq. 12.2 leads to the following corollary of Theorem 12.17:

Corollary 12.21 *Let μ be a measure and K its support. Suppose that μ has dimension at most k and that $\mathrm{reach}_\alpha(d_K) \geq R$ 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:*

$$\mathrm{W}_2(\mu,\mu') \leq \frac{R\sqrt{m_0}}{5 + 4/\alpha^2} - C(\mu)^{-1/k} m_0^{1/k+1/2}$$

Figure 12.4 illustrates the reconstruction Theorem 12.17 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.

12.5 Bibliographical notes

Most of the chapter comes from [36] that introduces and studies stability properties of distance functions to a probability measure.

Wasserstein distances are closely related to the theory of optimal transportation (see e.g. [107]). The distance W_1 is also known as the earth-mover distance, and has been used in vision by [99] and in image retrieval by [102] and others.

General results about semiconcave functions can be found in [100, 27].

The convergence properties of empirical measure with respect to the Wasserstein metric have been widely studied and quantitative results can be found in [25].

The complete proofs of Propositions 12.7, 12.10 and 12.14 are given in [36].

The Günther-Bishop Theorem is stated in [73, 3.101] and the proof of Proposition 12.20 can be found in [36].

12.6 Exercises

Exercise 12.1 *Let $\mu_1, \cdots, \mu_k$ be measures on $\mathbb{R}^d$ and let $\lambda_1, \cdots, \lambda_k \in \mathbb{R}$. Show that if all the λ_i's are non negative then $\mu = \sum_{i=1}^k \mu_i$ is a measure. Show that the set of probability measures on $\mathbb{R}^d$ is convex.*

Exercise 12.2 *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 $x_1, \cdots, x_m$ and $y_1, \cdots, y_n$ 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

$$\sum_{i=1}^{n} \pi_{i,j} = c_j \quad \textit{and} \quad \sum_{j=1}^{m} \pi_{i,j} = d_i.$$

Exercise 12.3 *Let μ be a probability measure on $\mathbb{R}^d$ anf 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 12.4 *Let μ be a probability measure on $\mathbb{R}^d$ and let $0 < m_0 < 1$ and $x \in R^d$. Recall that we denote by $\mathcal{R}_{\mu,m_0}(x)$ the set of submeasures minimizing the right hand term of Equation (12.1). Give an example of 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 $\mu(\mathbb{S}(x, \delta_{\mu,m_0}))$ of the sphere of center x and radius δ_{μ,m_0} must be positive.

Chapter 13

Homology inference

13.1 Simplicial homology and Betti numbers

In this section we introduce the basic notion of simplicial homology that are necessary to define and study topological persistence. In particular we restrict to the homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$. In the sequel of this section, K denotes a finite d-dimensional simplicial complex.

13.1.1 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 $\{\sigma_1, \cdots, \sigma_p\}$ is the set of k-simplices of K any k-chain c can be uniquely written

$$c = \sum_{i=1}^{p} \varepsilon_i \sigma_i \quad \text{with} \quad \varepsilon_i \in \mathbb{Z}/2\mathbb{Z} = \{0, 1\}$$

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 13.1 *The space of k-chains is the set $C_k(K)$ of the simplicial k-chains of K with the above defined operations. This is a $\mathbb{Z}/2\mathbb{Z}$-vector space.*

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 13.1, $C_1(K)$ is the $\mathbb{Z}/2\mathbb{Z}$-vector space generated by the edges $e_1 = ab$, $e_2 = bc$, $e_3 = ca$, $e_4 = cd$, i.e.

$$C_1(K) = \{0, e_1, e_2, e_3, e_4, e_1 + e_2, e_1 + e_3, e_1 + e_4, e_2 + e_3, e_2 + e_4, e_3 + e_4, e_1 + e_2 + e_3, \cdots\}$$

Summing $e_1 + e_2$ with $e_2 + e_3 + e_4$ gives $e_1 + e_3 + e_4$ (note that 0 denotes the empty simplex).

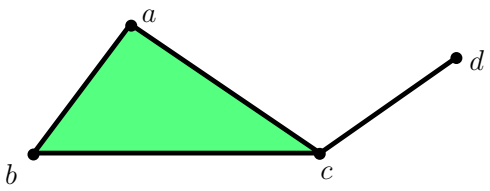

Figure 13.1: A very simple simplicial complex made of 4 vertices, 4 edges and 1 triangle.

Chains with coefficient in $\mathbb{Z}/2\mathbb{Z}$ have an obvious geometric interpretation: since any k-chain can be uniquely written as $c = \sigma_{i_1} + \sigma_{i_2} + \cdots \sigma_{i_m}$ where the σ_{i_j} are k-simplices, c can be considered as the union of the simplices σ_{i_j}. The sum of two k-chains is equal to their symetric difference.

13.1.2 The boundary operator and homology groups

Definition 13.2 *The boundary $\partial(\sigma)$ of a k-simplex σ is the sum of its $(k-1)$-faces. This is a $(k-1)$-chain.*

If $\sigma = [v_0, \cdots, v_k]$ is the k-simplex generated by the $(k+1)$-vertices $v_0, \cdots, v_k$, 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 generated by the sets of all the vertices of σ except v_i. The boundary operator, defined on the simplices of k, extends linearly to $C_k(K)$.

Definition 13.3 *The boundary operator is the linear map defined by*

$$\begin{array}{rccl} \partial: & C_k(K) & \to & C_{k-1}(K) \\ & c & \to & \partial c = \sum_{\sigma \in c} \partial(\sigma) \end{array}$$

Notice that one should denote ∂_k the above defined operator but to avoid heavy notations one usually omit the index notations.

Proposition 13.4

$$\partial\partial := \partial \circ \partial = 0$$

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 operators define a sequence of linear maps between spaces of chains.

Definition 13.5 *The chain complex associated to a complex K of dimension d is the following sequence of linear operators*

$$\emptyset \to C_d(K) \xrightarrow{\partial} C_{d-1}(K) \xrightarrow{\partial} \cdots C_{k+1}(K) \xrightarrow{\partial} C_k(K) \xrightarrow{\partial} \cdots C_1(K) \xrightarrow{\partial} C_0(K) \xrightarrow{\partial} \emptyset$$

For $k \in \{0, \cdots, d\}$, the set $Z_k(K)$ of *k-cycles* of K is the kernel of $\partial : C_k \to C_{k-1}$:

$$Z_k(K) := \ker(\partial : C_k \to C_{k-1}) = \{c \in C_k : \partial c = \emptyset\}$$

The image $B_k(K)$ of $\partial : C_{k+1} \to C_k$ is the set of k-chains bounding a $(k+1)$-chain:

$$B_k(K) := im(\partial : C_{k+1} \to C_k) = \{c \in C_k : \exists c' \in C_{k+1}, c = \partial c'\}$$

B_k and Z_k are subspaces of C_k and according to Proposition 13.4, one has

$$B_k(K) \subset Z_k(K) \subset C_k(K)$$

Definition 13.6 *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 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*; in this case there exists a boundary cycle $b \in B_k(K)$ such that $c' = c + b$, i.e. $c' - c = c' + c = b$.

13.2 An algorithm to compute Betti numbers

Let K be a finite simplicial complex of diemnsion d and

$$\mathbb{F} = \{\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K\}$$

a filtration of K such that for any $i = 0, \cdots m-1$,

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

By considering the evolution of the Betti numbers of the filtration as we add the simplices σ^i we get the following algorithm.

Algorithm 9 Betti numbers computation

Input: A filtration $\mathbb{F}$ of a d-dimensional simplicial complex K containing m simplices.
$\beta_0 \leftarrow 0;\ \beta_1 \leftarrow 0; \cdots\ \beta_d \leftarrow 0$
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 $\beta_0, \beta_1, \cdots, \beta_d$ of K.

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 let add the simplex σ^i of dimension $k+1$ to get K^i. Remark that according to Lemma 4.6, σ^i cannot be in 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:

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

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 sub-space 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 $B_k(K^{i-1}) \subset B_k(K^i)$, $\beta_{k+1}(K^i) \leq \beta_{k+1}(K^{i-1}) + 1$.

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 a distinction between the simplices of the filtration of K that is going to play an important role in the definition of topological persistence.

Definition 13.7 *Let K be a d-dimensional simplicial complex and let*

$$\mathbb{F} = \{\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K\}$$

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 13.1 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 require 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 13.5. 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.

13.3 Singular homology and topological invariance

The homology groups and Betti numbers are topological invariants: if K and K' are two simplicial complexes with homeomorphic supports then their homology groups are isomorphic and their Betti numbers are equal. This result is still true is the supports $|K|$ 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 and requires the notion of *singular homology* (see [94, 80] for details).

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 simplex spanned by the vertices x_i, $i = 1, \cdots k+1$, 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 $\sigma : \Delta_k \to X$. 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. The Proposition 13.4 still holds for the (singular) boundary operator and the k^{th} singular homolgy group of X is similarly defined 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 and thus they have the same Betti numbers that are, indeed, the ones of X. As a consequence, in the

following 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 $f : X \to Y$ is a continuous map between two topological spaces and if $\sigma : \Delta_k \to X$ is a singular simplex in X, then $f \circ \sigma : \Delta_k \to Y$ 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 $f_\star : H_k(X) \to H_k(Y)$. Moreover if f is an homeomorphism, $f_\star$ is an isomorphism and $f_\star^{-1} = (f^{-1})_\star$. Indeed, even if f is an homotopy equivalence with homotopic inverse $g : Y \to X$ [1], then $f_\star$ is an isomorphism with inverse $g_\star$. As a consequence, two spaces that are homotopy equivalent have the same Betti numbers (take care that, when X is not homotopy equivalent to a finite simplicial complex, its Betti numbers might not be finite).

13.4 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 11 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 $\mathrm{wfs}(K) > 0$ 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 11.5, all the r-offsets K^r of K, for $0 < r < \mathrm{wfs}(K)$, are homeomorphic and thus have isomorphic homology groups. The goal of this section is to provide an effective method to compute the Betti numbers $\beta_k(K^r)$, $0 < r < \mathrm{wfs}(K)$, from P.

Theorem 13.8 *Assume that* $\mathrm{wfs}(K) > 4\varepsilon$. *For* $\alpha > 0$ *such that* $4\varepsilon + \alpha < \mathrm{wfs}(K)$, *let* $i : P^{\alpha+\varepsilon} \hookrightarrow P^{\alpha+3\varepsilon}$ *be the canonical inclusion. Then for any non negative integer* k *and any* $0 < r < \mathrm{wfs}(K)$,

$$H_k(K^r) \cong im(i_\star : H_k(P^{\alpha+\varepsilon}) \to H_k(P^{\alpha+3\varepsilon})$$

[1] g is called a homotopic inverse of f if both $f \circ g$ and $g \circ f$ are homotopic to the indentity maps in Y and X respectively

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

$$H_k(K^\alpha) \to H_k(P^{\alpha+\varepsilon}) \to H_k(K^{\alpha+2\varepsilon}) \to H_k(P^{\alpha+3\varepsilon}) \to H_k(K^{\alpha+4\varepsilon}).$$

Since $\mathrm{wfs}(K) > \alpha + 4\varepsilon$, it follows from the Isotopy Lemma 11.5 that the homomorphisms $H_k(K^\alpha) \to H_k(K^{\alpha+2\varepsilon})$ and $H_k(K^{\alpha+2\varepsilon}) \to H_k(K^{\alpha+4\varepsilon})$ induced by the inclusion maps are indeed isomorphisms. Notice that the above sequence implies that their rank is finite (see Exercise 13.4). It immediately follows that the rank of $H_k(P^{\alpha+\varepsilon}) \to H_k(P^{\alpha+3\varepsilon})$ is equal to the rank of these isomorphisms which is equal to $\beta_k(K^\alpha)$. □

Theorem 13.8 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 4.4 implies that for any $r \geq 0$, P^r is homotopy equivalent to $\check{\mathrm{C}}\mathrm{ech}(P,r)$. As a consequence $H_k(P^r)$ and $H_k(\check{\mathrm{C}}\mathrm{ech}(P,r))$ are isomorphic. Moreover, one can show that the isomorphisms can be chosen to commute with the ones induced by inclusions maps, making the following diagram commutative

$$\begin{array}{ccc} H_k(P^r) & \to & H_k(P^{r'}) \\ \uparrow & & \uparrow \\ H_k(\check{\mathrm{C}}\mathrm{ech}(P,r)) & \to & H_k(\check{\mathrm{C}}\mathrm{ech}(P,r')) \end{array} \tag{13.1}$$

We immediately obtain the following result.

Proposition 13.9 *Assume that* $\mathrm{wfs}(K) > 4\varepsilon$. *For* $\alpha > 0$ *such that* $4\varepsilon + \alpha < \mathrm{wfs}(K)$, *let* $i : \check{\mathrm{C}}\mathrm{ech}(P, \alpha+\varepsilon) \hookrightarrow \check{\mathrm{C}}\mathrm{ech}(P, \alpha+3\varepsilon)$ *be the canonical inclusion. Then for any non negative integer* k *and any* $0 < r < \mathrm{wfs}(K)$,

$$H_k(K^r) \cong im(i_\star : H_k(\check{\mathrm{C}}\mathrm{ech}(P,\alpha+\varepsilon)) \to H_k(\check{\mathrm{C}}\mathrm{ech}(P,\alpha+3\varepsilon)).$$

Thanks to the previous proposition, inferring the Betti numbers of K^r now boils down to homology computation on Čech complexes that are finite. However, 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 4.7, we obtain the following theorem.

Theorem 13.10 *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*

$$\beta_k(K^r) = rk(H_k(\text{Rips}(P,\alpha) \to H_k(\text{Rips}(P,4\alpha))$$

where $rk(H_k(\text{Rips}(P,\alpha) \to H_k(\text{Rips}(P,4\alpha))$ *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 subcomplexes of the $Vietoris-Rips$ complex. Its answer will be provided by the persistence computation algorithm in next Section. 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 11.3. 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.
Let $L := \emptyset$, $\delta := +\infty$;
while $L \subsetneq P$ **do**
  Let $p := \text{argmax}_{w \in P} \min_{v \in L} \|w - v\|$; *// arbitr. chose p if* $L = \emptyset$
  $L \leftarrow L \cup \{p\}$;
  $\delta \leftarrow \max_{w \in P} \min_{v \in L} \|w - v\|$;
  Update $\text{Rips}(L, 4\delta)$ and $\text{Rips}(L, 16\delta)$;
  Compute $\beta_{k,4\delta}^{16\delta} = rk(i_\star : H_k(\text{Rips}(L,4\delta)) \to H_k(\text{Rips}(L,16\delta)))$;
Output: diagram showing the evolution of persistent Betti numbers, i.e. the ranks of $i_\star$ versus δ.

When applied to a point cloud sampled around a compact subset of $\mathbb{R}^d$ with positive weak feature size, the algorithm provides diagrams of persistent Betti numbers $\beta_{k,4\delta}^{16\delta}$ that are constant on some intervals of values δ

as illustrated on Figure 13.2. Identifying these intervals allows the user to determine the scales at which the topological features of K^r can be infered. 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 13.3.

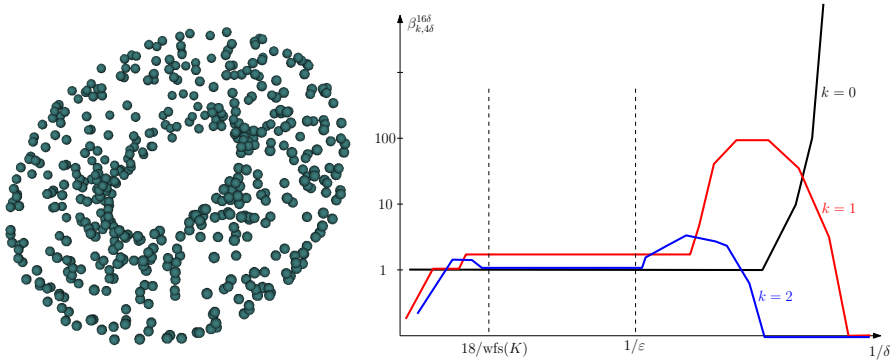

Figure 13.2: An example of persistent Betti numbers curves obtained from a point cloud sampled on a torus. They are plotted as functions of $1/\delta$. These numbers are constant on an interval containing $[18/\mathrm{wfs}(K), 1/\varepsilon]$ and correspond to the three first Betti numbers of the torus: $\beta_0 = 1$, $\beta_1 = 2$ and $\beta_2 = 1$.

The previous algorithm comes with the following theoretical guarantees justifying the existence of intervals of constant persistent Betti numbers.

Theorem 13.11 *Let $K \subset \mathbb{R}^d$ be a compact set with $\mathrm{wfs}(K) > 0$ 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 $\mathrm{wfs}(K) > 18\varepsilon$. Then at each iteration of the algorithm such that $\varepsilon < \delta < \frac{1}{18}\varepsilon$,*

$$\beta_k(K^r) = beta_{k,4\delta}^{16\delta}$$

for any $r \in (0, \mathrm{wfs}(K))$ and any non negative integer k.

Regarding complexity, it can be shown that the worst case running time of the algorithm for $K \subset \mathbb{R}^d$ is $O(8^{33^d}|P|^5)$. It depends exponentially on the dimension d but if, moreover, K is a smooth submanifold of $\mathbb{R}^d$ of dimension n, then the worst case running time of the algorithm drops to $O(8^{35^n}|P|)$. So, when K is smooth, the complexity is no longer exponential in the ambiant dimension d but only in the intrinsic dimension of K. The

example of Figure 13.3 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 computations if the complexity was exponential in $d = 1000$!

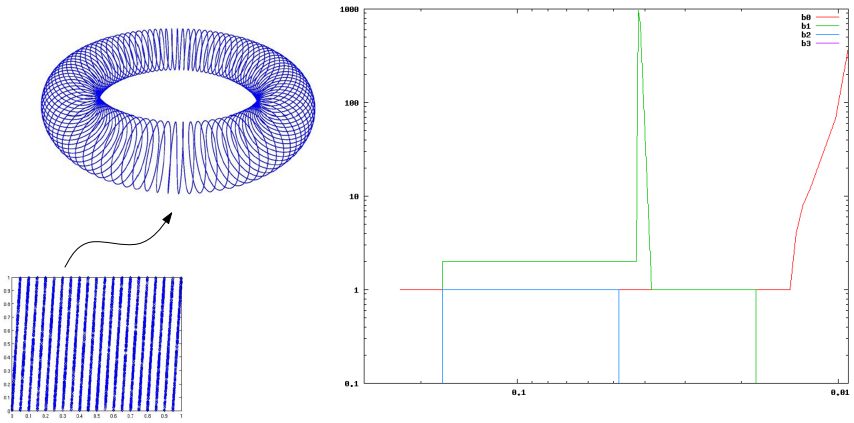

Figure 13.3: The persistent Betti numbers curves obtained from a point cloud sampled on a circular curve embedded on a torus which was itself non linearly embedded into $\mathbb{R}^{1000}$. They are plotted as functions of $1/\delta$. We distinguish two intervals on which the persistent Betti numbers are constant. The right most one (corresponding to the smaller range of δ) exhibits the Betti numbers of the circle curve $(1, 1, 0)$, while left most one (intuitively corresponding to a larger scale) exhibits the Betti numbers of the torus $(1, 2, 0)$.

13.5 Topological persistence

The algorithm of Section 13.2 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 topological persistence is to keep track of all this information and to pair the creation and destruction time of cycles appearing during the process.

13.5.1 A simple example

Before formally indroducing topological persistence, we first consider a very simple example. Let $f : [0, 1] \to \mathbb{R}$ be the function whose graph is rep-

resented on Figure 13.4. We are interested in studying the evolution of the topology of the sublevel sets $\{f \leq t\}$ 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 $a \leq t < b$, $\{f \leq t\}$ is a connected set (an interval). When t passes through the critical value b a second connected component appears and for $b \leq t < c$, $\{f \leq t\}$ 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 topological persistence framework, this pairing is either represented by the interval on the left of the graph of f on Figure 13.4 or by the point with coordinates (b, c) in the plane on the right of Figure 13.4. The length $c-b$ of the interval (b, c) represents the lifespan of the component created at b. Intuitively, the larger is the interval, the more relevant is the corresponding component. Now, continuying 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 of the sweep of the t values, 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 13.4 on the right). For technical reasons that will be become clear later in this chapter the diagonal $\{y = x\}$ is added to the diagram.

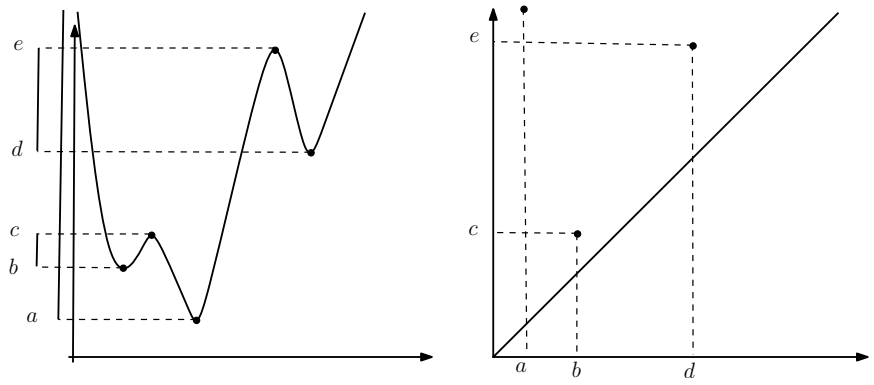

Figure 13.4: The persistence diagram of a real valued function.

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, as for the above example with connectedness changes, we can define pairs and persistence diagrams for each dimension.

Now, replacing the function f by the function g on the figure 13.5 which is "close" to f, we see that the number of pairs of g is much larger than the one of f. However most of these pairs correspond to short length intervals (points close to the diagonal) while the pairs corresponding to long interval are close to the ones of f (see figure 13.5). 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 diagram are of fundamental importance for using topological persistence in topological data analysis.

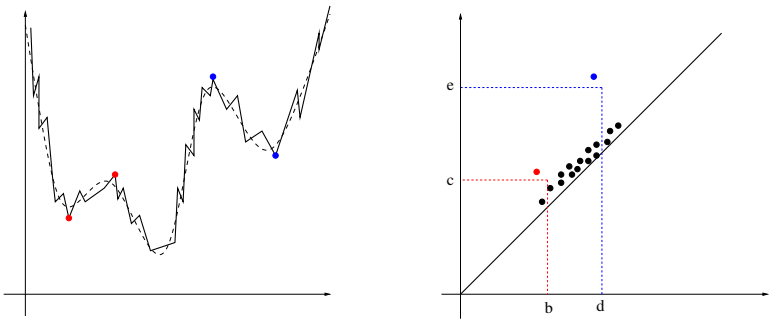

Figure 13.5: An approximation g of f and its persistence diagram

13.5.2 Topological persistence 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

$$\mathbb{F} = \{\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K\}$$

be a filtration of K such that for any $i = 0, \cdots, m-1$, $K^{i+1} = K^i \cup \sigma^{i+1}$ 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 $\partial : C_k^n \to C_{k-1}^{n-1}$, 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

$$Z_k^0 \subset Z_k^1 \subset \cdots \subset Z_k^n \cdots \subset Z_k^m = Z_k(K)$$

$$B_k^0 \subset B_k^1 \subset \cdots \subset B_k^n \cdots \subset B_k^m = B_k(K)$$

Definition 13.12 *For $p \in \{0, \cdots, m\}$ et $l \in \{0, \cdots, m-p\}$, 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 independant 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 13.2 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. Topological persistence 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 its paired negative simplex is added.

Cycle associated to a positive simplex

Lemma 13.13 *Let $\sigma = \sigma^i$ be a positive k-simplex in the filtration $\mathbb{F}$ of K. There exists a k-cycle c which is not a boundary in K^i, which contains σ and which 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, which is not a boundary, that contains this simplex and no other positive k-simplex. Since σ is positive, there exists a k-cycle, which is not a boundary, d in K^i containing σ. Let σ_{i_j}, $j = 1, \cdots, p$ be the positive k-simplices different from σ contained in d and let c_j be the not-boundary cycles containing them and not containing any other positive simplices. Then

$$c = d + c_1 + \cdots + c_p$$

is a k-cycle whose 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. □

Persistent homology basis and persistent pairs

The k-cycles associated to the positive k-simplices in Lemma 13.13 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 $c^{i_1}, \cdots, c^{i_p}$ be the cycles associated to the positive simplices $\sigma^{i_1}, \cdots \sigma^{i_p}$ whose homology classes form a basis of H_k^{j-1}. The boundary $d = \partial\sigma^j$ de 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} to obtain a basis of H_k^j.

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}, $i_k \neq l(j)$ in Z_k^j which is equivalent to the above decomposition of d.

Definition 13.14 *The simplices $\sigma^{l(j)}$ and σ^j are paired in a persistent pair $(\sigma^{l(j)}, \sigma^j)$. 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 $\mathbb{F}$ of a simplicial complex K of dimension d with m simplices.

Algorithm 11 Persistent pairs computation

Input: A filtration $\mathbb{F}$ of a d-dimensional simplicial complex K containing m simplices.
$L_0 = L_1 = \cdots = L_{d-1} = \emptyset$
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 $L_0, L_1, \cdots, L_{d-1}$;

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 sof dimension d and

$$\mathbb{F} = \{\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K\}$$

be a filtration of K such that for any $i = 0, \cdots m-1$,

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

Let $M = (m_{i,j})_{i,j=1,\cdots m}$ be the matrix with coefficient in $\mathbb{Z}/2\mathbb{Z}$ of the boundary operator defined by:

$$m_{i,j} = 1 \text{ if } \sigma^i \leq \sigma^j \quad \text{and} \quad m_{i,j} = 0 \text{ otherwise}$$

So if σ^j is a $(k+1)$-simplexe the j-th column of M represents the set of the k-dimensional faces of the boundary of σ^j. Since $\mathbb{F}$ is a 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:

$$(i = l(j)) \Leftrightarrow (m_{i,j} = 1 \text{ and } m_{i',j} = 0 \ \forall i' > i)$$

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 $\mathbb{F}$ of a d-dimensional simplicial complex K containing m simplices.
for $j = 0$ to m **do**
  while there exists $j' < j$ such that $l(j') == l(j)$ **do**
    $C_j = C_j + C_{j'} \ \ mod(2)$;
Output: Return the pairs $(l(j), j)$;

Proposition 13.15 *The previous algorithm computes the persistent pairs of the filtration* $\mathbb{F} = \{\emptyset = K^0 \subset K^1 \subset \cdots \subset K^m = K\}$.

Proof The result immediately follows from a sequence of elementary remarks.
Rem 1 : At each step of the algorithm, the column C_j represents a chain of the following form

$$\partial \left(\sigma^j + \sum_{i<j} \varepsilon_i \sigma^i \right) \text{ with } \varepsilon_i \in \{0, 1\}.$$

This is proven by an immediate induction.

Rem 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

$$\sigma^{l(j)} + \sum_{p<l(j)} \eta_p \sigma^p \quad \text{with} \quad \eta_p \in \{0,1\},$$

but according to the remarque 1 the column C_j also represent 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.

Rem 3 : If at the end of the algorithm the column C_j only contains zero terms, then σ^j is positive.

Indeed, according to remark 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.

Rem 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 remarks 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 so it was added to the persistent homology basis at time $l(j)$ and has not been paired before time j. As a consequence, $(\sigma^{l(j)}, \sigma^j)$ is a persistence pair. □

Remark 13.16 *It is useful to remark that the matrix M of the boundary operator is a sparse matrix. The complexes used in practice usually contains a large number of simplices (typically $m = O(10^6)$ or $m = O(10^7)$). It is thus mandatory not to represent M as an (m, m) standard matrix with m^2 entries!*

Remark 13.17 *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 $O(m \log m)$).*

13.5.3 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 $\mathbb{F}$. More generally, if the filtration $\mathbb{F}$ is indexed by a non decreasing sequence of real numbers (as in the case of a filtration associated to the sublevel sets of a function)

$$\mathbb{F} = \{\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \cdots \subset K^{\alpha_m} = K\} \text{ with } \alpha_0 \leq \alpha_1 \cdots \leq \alpha_m$$

a persistent pair of simplicices (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 only non decreasing, several pairs can be associated to the same point in the plane. The persistence diagrams 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 $\overline{\mathbb{R}} = \mathbb{R} \cup +\infty$.

******** Examples of persistence diagrams are given on the figures ******* ADD FIGURES/EXAMPLES HERE **********

Persistence diagrams can be compared using a matching distance called the *bottleneck distance*.

Definition 13.18 *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 $\|p - q\|_\infty = \max(|\ x_p - x_q\ |, |\ y_p - y_q\ |)$.

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 13.6).

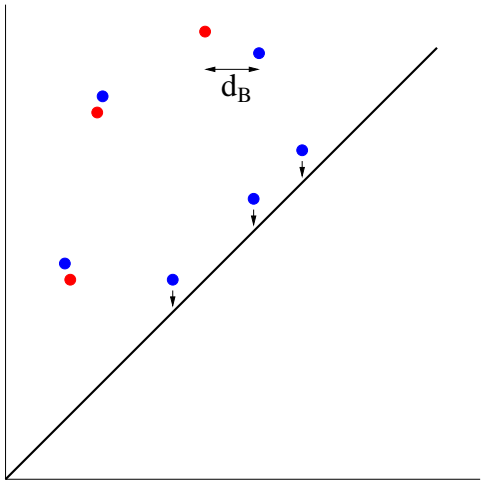

Figure 13.6: The bottleneck distance between two diagrams

13.5.4 Persistence modules, interleaving and stability

The proofs of the results presented in this section are beyond the scope of this book but the theoretical framework introduced in this section allows to efficiently prove the results of next Sections.

The notion of persistence can be extended to more general sequences of vector spaces in the following way.

Definition 13.19 *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 $(v_a^b : V_a \to V_b \mid a \leq b)$ 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.*

The persistence module $\mathbb{V}$ *is said to be q-tame if* $rk(v_a^b) < +\infty$ *whenever* $a < b$.

The sequence of homology groups of the subcomplexes of the filtration of K in the previous section together with the homomorphisms induced by the inclusion is a persistence module indexed over the set $I = \{0, 1, \cdots, m\}$. It can be shown that the persistence diagram of a filtration $\{\emptyset = K^{\alpha_0} \subset K^{\alpha_1} \subset \cdots \subset K^{\alpha_m} = K\}$ is completely determined by the rank of the homomorphisms $H_k(K^{\alpha_i}) \to H_k(K^{\alpha_j})$ for any $i < j$. This property extends to q-tame modules.

Theorem 13.20 *If a persistence module* $\mathbb{V}$ *is q-tame, then it has a well-defined persistence diagram* $\mathsf{dgm}(\mathbb{V}) \subset \overline{\mathbb{R}}^2$. *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 13.5.3.*

To avoid technical difficulties in the sequel of this section, we assume that all the considered persistence modules are indexed by $\mathbb{R}$.

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

$$(\phi_a : U_a \to V_{a+\varepsilon} \mid a \in \mathbb{R})$$

such that $v_{a+\varepsilon}^{b+\varepsilon} \circ \phi_a = \phi_b \circ u_a^b$ for all $a \leq b$. We write

$$\begin{aligned} \mathrm{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}) &= \{\text{homomorphisms } \mathbb{U} \to \mathbb{V} \text{ of degree } \varepsilon\}, \\ \mathrm{End}^\varepsilon(\mathbb{V}) &= \{\text{homomorphisms } \mathbb{V} \to \mathbb{V} \text{ of degree } \varepsilon\}. \end{aligned}$$

Composition is defined in the obvious way. For $\varepsilon \geq 0$, the most important degree-ε endomorphism is the shift map

$$1_{\mathbb{V}}^\varepsilon \in \mathrm{End}^\varepsilon(\mathbb{V}),$$

which is the collection of maps $(v_a^{a+\varepsilon})$ from the persistence structure on $\mathbb{V}$.

If Φ is a homomorphism $\mathbb{U} \to \mathbb{V}$ of any degree, then by definition $\Phi 1_{\mathbb{U}}^\varepsilon = 1_{\mathbb{V}}^\varepsilon \Phi$ for all $\varepsilon \geq 0$.

Definition 13.21 *Two persistence modules $\mathbb{U}, \mathbb{V}$ are said to be* ε**-interleaved** *if there are maps*

$$\Phi \in \mathrm{Hom}^{\varepsilon}(\mathbb{U}, \mathbb{V}), \qquad \Psi \in \mathrm{Hom}^{\varepsilon}(\mathbb{V}, \mathbb{U})$$

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 13.22 (Persistence Stability) *Let $\mathbb{U}, \mathbb{V}$ be two q-tame persistent modules that are ε-interleaved for some $\varepsilon \geq 0$. Denoting by $\mathsf{dgm}(\mathbb{U})$ and $\mathsf{dgm}(\mathbb{V})$ their persistence diagrams, we have*

$$d_B(\mathsf{dgm}(\mathbb{U}), \mathsf{dgm}(\mathbb{V})) \leq \varepsilon.$$

In the next sections, we apply the stability to different settings.

13.5.5 Persistence stability for functions

Let $f : X \to \mathbb{R}$ 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 $H_k(F_\alpha) \to H_k(F_\beta)$. So, the sublevel sets filtration of f induces a persistence module $\mathbb{F}_k$.

Proposition 13.23 *Let $f, g : X \to \mathbb{R}$ 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 $\alpha \in \mathbb{R}$, $F_\alpha \subseteq G_{\alpha+\varepsilon} \subseteq F_{\alpha+2\varepsilon} \subseteq G_{\alpha+3\varepsilon} \subseteq \cdots$ These inclusions induce homomorphisms $H_k(F_\alpha) \to H_k(G_{\alpha+\varepsilon})$ and $H_k(G_\alpha) \to H_k(F_{\alpha+\varepsilon})$ for all $\alpha \in \mathbb{R}$. The sets of these homomorphisms define an ε-interleaving between $\mathbb{F}_k$ and $\mathbb{G}_k$. □

In general, $\mathbb{F}_k$ and $\mathbb{G}_k$ are not q-tame. However, the following result provide sufficient conditions for f and g to be q-tame.

Proposition 13.24 *If X is homeomorphic to a finite simplicial complex and $f : X \to \mathbb{R}$ is continuous, then $\mathbb{F}_k$ is q-tame for any non negative integer k. In particular, $\mathsf{dgm}(\mathbb{F}_k)$ is well-defined.*

When $\mathbb{F}_k$ is q-tame for any non negative integer k, we say that the function f is q-tame. In the sequel, when there is no ambiguity, the notation $\mathsf{dgm}(f)$ denotes the persistence diagram of $\mathbb{F}_k$, for any k.

Applying the Persistence Stability Theorem 13.22, we immediately obtain the following stability result for functions.

Theorem 13.25 *Let X be a topological space homeomorphic to a finite simplicial complex and let $f, g : X \to \mathbb{R}$ be two continuous functions. Then*

$$d_B(\mathsf{dgm}(\mathbb{F}_k), \mathsf{dgm}(\mathbb{G}_k)) \leq \|f - g\|_\infty.$$

From a practical point of view, the above theorem provides a rigorous way to approximate persistence diagrams of continuous functions defined on a triangulated space. For example, assume that X is a triangulated surface in $\mathbb{R}^d$ where the diameter of each triangle is upper bounded by some $\delta > 0$ [2] and assume $f : X \to \mathbb{R}$ to be c-Lipschtiz 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 diagrams of the filtered complex induced by the values of f at the vertices of the triangulation (see Section **??**) is upper bounded by $c\delta$ (see Exercise 13.7).

13.5.6 Persistence stability for compact sets and complexes built on top of point clouds

Proposition 13.26 *Let $X \subset \mathbb{R}^d$ be a compact set. The distance function $d_X : \mathbb{R}^d \to \mathbb{R}$ is q-tame.*

Proof Given $0 \leq \alpha < \beta$ and a non negative integer k, we just need to prove that the homomorphism $H_k(X^\alpha) \to H_k(X^\beta)$ induced by the inclusion of the offsets, $X^\alpha \subset X^\beta$, has finite rank. Denoting $\varepsilon = (\beta - \alpha)/2 > 0$, since X is

[2] Notice that by doing subdivisions of the triangulation one can make δ arbitrarily small

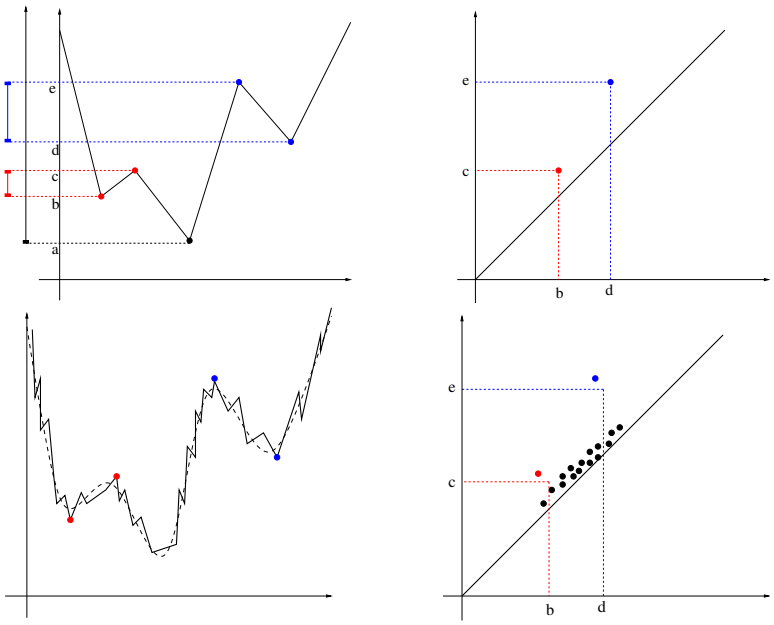

Figure 13.7: Comparing the persistence diagrams of two close functions defined on a segment

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

$$H_k(X^\alpha) \to H_k(P^{\alpha+\varepsilon}) \to H_k(X^\beta).$$

Now, it follows from the Nerve lemma that $P^{\alpha+\varepsilon}$ is a finite union of balls homotopy equivalent to $\check{\text{C}}\text{ech}(P, \alpha + \varepsilon)$ which is a finite simplicial complex. As a consequence $\dim H_k(P^{\alpha+\varepsilon}) < +\infty$ and $rk(H_k(X^\alpha) \to H_k(X^\beta)) \leq \dim H_k(P^{\alpha+\varepsilon}) < +\infty$. $\square$

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 13.27 *Let $X, Y \subset \mathbb{R}^d$ be compact. Then*

$$d_B(\mathsf{dgm}(d_X), \mathsf{dgm}(d_Y)) \leq d_H(X, Y).$$

In particular, if $P, Q \subset \mathbb{R}^d$ are finite point clouds, then for any non negative integer k,

$$d_B(\mathsf{dgm}(H_k(\check{\mathbb{C}}\text{ech}(P))), \mathsf{dgm}(H_k(\check{\mathbb{C}}\text{ech}(Q)))) \leq d_H(P, Q)$$

where $\check{\mathbb{C}}$ech(.) denotes the Čech filtration.

The second part of the corollary follows from the Nerve Lemma 4.4 and the paragraph before Proposition 13.9 showing that P^α and $\check{\mathrm{C}}\mathrm{ech}(P, \alpha)$ 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 $\mathcal{A}(P, \alpha)$ (see Exercise **??**) the same result also holds when $\check{\mathbb{C}}\mathrm{ech}(P)$ 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 13.27 less interesting. To overcome this issue one can consider the *Gromov-Hausdorff distance* defined in the following way.

Definition 13.28 *Let $X, Y \subset \mathbb{R}^d$ be two compact sets and let $\varepsilon \geq 0$. An ε-correspondence between P and Q 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

$$d_{GH}(X,Y) = \inf\{\varepsilon \geq 0 : \textit{there exists an } \varepsilon\textit{-correspondence between } X \textit{ and } Y\}.$$

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 13.27 has the following generalization.

Theorem 13.29 *Let $P, Q \subset \mathbb{R}^d$ be finite point clouds and let* Filt(.) *be any of the Čech , $Vietoris - Rips$ or alpha-shape filtered complexes. Then for any non negative integer k,*

$$d_B(\mathsf{dgm}(H_k(\mathrm{Filt}(P))), \mathsf{dgm}(H_k(\mathrm{Filt}(Q)))) \leq d_{GH}(P,Q).$$

This theorem can be extended to point clouds in non Euclidean metric spaces (except for the alpha-shape filtration which is non longer defined). In particular, to define the $Vietoris - Rips$ complex, one just need 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 13.29, 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}. But take care that the reverse is not true...

13.6 Bibliographical notes

A detailed introduction to algebraic topology and simplicial and singular homology can be found [94, 80].

Topological persistence has been independently introduced by different authors [63, 69, 101] and has know important developments during the last decade ****** put refs here *****.

The results of Section 13.4 are derived from [41] where complete proofs are given.

******** Put some ref about recent results on algos to compute persistence. **********

The stability of persistence diagrams has been initially proven by [51] for tame continuous functions defined on triagulable spaces. It was then extended and generalized by [30] and [37] to a more algebraic framework that appeared of fundamental importance in topological data analysis.

Theorem 13.29 and some of its extensions have found applications in shape classification [31] and in statistical analysis of data [43].

13.7 Exercises

Exercise 13.1 *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 4.2.

Exercise 13.2 (Difficult) *Let P be a finite set of points in $\mathbb{R}^2$. Prove that for any $\alpha \geq 0$*

$$\text{Rips}(P, \alpha) \subseteq \check{\text{C}}\text{ech}(P, \alpha\sqrt{\frac{d}{2(d+1)}}) \subseteq \text{Rips}(P, 2\alpha\sqrt{\frac{d}{2(d+1)}})$$

Hint: see [53], Theorem 2.5.

Exercise 13.3 *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 symetric difference* [3].

Exercise 13.4 *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.

[3] The symetric difference of two sets A and B is defined by $A\Delta B = (A \cup B) \setminus (A \cap B)$.

Exercise 13.5 *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 13.6 *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 13.13 is uniquely defined.

Exercise 13.7 *Let X be a (finitely) triangulated subset of $\mathbb{R}^d$ anf let $f : X \to \mathbb{R}$ 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,*

$$d_B(\mathsf{dgm}(K_f), \mathsf{dgm}(f)) \leq c\delta.$$

Bibliography

[1] Geometric range searching and its relative, booktitle=Advances in Discrete and Computational Geometry, author=Agarwal, Pankaj K and Erickson, Jeff, publisher= American Mathematical Society Press, series=Contemporary Mathematics, volume=223, number=4, pages=1–56, year=1999, editor=Chazelle, Bernard and Goodman, Jacob E. and Pollack Richard.

[2] Nir Ailon and Bernard Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 557–563. ACM, 2006.

[3] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, New York, USA, 2nd edition, 2008.

[4] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.

[5] Sunil Arya, David M Mount, Nathan S Netanyahu, Ruth Silverman, and Angela Y Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM (JACM)*, 45(6):891–923, 1998.

[6] 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.

[7] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16:78–96, 1987.

[8] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational*

Geometry, pages 201–290. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[9] M. Belkin, J. Sun, and Y. Wang. Constructing laplace operator from point clouds in r^d. In *Proc. of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1031–1040, 2009.

[10] M. Berger. *Géométrie (vols. 1-5)*. Fernand Nathan, Paris, 1977.

[11] J-D. Boissonnat and J. Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004.

[12] J-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential delaunay complexes. In *Proc. 26th Annual Symposium on Computational Geometry*, 2010.

[13] J-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. ACM Symp. on Computational Geometry*, pages 270–277, 2008.

[14] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Constructing stable Delaunay triangulations. Technical Report 8275, INRIA, 2013.

[15] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay stability via perturbations. Technical Report 8276, INRIA, 2013.

[16] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay triangulation of manifolds. Rapport de recherche RR-8389, INRIA, October 2013.

[17] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay triangulation of manifolds. Technical Report XXXX, INRIA, 2013.

[18] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. The Stability of Delaunay Triangulations. *Int. J. on Comp. Geom. (IJCGA)*, 23(4 & 5):303–333, 2014.

[19] Jean-Daniel Boissonnat and Arijit Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete and computational Geometry*, (November), 2013.

[20] Jean-Daniel Boissonnat, Frank Nielsen, and Richard Nock. Bregman Voronoi diagrams. *Discrete and Computational Geometry*, 44(2), 2010.

[21] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67:405–451, 2005.

[22] 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.

[23] Jean-Daniel Boissonnat and Mariette Yvinec. *Géométrie Algorithmique*. Ediscience international, Paris, 1995.

[24] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

[25] 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.

[26] S. S. Cairns. A simple triangulation method for smooth maniolds. *Bull. Amer. Math. Soc.*, 67:380–390, 1961.

[27] P. Cannarsa and C. Cinestrari. *Semiconcave Functions, Hamilton-Jacobi Equations, and Optimal Control*, volume 58. Brikhauser, 2004.

[28] R. Chaine. A geometric convection approach of 3d-reconstruction. In *1st Symposium on Geometry Processing*, pages 218–229, 2003.

[29] I. Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006.

[30] 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.

[31] 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.

[32] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 319–326, 2006.

[33] 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.

[34] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert. Shape smoothing using double offsets. In *Proc. ACM solid and Physical Modeling*, 2007.

[35] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert. Stability of Curvature Measures. *Computer Graphics Forum (proc. SGP 2009)*, pages 1485–1496, 2009.

[36] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.

[37] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. The structure and stability of persistence modules. arXiv:1207.3674 [math.AT], 2012.

[38] F. Chazal and A. Lieutier. The λ-medial axis. *Graphical Models*, 67(4):304–331, July 2005.

[39] 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.

[40] 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.

[41] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 232–241, 2008.

[42] 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.

[43] Frédéric Chazal, Marc Glisse, Catherine Labruère, and Bertrand Michel. Optimal rates of convergence for persistence diagrams in topological data analysis. *arXiv preprint arXiv:1305.6239*, 2013.

[44] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.

[45] 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.

[46] J. Cheeger, W. Müller, and R. Schrader. On the curvature of piecewise flat spaces. *Comm. Mth. Phys.*, 92:405–454, 1984.

[47] 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.

[48] F. H. Clarke. *Optimization and Nonsmooth Analysis*. Wiley-Interscience, New-York, 1983.

[49] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[50] Kenneth L Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17(4):830–847, 1988.

[51] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proc. 21st ACM Sympos. Comput. Geom.*, pages 263–271, 2005.

[52] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

[53] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7:339–358, 2007.

[54] Vin de Silva. A weak characterisation of the delaunay triangulation. *Geometriae Dedicata*, 135(1):39–64, 2008.

[55] David Dobkin and Richard J Lipton. Multidimensional searching problems. *SIAM Journal on Computing*, 5(2):181–186, 1976.

[56] R. A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.*, 6:343–367, 1991.

[57] 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.

[58] H. Edelsbrunner. Weighted alpha shapes. Technical Report UIUCDCS-R-92-1760, Dept. Comput. Sci., Univ. Illinois, Urbana, IL, 1992.

[59] H. Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13(1):415–440, 1995.

[60] 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.

[61] 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.

[62] 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.

[63] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.

[64] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. In *Proceedings of the 1992 workshop on Volume visualization*, pages 75–82. ACM, 1992.

[65] Herbert Edelsbrunner. Shape reconstruction with the Delaunay complex. In *LATIN'98: Theoretical Informatics*, volume 1380 of *Lecture Notes Comput. Sci.*, pages 119–132, 1998.

[66] Herbert Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, 2001.

[67] Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.

[68] D Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002.

[69] P. Frosini and C. Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9:596–603, 1999.

[70] J. H. Fu. Convergence of curvatures in secant approximations. *J. Differential Geometry*, 37:177–190, 1993.

[71] J. H. G. Fu. Tubular neighborhoods in Euclidean spaces. *Duke Mathematical Journal*, 52(4):1025–1046, 1985.

[72] W. Fulton. *Algebraic Topology: a First Course*. Springer-Verlag, 1995.

[73] S. Gallot, D. Hulin, and J. Lafontaine. *Riemannian Geometry*. Springer-Verlag, 1990.

[74] 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.

[75] Joachim Giesen and Uli Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete & Computational Geometry*, 32:245–267, 2004.

[76] K. Grove. Critical point theory for distance functions. In *Proc. of Symposia in Pure Mathematics*, volume 54, 1993. Part 3.

[77] Sariel Har-Peled. A replacement for voronoi diagrams of near linear size. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 94–103. IEEE, 2001.

[78] Sariel Har-Peled. *Geometric approximation algorithms*, volume 173. Amer Mathematical Society, 2011.

[79] Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory OF Computing*, 8:321–350, 2012.

[80] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[81] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, NY, 1976.

[82] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. ACM, 1998.

[83] William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[84] I. T. Jolliffe. *Principal component analysis*. Springer, 2002.

[85] X-Y. Li. Generating Well-Shaped d-dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.

[86] X-Y. Li. Generating Well-Shaped d-dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.

[87] 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.

[88] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as it medial axis. *Computer-Aided Design*, 36(11):1029–1046, September 2004.

[89] B. Mederos, N. Amenta, L. Velho, and L. H. de Figueiredo. Surface reconstruction for noisy point clouds. In *Proc. 3rd Sympos. on Geometry Processing*, pages 53–62, 2005.

[90] Stefan Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106(2):286–303, 1993.

[91] Robin A. Moser and Gábor Tardos. A constructive proof of the generalized Lovász lemma. *Journal of the ACM*, (2), 2010.

[92] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.

[93] J. R. Munkres. *Elementary differential topology*. Princeton University Press, 1966.

[94] J. R. Munkres. *Elements of algebraic topology*. Addison-Wesley, Redwood City, California, 1984.

[95] 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.

[96] Atsuyuki Okabe, Barry Boots, and Kokichi Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Chichester, UK, 1992.

[97] Richard B Paris and David Kaminski. *Asymptotics and Mellin-Barnes Integrals*, volume 85. Cambridge University Press, 2001.

[98] D. Pedoe. *Geometry, a comprehensive course*. Dover Publications, New York, 1970.

[99] 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.

[100] A. Petrunin. Semiconcave functions in alexandrov's geometry. In *Surveys in Differential Geometry: Metric and Comparison Geometry, volume XI*. Internationational Press of Boston, 2007.

[101] V. Robins. Toxard computing homology from finite approximations. *Topology Proceedings*, 24:503–532, 1999.

[102] 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.

[103] R. Seidel. The upper bound theorem for polytopes: an easy proof of its asymptotic version. *Comput. Geom. Theory Appl.*, 5:115–116, 1995.

[104] 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.

[105] R. Sibson. A vector identity for the Dirichlet tesselation. *Math. Proc. Camb. Phil. Soc.*, 87:151–155, 1980.

[106] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnet, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.

[107] C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.

[108] J. H. C. Whitehead. On C^1-complexes. *Ann. of Math.*, 41:809–824, 1940.

[109] H. Whitney. *Geometric integration theory*. Princeton University Press, 1957.

[110] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, 1994.

Index