



Piecewise linear reconstruction and meshing of submanifolds of Euclidean space

Arijit Ghosh

► To cite this version:

Arijit Ghosh. Piecewise linear reconstruction and meshing of submanifolds of Euclidean space. Computational Geometry [cs.CG]. Université de Nice Sophia Antipolis, 2012. English. NNT: . hal-01095861

HAL Id: hal-01095861

<https://inria.hal.science/hal-01095861>

Submitted on 18 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ DE NICE-SOPHIA ANTIPOLIS
ÉCOLE DOCTORALE STIC
SCIENCES ET TECHNOLOGIES DE L'INFORMATION
ET DE LA COMMUNICATION

THÈSE

pour l'obtention du grade de

Docteur en Sciences

de l'Université de Nice - Sophia Antipolis

Mention : INFORMATIQUE

présentée et soutenue par

Arijit GHOSH

Reconstruction et Maillages de Sous-Variétés

Thèse dirigée par Jean-Daniel BOISSONNAT
préparée au INRIA Sophia Antipolis, Team GÉOMÉTRICA
soutenue le 30 mai 2012

Jury :

<i>Rapporteurs :</i>	Dominique ATTALI	- CNRS GIPSA-LAB
	Tamal DEY	- The Ohio State University
	Gert VEGTER	- University of Groningen
<i>Directeur :</i>	Jean-Daniel BOISSONNAT	- INRIA Sophia Antipolis
<i>Président :</i>	Pierre PANSU	- University of Paris-Sud
<i>Examineurs :</i>	Jean-Marie MORVAN	- University Lyon 1
	Steve Y. OUDOT	- INRIA Saclay

UNIVERSITY OF NICE - SOPHIA ANTIPOLIS
DOCTORAL SCHOOL STIC
SCIENCES ET TECHNOLOGIES DE L'INFORMATION
ET DE LA COMMUNICATION

PHD THESIS

to obtain the title of

PhD of Science

of the University of Nice - Sophia Antipolis

Specialty : COMPUTER SCIENCE

Defended by

Arijit GHOSH

Piecewise linear reconstruction and meshing of submanifolds of Euclidean space

Thesis Advisor: Jean-Daniel BOISSONNAT

prepared at INRIA Sophia Antipolis, GEOMETRICA Team

defended on May 30, 2012

Jury :

<i>Reviewers :</i>	Dominique ATTALI	- CNRS GIPSA-LAB
	Tamal DEY	- The Ohio State University
	Gert VEGTER	- University of Groningen
<i>Advisor :</i>	Jean-Daniel BOISSONNAT	- INRIA Sophia Antipolis
<i>President :</i>	Pierre PANSU	- University of Paris-Sud
<i>Examinators :</i>	Jean-Marie MORVAN	- University Lyon 1
	Steve Y. OUDOT	- INRIA Saclay

Abstract

In this thesis we address some of the problems in the field of piecewise linear approximation of k -dimensional smooth submanifolds of Euclidean space $\mathbb{R}^d$. The main goal of this thesis was to develop algorithms that solve these problems with *theoretical guarantees*, i.e. the output being homeomorphic to the submanifold, and also have *intrinsic dimension sensitive complexity*, i.e. time and space complexity depend exponentially on the intrinsic dimension k of the submanifold and linearly on the ambient Euclidean dimension d .

The two standard questions in this field are the following:

- **Manifold reconstruction.** From a dense point sample $P \subset \mathbb{R}^d$, from an unknown smooth k -dimensional submanifold $\mathcal{M}$ of $\mathbb{R}^d$, we want to build a simplicial approximation $\hat{\mathcal{M}} \subset \mathbb{R}^d$ of $\mathcal{M}$ with theoretical guarantees.
- **Sampling and meshing manifolds.** For a given parameter ε and a k -dimensional smooth submanifold, known through some standard oracles, we want to generate a dense sample $P \subset \mathcal{M}$, according to the prescribed parameter ε , and build a simplicial approximation $\hat{\mathcal{M}}$ of $\mathcal{M}$ on top of the sample P with theoretical guarantees.

In this thesis we try to chip away at both these problems with the following results:

- For a dense point sample P of a smooth submanifold $\mathcal{M}$ of $\mathbb{R}^d$ we give sufficient conditions under which the *tangential Delaunay complex*, defined in [BF04, Flö03, Fre02], build using the point sample P is homeomorphic and a close geometric approximation of $\mathcal{M}$.
- We give an algorithm, whose complexity is intrinsic dimension sensitive, to reconstruct smooth k -dimensional submanifolds of $\mathbb{R}^d$ from a dense point sample P using tangential Delaunay complexes. We show, using the above result, that the output is homeomorphic and a close geometric approximation of $\mathcal{M}$. To the best of our knowledge, this is the first certified algorithm for manifold reconstruction whose complexity is intrinsic dimension sensitive.
- We give an algorithm to sample and mesh a k -dimensional smooth submanifold $\mathcal{M}$ of $\mathbb{R}^d$. According to the prescribed parameter ε , the algorithm generates a dense sample of $\mathcal{M}$ and a mesh with theoretical guarantees. The algorithm uses only simple numerical operations. We show that the size of the sample is $O(\varepsilon^{-k})$ and the asymptotic complexity of the algorithm is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (for fixed $\mathcal{M}$, d and k).
- We provide a counterexample to the result announced by Liebon and Letscher [LL00]. We show that density of the sample points on a manifold $\mathcal{M}$

alone cannot guarantee that the nerve of the intrinsic Voronoi diagram, i.e. the intrinsic Delaunay triangulation, is homeomorphic to the manifold $\mathcal{M}$.

- We introduce a parameterized notion of δ -generic point set for Delaunay triangulations. We show that Delaunay triangulation of a δ -generic point sample is (1) combinatorially stable under small perturbation of the underlying metric and vertex positions, and (2) simplices of Delaunay triangulation are *well shaped*.
- Using the stability results of Delaunay triangulations of δ -generic point set, we show that, for any sufficiently regular submanifold of Euclidean space, and appropriate ε and δ , any sample set which meets a localized δ -generic ε -dense sampling criteria, intrinsic Delaunay triangulation is equal to restricted Delaunay triangulation and tangential Delaunay triangulation, and intrinsic Delaunay triangulation is homeomorphic to the submanifold. We also give a refinement algorithm for generating intrinsic Delaunay triangulations of submanifolds.

Keywords. Delaunay complex, intrinsic Delaunay complex, manifold reconstruction, meshing, slivers, stability of Delaunay triangulation, Voronoi diagram, and weighted points.

Acknowledgments

First and foremost, I would like to thank Jean-Daniel for his guidance, support and encouragement during the whole period of my graduate studies. He gave me the freedom to work on problems I found interesting and always treated me as his equal rather than a student. He was always, and I hope still is, available to his students. I never had to fix an appointment to meet him. I could directly barge into his office and started talking, and he would always stop what he was doing and listen to me.

I want to thank the members of my thesis committee: Dominique Attali, Tamal Dey, Gert Vegter, Pierre Pansu, Jean-Marie Morvan, and Steve Y. Oudot. Their feedback, and insights have been very helpful.

I would like to thank Ramsay Dyer who has been my close collaborator and a very good friend. If Ramsay had not join Geometrica during my graduate studies then the second half of my thesis would have been completely different.

Steve other than being a member of my thesis committee was also a person who I collaborated with during my graduate studies. I learned a lot from him during our many interactions.

I would like to thank Mariette Yvunec for carefully reading the early versions of my papers. She not only helped us with the presentations but also gave important insights on the problems we have been working on.

It has been a great privilege to do my graduate studies in Geometrica. It is beyond doubt the best computational geometry group in the world. I would specially like to thank the permanent members, postdoctoral researchers and my fellow PhD students from the group: David Cohen- Steiner, Olivier Devillers, Monique Teillaud, Frederic Cazals, Pierre Alliez, Sylvain Pion, Julie Digne, Michael Hemmer, Samuel Hornus, Mikhail Bogdanov, Manuel Caroli, Ross Hemsley, Clement Maria, Pooran Memari, Quentin Merigot, Bertrand Pellenard, Nader Salman, Jane Tournois and Pedro Machado Manhães de Castro.

I also want to thank Geometrica secretaries, Caroline French and Nelly Bessega, for their administrative support throughout my studies. Without their help I would have long been deported from France.

I would specially like to thank Aurobrata Ghosh and Cheryl Northcott for being wonderful friends. They, together with Ramsay, helped to preserve my sanity in the final weeks during which period I was writing my thesis.

Finally, I want to end by thanking my family for their love and support.

Contents

I	Introduction and Background	1
1	Introduction	3
1.1	Manifold reconstruction	4
1.2	Triangulating manifold	5
1.3	Stability of Delaunay structures and intrinsic Delaunay triangulations . .	6
2	Backgrounds and notations	9
2.1	General notations	9
2.2	Sampling conditions	10
2.3	Simplexes	11
2.4	Weighted Delaunay triangulation	15
2.4.1	Properties of weighted Delaunay triangulation	17
2.5	Definitions and results from topology	20
II	Tangential Delaunay complexes	23
3	Manifold reconstruction	25
3.1	Introduction	25
3.2	Tangential complex and inconsistent configurations	27
3.3	Manifold reconstruction	31
3.3.1	Algorithm	31
3.4	Analysis of the algorithm	33
3.4.1	Properties of the tangential Delaunay complex	34
3.4.2	Properties of inconsistent configurations	36
3.4.3	Number of local neighbors	40
3.4.4	Correctness of the algorithm, and theoretical guarantees	41
3.4.5	Time and space complexity	43
3.5	Summary	45
4	Topological and geometric guarantees	47
4.1	Tangent space approximation	49
4.2	Piecewise-linear k -manifold	49
4.3	Homeomorphism	50
4.4	Pointwise approximation	67

4.5	Isotopy	67
5	Sampling and meshing of submanifolds	69
5.1	Introduction	69
5.2	Definitions and preliminaries	71
5.2.1	Revisiting tangential Delaunay, unweighted	71
5.3	Algorithm	73
5.3.1	Primitive operations	74
5.3.2	Computing the initial sample P_0	74
5.3.3	Good simplices and slivers	74
5.3.4	Picking region and good points	75
5.3.5	Refinement Algorithm	76
5.4	Analysis of the algorithm	77
5.4.1	Forbidden regions	78
5.4.2	Proof of termination	80
5.4.3	Combinatorial complexity analysis	86
5.5	Topological and geometric guarantees	93
5.6	Summary	94
III	Stability of Delaunay Triangulation	95
6	An obstruction to intrinsic Delaunay triangulations	97
6.1	Delaunay complex and Delaunay triangulation	97
6.2	A qualitative argument	98
6.3	An obstruction to intrinsic Delaunay triangulations	100
6.3.1	Sampling density alone is insufficient	100
6.3.2	Discussion	104
7	Stability of Delaunay triangulations	105
7.1	Introduction	105
7.2	Background	106
7.2.1	Sampling parameters and perturbations	107
7.2.2	Simplices	107
7.2.3	Complexes	111
7.3	Parameterized genericity	114
7.3.1	The Delaunay complex	114
7.3.2	Protection	115
7.3.3	Thickness from protection	120
7.4	Delaunay stability	122
7.4.1	Perturbations and circumcentres	123
7.4.2	Perturbations and protection	127
7.4.3	Perturbations and Delaunay stability	129
7.5	Summary	131

8	Constructing intrinsic Delaunay triangulations	133
8.1	Introduction	133
8.2	Background	134
8.2.1	Notations from Chapter 7	134
8.2.2	Simplex perturbation	134
8.2.3	Flakes	137
8.2.4	The Delaunay complex and protection	138
8.2.5	The Delaunay complex in other metrics	139
8.2.6	The Voronoi diagram	139
8.2.7	Background results for manifolds	140
8.3	Equating Delaunay structures	142
8.3.1	Delaunay structures on manifolds	142
8.3.2	Choice of local Euclidean metric	145
8.3.3	The protected tangential complex	147
8.4	Algorithm	150
8.4.1	Components of the algorithm	150
8.4.2	The refinement algorithm	154
8.5	Analysis of the algorithm	156
8.5.1	Termination of the algorithm	157
8.5.2	Output quality	164
8.6	Summary	167
9	Conclusion	169
A	Appendix for Chapter 2	173
A.1	Whitney's proof of Lemma 2.3.2	173
B	Appendix for Chapter 5	175
B.1	Proof of Lemmas 5.2.3 and 5.2.5	175
B.2	Proof of Lemma 5.4.1	176
B.2.1	Geodesic curves and balls	176
B.2.2	Injectivity radius and reach	177
B.2.3	Proof of Lemma 5.4.1	179
C	Appendix for Chapter 8	181
C.1	Forbidden volume calculation	181
	Bibliography	187

Part I

Introduction and Background

Chapter 1

Introduction

In this thesis, we will investigate the following problems

- **Manifold reconstruction**
- **Sampling and meshing manifolds**

One of the main goals of the thesis was to come up with algorithms whose complexity is *intrinsic dimension sensitive*, i.e. the complexity of the algorithm depends exponentially on the intrinsic dimension k of the manifold rather than the ambient dimension d . The state of affairs as of now is that the complexity of the above problems is exponentially in d , and whether the complexity can be made only polynomial in d (while still exponential in k) has been an open question.

Using the ideas from the papers [She05, BWY08], which can be summarized as *building locally and fitting globally*, the notion of tangential Delaunay complex introduced in [BF04, Flö03, Fre02], and the techniques of sliver removal by weighting the sample points and sliver refinement developed in [CDE⁺00a, Li03a], we were able to give new algorithms with intrinsic dimension sensitive complexities in:

- **Manifold reconstruction (Chapter 3).** We give an intrinsic dimension sensitive (complexity) and provably correct algorithm to reconstructs smooth submanifolds of $\mathbb{R}^d$ from a dense point sample.
- **Triangulating manifold (Chapter 5).** We give an algorithm to sample and mesh smooth submanifolds $\mathcal{M}$ of $\mathbb{R}^d$ according to the prescribed sampling parameter ε . If the intrinsic dimension of the manifold is k , then we show that the size of the sample is $O(\varepsilon^{-k})$ and the asymptotic time complexity of the algorithm is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (for fixed $\mathcal{M}$, d and k). We also show that the output mesh is homeomorphic and a close geometric approximation of $\mathcal{M}$.

To show that the output returned by our algorithms is homeomorphic to $\mathcal{M}$ we will use the following result:

- **Properties of Tangential Delaunay complex (Chapter 4).** We give sufficient conditions under which tangential Delaunay complex is homeomorphic and a close geometric approximation of the underlying manifold $\mathcal{M}$. We also show that under these conditions tangential Delaunay complex is isotopic to $\mathcal{M}$.

We also study the stability of Delaunay triangulations with the view towards better understanding intrinsic Delaunay triangulations on manifolds:

- **Obstruction to intrinsic Delaunay triangulation (Chapter 6).** We give a counterexample, to the announced result of Liebon and Letscher [LL00], showing that density of the sample points on a manifold alone cannot guarantee that the nerve of the intrinsic Voronoi diagram, i.e. the intrinsic Delaunay complex, is homeomorphic to the manifold.
- **Stability of Delaunay triangulations (Chapter 7).** We introduce a parameterized notion of δ -generic point set for Delaunay triangulations. We show that Delaunay triangulation of a δ -generic point sample is (1) combinatorially stable under small perturbation of the underlying metric and vertex positions, and (2) simplices of Delaunay triangulation are *well shaped*.
- **Constructing intrinsic Delaunay triangulations (Chapter 8).** We show that, for any sufficiently regular submanifold of Euclidean space, and appropriate ε and δ , any sample set which meets a localized δ -generic ε -dense sampling criteria, intrinsic Delaunay triangulation is equal to restricted Delaunay triangulation and tangential Delaunay triangulation, and intrinsic Delaunay triangulation is homeomorphic to the submanifold. We also give an algorithm for generating δ -generic point sets.

1.1 Manifold reconstruction

Manifold reconstruction consists of computing a piecewise linear approximation of an unknown manifold $\mathcal{M} \subset \mathbb{R}^d$ from a finite sample of unorganized points P lying on $\mathcal{M}$ or close to $\mathcal{M}$. The special case of reconstruction of two-dimensional surfaces embedded in $\mathbb{R}^3$ have been studied extensively in the fields of Computational Geometry, Computer Graphics and Computer Vision. Refer to [CG06, Dey06] for recent results. The output of those methods is a triangulated surface that approximates $\mathcal{M}$.

The general problem in higher dimensions is also of great practical interest in data analysis and machine learning. Well-known global techniques have been developed for approximating linear manifolds, like principal component analysis (PCA) or multi-dimensional scaling (MDS). When the manifold is nonlinear, more local techniques have attracted more attention. Among the prominent algorithms are Isomap [TdSL00], LLE [RS00], Laplacian eigenmaps [BN02], Hessian eigenmaps [DG03], diffusion maps [LL06, NLCK05], principal manifolds [ZZ04]. Most of those methods reduce to computing an eigendecomposition of some connection matrix. In all cases, the output is a mapping of the original data points into $\mathbb{R}^k$ where k is the estimated intrinsic dimension of $\mathcal{M}$. These methods come with no or very limited guarantees. To be able to better approximate the sampled manifold, another approach is to extend the work on surface reconstruction and to construct a piecewise linear approximation of $\mathcal{M}$ from the sample in such a way that, under appropriate sampling conditions, the quality of the approximation can be guaranteed. All attempts to extend those geometric approaches from surface reconstruction to more general manifolds have led to algorithms whose complexities depend exponentially on d [BGO09, CL08, CDR05a, NSW08a].

In Chapter 3, we extend the geometric techniques developed in small dimensions and propose an algorithm that can reconstruct smooth k -manifolds of arbitrary topology

while avoiding the computation of data structures in the ambient space. The complexity of the algorithm is linear in d , quadratic in the size n of the sample, and exponential in k . Our work builds on [BGO09] and [CDR05a] but dramatically reduces the dependence on d . To the best of our knowledge, this is the first certified algorithm, with theoretical guarantees (proved using results from Chapter 4), for manifold reconstruction whose complexity depends only linearly on the ambient dimension d .

We based our approach on two key ideas: the notion of *tangential Delaunay complex* introduced in [BF04, Flö03, Fre02], and the technique of *sliver removal* by weighting the sample points [CDE⁺00a]. The tangential complex is obtained by gluing stars, extracted from local (Delaunay) triangulations, at each sample point. The tangential complex is a subcomplex of the d -dimensional Delaunay triangulation of the sample points but the stars can be computed using mostly operations in the k -dimensional tangent spaces at the sample points. This is the key reason why the complexity of the algorithm depends exponentially on k rather than d . However, due to the presence of so-called *inconsistencies*, the local triangulations may not form a triangulated manifold. The idea of removing inconsistencies among the stars that have been computed independently has already been used for maintaining dynamic meshes [She05] and generating anisotropic meshes [BWY08]. Our approach is close in spirit to the one in [BWY08]. We also show that, under appropriate sample conditions, we can remove inconsistencies by weighting the sample points. We can then prove that the approximation returned by our algorithm is ambient isotopic to $\mathcal{M}$, and a close geometric approximation of $\mathcal{M}$.

1.2 Triangulating manifold

In Chapter 5, we study the algorithmic problem of sampling and meshing a k -manifold $\mathcal{M}$ of positive reach embedded in $\mathbb{R}^d$. Manifolds of positive reach have been introduced by Federer [Fed59, Fed69] and include in particular C^2 -manifolds. By mesh, we mean simplicial approximation of $\mathcal{M}$. We are especially interested in the case where the dimension k of $\mathcal{M}$ is much smaller than d , and intend to design an algorithm whose complexity depends on k rather than on d . Applications can be found in scientific computing for solving partial differential equations where the domain of interest has the structure of a manifold, in dynamical systems for computing the topology of space attractors, and in statistics and machine learning to approximate statistical manifolds.

The problem of triangulating manifolds has a long history in the mathematical literature. In differential topology, seminal contributions are due to Whitney [Whi57a], Cairns [Cai61], Munkres [Mun66], Whitehead [Whi40] to name a few. Although these papers are not of an algorithmic nature, they introduce and study several interesting concepts that have been extensively used in Computational Geometry recently such as Voronoi diagrams restricted to a manifold, ε -sample of a manifold, fat (or thick) triangulations. However, these papers do not discuss the geometric quality of the approximation nor the size of the sample. The optimal sampling and approximation of convex bodies is also a long standing problem in convex optimization with major contributions by Gruber [Gru93, Gru04] and Dudley [Dud74]. Recently, Clarkson [Cla06] extended this line of work to non-convex smooth manifolds of arbitrary dimensions. However, his algorithm follows an intrinsic point of view which makes it difficult to use in practice since it requires to compute geodesic distances on the manifold which may be quite

complicated in practice [PC05]. Other, more practical algorithms for approximating convex bodies, including the well-known sandwich algorithm, have been analyzed by Kamenev [Kam08]. We are not aware of similar studies for non convex manifolds except for the case of surfaces embedded in $\mathbb{R}^3$ which has been extensively studied in the Computational Geometry literature. See [CG06] for a recent survey. As in the case for surface reconstruction algorithms mentioned earlier, these methods start by computing some subdivision of the embedding space (such as a grid or a triangulation of the sample points) and their direct extension to higher dimensions would face an exponential dependence on d . A step in this direction is the extension of the celebrated Marching Cube algorithm to manifolds of higher dimensions [BWC02, Min03]. Continuation methods [Hen02] do not use any subdivision of the ambient space and are close in spirit to our approach but they lack theoretical guarantees.

We follow the extrinsic approach but show that we can avoid using any d -dimensional data structure (except in the initialization step). The algorithm starts with a sufficiently dense sample of $\mathcal{M}$ and then refines the sample and builds a mesh that approximates $\mathcal{M}$ so as to satisfy a prescribed sampling rate ε . The size of the initial sample does not depend on ε but only on $\mathcal{M}$.

In the same spirit as the manifold reconstruction algorithm in Chapter 3, we build the mesh by glueing local stars, extracted from local (Delaunay) triangulations, at each sample point. These stars can be computed using mostly operations in the tangent spaces at the sample points. As discussed in the case of manifold reconstruction, these stars do not glue coherently to output a piecewise linear manifold due to the presence of so called inconsistencies. The crucial observation is that by refining the sample, we can ensure that all the stars become coherent leading to a k -dimensional mesh $\hat{\mathcal{M}}$. For ε small enough, we show that the size of the sample is $O(\varepsilon^{-k})$ and that $\hat{\mathcal{M}}$ (output) is a good approximation of $\mathcal{M}$. Our bound on the Hausdorff distance between $\hat{\mathcal{M}}$ and $\mathcal{M}$ matches the lower bound of Clarkson [Cla06] (up to a multiplicative constant that depends on $\mathcal{M}$).

To refine the mesh according to a sampling parameter ε , we need an *oracle* to query the manifold and to compute new points on $\mathcal{M}$. This is a critical issue with respect to practical efficiency. In our algorithm, we only need to compute a point in the (0-dimensional) intersection of $\mathcal{M}$ with a $(d - k)$ -flat. The asymptotic complexity of the algorithm is $O(\varepsilon^{-k^2-k})$ for fixed k , d , and $\mathcal{M}$. Hence, while our approach is extrinsic, the ambient dimension appears only in the constant hidden in the big- O .

1.3 Stability of Delaunay structures and intrinsic Delaunay triangulations

Delaunay triangulations, introduced by B. Delaunay in 1934, have been extensively studied in Computational Geometry and have found applications in many domains of science. Rather surprisingly though, the stability of those structures has not been studied in a systematic way. Our motivation comes from the recent attempts to extend Delaunay triangulations beyond Euclidean spaces. A first attempt in that direction is the generation of anisotropic meshes where a metric tensor field is given that varies over a domain of $\mathbb{R}^d$ we want to mesh. A related (and more general) question is to define intrinsic Delaunay triangulations on a Riemannian manifold [LL00]. We might expect that, when the density of points is dense enough, all these Delaunay-like structures are similar. In fact this is the

type of result that can be found in [LS03] and in [LL00]. However, the result of Labelle and Shewchuk [LS03] is limited to the 2-dimensional case and the paper of Leibon and Letscher [LL00] contains a flaw. We give a counterexample to the claimed results of Leibon and Letscher [LL00] in Chapter 6. These papers in fact miss an important fact that the general (non-Euclidean) Delaunay triangulations are not well behaved when the points are roughly cospherical, even if they are not exactly cospherical, when the dimension is greater than two.

In Chapter 7, we introduce a parametrized notion of genericity for Delaunay triangulations which, in particular, implies that the Delaunay simplices of δ -generic point sets are well shaped. Using the above framework, in Chapter 8, we study the stability of Delaunay triangulations under perturbations of the metric and of the vertex positions. We prove that for a regular submanifold of Euclidean space, and appropriate ϵ and δ , any sample set which meets a localized δ -generic ϵ -dense sampling criteria yields an intrinsic Delaunay triangulation homeomorphic to the manifold. Finally in Chapter 8, we give a provably correct algorithm to produce intrinsic Delaunay triangulations of submanifolds.

Chapter 2

Backgrounds and notations

This chapter gives the notations, basic definitions, and variants and generalization of the standard results from [BG11, ELS⁺00, Li00, Li03a, LT01, Whi57a], which will be used in the thesis.

2.1 General notations

In this thesis, unless stated otherwise, $\mathcal{M}$ denotes a compact smooth k -dimensional submanifold of $\mathbb{R}^d$ without boundary, and P a finite set of points on $\mathcal{M}$. The tangent space at $x \in \mathcal{M}$ is denoted by $T_x\mathcal{M}$ and the normal space by $N_x\mathcal{M}$.

Within the context of the standard d -dimensional Euclidean space $\mathbb{R}^d$, we denote the the standard Euclidean norm by $\|\cdot\|$, i.e. for all $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ in $\mathbb{R}^d$,

$$\|x - y\| = \sqrt{\sum_{i=1}^d (x_i - y_i)^2}.$$

For a point p in $\mathbb{R}^d$ and $r \geq 0$, $B(p, r)$ ($\bar{B}(p, r)$) denotes the d -dimensional open (topological closure of $B(p, r)$) ball centered at p of radius r with respect to the standard Euclidean norm, i.e $B(p, r) = \{x \in \mathbb{R}^d \mid \|p - x\| < r\}$ and $\bar{B}(p, r) = \{x \in \mathbb{R}^d \mid \|p - x\| \leq r\}$.

Generally, we denote the topological closure of a set X by $\bar{X}$, the interior by $\text{int}(X)$, and the boundary by ∂X . The convex hull is denoted $\text{conv}(X)$, and the affine hull is $\text{aff}(X)$.

For a given point p in the point set P , $\text{nn}(p)$ denotes the distance of p to its nearest neighbor in $P \setminus \{p\}$, i.e.

$$\text{nn}(p) = \min_{x \in P, x \neq p} \|x - p\|.$$

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

$$\angle(U, V) = \max_{u \in U} \min_{v \in V} \angle(u, v) = \arccos \inf_{u \in U} \sup_{v \in V} \frac{u^\top v}{\|u\| \|v\|},$$

where u and v are vectors in U and V respectively. This is the largest principal angle between U and V . The angle between affine subspaces is defined as the angle between the corresponding parallel vector subspaces.

The following lemma directly follows from the definition of the angle between affine spaces.

Lemma 2.1.1 *Let U and V be affine spaces of $\mathbb{R}^d$ with $\dim(U) \leq \dim(V)$, and let $U^\perp$ and $V^\perp$ be affine spaces of $\mathbb{R}^d$ with $\dim(U^\perp) = d - \dim(U)$ and $\dim(V^\perp) = d - \dim(V)$.*

1. *If $U^\perp$ and $V^\perp$ are the orthogonal complements of U and V in $\mathbb{R}^d$, then $\angle(U, V) = \angle(V^\perp, U^\perp)$.*
2. *If $\dim(U) = \dim(V)$ then $\angle(U, V) = \angle(V, U)$.*

Proof Without loss of generality we assume that the affine spaces U , V , $U^\perp$ and $V^\perp$ are vector subspaces of $\mathbb{R}^d$, i.e. they all pass through the origin.

1. Suppose $\angle(U, V) = \alpha$. Let $v_* \in V^\perp$ be a unit vector. There are unit vectors $u \in U$, and $u_* \in U^\perp$ such that $v_* = au + bu_*$. We will show that $\angle(v_*, u_*) \leq \alpha$. First note that this angle is complementary to $\angle(v_*, u)$, i.e.,

$$\angle(v_*, u_*) = \frac{\pi}{2} - \angle(v_*, u). \quad (2.1)$$

There is a unit vector $v \in V$ such that $\angle(u, v) = \alpha_0 \leq \alpha$. Viewing angles between unit vectors as distances on the unit sphere, we exploit the triangle inequality: $\angle(v_*, v) \leq \angle(v_*, u) + \angle(u, v)$, from whence

$$\angle(v_*, u) \geq \frac{\pi}{2} - \alpha_0.$$

Using this expression in (2.1), we find

$$\angle(v_*, u_*) \leq \alpha_0 \leq \alpha,$$

which implies, since v_* was chosen arbitrarily, that $\angle(V^\perp, U^\perp) \leq \angle(U, V)$.

Since $\dim V^\perp \leq \dim U^\perp$, and the orthogonal complement is a symmetric relation on subspaces, the same argument yields the reverse inequality.

2. Let $\angle(U, V) = \alpha$, and let $P : U \rightarrow V$ denotes the projection map of the vector space U on V .

Case a. $\alpha \neq \pi/2$. Since $\alpha \neq \pi/2$ and $\dim(U) = \dim(V)$, the projection map P is an isomorphism between vector spaces U and V . Therefore, for any unit vector $v \in V$ there exist a vector $u \in U$ such that $P(u) = v$. From the definition of angle between affine space and the linear map P , we have $\angle(u, v) (= \angle(v, u))$. This implies, $\angle(V, U) \leq \angle(U, V) = \alpha$. Using the same arguments we can show that $\angle(U, V) \leq \angle(V, U)$ hence $\angle(U, V) = \angle(V, U)$.

Case b. $\alpha = \pi/2$. We have $\angle(V, U) = \pi/2$, otherwise using the same arguments as in Case 1 we can show that $\alpha = \angle(U, V) \leq \angle(V, U) < \pi/2$. $\square$

2.2 Sampling conditions

Medial axis and local feature size. The *medial axis* of $\mathcal{M}$ is the closure of the set of points of $\mathbb{R}^d$ that have more than one nearest neighbor on $\mathcal{M}$. The *local feature size* of $x \in \mathcal{M}$, $\text{lfs}(x)$, is the distance of x to the medial axis of $\mathcal{M}$ [AB99]. As is well known and can be easily proved, lfs is *Lipschitz continuous*, i.e., $\text{lfs}(x) \leq \text{lfs}(y) + \|x - y\|$.

The infimum of lfs over $\mathcal{M}$ is called the *reach* of $\mathcal{M}$. In this thesis, we assume that the reach of $\mathcal{M}$ is (strictly) positive.

Sampling definitions. A point sample P is said to be a (ε, δ) -lfs sample (where $0 < \delta < \varepsilon < 1$) if

- (C1) for any point $x \in \mathcal{M}$, there exists a point $p \in P$ such that $\|x - p\| \leq \varepsilon \text{lfs}(x)$, and
- (C2) for any two distinct points $p, q \in P$, $\|p - q\| \geq \delta \text{lfs}(p)$.

The parameter ε is called the *sampling rate*, δ the *sparsity*, and ε/δ the *sampling ratio* of the sample P . If P satisfy only condition (C1) or (C2) then P is called ε -lfs sample or δ -lfs sparse sample of P respectively.

The following lemma, proved in [GW04a], states basic properties of manifold samples. As before, we write $\text{nn}(p)$ for the distance between $p \in P$ and its nearest neighbor in $P \setminus \{p\}$.

Lemma 2.2.1 *Given a (ε, δ) -lfs sample P of $\mathcal{M}$, we have*

1. $\delta \text{lfs}(p) \leq \text{nn}(p) \leq \frac{2\varepsilon}{1-\varepsilon} \text{lfs}(p)$.
2. For any two points $p, q \in \mathcal{M}$ such that $\|p - q\| = t \text{lfs}(p)$, $0 < t < 1$, $\sin \angle(pq, T_p) \leq t/2$.
3. Let p be a point in $\mathcal{M}$. Let x be a point in T_p such that $\|p - x\| \leq t \text{lfs}(p)$ for some $0 < t \leq 1/4$. Let x' be the point on $\mathcal{M}$ closest to x . Then $\|x - x'\| \leq 2t^2 \text{lfs}(p)$.

A point sample P is said to be a (ε, δ) -rch sample of $\mathcal{M}$, if P satisfies conditions (C1) and (C2) are satisfied when we replace lfs by rch in those conditons. Similarly, we can define ε -rch sample and δ -rch sparse sample of $\mathcal{M}$.

The Lemma 2.2.1 (2) and (3) holds exactly if we replace $\text{lfs}(p)$ with $\text{rch}(\mathcal{M})$.

- Lemma 2.2.2**
1. For any point $q \in \mathcal{M}$ such that $\|p - q\| = \varepsilon \text{rch}(\mathcal{M})$ for some $0 < \varepsilon < 1$, $\sin \angle(pq, T_p \mathcal{M}) \leq \varepsilon/2$.
 2. Let q be a point in $T_p \mathcal{M}$ such that $\|p - q\| = \varepsilon \text{rch}(\mathcal{M})$ for some $0 < \varepsilon \leq 1/4$. Let q' be the point on $\mathcal{M}$ closest to q . Then $\|q - q'\| \leq 2\varepsilon \|p - q\|$.

2.3 Simplicies

A j -dimensional simplex (or j -simplex for short) σ is the convex hull of $j + 1$ affinely independent points $p_0, \dots, p_j$. We write $\sigma = [p_0, \dots, p_j]$ and, for convenience, we may confound a simplex and the set of its vertices. We write c_σ for the circumcenter of σ (i.e. the center of its minimum enclosing d -ball), $\text{aff}(\sigma)$ for the j -dimensional affine hull of σ , N_σ for the $(d - j)$ -dimensional affine space normal to $\text{aff}(\sigma)$ and passing through c_σ (which lies in $\text{aff}(\sigma)$).

For any j -simplex σ , we denote by R_σ the circumradius of σ (i.e. the radius of its minimum enclosing d -ball), by L_σ (or Δ_σ) the length of its shortest (or longest) edge, by $\rho_\sigma = R_\sigma/L_\sigma$ the *radius-edge ratio* of σ , and by $\text{vol}(\sigma)$ the j -dimensional volume of σ . For a vertex p of σ , we write $\sigma_p = \sigma \setminus \{p\}$ for the $(j - 1)$ -dimensional face of σ opposite to p , and $D_\sigma(p)$ for the distance from p to the affine hull $\text{aff}(\sigma_p)$ of σ_p . We will call $D_p(\sigma)$ the *altitude* of p in σ . In addition, we define the *fatness* of σ as

$$\Theta_\sigma = \begin{cases} 1 & \text{if } j = 0 \\ \text{vol}(\sigma)/\Delta_\sigma^j & \text{if } j > 0 \end{cases} \quad (2.2)$$

Part (1) and (2) of the following lemma follows directly from the definition of fatness, and part (3) is a generalization of Torus Lemma from [ELS⁺00].

Lemma 2.3.1 (Properties of simplices) *Let $\sigma = [p_0, \dots, p_j]$ be a j -simplex and p be a vertex of σ .*

1. $\Theta_\sigma \leq \frac{1}{j!}$
2. $j! \Theta_\sigma \leq \frac{D_\sigma(p)}{\Delta_\sigma} \leq j 2^{j-1} \rho_\sigma^{j-1} \times \frac{\Theta_\sigma}{\Theta_{\sigma_p}}$
3. *The distance of p from the $(j-1)$ -sphere $\partial B(c_{\sigma_p}, R_{\sigma_p}) \cap \text{aff}(\sigma_p)$ is less than $b(\sigma) D_\sigma(p)$, where*

$$b(\rho_\sigma) = 1 + \frac{1}{1 - \sqrt{1 - \frac{1}{4\rho_\sigma^2}}}.$$

Proof 1. Without loss of generality we assume that $\sigma = [p_0, \dots, p_j]$ is embedded in $\mathbb{R}^j$. From the definition of fatness we have

$$\Delta_\sigma^j \Theta_\sigma = \text{vol}(\sigma) = \frac{|\det(p_1 - p_0 \dots p_j - p_0)|}{j!} \leq \frac{\Delta_\sigma^j}{j!}.$$

2. Using the bound from Lemma 2.3.1 (1) and the definition of fatness, we get

$$\begin{aligned} D_\sigma(p) &= \frac{j \text{vol}(\sigma)}{\text{vol}(\sigma_p)} \\ &\geq \frac{j \Theta_\sigma \Delta_\sigma^j}{\Delta_{\sigma_p}^{j-1} / (j-1)!} \\ &\geq j! \Theta_\sigma \Delta_\sigma. \end{aligned}$$

We deduce, using $R_\sigma / \rho_\sigma = L_\sigma \leq \Delta_\sigma \leq 2R_\sigma$,

$$\begin{aligned} \frac{D_\sigma(p)}{\Delta_\sigma} &= \frac{j \text{vol}(\sigma)}{\Delta_\sigma \text{vol}(\sigma_p)} = j \frac{\Theta_\sigma \Delta_\sigma^{j-1}}{\Theta_{\sigma_p}^{j-1} \Delta_{\sigma_p}^{j-1}} \\ &\leq j \frac{\Theta_\sigma \Delta_\sigma^{j-1}}{\Theta_{\sigma_p} L_{\sigma_p}^{j-1}} \\ &\leq j 2^{j-1} \rho_\sigma^{j-1} \times \frac{\Theta_\sigma}{\Theta_{\sigma_p}}. \end{aligned}$$

3. Let p^* be the point closest to p on $\partial B(c_{\sigma_p}, R_{\sigma_p}) \cap \text{aff}(\sigma_p)$ and p' be the point closest to p on $\text{aff}(\sigma_p)$. We denote by H the distance of c_σ from $\text{aff}(\sigma_p)$, $Q = \|p' - p^*\|$, and by t the angle between $q c_{\sigma_p}$ and $q c_\sigma$, where q is a vertex of σ_p (see Figure 2.1). Then $D_\sigma(p) = \|p - p'\|$ and $R_{\sigma_p} = R_\sigma \cos t$, which implies that

$$\cos t = \frac{R_{\sigma_p}}{R_\sigma} \geq \frac{L_{\sigma_p}}{2R_\sigma} \geq \frac{L_\sigma}{2R_\sigma} = \frac{1}{2\rho_\sigma}.$$

We also have $H = R_\sigma \sin t = R_{\sigma_p} \tan t$. Note that the points $c_\sigma, c_{\sigma_p}, p, p^*$ and p' lie on a 2-dimensional affine space. We have to consider the following cases:




Figure 2.1: Refer to the proof of Lemma 2.3.1 (3), Case (a).

(a) $p' \in B(c_\sigma, R_\sigma)$, and c_σ and p lie on opposite sides of $\text{aff}(c_{\sigma_p} p^*)$. We have

$$\frac{D_\sigma(p)}{Q} \geq \frac{R_\sigma - H}{R_{\sigma_p}} \geq \frac{R_\sigma - H}{R_\sigma} \geq 1 - \sin t,$$

see Figure 2.1. The distance from p to p^* is less than

$$D_\sigma(p) + Q \leq \left(1 + \frac{1}{1 - \sin t}\right) D_\sigma(p).$$




Figure 2.2: Refer to the proof of Lemma 2.3.1 (3), Case (b).

(b) $p' \in B(c_\sigma, R_\sigma)$, and c_σ and p lie on the same side of $\text{aff}(c_{\sigma_p} p^*)$. We have

$$\frac{D_\sigma(p)}{Q} \geq \frac{R_\sigma + H}{R_{\sigma_p}} \geq \frac{R_\sigma}{R_{\sigma_p}} \geq 1,$$

see Figure 2.2. The distance from p to p^* is less than

$$D_\sigma(p) + Q \leq 2D_\sigma(p).$$




Figure 2.3: Refer to the proof of Lemma 2.3.1 (3), Case (c).

(c) $p' \notin B(c_\sigma, R_\sigma)$. From Figure 2.3 we can see that p should lie on the right hand side of $\text{aff}(p^* \bar{p})$. From Intersecting Secants Theorem, we have

$$Q \times (2R_{\sigma_p} + Q) = \|p' - p\| \times \|p' - p''\| \quad (2.3)$$

Using Eq. (2.3) the facts that $\|p' - p''\| \leq 2H$, and $H = R_{\sigma_p} \tan t$, we have

$$Q \leq \frac{Q(2R_{\sigma_p} + Q)}{2R_{\sigma_p}} = \frac{\|p' - p\| \|p' - p''\|}{2R_{\sigma_p}} \leq D_\sigma(p) \tan t$$

The distance from p to p^* is less than

$$D_\sigma(p) + Q \leq (1 + \tan t) D_\sigma(p).$$

The lemma follows by observing that

$$\begin{aligned} 1 + \frac{1}{1 - \sin t} &\geq 1 + \frac{\sin t (1 + \sin t)}{\cos^2 t} = 1 + \frac{\tan t (1 + \sin t)}{\cos t} \\ &\geq 1 + \tan t \end{aligned}$$

and

$$\sin^2 t = 1 - \cos^2 t \leq 1 - 1/4\rho_\sigma^2$$

□

The following lemma is due to Whitney [Whi57a]. The proof is included in Appendix A for completeness.

Lemma 2.3.2 (Whitney's angle bound) *Let $\sigma = [p_0, \dots, p_j]$ be a j -dimensional simplex and let H be an affine flat such that σ is contained in the offset of H by η (i.e. any point of σ is at distance at most η from H). If u is a unit vector in $\text{aff}(\sigma)$, then there exists a unit vector u_H in H such that*

$$\sin \angle(u, u_H) \leq \frac{2\eta}{(j-1)! \Theta_\sigma L_\sigma}.$$

We deduce from the above lemma the following important corollary that bounds the angle between a simplex and the tangent space at a vertex of the simplex. See also Lemma 1 in [Fu93] and Lemma 16 in [CDR05a].

Corollary 2.3.3 (Tangent approximation) *Let σ be a j -simplex, $j \leq k$, with vertices on $\mathcal{M}$, and let p be vertex of σ . Assuming that $\Delta_\sigma < \text{lfs}(p)$, we have*

$$\sin \angle(\text{aff}(\sigma), T_p \mathcal{M}) \leq \frac{\Delta_\sigma^2}{\Theta_\sigma L_\sigma \text{lfs}(p)} \leq \frac{2\rho_\sigma \Delta_\sigma}{\Theta_\sigma \text{lfs}(p)}.$$

Proof It suffices to take $H = T_p \mathcal{M}$ and to use $\eta = \Delta_\sigma^2/2 \text{lfs}(p)$ (from Lemma 2.2.1 (2)) and $R_\sigma/\rho_\sigma = L_\sigma \leq \Delta_\sigma \leq 2R_\sigma$. Hence

$$\sin \angle(\text{aff}(\sigma), T_p \mathcal{M}) \leq \frac{2\eta}{(j-1)! \Theta_\sigma L_\sigma} \leq \frac{2\eta}{\Theta_\sigma L_\sigma} \leq \frac{\Delta_\sigma^2}{\Theta_\sigma L_\sigma \text{lfs}(p)} \leq \frac{2\rho_\sigma \Delta_\sigma}{\Theta_\sigma \text{lfs}(p)}$$

□

The property of a simplex having a *good (or bad) radius-edge ratio* is defined in terms of a parameter ρ_0 . The value of ρ_0 will be fixed in the respective chapters.

Definition 2.3.4 (Good/bad radius-edge ratio) *Given a positive parameter ρ_0 , a simplex σ is said to have good (bad) radius-edge ratio if $\rho_\sigma \leq \rho_0$ ($\rho_\sigma > \rho_0$).*

A *sliver* is a special type of flat simplex. The property of being a sliver is defined in terms of a parameter Θ_0 . The value to Θ_0 , to be fixed later.

The following definition is a variant of a definition given in [Li03a].

Definition 2.3.5 (Θ_0 -fat simplices and Θ_0 -slivers) *Given a positive parameter Θ_0 , a simplex σ is said to be a Θ_0 -fat if for all subsimplex σ' of σ , we have $\Theta_{\sigma'} \geq \Theta_0^k$ where the dimension of σ' is k .*

A simplex σ of dimension 2 is a Θ_0 -sliver if all the proper subsimplices of σ is Θ_0 -fat, and $\Theta(\sigma) < \Theta_0^k$ where k is the dimension of the simplex σ .

2.4 Weighted Delaunay triangulation

Weighted points. A weighted point is a pair consisting of a point p of $\mathbb{R}^d$, called the *center* of the weighted point, and a non-negative real number $\omega(p)$, called the *weight* of the weighted point. It might be convenient to identify a weighted point $(p, \omega(p))$ and the hypersphere (we will simply say sphere in the sequel) centered at p of radius $\omega(p)$.

Two weighted points (or spheres) $(p, \omega(p))$ and $(q, \omega(q))$ are called *orthogonal* when $\|p - q\|^2 = \omega(p)^2 + \omega(q)^2$, *further than orthogonal* when $\|p - q\|^2 > \omega(p)^2 + \omega(q)^2$, and *closer than orthogonal* when $\|p - q\|^2 < \omega(p)^2 + \omega(q)^2$.

Given a point set $P = \{p_1, \dots, p_n\} \subseteq \mathbb{R}^d$, a *weight function* on P is a function ω that assigns to each point $p_i \in P$ a non-negative real weight $\omega(p_i)$: $\omega(P) = (\omega(p_1), \dots, \omega(p_n))$. We write $p_i^\omega = (p_i, \omega(p_i))$ and $P^\omega = \{p_1^\omega, \dots, p_n^\omega\}$.

We define the *relative amplitude* of ω as

$$\tilde{\omega} = \max_{p \in P, q \in P \setminus \{p\}} \frac{\omega(p)}{\|p - q\|}. \quad (2.4)$$

Unless stated otherwise, we will assume the following hypothesis.

Hypothesis 2.4.1 $\tilde{\omega} \leq \omega_0$, for some constant $\omega_0 \in [0, 1/2)$

Observe that, under this hypothesis, all the balls bounded by weighted spheres are disjoint.

Given a subset σ of $d + 1$ weighted points whose centers are affinely independent, there exists a unique sphere orthogonal to the weighted points of σ . The sphere is called the *orthosphere* of σ and its center o_σ and radius Φ_σ are called the *orthocenter* and the *orthoradius* of σ . If the weights of the vertices of σ are 0 (or all equal), then the orthosphere is simply the *circumscribing sphere* of σ whose center and radius are respectively called *circumcenter* and *circumradius*. If σ is a j -simplex, $j < d$, the orthosphere of σ is the smallest sphere that is orthogonal to the (weighted) vertices of σ . Its center o_σ lies in $\text{aff}(\sigma)$.

A finite set of weighted points P^ω is said to be in *general position* if there exists no sphere orthogonal to $d + 2$ weighted points of P^ω .

Weighted Voronoi diagram and Delaunay triangulation. Let ω be a weight function defined on P . We define the weighted Voronoi cell of $p \in P$ as

$$\text{Vor}^\omega(p) = \{x \in \mathbb{R}^d : \|p - x\|^2 - \omega(p)^2 \leq \|q - x\|^2 - \omega(q)^2, \forall q \in P\}. \quad (2.5)$$

The weighted Voronoi cells and their k -dimensional faces, $0 \leq k \leq d$, form a cell complex that decomposes $\mathbb{R}^d$ into convex polyhedral cells. This cell complex is called the weighted Voronoi diagram or power diagram of P [Aur87].

Let σ be a subset of points of P and write $\text{Vor}^\omega(\sigma) = \bigcap_{x \in \sigma} \text{Vor}^\omega(x)$. We will assume that the points of P are in general position. Then, $\text{Vor}^\omega(\sigma) = \emptyset$ when $|\sigma| > d + 1$, and the collection of all simplices $\text{conv}(\sigma)$ such that $\text{Vor}^\omega(\sigma) \neq \emptyset$ constitutes a triangulation called the weighted Delaunay triangulation $\text{Del}^\omega(P)$. The mapping that associates to the face $\text{Vor}^\omega(\sigma)$ of $\text{Vor}^\omega(P)$ the face $\text{conv}(\sigma)$ of $\text{Del}^\omega(P)$ is a *duality*, i.e. a bijection that reverses the inclusion relation.

Alternatively, a d -simplex σ is in $\text{Del}^\omega(P)$ if the orthosphere o_σ of σ is further than orthogonal from all weighted points in $P^\omega \setminus \sigma$.

Observe that the definition of weighted Voronoi diagrams makes sense if, for some $p \in P$, $\omega(p)^2 < 0$, i.e. some of the weights are imaginary. In fact, since adding a same positive quantity to all $\omega(p)^2$ does not change the diagram, handling imaginary weights is as easy as handling real weights. In the sequel, we will only consider real positive weights, except in Lemma 2.4.2.

The weighted Delaunay triangulation of a set of weighted points can be computed efficiently in small dimensions and has found many applications, see e.g. [Aur87]. We use weighted Delaunay triangulations for two main reasons. The first one is that the restriction of a d -dimensional weighted Voronoi diagram to an affine space of dimension k is a k -dimensional weighted Voronoi diagram that can be computed without computing the d -dimensional diagram (see Lemma 2.4.2). The other main reason is that some flat simplices named *slivers* can be removed from a Delaunay triangulation by weighting the vertices (see [BGO09, CDE⁺00a, CDR05a] and Section 3.3).

Lemma 2.4.2 *Let H be a k -dimensional affine space of $\mathbb{R}^d$. The restriction of the weighted Voronoi diagram of $P = \{p_0, \dots, p_m\}$ to H is identical to the k -dimensional weighted Voronoi*




Figure 2.4: Refer to Lemma 2.4.2. The grey line denotes the k -dimensional plane H and the black line denotes $\text{Vor}^\omega(p_i p_j)$.

diagram of $P' = \{p'_0, \dots, p'_m\}$ in H , where p'_i denotes the orthogonal projection of p_i onto H and the squared weight of p'_i is

$$\xi(p'_i)^2 = \omega(p_i)^2 - \|p_i - p'_i\|^2 + \lambda^2$$

where $\lambda = \max_{p_j \in P} \|p_j - p'_j\|$ is used to have all weights non-negative. In other words,

$$\text{Vor}^\omega(p_i) \cap H = \text{Vor}^\xi(p'_i)$$

where

$$\text{Vor}^\xi(p'_i) = \{x \in H : \|x - p'_i\|^2 - \xi(p'_i)^2 \leq \|x - p'_j\|^2 - \xi(p'_j)^2, \forall p'_j \in P'\}.$$

Proof By Pythagoras theorem, we have $\forall x \in H \cap \text{Vor}^\omega(p_i), \|x - p_i\|^2 - \omega(p_i)^2 \leq \|x - p_j\|^2 - \omega(p_j)^2 \Leftrightarrow \|x - p'_i\|^2 + \|p_i - p'_i\|^2 - \omega(p_i)^2 \leq \|x - p'_j\|^2 + \|p_j - p'_j\|^2 - \omega(p_j)^2$, where p'_i denotes the orthogonal projection of $p_i \in P$ onto H . See Figure 2.4. Hence the restriction of $\text{Vor}^\omega(P)$ to H is the weighted Voronoi diagram $\text{Vor}^\xi(P')$ of the points P' with the weight function: $\xi : P' \rightarrow [0, \infty)$, with

$$\xi(p'_i)^2 = -\|p_i - p'_i\|^2 + \omega(p_i)^2 + \lambda^2$$

where $\lambda = \max_{p_j \in P} \|p_j - p'_j\|$. □

2.4.1 Properties of weighted Delaunay triangulation

Circumradius and orthoradius. The following lemma states some basic facts about weighted Voronoi diagrams when the relative amplitude of the weighting function is bounded. Similar results were proved in [CDE⁺00a].

The following lemma states some basic facts about weighted Voronoi diagrams when the relative amplitude of the weighting function is bounded. Similar results were proved in [CDE⁺00a].

Lemma 2.4.3 Assume that Hypothesis 2.4.1 is satisfied. If τ is a simplex of $\text{Del}^\omega(P)$ and p and q are any two vertices of τ , then

1. $\forall z \in \text{aff}(\text{Vor}^\omega(\tau)), \|q - z\| \leq \frac{\|p - z\|}{\sqrt{1 - 4\omega_0^2}}.$
2. $\forall z \in \text{aff}(\text{Vor}^\omega(\tau)), \sqrt{\|z - p\|^2 - \omega^2(p)} \geq \Phi_\tau.$
3. $\forall \sigma \subseteq \tau, \Phi_\sigma \leq \Phi_\tau.$




Figure 2.5: For the proof of Lemma 2.4.3.

Proof Refer to Figure 2.5.

1. If $\|z - q\| \leq \|z - p\|$, then the lemma is proved since $0 < \sqrt{1 - 4\omega_0^2} \leq 1$. Hence assume that $\|z - q\| > \|z - p\|$. Since $z \in \text{aff}(\text{Vor}^\omega(\tau))$

$$\begin{aligned}
 \|z - p\|^2 &= \|z - q\|^2 + \omega(p)^2 - \omega(q)^2 \\
 &\geq \|z - q\|^2 - \omega(q)^2 \\
 &\geq \|z - q\|^2 - \omega_0^2 \|p - q\|^2 \\
 &\geq \|z - q\|^2 - \omega_0^2 (\|z - p\| + \|z - q\|)^2 \\
 &> \|z - q\|^2 - 4\omega_0^2 \|z - q\|^2 = (1 - 4\omega_0^2) \|z - q\|^2.
 \end{aligned}$$

2. We know that $o_\tau = \text{aff}(\text{Vor}^\omega(\tau)) \cap \text{aff}(\tau)$. Therefore, using Pythagoras theorem,

$$\begin{aligned}
 \|z - p\|^2 - \omega(p)^2 &= \|p - o_\tau\|^2 + \|o_\tau - z\|^2 - \omega(p)^2 \\
 &= \Phi_\tau^2 + \|o_\tau - z\|^2 \geq \Phi_\tau^2.
 \end{aligned}$$

3. The result directly follows from part 2 and the fact that $\text{aff}(\text{Vor}^\omega(\tau)) \subseteq \text{aff}(\text{Vor}^\omega(\sigma))$ (since $\sigma \subseteq \tau$). $\square$

Excentricity. Let σ be a simplex and p be a vertex of σ . We define the *excentricity* $H_\sigma(p, \omega(p))$ of σ with respect to p as the signed distance from o_σ to $\text{aff}(\sigma_p)$. Hence, $H_\sigma(p, \omega(p))$ is positive if o_σ and p lie on the same side of $\text{aff}(\sigma_p)$ and negative if they lie on different sides of $\text{aff}(\sigma_p)$. The following lemma bounds the excentricity of a simplex.

The following lemma is a generalization of Claim 13 from [CDE⁺00a]. It bounds the excentricity of a simplex σ with respect to a vertex $p \in \sigma$ as a function of the weight $\omega(p)$.

Lemma 2.4.4 *Let σ be a simplex of $\text{Del}^\omega(\mathcal{P})$ and let p be any vertex of σ . We have*

$$H_\sigma(p, \omega(p)) = H_\sigma(p, 0) - \frac{\omega(p)^2}{2D_\sigma(p)}.$$




Figure 2.6: For the proof of Lemma 2.4.4.

Proof Refer to Figure 2.6. For convenience, we write $R = \Phi_\sigma$ for the orthoradius of σ and $r = \Phi_{\sigma_p}$ for the orthoradius of σ_p . The orthocenter o_{σ_p} of σ_p is the projection of o_σ onto $\text{aff}(\sigma_p)$. When the weight $\omega(p)$ varies while the weights of other points remain fixed, o_σ moves on a (fixed) line L that passes through o_{σ_p} . Now, let p' and p'' be the projections of p onto L and $\text{aff}(\sigma_p)$ respectively. Write $\delta = \|p - p'\|$ for the distance from p to L . Since p and L (as well as all the objects of interest in this proof) belong to $\text{aff}(\sigma)$, $\|p' - o_{\sigma_p}\| = \|p - p''\| = D_\sigma(p)$.

We have $R^2 + \omega(p)^2 = (H_\sigma(p, \omega(p)) - D_\sigma(p))^2 + \delta^2$. We also have $R^2 = H_\sigma(p, \omega(p))^2 + r^2$ and therefore $H_\sigma(p, \omega(p))^2 = (H_\sigma(p, \omega(p)) - D_\sigma(p))^2 + \delta^2 - \omega(p)^2 - r^2$. We deduce that

$$H_\sigma(p, \omega(p)) = \frac{D_\sigma(p)^2 + \delta^2 - r^2}{2D_\sigma(p)} - \frac{\omega(p)^2}{2D_\sigma(p)}.$$

The first term on the right side is $H_\sigma(p, 0)$ and the second is the displacement of o_σ when we change the squared weight of p to $\omega(p)^2$. $\square$

2.5 Definitions and results from topology

In this section, we will give the standard definitions and results from topology used in this thesis. See, for example, [Bre94, Dug66, Mun00].

Definition 2.5.1 (Homeomorphism) A function $f : X \rightarrow Y$ between topological spaces is called a homeomorphism if $f^{-1} : Y \rightarrow X$ (i.e., f is bijective) and both f and f^{-1} are continuous.

Definition 2.5.2 (Hausdorff space) A topological space X is a Hausdorff space if for any two points x, y ($x \neq y$) there are disjoint open spaces U and V with $x \in U$ and $y \in V$.

Definition 2.5.3 (Covering) A covering of a topological space X is a collection of sets whose union is X . It is an open covering if the sets are open. A subcover is a subset of this collection which still covers the space.

Definition 2.5.4 (Compact) A topological space X is said to be compact if every open covering of X has a finite subcover.

Theorem 2.5.5 If X is a Hausdorff space, then any compact subset of X is closed.

Theorem 2.5.6 If X is compact and $f : X \rightarrow Y$ is continuous, then $f(X)$ is compact.

Theorem 2.5.7 If X is compact, Y is a Hausdorff space and $f : X \rightarrow Y$ is continuous, injective, surjective, then f is a homeomorphism.

Definition 2.5.8 (Connected) A topological space is connected if it is not the disjoint union of two nonempty open sets.

Proposition 2.5.9 If $f : X \rightarrow Y$ is continuous and X is connected, then $f(X)$ is connected.

Lemma 2.5.10 If $\{Y_i\}$ is a collection of connected sets in a topological space X and if no two of the Y_i are disjoint, then $\bigcup Y_i$ is connected.

Lemma 2.5.11 The relation “ p and q belong to a connected subset of X ” is an equivalence relation.

Definition 2.5.12 (Components) The equivalence classes of the equivalence relation in Lemma 2.5.11 are called the components of X .

Lemma 2.5.13 For a topological space X :

1. Components of space X are connected and closed.
2. Each connected subset is contained in a component of X .
3. Components are either equal or disjoint, and fill out X .

Definition 2.5.14 (Open map) A function $f : X \rightarrow Y$ is an open map if for any open set U in X , the image $f(U)$ is open in Y .

Lemma 2.5.15 *Let $f : X \rightarrow Y$ is a continuous map which is also a open map. If f is a bijection then, then f is a homeomorphism.*

Theorem 2.5.16 (Invariance of domain generalized) *If M and N are topological k -manifold without boundary and $f : M \rightarrow N$ is a continuous function which is locally one-one, i.e., for every point $p \in M$ there exists a open set $U_p \subset M$ with $p \in U_p$ and f restricted to U_p is injective, then f is an open map.*

Definition 2.5.17 (C^r -diffeomorphism) *Let $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$. A bijective C^r -function $f : U \rightarrow V$ is called C^r -diffeomorphism if f^{-1} map is a C^r -function.*

Definition 2.5.18 (Isotopy) *Let X, Y be topological space. The map $F : X \times [0, 1] \rightarrow Y$ is called an isotopy of X if the family of maps*

$$F_t : X \rightarrow Y, \quad x \mapsto F(x, t)$$

are homeomorphism between X and $F_t(X)$.

We will now recall the definition of covering space. See, e.g. [Hat02, Mas67].

Definition 2.5.19 (Covering space) *Let X be a topological space. A covering space of X is a space $\tilde{X}$ together with a continuous surjective map $f : \tilde{X} \rightarrow X$ satisfying the following condition: There exist an open cover $\{U_\alpha\}$ of X such that for each α , $f^{-1}(U_\alpha)$ is a disjoint union of open sets in $\tilde{X}$, each of which is mapped homeomorphically onto U_α by p .*

The following lemma follows directly from the above definition. See, e.g. [Hat02, Mas67].

Lemma 2.5.20 *Let $f : \tilde{X} \rightarrow X$ be a covering map. If X is connected, then the cardinality of $f^{-1}(x)$ is constant for all $x \in X$.*

For a given simplicial complex $\mathcal{K}$, let $\mathcal{K}^j$ denotes the the subcomplex containing all i -dimensional simplices of $\mathcal{K}$ with $i \leq j$. The following lemma is a special case¹ of a standard result from piecewise linear topology. See, e.g, Appendix II of [Whi57a, Lemma 15a].

Lemma 2.5.21 *Let $\mathcal{K}$ be a j -dimensional piecewise linear manifold with oriented j -dimensional simplices of $\mathcal{K}$ be oriented, and let the continuous map $f : \mathcal{K} \rightarrow \mathbb{R}^j$ be a simplexwise positive map for all the j -simplices in $\mathcal{K}$. Then for any connected open subset R of $\mathbb{R}^j \setminus f(\partial\mathcal{K})$, any two points of R not in $f(\mathcal{K}^{j-1})$ are covered the same number of times. If this number is 1, then f , restricted to the open subset $R' = f^{-1}(R)$ of $\mathcal{K}$, is injective.*

¹The Lemma 15a from Appendix II of [Whi57a] holds for *oriented pseudo-manifolds* which are a generalization of oriented piecewise linear manifolds.

Part II

Tangential Delaunay complexes

Chapter 3

Manifold reconstruction

In this chapter we give a provably correct algorithm to reconstruct a k -dimensional manifold embedded in d -dimensional Euclidean space. The input to our algorithm is a point sample coming from an unknown manifold. Our approach is based on two main ideas : the notion of tangential Delaunay complex defined in [BF04, Flö03, Fre02], and the technique of sliver removal by weighting the sample points [CDE⁺00a]. Differently from previous methods, we do not construct any subdivision of the d -dimensional ambient space. As a result, the running time of our algorithm depends only linearly on the extrinsic dimension d while it depends quadratically on the size of the input sample, and exponentially on the intrinsic dimension k . To the best of our knowledge, this is the first certified algorithm for manifold reconstruction whose complexity depends linearly on the ambient dimension. Using the results from Chapter 3, we will show that for a dense enough sample the output of our algorithm is ambient isotopic to the manifold and a close geometric approximation of the manifold.

3.1 Introduction

Manifold reconstruction consists in computing a piecewise linear approximation of an unknown manifold $\mathcal{M} \subset \mathbb{R}^d$ from a finite sample of unorganized points P lying on $\mathcal{M}$ or close to $\mathcal{M}$. When the manifold is a two-dimensional surface embedded in $\mathbb{R}^3$, the problem is known as the surface reconstruction problem. Surface reconstruction is a problem of major practical interest which has been extensively studied in the fields of Computational Geometry, Computer Graphics and Computer Vision. In the last decade, solid foundations have been established and the problem is now pretty well understood. Refer to Dey's book [Dey06], and the survey by Cazals and Giesen in [CG06] for recent results. The output of those methods is a triangulated surface that approximates $\mathcal{M}$. This triangulated surface is usually extracted from a 3-dimensional subdivision of the ambient space (typically a grid or a triangulation). Although rather inoffensive in 3-dimensional space, such data structures depend exponentially on the dimension of the ambient space, and all attempts to extend those geometric approaches to more general manifolds have led to algorithms whose complexities depend exponentially on d [BGO09, CL08, CDR05a, NSW08a].

The problem in higher dimensions is also of great practical interest in data analysis and machine learning. In those fields, the general assumption is that, even if the data are represented as points in a very high dimensional space $\mathbb{R}^d$, they in fact live on a manifold

of much smaller intrinsic dimension [SL00]. If the manifold is linear, well-known global techniques like principal component analysis (PCA) or multi-dimensional scaling (MDS) can be efficiently applied. When the manifold is highly nonlinear, several more local techniques have attracted much attention in visual perception and many other areas of science. Among the prominent algorithms are Isomap [TdSL00], LLE [RS00], Laplacian eigenmaps [BN02], Hessian eigenmaps [DG03], diffusion maps [LL06, NLCK05], principal manifolds [ZZ04]. Most of those methods reduce to computing an eigendecomposition of some connection matrix. In all cases, the output is a mapping of the original data points into $\mathbb{R}^k$ where k is the estimated intrinsic dimension of $\mathcal{M}$. Those methods come with no or very limited guarantees. For example, Isomap provides a correct embedding only if $\mathcal{M}$ is isometric to a convex open set of $\mathbb{R}^k$ and LLE can only reconstruct topological balls. To be able to better approximate the sampled manifold, another route is to extend the work on surface reconstruction and to construct a piecewise linear approximation of $\mathcal{M}$ from the sample in such a way that, under appropriate sampling conditions, the quality of the approximation can be guaranteed. First investigation on this problem can be found in the work of Cheng, Dey and Ramos [CDR05a], followed by Boissonnat, Guibas and Oudot [BGO09]. In both cases, however, the complexity of the algorithms is exponential in the ambient dimension d , which highly reduces their practical relevance.

In this chapter, we extend the geometric techniques developed in small dimensions and propose an algorithm that can reconstruct smooth manifolds of arbitrary topology while avoiding the computation of data structures in the ambient space. We assume that $\mathcal{M}$ is a smooth manifold of known dimension k and that we can compute the tangent space to $\mathcal{M}$ at any sample point. Under those conditions, we propose a provably correct algorithm that constructs a simplicial complex of dimension k that approximates $\mathcal{M}$. The complexity of the algorithm is linear in d , quadratic in the size n of the sample, and exponential in k . Our work builds on [BGO09] and [CDR05a] but dramatically reduces the dependence on d . To the best of our knowledge, this is the first certified algorithm for manifold reconstruction whose complexity depends only linearly on the ambient dimension. In the same spirit, Chazal and Oudot [CO08] have devised an algorithm of intrinsic complexity to solve the easier problem of computing the homology of a manifold from a sample.

Our approach is based on two main ideas : the notion of *tangential Delaunay complex* introduced in [BF04, Flö03, Fre02], and the technique of sliver removal by weighting the sample points [CDE⁺00a]. The tangential complex is obtained by gluing local (Delaunay) triangulations around each sample point. The tangential complex is a subcomplex of the d -dimensional Delaunay triangulation of the sample points but it can be computed using mostly operations in the k -dimensional tangent spaces at the sample points. Hence the dependence on k rather than d in the complexity. However, due to the presence of so-called inconsistencies, the local triangulations may not form a triangulated manifold. Although this problem has already been reported [Fre02], no solution was known except for the case of curves ($k = 1$) [Flö03]. The idea of removing inconsistencies among local triangulations that have been computed independently has already been used for maintaining dynamic meshes [She05] and generating anisotropic meshes [BWY08]. Our approach is close in spirit to the one in [BWY08]. We show that, under appropriate sample conditions, we can remove inconsistencies by weighting the sample points. We can then prove that the approximation returned by our algorithm is ambient isotopic to

$\mathcal{M}$, and a close geometric approximation of $\mathcal{M}$.

Our algorithm can be seen as a *local* version of the cocone algorithm of Cheng et al. [CDR05a]. By local, we mean that we do not compute any d -dimensional data structure like a grid or a triangulation of the ambient space. Still, the tangential complex is a subcomplex of the weighted d -dimensional Delaunay triangulation of the (weighted) data points and therefore implicitly relies on a global partition of the ambient space. This is a key to our analysis and distinguishes our method from other local algorithms that have been proposed in the surface reconstruction literature [CSD04, GKS00].

Organization of the chapter. In Section 3.2, we define the two main notions of this chapter: the tangential complex and inconsistent configurations. The algorithmic part is given in Section 3.3. The main structural results are given in Section 3.4. Under some sampling condition, we bound the shape measure of the simplices of the tangential complex in Section 3.4.1 and of inconsistent configurations in Section 3.4.2. A crucial fact is that inconsistent configurations cannot be fat. We also bound the number of simplices and inconsistent configurations that can be incident on a point in Section 3.4.3. In Sections 3.4.4 and 3.4.5, we prove the correctness of the algorithm and theoretical guarantees on the output, and space and time complexity respectively. Finally, in Section 3.5, we conclude with some possible extensions.

3.2 Tangential complex and inconsistent configurations




Figure 3.1: $\mathcal{M}$ is the red curve. The sample P is the set of small blue circles. The tangent space at u is denoted by $T_u \mathcal{M}$. The Voronoi diagram of the sample is in grey. The edges of the Delaunay triangulation $\text{Del}(P)$ are the blue line segments between small circles. In bold, $\text{star}(u) = \{uv, uw\}$.

Let u be a point of P . We denote by $\text{Del}_u^\omega(P)$ the weighted Delaunay triangulation of P restricted to the tangent space $T_u \mathcal{M}$. Equivalently, the simplices of $\text{Del}_u^\omega(P)$ are the simplices of $\text{Del}^\omega(P)$ whose Voronoi dual faces intersect $T_u \mathcal{M}$, i.e., $\sigma \in \text{Del}_u^\omega(P)$ iff $\text{Vor}^\omega(\sigma) \cap T_u \mathcal{M} \neq \emptyset$. Observe that $\text{Del}_u^\omega(P)$ is in general a k -dimensional triangulation. Since this situation can always be ensured by applying some infinitesimal perturbation on P , we will assume, in the rest of the chapter, that all $\text{Del}_u^\omega(P)$ are k -dimensional triangulations. Finally, write $\text{star}(u)$ for the *star* of u in $\text{Del}_u^\omega(P)$, i.e. the set of simplices that are incident to u in $\text{Del}_u^\omega(P)$ (see Figure 3.1).




Figure 3.2: Stitching the stars $\text{star}(u)$ to get tangential Delaunay complex $\text{Del}_{T\mathcal{M}}^\omega(P)$. The above figure is taken from [She05].

We denote by *tangential Delaunay complex* or *tangential complex* for short, the simplicial complex $\{\sigma : \sigma \in \text{star}(u), u \in P\}$, i.e.,

$$\text{Del}_{T\mathcal{M}}^\omega(P) = \bigcup_{u \in P} \text{star}(u).$$

See Figure 3.2. We denote it by $\text{Del}_{T\mathcal{M}}^\omega(P)$. By our assumption above, $\text{Del}_{T\mathcal{M}}^\omega(P)$ is a k -dimensional subcomplex of $\text{Del}^\omega(P)$.

By duality, computing $\text{star}(u)$ is equivalent to computing the restriction to $T_u\mathcal{M}$ of the (weighted) Voronoi cell of u , which, by Lemma 2.4.2, reduces to computing a cell in a k -dimensional weighted Voronoi diagram embedded in $T_u\mathcal{M}$. To compute such a cell, we need to compute the intersection of $|P| - 1$ halfspaces of $T_u\mathcal{M}$ where $|P|$ is the cardinality of P . Each halfspace is bounded by the bisector consisting of the points of $T_u\mathcal{M}$ that are at equal weighted distance from u^ω and some other point in P^ω . This can be done in optimal time [Cha93, CS89]. It follows that the tangential complex can be computed without constructing any data structure of dimension higher than k , the intrinsic dimension of $\mathcal{M}$.




Figure 3.3: The stars of the points u and v are drawn in red and blue respectively. The inconsistent simplex $\sigma \in \text{star}(u)$ ($\sigma \notin \text{star}(v)$) is drawn in blue.

The tangential Delaunay complex is *not* in general a triangulated manifold and therefore not a good approximation of $\mathcal{M}$. This is due to the presence of so-called *inconsistencies*. Consider a k -simplex σ of $\text{Del}_{T\mathcal{M}}^\omega(P)$ with two vertices u and v such that σ is in $\text{star}(u)$ but not in $\text{star}(v)$ (refer to Figure 3.3). The k -dimensional simplex σ is called an *inconsistent simplex*. We write $B_u(\sigma)$ (and $B_v(\sigma)$) for the open ball centered on




Figure 3.4: An inconsistent configuration in the unweighted case. Edge $\sigma = [u, v]$ is in $\text{Del}_u(\mathcal{P})$ but not in $\text{Del}_v(\mathcal{P})$ since $\text{Vor}(uv)$ intersects $T_u\mathcal{M}$ but not $T_v\mathcal{M}$. This happens because the line segment $[m_u(\sigma)m_v(\sigma)]$ penetrates (at i_ϕ) the Voronoi cell of a point $w \neq u, v$, therefore creating an inconsistent configuration $\phi = [u, v, w]$. Also note that i_ϕ is the center of an empty sphere circumscribing simplex ϕ .

$T_u\mathcal{M}$ (and $T_v\mathcal{M}$) that is orthogonal to the (weighted) vertices of σ , and denote by $m_u(\sigma)$ (and $m_v(\sigma)$), or m_u (m_v) for short, its center. According to our definition, σ is inconsistent simplex iff $B_u(\sigma)$ is further than orthogonal from all weighted points in $\mathcal{P}^\omega \setminus \sigma$ while there exists a weighted point in $\mathcal{P}^\omega \setminus \sigma$ that is closer than orthogonal from $B_v(\sigma)$. We deduce from the above discussion that the line segment $[m_u m_v]$ has to penetrate the interior of $\text{Vor}^\omega(w)$, where w^ω is a weighted point in $\mathcal{P}^\omega \setminus \sigma$. Refer to Figure 3.4.

For a given constant Θ_0 , we now formally define a Θ_0 -inconsistent configuration as follows.

Definition 3.2.1 (Inconsistent configuration) Let $\phi = [p_0, \dots, p_{k+1}]$ be a $(k+1)$ -simplex, and let u, v , and w be three vertices of ϕ . We say that ϕ is a Θ_0 -inconsistent (or inconsistent for short) configuration of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ witnessed by the triplet (u, v, w) if

- The k -dimensional simplex $\sigma = \phi \setminus \{w\}$ is an inconsistent simplex with σ is in $\text{star}(u)$ but not in $\text{star}(v)$.
- $\text{Vor}^\omega(w)$ is one of the first weighted Voronoi cells of $\text{Vor}^\omega(\mathcal{P})$, other than the weighted Voronoi cells of the vertices of σ , that is intersected by the line segment $[m_u m_v]$ oriented from m_u to m_v . Here $m_u = T_u\mathcal{M} \cap \text{Vor}^\omega(\sigma)$ and $m_v = T_v\mathcal{M} \cap \text{aff}(\text{Vor}^\omega(\sigma))$. Let i_ϕ denote the point where the oriented segment $[m_u m_v]$ first intersects $\text{Vor}^\omega(w)$.
- σ is a Θ_0 -fat simplex.




Figure 3.5: In Figure (a), $\mathcal{M}$ is the black curve, the sample P is the set of small circles, the tangent space at a point $x \in P$ is denoted by $T_x \mathcal{M}$ and the Voronoi diagram of the sample is in grey and $\text{Del}_{T_M}(P)$ is the line segments between the sample points. In dashed lines, are the inconsistent simplices in $\text{Del}_{T_M}(P)$. In Figure (b), the grey triangles denote the inconsistent configurations corresponding to the inconsistent simplices in Figure (a).

Note that i_ϕ is the center of a sphere that is orthogonal to the weighted vertices of σ and also to w^ω , and further than orthogonal from all the other weighted points of P^ω . Equivalently, i_ϕ is the point on $[m_u m_v]$ that belongs to $\text{Vor}^\omega(\phi)$.

An inconsistent configuration is therefore a $(k+1)$ -simplex of $\text{Del}^\omega(P)$. However, an inconsistent configuration does not belong to $\text{Del}_{T_M}^\omega(P)$ since $\text{Del}_{T_M}^\omega(P)$ has no $(k+1)$ -simplex under our general position assumption. Moreover, the lower dimensional faces of an inconsistent configuration do not necessarily belong to $\text{Del}_{T_M}^\omega(P)$.

Since the inconsistent configurations are $k+1$ -dimensional simplices, we will use the same notations for inconsistent configurations as for simplices, e.g. R_ϕ and c_ϕ for the circumradius and the circumcenter of ϕ , ρ_ϕ and Θ_ϕ for its radius-edge ratio and fatness respectively.

We write $\text{Inc}^\omega(P)$ for the subcomplex of $\text{Del}(P)$ consisting of all the Θ_0 -inconsistent configurations of $\text{Del}_{T_M}^\omega(P)$ and their subfaces. We also define the *completed complex* as $C^\omega(P) = \text{Del}_{T_M}^\omega(P) \cup \text{Inc}^\omega(P)$. Refer to Figure 3.5.

An important observation, stated as Lemma 3.4.6 in Section 3.4.2, is that, if ε is sufficiently small with respect to Θ_0 , then the fatness of ϕ is less than Θ_0 . Hence, if the subfaces of ϕ are Θ_0 -fat simplices, ϕ will be a Θ_0 -sliver. This observation is at the core of our reconstruction algorithm.

3.3 Manifold reconstruction

The algorithm removes all Θ_0 -slivers from $C^\omega(P)$ by weighting the points of P . By the observation just above, all inconsistencies in the tangential complex will then also be removed. All simplices being consistent, the resulting weighted tangential Delaunay complex $\hat{\mathcal{M}}$ output by the algorithm will be a simplicial k -manifold that approximates $\mathcal{M}$ well, as will be shown in Theorem 3.4.13 from Section 3.4.4.

In this section, we describe the algorithm. Its analysis is deferred to Section 3.4.

3.3.1 Algorithm

Let $\mathcal{M}$ be a compact smooth submanifold of positive reach and without boundaries, and let P be an (ε, δ) -lfs sample of $\mathcal{M}$. $\mathcal{M}$, ε , δ are unknown and the input to the algorithm consists only of the sample P and an upper bound η_0 on the sampling ratio (ε/δ) of P . As shown in [CWW08, GW04a], we can estimate the tangent space $T_p\mathcal{M}$ at each sample point p and also the dimension k of the manifold from P and η_0 . We assume now that k and $T_p\mathcal{M}$, for any point $p \in P$, are known.

The algorithm fixes ω_0 , the bound on the relative amplitude of the weight assignment, in the interval $[0, 1/2)$ (Hypothesis 2.4.1). The algorithm also fixes Θ_0 to a constant defined in Theorem 3.4.12, that depends on k , ω_0 and η_0 .

We define the *local neighborhood* of $p \in P$ as

$$LN(p) = \{q \in P : |B(p, \|p - q\|) \cap P| \leq N\}. \quad (3.1)$$

where N is a constant that depends on k , ω_0 and η_0 to be defined in Section 3.4.3. We will show in Lemma 3.4.9, that $LN(p)$ includes all the points of P that can form an edge with p in $C^\omega(P)$. In fact, the algorithm can use instead of $LN(p)$ any subset of P that contains $LN(p)$. This will only affect the complexity of the algorithm, not the output.

We will also need the definition of $m(p)$ defined in Section 2.1.

Outline of the algorithm. Initially, all the sample points in P are assigned zero weights, and the completed complex $C^\omega(P)$ is built for this zero weight assignment. Then the algorithm processes each point $p_i \in P = \{p_1, \dots, p_n\}$ in turn, and assigns a new weight to p_i . The new weight is chosen so that all the simplices of all dimensions in $C^\omega(P)$ are Θ_0 -fat. See Algorithm 1.

We now give the details of the functions used in the manifold reconstruction algorithm. The function **update_completed_complex**(Q, ω) is described as Algorithm 2. It makes use of two functions, **build_star**(p) and **build_inconsistent_configurations**(p, σ).

The function **build_star**(p) calculates the weighted Voronoi cell of p , which reduces to computing the intersection of the halfspaces of $T_p\mathcal{M}$ bounded by the (weighted) bisectors between p and other points in $LN(p)$.

The function **build_inconsistent_configurations**(u, σ) adds to $C^\omega(P)$ all the inconsistent configurations of the form $\phi = \sigma \cup \{w\}$ where σ is an inconsistent simplex of $\text{star}(u)$. More precisely, for each vertex $v \neq u$ of σ such that $\sigma \notin \text{star}(v)$, we calculate the points $w \in LN(p)$, such that (u, v, w) witnesses the inconsistent configuration $\phi = \sigma \cup \{w\}$. Specifically, we compute the restriction of the Voronoi diagram of the points in $LN(u)$ to the line segment $[m_u, m_v]$, where $m_u = T_u\mathcal{M} \cap \text{aff}(\text{Vor}^\omega(\sigma))$ and $m_v = T_v\mathcal{M} \cap \text{aff}(\text{Vor}^\omega(\sigma))$.

Algorithm 1 **Manifold_reconstruction**($P = \{p_0, \dots, p_n\}, \eta_0$)

```

// Initialization
for  $i = 1$  to  $n$  do
    calculate the local neighborhood  $LN(p_i)$ 
for  $i = 1$  to  $n$  do
     $\omega(p_i) \leftarrow 0$ 
// Build the full unweighted completed complex  $C^\omega(P)$ 
 $C^\omega(P) \leftarrow \text{update\_completed\_complex}(P, \omega)$ 
// Weight assignment to remove inconsistencies
for  $i = 1$  to  $n$  do
     $\omega(p_i) \leftarrow \text{weight}(p_i, \omega)$ 
    update\_completed\_complex( $LN(p_i), \omega$ )
// Output
output :  $\hat{\mathcal{M}} \leftarrow \text{Del}_{T\mathcal{M}}^\omega(P)$ 

```

Algorithm 2 Function **update_completed_complex**(Q, ω)

```

for each point  $q \in Q$  do
    build_star( $q$ )
for each  $q \in Q$  do
    for each  $k$ -simplex  $\sigma$  in  $\text{star}(q)$  do
        if  $\sigma$  is  $\Theta_0$ -fat and  $\exists v \in \sigma, \sigma \notin \text{star}(v)$  then
            //  $\sigma$  is inconsistent
            build_inconsistent_configurations( $q, \sigma$ )

```

Algorithm 3 Function **weight**(p, ω)

```

 $S(p) \leftarrow \text{candidate\_slivers}(p, \omega)$ 
//  $J(p)$  is the set of squared weights of  $p$  such that  $C^\omega(P)$  contains
// no  $\Theta_0$ -sliver incident to  $p$ 
 $J(p) \leftarrow [0, \omega_0^2 \text{mn}(p)^2] \setminus \bigcup_{\sigma \in S(p)} W(\sigma)$ 
 $\omega(p)^2 \leftarrow$  a squared weight from  $J(p)$ 
return  $\omega(p)$ 

```

According to the definition of an inconsistent configuration, w is one of the sites whose (restricted) Voronoi cell is the first to be intersected by the line segment $[m_u m_v]$, oriented from m_u to m_v . We add inconsistent configuration $\phi = \sigma \cup \{w\}$ to the completed complex.

We now give the details of function **weight**(p, ω) that computes $\omega(p)$, keeping the other weights fixed (see Algorithm 3). This function extends a similar subroutine introduced in [CDE⁺00a] for removing slivers in $\mathbb{R}^3$. We need first to define candidate simplices. A *candidate simplex* of p is defined as a simplex of $C^\omega(P)$ that becomes incident to p when the weight of p is varied from 0 to $\omega_0 \text{nn}(p)$, keeping the weights of all the points in $P \setminus \{p\}$ fixed. Note that a candidate simplex of p is incident to p for some weight $\omega(p)$ but does not necessarily belong to $\text{star}(p)$.

Let σ be a candidate simplex of p that is a Θ_0 -sliver. We associate to σ a forbidden interval $W(\sigma)$ that consists of all squared weights $\omega(p)^2$ for which σ appears as a simplex in $C^\omega(P)$ (the weights of the other points remaining fixed).

The function **candidate_slivers**(p, ω) varies the weight of p and computes all the candidate slivers of p and their corresponding weight intervals $W(\sigma)$. More precisely, this function follows the following steps.

1. We first detect all candidate j -simplices for all $2 \leq j \leq k + 1$. This is done in the following way. We vary the weight of p from 0 to $\omega_0 \text{nn}(p)$, keeping the weights of the other points fixed. For each new weight assignment to p , we modify the stars and inconsistent configurations of the points in $LN(p)$ and detect the new j -simplices incident to p that have not been detected so far. The weight of point p changes only in a finite number of instances $0 = P_0 < P_1 < \dots < P_{n-1} < P_n = \omega_0 \text{nn}(p)$.

2. We determine the next weight assignment of p in the following way. For each new simplex σ currently incident to p , we keep it in a priority queue ordered by the weight of p at which σ will disappear for the first time. Hence the minimum weight in the priority queue gives the next weight assignment for p . Since the number of points in $LN(p)$ is bounded, the number of simplices incident to p is also bounded, as well as the number of times we have to change the weight of p .

3. For each candidate sliver σ of p which is detected, we compute $W(\sigma)$ on the fly.

3.4 Analysis of the algorithm


The analysis of the algorithm relies on structural results that will be proved in Sections 3.4.1, 3.4.2 and 3.4.3. We will then prove that the algorithm is correct and analyze its complexity in Sections 3.4.4 and 3.4.5.

In Section 5.4 (Theorem 3.4.13), we will also show that the output $\hat{\mathcal{M}}$ of the reconstruction algorithm is a good approximation of $\mathcal{M}$.

For this section, the following hypothesis is assumed to be satisfied as well as Hypothesis 2.4.1.

Hypothesis 3.4.1 P is an (ε, δ) -lfs sample of $\mathcal{M}$ of sampling ratio $\varepsilon/\delta \leq \eta_0$ for some positive constant η_0 .

The bounds to be given in the lemmas of this section will depend on the dimension k of $\mathcal{M}$, the bound η_0 on the sampling ratio, and on a positive scalar Θ_0 that bounds the fatness and will be used to define slivers, fat simplices and inconsistent configurations.



It follows (see Figure 3.6), that $\angle ptx > \pi/2$, which implies that $\|x - p\|^2 - \|x - t\|^2 - \|p - t\|^2 > 0$.

Hence,

$$\begin{aligned} \|x - p\|^2 - \|x - t\|^2 - \omega^2(p) + \omega^2(t) &\geq \|p - t\|^2 - \omega^2(p) \\ &\geq \|p - t\|^2 - \omega_0^2 \|p - t\|^2 \\ &> 0 \quad (\text{since } \omega_0 < \tfrac{1}{2}) \end{aligned}$$

This implies $x \notin \text{Vor}^\omega(p)$, which contradicts our initial assumption. We conclude that $\text{Vor}^\omega(p) \cap T_p \mathcal{M} \subseteq B(p, C_1 \varepsilon \text{lfs}(p))$ if Eq. (3.2) is satisfied, which is true for $C_1 \stackrel{\text{def}}{=} 3 + \sqrt{2} \approx 4.41$ and $\varepsilon \leq \varepsilon_0 < 0.09$. $\square$

The following lemma states that, under Hypotheses 2.4.1 and 3.4.1, the simplices of $\text{Del}_p^\omega(\mathcal{P})$ are small, have a good radius-edge ratio and a small excentricity.

Lemma 3.4.3 *Assume that Hypotheses 2.4.1 and 3.4.1 are satisfied. There exists positive constants C_2 , C_3 and C_4 that depend on ω_0 and η_0 such that, if $\varepsilon < \frac{1}{2C_2}$, the following holds.*

1. *If pq is an edge of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, then $\|p - q\| < C_2 \varepsilon \text{lfs}(p)$.*
2. *If σ is a simplex of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, then $\Phi_\sigma \leq C_3 L_\sigma$ and $\Gamma_\sigma = \Delta_\sigma / L_\sigma \leq C_3$.*
3. *If σ is a simplex of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ and p a vertex of σ , the excentricity $|H_\sigma(p, \omega(p))|$ is at most $C_4 \varepsilon \text{lfs}(p)$.*

Proof 1a. Consider first the case where pq is an edge of $\text{Del}_p^\omega(\mathcal{P})$. Then $T_p \mathcal{M} \cap \text{Vor}^\omega(pq) \neq \emptyset$. Let $x \in T_p \mathcal{M} \cap \text{Vor}^\omega(pq)$. From Lemma 3.4.2, we have $\|p - x\| \leq C_1 \varepsilon \text{lfs}(p)$. By Lemma 2.4.3,

$$\|q - x\| \leq \frac{\|p - x\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C_1 \varepsilon \text{lfs}(p)}{\sqrt{1 - 4\omega_0^2}}.$$

Hence, $\|p - q\| \leq C'_1 \varepsilon \text{lfs}(p)$ where $C'_1 \stackrel{\text{def}}{=} C_1(1 + 1/\sqrt{1 - 4\omega_0^2})$.

1b. From the definition of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, there exists a vertex r of τ such that $pq \in \text{star}(r)$. From 1a, $\|r - p\|$ and $\|r - q\|$ are at most $C'_1 \varepsilon \text{lfs}(r)$. Using the fact that lfs is 1-Lipschitz, $\text{lfs}(p) \geq \text{lfs}(r) - \|p - r\| \geq (1 - C'_1 \varepsilon) \text{lfs}(r)$ (from part 1a), which yields $\text{lfs}(r) \leq \frac{\text{lfs}(p)}{1 - C'_1 \varepsilon}$. It follows that

$$\|p - q\| \leq \|p - r\| + \|r - q\| \leq \frac{2C'_1 \varepsilon \text{lfs}(p)}{1 - C'_1 \varepsilon}.$$

The first part of the lemma is proved by taking $C_2 \stackrel{\text{def}}{=} \frac{5C'_1}{2}$ and using $2C_2 \varepsilon < 1$.

2. Without loss of generality, let $\sigma_1 \in \text{star}(p)$ be a simplex incident to p and $\sigma \subseteq \sigma_1$. Let $z \in \text{Vor}^\omega(\sigma_1) \cap T_p \mathcal{M}$, and $r_z = \sqrt{\|z - p\|^2 - \omega^2(p)}$. The ball centered at z with radius r_z is orthogonal to the weighted vertices of σ_1 . From Lemmas 2.4.3 (2) and (3), we have $r_z \geq \Phi_{\sigma_1}$ and $\Phi_\sigma \leq \Phi_{\sigma_1}$. Hence it suffices to prove that there exists a constant C_3 such that $r_z \leq C_3 L_\sigma$. Since $z \in \text{Vor}^\omega(\sigma) \cap T_p \mathcal{M}$, we deduce from Lemma 3.4.2 that $\|z - p\| \leq C_1 \varepsilon \text{lfs}(p)$. Therefore

$$r_z = \sqrt{\|z - p\|^2 - \omega^2(p)} \leq \|z - p\| \leq C_1 \varepsilon \text{lfs}(p).$$

For any vertex q of σ , we have $\|p - q\| \leq C_2 \varepsilon \text{lfs}(p)$ (By part 1). Using $2C_2 \varepsilon < 1$ and the fact that lfs is 1-Lipschitz, $\text{lfs}(p) \leq 2\text{lfs}(q)$. Therefore, taking for q a vertex of the shortest edge of σ , we have, using Lemma 2.2.1 and Hypothesis 1,

$$r_z \leq C_1 \varepsilon \text{lfs}(p) \leq C_1 \left(\frac{\varepsilon}{\delta} \times \delta \right) \times 2\text{lfs}(q) \leq 2C_1 \eta_0 L_\sigma.$$

From part 1 of the lemma we have $\Delta_{\sigma_1} \leq 2C_2 \varepsilon \text{lfs}(p)$. Therefore

$$\Gamma_\sigma = \frac{\Delta_\sigma}{L_\sigma} \leq \frac{\Delta_{\sigma_1}}{L_\sigma} \leq \frac{2C_2 \varepsilon \text{lfs}(p)}{\delta \text{lfs}(q)} \leq 4C_2 \eta_0 \stackrel{\text{def}}{=} C_3.$$

The last inequality follows from the fact that $\text{lfs}(p) \leq 2\text{lfs}(q)$ and $\varepsilon/\delta \leq \eta_0$.

3. From the definition of $H_\sigma(p, \omega(p))$, we have for all vertices $q \in \sigma_p$

$$|H_\sigma(p, \omega(p))| = \text{dist}(o_\sigma, \text{aff}(\sigma_p)) = \|o_\sigma - o_{\sigma_p}\| \leq \|o_\sigma - q\|$$

Using the facts that $\Phi_\sigma \leq C_3 L_\sigma$ (from part 2), $L_\sigma \leq \|p - q\|$, $\omega(q) \leq \omega_0 \|p - q\|$, and $\|p - q\| \leq C_2 \varepsilon \text{lfs}(p)$ (from part 1), we have

$$\begin{aligned} \|o_\sigma - q\| &= \sqrt{\Phi_\sigma^2 + \omega(q)^2} \\ &\leq \sqrt{C_3^2 L_\sigma^2 + \omega_0^2 \|p - q\|^2} \\ &\leq \|p - q\| \sqrt{C_3^2 + \omega_0^2} \\ &\leq C_2 \sqrt{C_3^2 + \omega_0^2} \times \varepsilon \text{lfs}(p) \\ &\stackrel{\text{def}}{=} C_4 \varepsilon \text{lfs}(p) \end{aligned}$$

□

3.4.2 Properties of inconsistent configurations

We now give lemmas on inconsistent configurations which are central to the proof of correctness of the reconstruction algorithm given later in the chapter. The first lemma is the analog of Lemma 3.4.3 applied to inconsistent configurations. Differently from Lemma 3.4.3, we need to use Corollary 2.3.3 to control the orientation of the facets of $\text{Del}_{\mathcal{M}}^\omega(P)$ and require the following additional hypothesis relating the sampling rate ε and the fatness bound Θ_0 .

Hypothesis 3.4.4 $2A\varepsilon < 1$ where $A \stackrel{\text{def}}{=} 2C_2 C_3 / \Theta_0^k$, and C_2, C_3 are the constants defined in Lemma 3.4.3.

Lemma 3.4.5 Assume that Hypotheses 2.4.1, 3.4.1 and 3.4.4 are satisfied. Let $\phi \in \text{Inc}^\omega(P)$ be an inconsistent configuration witnessed by (u, v, w) . There exist positive constants $C'_2 > C_2$, $C'_3 > C_3$ and $C'_4 > C_4$ that depend on ω_0 and η_0 s.t., if $\varepsilon < 1/C'_2$, then

1. $\|p - i_\phi\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(p)$ for all vertices p of ϕ .

2. If pq is an edge of ϕ then $\|p - q\| \leq C'_2 \varepsilon \text{lfs}(p)$.
3. If $\sigma \subseteq \phi$, then $\Phi_\sigma \leq C'_3 L_\sigma$ and $\Gamma_\sigma = \Delta_\sigma / L_\sigma \leq C'_3$.
4. If $\sigma \subseteq \phi$ and p is any vertex of σ , $|H_\sigma(p, \omega(p))|$ of σ is at most $C'_4 \varepsilon \text{lfs}(p)$.

Proof From the definition of inconsistent configurations, the k -dimensional simplex $\tau = \phi \setminus \{w\}$ belongs to $\text{Del}_u^\omega(\mathcal{P})$. We first bound $\text{dist}(i_\phi, \text{aff}(\tau)) = \|o_\tau - i_\phi\|$ where o_τ is the orthocenter of τ . Let $m_u \in \text{Vor}^\omega(\tau) \cap T_u \mathcal{M}$ denote, as in Definition 3.2.1, the point of $T_u \mathcal{M}$ that is the center of the ball orthogonal to the weighted vertices of τ . By definition, m_u is further than orthogonal to all other weighted points of $\mathcal{P} \setminus \tau$. Observe that $\|u - o_\tau\| \leq \|u - m_u\|$, since o_τ belongs to $\text{aff}(\tau)$ and therefore is the closest point to u in $\text{aff}(\text{Vor}^\omega(\tau))$. Moreover, by Lemma 3.4.2, $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$. Then, by Lemma 2.4.3, we have for all vertices $p \in \tau$

$$\|p - o_\tau\| \leq \frac{\|u - o_\tau\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{\|u - m_u\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C_1 \varepsilon \text{lfs}(u)}{\sqrt{1 - 4\omega_0^2}}. \quad (3.3)$$

Let $C'_2 \stackrel{\text{def}}{=} C_2 / \sqrt{1 - 4\omega_0^2} > C_2$. We now use the facts that the k -dimensional simplex τ is a Θ_0 -fat simplex (by definition of an inconsistent configuration), $\Delta_\tau \leq 2C_2 \varepsilon \text{lfs}(p)$ (Lemma 3.4.3 (1) and the Triangle Inequality), and $\Gamma_\tau \leq C_3$ (Lemma 3.4.3 (2)). Then Corollary 2.3.3 yields

$$\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq \frac{\Gamma_\tau \Delta_\tau}{\Theta_0^k \text{lfs}(p)} \leq \frac{2C_2 C_3 \varepsilon}{\Theta_0^k} = A \varepsilon \quad (3.4)$$

for all vertices p of τ . Which implies, together with $2A \varepsilon < 1$ (Hypothesis 3.4.4),

$$\tan^2 \angle(\text{aff}(\tau), T_u \mathcal{M}) \leq \frac{A^2 \varepsilon^2}{1 - A^2 \varepsilon^2} < 4A^2 \varepsilon^2, \quad (3.5)$$

Observing again that $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$ (from Lemma 3.4.2) and Eq. 3.4, we deduce

$$\|m_u - o_\tau\| \leq \|m_u - u\| \sin \angle(\text{aff}(\tau), T_u \mathcal{M}) \leq A C_1 \varepsilon^2 \text{lfs}(u). \quad (3.6)$$

We also have, $\|v - o_\tau\| \leq \|v - m_u\|$ as o_τ is the closest point to v in $\text{aff}(\text{Vor}^\omega(\tau))$. Hence we have, using Eq. (3.3) and (3.5),

$$\|m_v - o_\tau\| \leq \|v - o_\tau\| \tan \angle(\text{aff}(\tau), T_v \mathcal{M}) < \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \text{lfs}(u) \quad (3.7)$$

Let i_ϕ denote, as in Definition 3.2.1, the first point of the line segment $[m_u m_v]$ that is in $\text{Vor}^\omega(\phi)$. We get from Eq. (3.6) and (3.7) that

$$\|o_\tau - i_\phi\| \leq \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \text{lfs}(u).$$

1. Using Lemma 3.4.2, and the facts that $2A\varepsilon < 1$ and $\|u - o_\tau\| \leq \|u - m_u\|$, we get

$$\begin{aligned}
\|u - i_\phi\| &\leq \|u - o_\tau\| + \|o_\tau - i_\phi\| \\
&\leq \|u - m_u\| + \|o_\tau - i_\phi\| \\
&\leq \left(C_1\varepsilon + \frac{2AC_1\varepsilon^2}{\sqrt{1-4\omega_0^2}} \right) \text{lfs}(u) \\
&\leq \left(C_1\varepsilon + \frac{C_1\varepsilon}{\sqrt{1-4\omega_0^2}} \right) \text{lfs}(u) \\
&\leq \frac{C_2}{4} \varepsilon \text{lfs}(u)
\end{aligned} \tag{3.8}$$

where C_2 is the constant introduced in Lemma 3.4.3. Eq. (3.8), together with Lemma 2.4.3 and $C'_2 = C_2/\sqrt{1-4\omega_0^2}$, yields

$$\|p - i_\phi\| \leq \frac{\|u - i_\phi\|}{\sqrt{1-4\omega_0^2}} \leq \frac{C'_2}{4} \varepsilon \text{lfs}(u) \tag{3.9}$$

for all vertices p of ϕ . We deduce $\|p - u\| \leq \|p - i_\phi\| + \|u - i_\phi\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(u)$.

We now express $\text{lfs}(u)$ in terms of $\text{lfs}(p)$ using the fact that lfs is 1-Lipschitz and using $C'_2\varepsilon < 1$:

$$|\text{lfs}(p) - \text{lfs}(u)| \leq \|p - u\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(u) \leq \frac{1}{2} \text{lfs}(u). \tag{3.10}$$

We deduce that $\text{lfs}(u) \leq 2\text{lfs}(p)$ and $\|p - i_\phi\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(p)$

2. Using Eq. (3.9) and (3.10), from part 1 of this lemma, we have

$$\|p - q\| \leq \|p - i_\phi\| + \|q - i_\phi\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(u) \leq C'_2 \varepsilon \text{lfs}(p).$$

3. If σ belongs to $\text{Del}_{T\mathcal{M}}^\omega(P)$, the result has been proved in Lemma 3.4.3 (3) with $C'_3 = C_3$. Let $r_\phi = \sqrt{\|i_\phi - u\|^2 - \omega(u)^2}$. Since $i_\phi \in \text{Vor}^\omega(\sigma)$, the sphere centered at i_ϕ with radius r_ϕ is orthogonal to the weighted vertices of σ . From Lemmas 2.4.3 (2) and (3), we have $r_\phi \geq \Phi_\phi$ and $\Phi_\phi \geq \Phi_\sigma$ respectively. Hence it suffices to show that there exists a constant C'_3 such that $r_\phi \leq C'_3 L_\sigma$. Using Eq. (3.8), we get

$$r_\phi = \sqrt{\|i_\phi - u\|^2 - \omega(u)^2} \leq \|i_\phi - u\| \leq \frac{C_2}{4} \varepsilon \text{lfs}(u).$$

Let q be a vertex of a shortest edge of σ . We have, from part 2 of this lemma, $\|u - q\| \leq C'_2 \varepsilon \text{lfs}(q) < \text{lfs}(q)$. From which we deduce that $\text{lfs}(u) \leq 2\text{lfs}(q)$. Therefore, using Hypothesis 3.4.1,

$$\Phi_\sigma \leq r_\phi \leq \frac{C_2}{2} \varepsilon \text{lfs}(q) = \frac{C_2}{2} \left(\frac{\varepsilon}{\delta} \times \delta \right) \times \text{lfs}(q) \leq \frac{C_2}{2} \eta_0 L_\sigma.$$

From part 2 of the lemma we have $\Delta_\phi \leq 2C'_2 \varepsilon \text{lhs}(q)$, and using the facts that $L_\sigma \geq \delta \text{lhs}(q)$ and $\varepsilon/\delta \leq \eta_0$, this implies

$$\Gamma_\sigma = \frac{\Delta_\sigma}{L_\sigma} \leq \frac{\Delta_\phi}{L_\sigma} \leq 2C'_2 \eta_0 \leq C'_3 \stackrel{\text{def}}{=} \max\{4C_2 \eta_0, 2C'_2 \eta_0\}$$

4. Using the same arguments as in the proof of Lemma 3.4.3 (3) we have for all vertices q of σ_p

$$|H_\sigma(p, \omega(p))| \leq \|o_\sigma - p\|$$

Using the facts that $\Phi_\sigma \leq C'_3 L_\sigma$ (from part 3), $L_\sigma \leq \|p - q\|$, $\omega(q) \leq \omega_0 \|p - q\|$, and $\|p - q\| \leq 2C'_2 \varepsilon \text{lhs}(p)$ (from part 2), we have

$$\begin{aligned} \|o_\sigma - q\| &= \sqrt{\Phi_\sigma^2 + \omega(q)^2} \\ &\leq \sqrt{C_3'^2 L_\sigma^2 + \omega_0^2 \|p - q\|^2} \\ &\leq \|p - q\| \sqrt{C_3'^2 + \omega_0^2} \\ &\leq C'_2 \sqrt{C_3'^2 + \omega_0^2} \times \varepsilon \text{lhs}(p) \\ &\stackrel{\text{def}}{=} C'_4 \varepsilon \text{lhs}(p) \end{aligned}$$

□

The next crucial lemma bounds the fatness of inconsistent configurations.

Lemma 3.4.6 *Assume Hypotheses 2.4.1, 3.4.1 and 3.4.4, and $\varepsilon < 1/2C'_2$. The fatness Θ_ϕ of an inconsistent configuration ϕ is at most*

$$\frac{C'_2 \varepsilon}{(k+1)!} \left(1 + \frac{2C_3}{\Theta_0^k} \right)$$

Proof Let ϕ be witnessed by (u, v, w) . From the definition of inconsistent configurations, the k -dimensional simplex $\sigma = \phi \setminus \{w\}$ belongs to $\text{star}(u)$ and σ is a Θ_0 -fat simplex. Using Lemma 3.4.5 (2), we have $\Delta_\phi \leq 2C'_2 \varepsilon \text{lhs}(u)$. As in the proof of Lemma 3.4.5, we have $\sin \angle(T_u \mathcal{M}, \text{aff}(\sigma)) \leq \frac{\Gamma_\sigma \Delta_\sigma}{\Theta_\sigma \text{lhs}(u)}$ (refer to Eq. (3.4)) by Corollary 2.3.3 and the fact that $\Delta_\phi < \text{lhs}(u)$. Also, using the fact that $\|u - w\| < C'_2 \varepsilon \text{lhs}(u) < \text{lhs}(u)$ (from Lemma 3.4.3 (2) and $\varepsilon < 1/2C'_2$) and Lemma 2.2.1 (2), $\sin \angle(uw, T_u \mathcal{M}) \leq \frac{\|u-w\|}{2\text{lhs}(u)} \leq \frac{\Delta_\phi}{2\text{lhs}(u)}$. We can bound the altitude $D_\phi(w)$ of w in ϕ

$$\begin{aligned} D_\phi(w) &= \text{dist}(w, \text{aff}(\sigma)) \\ &= \sin \angle(uw, \text{aff}(\sigma)) \times \|u - w\| \\ &\leq (\sin \angle(uw, T_u \mathcal{M}) + \sin \angle(\text{aff}(\sigma), T_u \mathcal{M})) \times \Delta_\phi \\ &\leq \left(\frac{\Delta_\phi}{2\text{lhs}(u)} + \frac{\Gamma_\sigma \Delta_\sigma}{\Theta_\sigma \text{lhs}(u)} \right) \Delta_\phi \\ &\leq \frac{\Delta_\phi^2}{2\text{lhs}(u)} \left(1 + \frac{2\Gamma_\sigma}{\Theta_\sigma} \right). \end{aligned} \tag{3.11}$$

From the definition of fatness of a simplex and Lemma 2.3.1 (1), we get

$$\text{vol}(\sigma) = \Theta_\sigma \Delta_\sigma^k \leq \frac{\Delta_\sigma^k}{k!}. \quad (3.12)$$

We deduce

$$\begin{aligned} \Theta_\phi &= \frac{\text{vol}(\phi)}{\Delta_\phi^{k+1}} \\ &= \frac{D_\phi(w) \text{vol}(\sigma)}{(k+1)} \times \frac{1}{\Delta_\phi^{k+1}} \\ &\leq \frac{\Delta_\phi^2}{2\text{lfs}(u)} \left(1 + \frac{2\Gamma_\sigma}{\Theta_\sigma}\right) \times \frac{\Delta_\sigma^k}{(k+1)! \Delta_\phi^{k+1}} \quad \text{using Eq. (3.11) and (3.12)} \\ &\leq \frac{C'_2 \varepsilon}{(k+1)!} \left(1 + \frac{2C_3}{\Theta_0^k}\right). \end{aligned}$$

The last inequality comes from the facts that σ is Θ_0 -fat, $\Delta_\phi \leq 2C'_2 \varepsilon \text{lfs}(u)$ (from Lemma 3.4.5 (2)) and $\Gamma_\tau \leq C_3$ (from Lemma 3.4.3 (3)). $\square$

A consequence of the lemma is that, if the subfaces of ϕ are Θ_0 -fat simplices and if the following hypothesis

Hypothesis 3.4.7 $\frac{C'_2 \varepsilon}{(k+1)!} \left(1 + \frac{2C_3}{\Theta_0^k}\right) < \Theta_0^{k+1}$

is satisfied, then ϕ is a Θ_0 -sliver. Hence, techniques to remove slivers can be used to remove inconsistent configurations.

In the above lemmas, we assumed that ε is small enough. Specifically in addition to Hypotheses 2.4.1, 3.4.1 and 3.4.4, we assumed that $2C_2 \varepsilon < 1$ in Lemma 3.4.3, $C'_2 \varepsilon < 1$ in Lemma 3.4.5 and $2C'_2 \varepsilon < 1$ in Lemma 3.4.6. We will make another hypothesis that subsumes these two previous conditions.

Hypothesis 3.4.8 $C'_2(1 + C'_2 \eta_0) \varepsilon < 1/2$.

Observe that this hypothesis implies $C'_2(1 + C'_2) \varepsilon < 1/2$ since $\eta_0 > 1$.

3.4.3 Number of local neighbors

We will use the result from this section for the analysis of the algorithm, and also for calculating its time and space complexity.

Let $N \stackrel{\text{def}}{=} (4C'_2 \eta_0 + 6)^k$, where the constant C'_2 is defined in Lemma 3.4.5.

Lemma 3.4.9 Assume Hypotheses 2.4.1, 3.4.1, 3.4.4 and 3.4.8. The set

$$LN(p) = \{q \in P : |B(p, \|p - q\|) \cap P| \leq N\},$$

where $N = 2^{O(k)}$ and the constant in big-O depends on ω_0 and η_0 , includes all the points of P that can form an edge with p in $C^\omega(P)$.

Proof Lemmas 3.4.3 and 3.4.5 show that, in order to construct $\text{star}(p)$ and search for inconsistencies involving p , it is enough to consider the points of P that lie in ball $B_p = B(p, C'_2 \varepsilon \text{lfs}(p))$. Therefore it is enough to count the number of points in $B_p \cap P$.

Let x and y be two points of $B_p \cap P$. Since $\text{lfs}()$ is a 1-Lipschitz function, we have

$$\text{lfs}(p)(1 - C'_2 \varepsilon) \leq \text{lfs}(x), \text{lfs}(y) \leq \text{lfs}(p)(1 + C'_2 \varepsilon). \quad (3.13)$$

By definition of an (ε, δ) -lfs sample of $\mathcal{M}$, the two balls $B_x = B(x, r_x)$ and $B_y = B(y, r_y)$, where $r_x = \delta \text{lfs}(x)/2$ and $r_y = \delta \text{lfs}(y)/2$, are disjoint. Moreover, both balls are contained in the ball $B_p^+ = B(p, r^+)$, where $r^+ = C'_2 \varepsilon \text{lfs}(p) + (1 + C'_2 \varepsilon) \delta \text{lfs}(p)$.

Let $\bar{B}_x = B_x \cap T_p \mathcal{M}$, $\bar{B}_y = B_y \cap T_p \mathcal{M}$ and $\bar{B}_p^+ = B_p^+ \cap T_p \mathcal{M}$. From Lemma 2.2.1 (2), the distance from x to $T_p \mathcal{M}$ is

$$\text{dist}(x, T_p \mathcal{M}) = \|p - x\| \times \sin(\angle px, T_p \mathcal{M}) \leq C_2'^2 \varepsilon^2 \text{lfs}(p)/2. \quad (3.14)$$

Using Eq.s (3.13), (3.14) and the fact that $\varepsilon/\delta \leq \eta_0$, we see that $\bar{B}_x$ is a k -dimensional ball of squared radius

$$\begin{aligned} \delta^2 \text{lfs}^2(x)/4 - \text{dist}(x, T_p \mathcal{M})^2 &\geq \delta^2 \text{lfs}^2(p)(1 - C'_2 \varepsilon)^2/4 - C_2'^4 \varepsilon^4 \text{lfs}(p)/4 \\ &\geq \delta^2 \text{lfs}^2(p)/4 \times \left((1 - C'_2 \varepsilon)^2 - C_2'^4 \eta_0^2 \varepsilon^2 \right) \stackrel{\text{def}}{=} (r^-)^2. \end{aligned}$$

We can now use a packing argument. Since the balls $\bar{B}_x$, x in $B_p \cap P$, are disjoint and all contained in B_p^+ , the number of points of $B_p \cap P$ is at most

$$\begin{aligned} \left(\frac{r^+}{r^-} \right)^k &= \left(\frac{(C'_2 \varepsilon + (1 + C'_2 \varepsilon) \delta)^2}{\delta^2/4 \times \left((1 - C'_2 \varepsilon)^2 - C_2'^4 \eta_0^2 \varepsilon^2 \right)} \right)^{k/2} \\ &\leq \left(\frac{4(C'_2 \eta_0 + (1 + C'_2 \varepsilon))^2}{(1 - C'_2 \varepsilon)^2 - C_2'^4 \eta_0^2 \varepsilon^2} \right)^{k/2} \\ &= \left(\frac{4(C'_2 \eta_0 + (1 + C'_2 \varepsilon))^2}{(1 - C'_2 \varepsilon - C_2'^2 \eta_0 \varepsilon)(1 - C'_2 \varepsilon + C_2'^2 \eta_0 \varepsilon)} \right)^{k/2} \\ &\leq (4C'_2 \eta_0 + 6)^k \stackrel{\text{def}}{=} N \quad (\text{using Hypothesis 3.4.8}) \end{aligned}$$

And the result follows. $\square$

3.4.4 Correctness of the algorithm, and theoretical guarantees

Definition 3.4.10 (Sliverity range) Let ω be a weight assignment satisfying Hypothesis 2.4.1. The weight of all the points in $P \setminus \{p\}$ are fixed and the weight $\omega(p)$ of p is varying. The sliverity range $\Sigma(p)$ of a point $p \in P$ is the measure of the set of all squared weights $\omega(p)^2$ for which p is a vertex of a Θ_0 -sliver in $C^\omega(P)$.

Lemma 3.4.11 Under Hypotheses 2.4.1, 3.4.1, 3.4.4, 3.4.7 and 3.4.8, the sliverity range satisfies

$$\Sigma(p) < 2N^{k+1} C_5 \Theta_0 \text{nn}(p)^2$$

for some constant C_5 that depends on k , ω_0 and η_0 but not on Θ_0 .

Proof Let σ be a j -dimensional simplex of $C^\omega(P)$ incident on p (with $2 \leq j \leq k+1$). assume that σ is a Θ_0 -sliver. If $\omega(p)$ is the weight of p , we write $H_\sigma(p, \omega(p))$ for the excentricity of σ with respect to p and parameterized by $\omega(p)$. From Lemma 3.4.5(4), we have

$$|H_\sigma(p, \omega(p))| \leq C'_4 \varepsilon \text{fs}(p) \stackrel{\text{def}}{=} D \quad (3.15)$$

Using Lemma 2.3.1 (2), we have

$$D_p(\sigma) \leq j \Gamma_\sigma^{j-1} \Delta_\sigma \times \frac{\Theta_\sigma}{\Theta_{\sigma_p}} \leq (k+1) C_3'^k \Theta_0 \Delta_\sigma \stackrel{\text{def}}{=} E \quad (3.16)$$

The last inequality follows from the facts that $j \leq k+1$, $\Gamma_\sigma \leq C_3'$ (from Lemmas 3.4.3 (2) and 3.4.5 (3)) and σ is a Θ_0 -sliver. Moreover, from Lemma 2.4.4,

$$H_\sigma(p, \omega(p)) = H_\sigma(p, 0) - \frac{\omega(p)^2}{2D_p(\sigma)}. \quad (3.17)$$

It then follows from Eq. (3.15), (3.16) and (3.17) that the set of squared weights of p for which σ belongs to $C^\omega(P)$ is a subset of the following interval

$$[2D_p(\tau)H_\sigma(p, 0) - \beta, 2D_p(\tau)H_\sigma(p, 0) + \beta],$$

where $\beta = 2DE$. Therefore, from Eq. (3.16) and (3.17), the measure of the set of weights for which σ belongs to $C^\omega(P)$ is at most

$$2\beta = 4DE = 4(k+1)C_3'^k \Theta_0 \Delta_\sigma C'_4 \varepsilon \text{fs}(p).$$

Let q_1 and q_2 be two vertices of τ such that $\Delta_\sigma = \|q_1 - q_2\|$. Using Lemmas 3.4.3 (1) and 3.4.5 (2), we get

$$\Delta_\sigma \leq \|p - q_1\| + \|p - q_2\| \leq 2C_2' \varepsilon \text{fs}(p).$$

Using this inequality, $\text{fs}(p) \leq \text{nn}(p)/\delta$ (Lemma 2.2.1) and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 3.4.1), the sliverity range of σ is at most $8(k+1)C_3'^k C_2' C_4' \Theta_0 \eta_0^2 \text{nn}(p)^2 = C_5 \Theta_0 \text{nn}(p)^2$ with $C_5 \stackrel{\text{def}}{=} 8(k+1)C_3'^k C_2' C_4' \eta_0^2$. By Lemma 3.4.9, the number of j -simplices that are incident to p is at most N^j . Hence, the sliverity range of p is at most

$$\Sigma(p) \leq \sum_{j=3}^{k+1} N^j C_5 \Theta_0 \text{nn}(p)^2 < 2N^{k+1} C_5 \Theta_0 \text{nn}(p)^2.$$

The last inequality follows from the fact that $\sum_{j=3}^{k+1} N^j < 2N^{k+1}$ as $N > 2$. $\square$

Theorem 3.4.12 *Let P be an (ε, δ) -lfs sample of $\mathcal{M}$, $\varepsilon/\delta \leq \eta_0$ and $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$. If Hypotheses 2.4.1, 3.4.1, 3.4.4, 3.4.7 and 3.4.8 are satisfied, the simplicial complex $\hat{\mathcal{M}}$ output by Algorithm 1 has no inconsistencies and its simplices are all Θ_0 -fat.*

Proof The sliverity range $\Sigma(p)$ of p is at most $2N^{k+1} C_5 \Theta_0 \text{nn}(p)^2$ from Lemma 3.4.11. Since $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$, $\Sigma(p)$ is less than the total range of possible squared weights $\omega_0^2 \text{nn}(p)^2$.

Hence, Function **weight** (p, ω) will always find a weight for any point $p \in P$ and any weight assignment of relative amplitude at most ω_0 for the points of $P \setminus \{p\}$.

Since the algorithm removes all the simplices of $C^\omega(P)$ that are not Θ_0 -fat, all the simplices of $\hat{M}$ are Θ_0 -fat.

By Lemma 3.4.6 and Hypothesis 3.4.7, all inconsistent configurations in $C^\omega(P)$ are Θ_0 -slivers. It follows that $\hat{M}$ has no inconsistency since, when the algorithm terminates, all simplices of $C^\omega(P)$ are Θ_0 -fat. $\square$

Using Theorem 4.0.3 from Chapter 4, we have the following guarantees on the output $\hat{M}$ of the algorithm.

Theorem 3.4.13 (Isotopy and geometric approximation) *Let P be an (ε, δ) -lfs sample of M , $\varepsilon/\delta \leq \eta_0$ and $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$. If Hypotheses 2.4.1, 3.4.1, 3.4.4, 3.4.7 and 3.4.8 are satisfied, the simplicial complex $\hat{M}$ output by Algorithm 1 satisfy the following geometric and topological properties:*

1. $\hat{M}$ is a piecewise linear (PL) k -manifold without boundary;
2. Let τ be a k -simplex in $\hat{M}$. For all vertices p of τ , we have $\sin \angle(T_p M, \text{aff}(\tau)) = \sin \angle(\text{aff}(\tau), T_p M) = O(\varepsilon)$, where the constant in the big-O depends on k , η_0 , ω_0 , and Θ_0 ;
3. There exists a homeomorphism $f : \hat{M} \rightarrow M$ between $\hat{M}$ and M ;
4. There exists an isotopy $F : \hat{M} \times [0, 1] \rightarrow \mathbb{R}^d$ such that the map $F(\cdot, 0)$ restricted to $\hat{M}$ is the identity map on $\hat{M}$ and $F(\hat{M}, 1) = M$;
5. For all x in M , $\text{dist}(x, f^{-1}(x)) = O(\varepsilon^2 \text{lfs}(x))$ where the constant in the big-O depends on k , η_0 , ω_0 , and Θ_0 .

Proof From Theorem 3.4.12, we know that $\hat{M}$ output by the algorithm will have no inconsistencies and all the simplices in $\hat{M}$ will be Θ_0 -fat.

This implies the conditions **C1** and **C2** of Theorem 4.0.3 are satisfied. Therefore from Theorem 4.0.3, if ε is sufficiently small, we get our result. $\square$

3.4.5 Time and space complexity

Theorem 3.4.14 *Assume that Hypotheses 2.4.1, 3.4.1, 3.4.4, 3.4.7 and 3.4.8 are satisfied. Then the space complexity of the algorithm is*

$$(O(d) + 2^{O(k^2)})|P|$$

and the time complexity is

$$O(d)|P|^2 + d 2^{O(k^2)}|P|.$$

Proof **1. Space Complexity :** For each point $p \in P$ we maintain $LN(p)$. The total space complexity for storing $LN(p)$ for each point $p \in P$ is thus $O(N|P|)$ by definition of $LN(p)$.

By Lemma 2.4.2, each $\text{star}(p)$, $p \in P$, has the same combinatorial complexity as a Voronoi cell in the k -dimensional flat $T_p M$. Since the sites needed to compute this

Voronoi cell all belong to $LN(p)$, there number is at most N by Lemma 3.4.9. From the Upper Bound Theorem of convex geometry, see e.g. [BY98], the combinatorial complexity of each star is therefore $O(N^{\lfloor k/2 \rfloor})$. Hence the total space complexity of the tangential Delaunay complex is $O(kN^{\lfloor k/2 \rfloor})|P|$.

For a given inconsistent Θ_0 -fat k -simplex in $\text{star}(p)$, we can have from Lemmas 3.4.5 (2) and 3.4.9, at most $k|LN(p)| \leq kN$ different inconsistent configurations. Hence, the number of inconsistent configurations to be stored in the completed complex $C^\omega(P)$ is at most $O(k^2 N^{\lfloor k/2 \rfloor + 1})|P|$.

With $N = O(2^k)$ (refer to Lemma 3.4.9), we conclude that the total space complexity of the algorithm is

$$O(k^2 N^{\lfloor k/2 \rfloor + 1} + d)|P| = (O(d) + 2^{O(k^2)})|P|.$$

2. Time complexity : In the initialization phase, the algorithm computes $LN(p)$ for all $p \in P$ and initializes the weights to 0. This can easily be done in time $O(d)|P|^2$.

Then the algorithm builds $C^\omega(P)$ for the zero weight assignment. The time to compute $\text{star}(p)$ is dominated by the time to compute the cell of p in the weighted k -dimensional Voronoi diagram of the projected points of $LN(p)$ onto $T_p\mathcal{M}$. Since, by definition, $|LN(p)| \leq N$, the time for building the star of p is the same as the time to compute the intersection of N halfspaces in $\mathbb{R}^k$,

$$O(kdN + k^3(N \log N + N^{\lfloor k/2 \rfloor})),$$

see e.g. [Cha93, BY98]. The factor $O(kd)$ appears in the first term because to calculate the projection of a point in $\mathbb{R}^d$ on a k -flat we have to do k inner products. The $O(k^3)$ factor comes from the fact that the basic operation we need to perform is to decide whether a point lies in the ball orthogonal to a k -simplex. This operation reduces to the evaluation of the sign of the determinant of a $(k+2) \times (k+2)$ matrix. The $N^{\lfloor k/2 \rfloor}$ term bounds the combinatorial complexity of a cell in the Voronoi diagram of N sites in a k -flat. Therefore the time needed to build the stars of all the points p in P is $O(kdN + k^3(N \log N + N^{\lfloor k/2 \rfloor}))|P|$.

Let $\sigma = [p_0, \dots, p_k]$ be a Θ_0 -fat k -simplex in $\text{star}(u)$. For each vertex $v (\neq u)$ of σ with $\sigma \notin \text{star}(v)$, we need to compute the inconsistent configurations of the form $\phi = [p_0, \dots, p_k, w]$ witnessed by (u, v, w) where $w \in LN(p) \setminus \sigma$. The number of such inconsistent configurations is therefore less than $|LN(p)| \leq N$. The time complexity to compute all the inconsistent configurations of the form $\phi = [p_0, \dots, p_k, w]$ witnessed by the triplet (u, v, w) is $O(dN)$. Since the number of choices of v is at most k , hence the time complexity for building all the inconsistent configurations of the form $\phi = [p_0, \dots, p_k, w]$ witnessed by (u, v, w) with $v (\neq u)$ being a vertex of σ and w a point in $LN(u) \setminus \sigma$ is

$$O(dkN) \tag{3.18}$$

The time complexity to build all the inconsistent configurations corresponding to $\text{star}(u)$ is $O(dkN^{\lfloor k/2 \rfloor + 1})$ since the number of k -simplices in the star of a point p is $O(N^{\lfloor k/2 \rfloor})$.

Hence, the time complexity for building the inconsistent configurations of $C^\omega(P)$ is $O(dN + kN^{\lfloor k/2 \rfloor + 1})|P|$. Therefore the total time complexity of the initialization phase is

$$O(dkN + k^3 N \log N + (dk + k^3)N^{\lfloor k/2 \rfloor + 1})|P|$$

Consider now the main loop of the algorithm. The time complexity of function **weight**(p, ω) is $O((d + k^3)N^{k+1})$ since we need to sweep over at most all $(k+1)$ -simplices

incident on p with vertices in $LN(p)$. The number of such simplices is at most N^{k+1} . We easily deduce from the above discussion that the time complexity of Function **update_complete_complex** $(LN(p), \omega)$ is

$$O(dkN + k^3N \log N + (dk + k^3)N^{\lfloor k/2 \rfloor + 1})N.$$

Since functions **weight** (p, Θ_0, ω) and **update_complete_complex** $(C^\omega(P), p, \omega)$ are called $|P|$ times, we conclude that the time complexity of the main loop of the algorithm is $O(dkN^2 + (k^3 + dk)N^{k+1})|P|$.

Combining the time complexities for all the steps of the algorithm and using $N = O(2^k)$ (refer to Lemma 3.4.9), we get the total time complexity of the algorithm

$$O(d)|P|^2 + O(dkN^2 + (dk + k^3)N^{k+1})|P| = O(d)|P|^2 + d2^{O(k^2)}|P|.$$

□

Observe that, since P is an (ε, δ) -lfs sample of $\mathcal{M}$ with $\varepsilon/\delta \leq \eta_0$, $|P| = O(\varepsilon^k)$.

3.5 Summary

We have given the first algorithm that is able to reconstruct a smooth closed manifold in a time that depends only linearly on the dimension of the ambient space. We believe that our algorithm is of practical interest when the dimension of the manifold is small, even if it is embedded in a space of high dimension. This situation is quite common in practical applications in machine learning. Unlike most surface reconstruction algorithms in $\mathbb{R}^3$, our algorithm does not need to orient normals (a critical issue in practical applications) and, in fact, works for non orientable manifolds. Note that the cocone algorithm [ACDL02b] for surface reconstruction does not require oriented normals.

The algorithm is simple. The basic ingredients we need are data structures for constructing weighted Delaunay triangulations in $\mathbb{R}^k$. We have assumed that the intrinsic dimension of $\mathcal{M}$ and tangent space at each sample point is known. If not, we can use algorithms given in [CWW08, GW04a] to estimate the dimension of $\mathcal{M}$ and the tangent space at each sample point. Moreover, our algorithm is easy to parallelize. One interesting feature of our approach is that it is robust and still works if we only have approximate tangent spaces at the sample points. We will report on experimental results in a forthcoming paper.

We have assumed that we know an upper bound on the sampling ratio η_0 of the input sample P . Ideas from [BGO09, FR02] may be useful to convert a sample to a subsample with a bounded sampling ratio.

We foresee other applications of the tangential complex and of our construction each time computations in the tangent space of a manifold are required, e.g. for dimensionality reduction and approximating the Laplace Beltrami operator [BSW08]. It easily follows from [BNN10] that our reconstruction algorithm can also be used in Bregman spaces where the Euclidean distance is replaced by any Bregman divergence, e.g. Kullback-Leibler divergence. This is of particular interest when considering statistical manifolds like, for example, spaces of images [CIdSZ08].

Chapter 4

Topological and geometric guarantees

In this chapter we give conditions under which $\text{Del}_{T\mathcal{M}}^\omega(P)$ is isotopic to and a close geometric approximation of the manifold $\mathcal{M}$. The results in this chapter will be used to prove the theoretical guarantees on the quality of the output by different algorithms in this thesis.

Sampling and general position assumption. For the rest of this chapter, we will assume that $\mathcal{M}$ is a compact smooth submanifold of $\mathbb{R}^d$ without boundaries, and $P \subset \mathcal{M}$ is an (ε, δ) -lfs sample of $\mathcal{M}$ with $\varepsilon < 0.09$ and $\varepsilon/\delta \leq \eta_0$. Additionally we also assume that $\omega : P \rightarrow [0, \infty)$ be a weight assignment with $\tilde{\omega} \leq \omega_0 \in [0, 1/2)$.

For $p \in P$, let $P_p \stackrel{\text{def}}{=} \{p'_0, \dots, p'_n\}$ where p'_j denotes the orthogonal projection of $p_j \in P$ onto $T_p\mathcal{M}$. Note that $p' = p$. We define

$$\xi_p : P_p \rightarrow [0, \infty), \text{ with } \xi_p(p'_j)^2 \stackrel{\text{def}}{=} \omega(p_j)^2 - \|p_j - p'_j\|^2 + \lambda^2$$

where $\lambda = \max_j \|p_j - p'_j\|$.

From Lemma 2.4.2, we get $\text{Vor}^\omega(P) \cap T_p\mathcal{M}$ to be identical to the weighted Voronoi diagram $\text{Vor}^{\xi_p}(P_p)$ of P_p in $T_p\mathcal{M}$.

Since $\varepsilon < 0.09$, $\text{Vor}^{\xi_p}(p)$ is bounded from Lemmas 2.4.2 and 3.4.2, and this implies $\dim \text{aff}(P_p) = k$. Let the weighted points $P_p^{\xi_p}$ be in *general position* on $T_p\mathcal{M}$, i.e., there exists no orthosphere centered on $T_p\mathcal{M}$ that is orthogonal to $k + 2$ weighted points of $P_p^{\xi_p}$. See Sections 2.4. As we have already discussed in Section 3.2, we can guarantee that the general position assumption is satisfied by infinitesimal perturbation of the point sample or the weight assignment ω .

General position assumption, together with $\dim \text{aff}(P_p) = k$, implies that $\text{Del}^{\xi_p}(P_p)$ is a triangulation of $\text{conv}(P_p)$. See, [Aur87, AE84]. For the rest of this chapter we will assume $\text{Del}^{\xi_p}(P_p)$ is a triangulation of $\text{conv}(P_p)$. Since $\text{Vor}^{\xi_p}(p)$ is bounded (from Lemma 3.4.2), p is an interior vertex of $\text{Del}^{\xi_p}(P_p)$. This implies the following result.

Lemma 4.0.1 *For all $p \in P$, $\text{star}(p)$ is isomorphic to the star of an interior vertex of a triangulated k -dimensional convex domain.*

Medial axis and the projection map. Let $\mathcal{O}$ denote the medial axis of $\mathcal{M}$, and let

$$\pi : \mathbb{R}^d \setminus \mathcal{O} \longrightarrow \mathcal{M}$$

denote the projection map that maps each point of $\mathbb{R}^d \setminus \mathcal{O}$ to its closest point on $\mathcal{M}$. The following lemma is a standard result from Federer [Fed69].

Lemma 4.0.2 *Let $\mathcal{M}$ be a smooth submanifold of $\mathbb{R}^d$ without boundary. Then, the projection map $\pi : \mathbb{R}^d \setminus \mathcal{O} \rightarrow \mathcal{M}$ is a C^1 -function in $\mathbb{R}^d \setminus \mathcal{O}$.*

1. *The map π is a C^1 -function.*
2. *For all $x \in \mathbb{R}^d \setminus \mathcal{O}$, the kernel of the linear map $d\pi(x) : \mathbb{R}^d \rightarrow T_{\pi(x)}\mathcal{M}$, where $d\pi(x)$ denotes the derivative of π at x , is parallel to $N_{\pi(x)}\mathcal{M}$ and has dimension $d - k$.*

Additional notations. Let $T_p\mathcal{M}^{1/16} = B(p, \frac{\text{fs}(p)}{16}) \cap T_p\mathcal{M}$. We define the map

$$\widetilde{\pi}_p : T_p\mathcal{M}^{1/16} \longrightarrow \mathcal{M}$$

as the restriction of the projection map π to $T_p\mathcal{M}^{1/16}$.

We define in addition

$$\pi_p^* : \mathbb{R}^d \setminus \mathcal{O} \longrightarrow T_p\mathcal{M}$$

as the map that maps each point $x \in \mathbb{R}^d \setminus \mathcal{O}$ to the point of intersection of $T_p\mathcal{M}$ and $N_{\pi(x)}\mathcal{M}$, the normal space of $\mathcal{M}$ at $\pi(x)$.

Main result. The main result of this chapter is the following theorem.

Theorem 4.0.3 *Let the weight assignment $\omega : \mathcal{P} \rightarrow [0, \infty)$, with $\tilde{\omega} \leq \omega_0 \in [0, 1/2)$, satisfy the following conditions:*

C1. $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ has no k -dimensional inconsistent simplices.

C2. All the simplices in $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ are Θ_0 -fat.

Then there exists ε_0 that depends only on k , ω_0 , η_0 , and Θ_0 such that for all $\varepsilon \leq \varepsilon_0$, the simplicial complex $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, also noted $\hat{\mathcal{M}}$ for convenience, and the map $\pi|_{\hat{\mathcal{M}}}$ satisfy the following properties:

P1. *Tangent space approximation: Let τ be a k -simplex in $\hat{\mathcal{M}}$. For all vertices p of τ , we have $\sin \angle(T_p\mathcal{M}, \text{aff}(\tau)) = \sin \angle(\text{aff}(\tau), T_p\mathcal{M}) = O(\varepsilon)$.*

P2. *PL k -manifold without boundary: $\hat{\mathcal{M}}$ is a piecewise linear k -manifold without boundary;*

P3. *Homeomorphism: The map $\pi|_{\hat{\mathcal{M}}}$ provides a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$;*

P4. *Pointwise approximation: $\forall x \in \mathcal{M}$, $\text{dist}(x, \pi|_{\hat{\mathcal{M}}}^{-1}(x)) = O(\varepsilon^2 \text{fs}(x))$;*

P5. *Isotopy: There exists an isotopy $F : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ such that the map $F(\cdot, 0)$ restricted to $\hat{\mathcal{M}}$ is the identity map on $\hat{\mathcal{M}}$ and $F(\hat{\mathcal{M}}, 1) = \mathcal{M}$.*

The constants in the big- O notations depend on k , ω_0 , η_0 and Θ_0 .

The rest of this chapter will be devoted to the proof of Theorem 4.0.3. Some definitions and results from topology we use in this chapter are recalled in Section 2.5.

4.1 Tangent space approximation

The property **P1** of Theorem 4.0.3 follows directly from Corollary 2.3.3.

Lemma 4.1.1 *Under conditions **C1** and **C2** of Theorem 4.0.3, let τ be a k -simplex in $\hat{\mathcal{M}}$. For all vertices p of τ , we have $\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) = \sin \angle(T_p \mathcal{M}, \text{aff}(\tau)) \leq A\varepsilon$. The constant A is defined in Hypothesis 3.4.4 and depends on k , ω_0 , η_0 and Θ_0 .*

Proof Since $\dim(\tau) = \dim(T_p \mathcal{M}) (= k)$, we have from Lemma 2.1.1, $\angle(\text{aff}(\tau), T_p \mathcal{M}) = \angle(T_p \mathcal{M}, \text{aff}(\tau))$.

All the simplices in $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ are Θ_0 -fat (condition **C2**), hence $\Theta_\tau \geq \Theta_0$. Using Corollary 2.3.3, and the facts that $\tau \in \text{star}(p)$ (condition **C1**), $\Delta_\tau \leq 2C_2\varepsilon \text{lfs}(p)$ (from Lemma 3.4.3 (1) and the Triangle inequality) and $\Gamma_\tau \leq C_3$ (from Lemma 3.4.3 (2)), we get

$$\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq \frac{\Gamma_\tau \Delta_\tau}{\Theta_0^k \text{lfs}(p)} \leq \frac{2C_2 C_3 \varepsilon}{\Theta_0^k} = A\varepsilon$$

□

4.2 Piecewise-linear k -manifold

For a given simplicial complex $\mathcal{K}$ and a simplex σ in $\mathcal{K}$, the *star* and *link* of a simplex σ in $\mathcal{K}$ $\sigma \in \mathcal{K}$ denotes the subcomplexes

$$\text{st}(\sigma, \mathcal{K}) = \{\tau : \text{for some } \tau_1 \in \mathcal{K}, \sigma, \tau \subseteq \tau_1\}$$

and

$$\text{lk}(\sigma, \mathcal{K}) = \{\tau \in \text{st}(\sigma, \mathcal{K}) : \tau \cap \sigma = \emptyset\}$$

respectively. The *open star* of σ in $\mathcal{K}$, $\mathring{\text{st}}(\sigma, \mathcal{K}) = \text{st}(\sigma, \mathcal{K}) \setminus \text{lk}(\sigma, \mathcal{K})$.

The following lemma is a direct consequence of condition **C1**.

Lemma 4.2.1 *Under condition **C1** of Theorem 4.0.3 and for a sufficiently small ε , $\text{star}(p) = \text{st}(p, \hat{\mathcal{M}})$ for all $p \in \mathcal{P}$.*

Proof Refer to the discussion on general position assumption at the beginning of this chapter. From Lemma 3.4.2, we get $\dim \text{aff}(\mathcal{P}_p) = k$. Moreover, since $\dim \text{aff}(\mathcal{P}_p) = k$ and $\text{Del}^{\xi_p}(\mathcal{P}_p)$ is a triangulation of $\text{conv}(\mathcal{P}_p)$, the maximal dimension of the simplices in $\text{st}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$ is k . This implies that dimension of maximal simplices in $\text{star}(p)$ is k . Since we have no k -dimensional inconsistent simplices in $\hat{\mathcal{M}}$, from condition **C1** of Theorem 4.0.3, we get $\text{star}(p) = \text{st}(p, \hat{\mathcal{M}})$. □

We call a simplicial complex $\mathcal{K}$ a PL $(j-1)$ -sphere (or j -ball) if there is a PL homeomorphism from $\mathcal{K}$ to $\partial \Delta^j$ (or Δ^j) where Δ^j denotes the regular j -dimensional simplex. We will use the following simple results of PL topology. See, Zeeman's Seminars in Combinatorial Topology [Zee66, Lem. 8 & 9 of Chap. 3].

Lemma 4.2.2 1. *Let $S \subset \mathbb{R}^m$ be a finite point set with $\dim \text{aff}(S) = j$, then $\text{conv}(S)$ is a PL j -ball.*

2. A simplicial complex $\mathcal{K} \subset \mathbb{R}^m$ is a PL j -manifold iff for all vertices v of $\mathcal{K}$, $\text{lk}(v, \mathcal{K})$ is either a PL $(j-1)$ -sphere or PL $(j-1)$ -ball depending on whether v belongs to $\int \mathcal{K}$ or $\partial \mathcal{K}$ respectively.

The following result is a direct consequence of the results from [Aur87] and Lemma 4.2.2.

Lemma 4.2.3 *Let $S \subset \mathbb{R}^j$ be a finite point set with $\dim \text{aff}(S) = j$, and $\omega : S \rightarrow [0, \infty)$ be a weight assignment. Additionally also assume that the weighted points S^ω are in general position (see the discussion on general position assumption given at the beginning of this chapter), i.e., there exists no sphere orthogonal to $j+2$ weighted points of S^ω . If v is a vertex of $\text{Del}^\omega(S)$ and v does not lie on the boundary of $\text{conv}(S)$, then $\text{lk}(v, \text{Del}^\omega(S))$ is a PL $(j-1)$ -sphere.*

Proof Since $\dim \text{aff}(S) = j$, $\text{conv}(S)$ is a PL j -ball according to Lemma 4.2.2. From [Aur87] and the fact that the weighted points S^ω are in general position, we get that $\text{Del}^\omega(S)$ is a triangulation of $\text{conv}(S)$. Therefore as the vertex v of $\text{Del}^\omega(S)$ is not on the boundary of $\text{conv}(S)$, $\text{lk}(v, \text{Del}^\omega(S))$ is a PL $(j-1)$ -sphere. $\square$

We can now prove that $\hat{\mathcal{M}}$ is a PL k -manifold without boundary (Property P2).

Lemma 4.2.4 (PL k -manifold without boundary) *Assume condition C1 of Theorem 4.0.3. If ε is sufficiently small, then $\hat{\mathcal{M}}$ is a PL k -manifold without boundary.*

Proof From Lemma 4.2.1, we have $\text{star}(p) = \text{st}(p, \hat{\mathcal{M}})$. By Lemma 4.0.1, $\text{st}(p, \hat{\mathcal{M}})$ is isomorphic to the star of an interior vertex of a k -dimensional triangulated convex domain. Hence, $\text{lk}(p, \hat{\mathcal{M}})$ is a PL $(k-1)$ -sphere from Lemma 4.2.3. This implies that $\hat{\mathcal{M}}$ is a PL k -manifold with no boundary by Lemma 4.2.2 (2). $\square$

Property P2 follows since the output $\hat{\mathcal{M}}$ of the Algorithm 1 has no inconsistency and therefore satisfies Condition C1.

4.3 Homeomorphism

The next lemma establishes Property P3 of Theorem 4.0.3.

Lemma 4.3.1 (Homeomorphism) *Assume conditions C1 and C2 of Theorem 4.0.3. For ε sufficiently small, the map π restricted to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$.*

Outline of the proof. The proof of the lemma is quite long and technical. Before we give the full details, we want to give a brief outline of the proof. The proof follows the following steps:

- S1.** We prove in Lemma 4.3.15 that the map π_p^* restricted to the open star $\text{st}(p, \hat{\mathcal{M}})$ of p is injective. Then using Lemma 4.3.15, we will show in Lemma 4.3.16 that the map π restricted to the open star $\text{st}(p, \hat{\mathcal{M}})$ of p is injective.
- S2.** We show in Lemma 4.3.17 that for all p in $\mathcal{P}$, $\pi|_{\hat{\mathcal{M}}}^{-1}(p) = \{p\}$.
- S3.** In Lemma 4.3.20, we prove that for all connected components $\mathcal{M}_i$ of $\mathcal{M}$, $\mathcal{P} \cap \mathcal{M}_i \neq \emptyset$.

S4. We show that $\pi|_{\hat{\mathcal{M}}}$ is an *open map* in Lemma 4.3.22, and, in Lemma 4.3.23, we show that, for each *connected component* $\hat{\mathcal{M}}_i$ of $\hat{\mathcal{M}}$, there exists a connected component $\mathcal{M}_j$ of $\mathcal{M}$ such that $\pi : \hat{\mathcal{M}}_i \rightarrow \mathcal{M}_j$ is surjective and a *covering map* of $\mathcal{M}_j$.

The open map property follows from the facts that $\hat{\mathcal{M}}$ is a PL (*topological* is good enough) k -manifold without boundaries, $\hat{\mathcal{M}} = \cup_{p \in P} \mathring{\text{st}}(p, \hat{\mathcal{M}})$ (from Lemma 4.3.19), and **S1** which proves that π restricted to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is injective for all $p \in P$.

The covering map property follows from the facts that $\hat{\mathcal{M}}$ and $\mathcal{M}$ are compact k -manifolds without boundaries, $\pi|_{\hat{\mathcal{M}}}$ is an open map, π restricted to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is injective, and Lemma 4.3.19 (a technical lemma about $\hat{\mathcal{M}}$).

S5. Using **S1**, **S2**, **S3** and **S4**, we show that $\pi|_{\hat{\mathcal{M}}} : \hat{\mathcal{M}} \rightarrow \mathcal{M}$ is both injective and surjective. Since $\hat{\mathcal{M}}$ is a compact space and $\mathcal{M}$ is a Hausdorff space, $\pi|_{\hat{\mathcal{M}}}$ provides a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$ from Theorem 2.5.7. This completes the proof of Lemma 4.3.1.

The broad outline of the proof is similar to the proof of homeomorphism given in [ACDL02b] but the technical details are quite different in most places. The difficulties that arose in getting results analogous to the ones in [ACDL02b] were handled using ideas from [Whi57a, Chap. II].

Details of the proof. Before we can prove the fundamental lemmas mentioned in the outline of the proof of Lemma 4.3.1, we will need some technical lemmas.

The following lemma is a generalization of [NSW08a, Proposition 6.2] which bounds the variation of the angle between tangent spaces between two points on the manifold $\mathcal{M}$.

Lemma 4.3.2 (Tangent variation) *Let $p, q \in \mathcal{M}$ and $\|p - q\| \leq \text{lfs}(p)/12$. Then,*

$$\sin \angle(T_q \mathcal{M}, T_p \mathcal{M}) = \sin \angle(T_p \mathcal{M}, T_q \mathcal{M}) < \frac{12\|p - q\|}{\text{lfs}(p)}.$$

Proof From Lemma 2.1.1, we have $\angle(T_p \mathcal{M}, T_q \mathcal{M}) = \angle(T_q \mathcal{M}, T_p \mathcal{M})$.

Let $t = \frac{\|p - q\|}{\text{lfs}(p)}$. Using the fact that lfs is 1-Lipschitz, we have

$$(1 - t)\text{lfs}(p) \leq \text{lfs}(q) \leq (1 + t)\text{lfs}(p) \quad (4.1)$$

We will show that for any unit vector u in $T_p \mathcal{M}$ there exists a unit vector v in $T_q \mathcal{M}$ such that $\sin \angle(u, v) \leq 12t$.

For a unit vector u in $T_p \mathcal{M}$, let p_u be a point in $T_p \mathcal{M}$ such that $p_u = p + t\text{lfs}(p) \cdot u$. We will use Lemma 2.3.2 on the edge $[p, p_u]$ and the tangent space $T_q \mathcal{M}$ to show that there exists a unit vector $v \in T_q \mathcal{M}$ such that $\sin \angle(u, v) < 12t$.

Let p'_u denote the point closest to p_u on $\mathcal{M}$. Then, from Lemma 2.2.1 (3) and Eq. (4.1), we have

$$\begin{aligned} \|q - p'_u\| &\leq \|q - p\| + \|p - p_u\| + \|p_u - p'_u\| \\ &\leq 2t(1 + t)\text{lfs}(p) \\ &\leq 2t\left(\frac{1 + t}{1 - t}\right)\text{lfs}(q). \end{aligned}$$

Using Lemma 2.2.1 (2) and Eq. (4.1), we have

$$\begin{aligned} \text{dist}(p, T_q \mathcal{M}) &\leq \|p - q\| \sin \angle(pq, T_q \mathcal{M}) \\ &\leq \frac{t^2}{2(1-t)^2} \text{lfs}(q) \end{aligned} \quad (4.2)$$

Using Lemmas 2.2.1 (2) and (3), and Eq. (4.1), we have

$$\begin{aligned} \text{dist}(p_u, T_q \mathcal{M}) &\leq \text{dist}(p'_u, T_q \mathcal{M}) + \|p_u - p'_u\| \\ &\leq 2t^2 \left(\frac{1+t}{1-t} \right)^2 \text{lfs}(q) + 2t^2 \text{lfs}(p) \\ &\leq 2t^2 \left(\frac{1+t}{1-t} \right)^2 \text{lfs}(q) + \frac{2t^2}{1-t} \text{lfs}(q) \\ &= \frac{2t^2}{1-t} \left(\frac{(1+t)^2}{(1-t)} + 1 \right) \text{lfs}(q) \end{aligned} \quad (4.3)$$

Let $\eta = \max\{\text{dist}(p, T_q \mathcal{M}), \text{dist}(p_u, T_q \mathcal{M})\}$. From Eq. (4.2) and (4.3), we have

$$\eta \leq \frac{2t^2}{1-t} \left(\frac{(1+t)^2}{(1-t)} + 1 \right) \text{lfs}(q)$$

From Lemma 2.3.2, there exists unit vector v in $T_q \mathcal{M}$ such that

$$\begin{aligned} \sin \angle(u, v) &\leq \frac{2\eta}{\Theta_{[p, p_u]} \|p - p_u\|} \\ &\leq \frac{4t^2}{1-t} \left(\frac{(1+t)^2}{(1-t)} + 1 \right) \text{lfs}(q) \times \frac{1}{\frac{t \text{lfs}(q)}{1+t}} \\ &= 4t \left(\frac{(1+t)^3}{(1-t)^2} + \frac{1+t}{1-t} \right) < 12t \end{aligned}$$

The above inequality follows from the facts that $\eta \leq \frac{2t^2}{1-t} \left(\frac{(1+t)^2}{(1-t)} + 1 \right) \text{lfs}(q)$, $\|p - p_u\| = t \text{lfs}(p)$, $\text{lfs}(p) \geq \frac{\text{lfs}(q)}{1+t}$ (from Eq. (4.1)), $\Theta_{[p, p_u]} = 1$ and $t \leq 1/12$. $\square$

We will show that the map $\tilde{\pi}_p$ is a diffeomorphism using Lemma 4.3.2.

Lemma 4.3.3 ($\tilde{\pi}_p$ is a C^1 -diffeomorphism) *The map $\tilde{\pi}_p$ is a C^1 -diffeomorphism between $T_p^{1/16}$ and $\tilde{\pi}_p(T_p^{1/16})$, and $B(p, \frac{14 \text{lfs}(p)}{16^2}) \cap \mathcal{M} \subseteq \tilde{\pi}_p(T_p^{1/16})$.*

Proof By Lemma 4.0.2, $\tilde{\pi}_p$ is a C^1 -function.

For all $x \in T_p \mathcal{M}^{1/16}$, we have from Lemma 2.2.1 (3)

$$\begin{aligned} \|p - \tilde{\pi}_p(x)\| &\leq \|p - x\| + \|x - \tilde{\pi}_p(x)\| \\ &\leq \frac{\text{lfs}(p)}{16} + \frac{2 \text{lfs}(p)}{16^2} = \frac{18 \text{lfs}(p)}{16^2} \end{aligned} \quad (4.4)$$

and using the fact that lfs is 1-Lipschitz we have

$$\text{lfs}(\tilde{\pi}_p(x)) \geq \text{lfs}(p) - \|p - \tilde{\pi}_p(x)\| \geq \left(1 - \frac{18}{16^2}\right) \text{lfs}(p), \quad (4.5)$$

the last inequality follows from Eq. (4.4).

From Eq. (4.4), and Lemmas 4.0.2 and 4.3.2, we have that $d\tilde{\pi}_p(x)$, the derivative of $\tilde{\pi}_p$, for all $x \in T_p^{1/16}$ is non-singular. Indeed for $x \in T_p^{1/16}$, we have from Lemmas 2.1.1 and 4.3.2, and Eq. (4.4)

$$\begin{aligned} \sin \angle(T_p \mathcal{M}, T_{\tilde{\pi}_p(x)} \mathcal{M}) &= \sin \angle(T_{\tilde{\pi}_p(x)} \mathcal{M}, T_p \mathcal{M}) \\ &\leq \frac{12\|p - \tilde{\pi}_p(x)\|}{\text{lfs}(p)} \leq \frac{27}{32} < 1. \end{aligned} \quad (4.6)$$

The map $\tilde{\pi}_p$ is injective. Otherwise there exists $x, y (x \neq y) \in T_p \mathcal{M}^{1/16}$ such that $\tilde{\pi}_p(x) = \tilde{\pi}_p(y)$. This implies the line segment $[x, y] \in T_p \mathcal{M}$ is orthogonal to $T_{\tilde{\pi}_p(x)} \mathcal{M}$. We have reached a contradiction since $\sin \angle(T_p \mathcal{M}, T_{\tilde{\pi}_p(x)} \mathcal{M}) < 1$ for all $x \in T_p \mathcal{M}^{1/16}$ from Eq. (4.6).

Since $\tilde{\pi}_p$ is injective and the derivative of $\tilde{\pi}_p$ is non-singular, $\tilde{\pi}_p$ is a diffeomorphism from the Inverse Function Theorem.

The fact that $B(p, \frac{14\text{lfs}(p)}{16^2}) \cap \mathcal{M} \subseteq \tilde{\pi}_p(T_p \mathcal{M}^{1/16})$ follows from Lemma 2.2.1 (3). $\square$

The following is a structural result on π , π_p and π_p^* .

Lemma 4.3.4 *Assume conditions C1 and C2 of Theorem 4.0.3. Let ε be sufficiently small and $x \in \text{st}(p, \hat{\mathcal{M}})$. There exists a constant C depending on Θ_0 , ω_0 and η_0 such that*

$$\max\{\|\pi_p(x) - x\|, \|\pi(x) - x\|, \|\pi_p^*(x) - x\|\} \leq C \varepsilon^2 \text{lfs}(p).$$

Proof Let τ be a k -simplex in $\text{star}(p)$ ($= \text{st}(p, \hat{\mathcal{M}})$ from Lemma 4.2.1). We will show that for all $x \in \tau$,

$$\max\{\|\pi_p(x) - x\|, \|\pi(x) - x\|, \|\pi_p^*(x) - x\|\} \leq C \varepsilon^2 \text{lfs}(p).$$

1. From Lemma 3.4.3 (1), we have for all vertices q of τ , $\|p - q\| \leq C_2 \varepsilon \text{lfs}(p)$. Therefore for all $x \in \tau$, $\|p - x\| \leq C_2 \varepsilon \text{lfs}(p)$.

From Lemma 4.1.1, we have

$$\|\pi_p(x) - x\| \leq \|p - x\| \sin \angle(T_p \mathcal{M}, \text{aff}(\tau)) \leq AC_2 \varepsilon^2 \text{lfs}(p). \quad (4.7)$$

2. Using the definition of the map π , the fact that $\|p - \pi_p(x)\| \leq C_2 \varepsilon \text{lfs}(p)$, Lemma 2.2.1 (3) and Eq. (4.7), we have

$$\begin{aligned} \|\pi(x) - x\| &\leq \|\pi(\pi_p(x)) - x\| \\ &\leq \|\pi(\pi_p(x)) - \pi_p(x)\| + \|\pi_p(x) - x\| \\ &\leq 2 \left(\frac{\|\pi_p(x) - p\|}{\text{lfs}(p)} \right)^2 \text{lfs}(p) + AC_2 \varepsilon^2 \text{lfs}(p) \\ &\leq (2C_2 + A) C_2 \varepsilon^2 \text{lfs}(p) \end{aligned} \quad (4.8)$$

3. We assume that ε is small enough such that $(2C_2 + A)C_2\varepsilon^2\text{fs}(p) < \frac{14\text{fs}(p)}{16^2}$. From Lemma 4.3.3, we have for all $x \in \tau$, $\tilde{\pi}_p^{-1}(\pi(x)) \in N_{\pi(x)}\mathcal{M} \cap T_p\mathcal{M}$. From Lemma 4.3.2, we have for all $x \in \tau$,

$$\sin \angle(T_p\mathcal{M}, T_{\pi(x)}\mathcal{M}) \leq \frac{12\|\pi(x) - p\|}{\text{fs}(p)} \leq \frac{12 \times 14}{16^2} < 1.$$

This implies that $|N_{\pi(x)}\mathcal{M} \cap T_p\mathcal{M}| = 1$. Therefore π_p^* restricted to τ is equal to $\tilde{\pi}_p^{-1} \circ \pi$, i.e. for all $x \in \tau$, $\pi_p^*(x) = \tilde{\pi}_p^{-1}(\pi(x))$.

For a given $x \in \tau$, let $y = \tilde{\pi}_p^{-1}(\pi(x))$ and $t = \frac{\|p - y\|}{\text{fs}(p)}$. From Lemma 4.3.3, we have $y \in T_p\mathcal{M}^{1/16} = B(p, \frac{\text{fs}(p)}{16}) \cap T_p\mathcal{M}$. Since $\tilde{\pi}_p$ is a restriction of π to $T_p\mathcal{M}$, we have

$$\|y - \pi(x)\| = \|y - \tilde{\pi}_p(y)\| \leq 2t^2\text{fs}(p). \quad (4.9)$$

The above inequality follows from Lemma 2.2.1 (3) and the fact that $\tilde{\pi}_p(y) = \pi(x)$.

From Eq. (4.8) and ε sufficiently small, we have $\|p - \pi(x)\| \leq C_2\varepsilon\text{fs}(p) + (2C_2 + A)C_2\varepsilon^2\text{fs}(p) \leq 2C_2\varepsilon\text{fs}(p)$.

Using the facts that $t \leq 1/16$, $t = \frac{\|p - y\|}{\text{fs}(p)}$, $\|y - \pi(x)\| \leq 2t^2\text{fs}(p)$ and $\|p - \pi(x)\| \leq 2C_2\varepsilon\text{fs}(p)$, we have

$$t/2 < t - 2t^2 \leq \|p - y\| - \|y - \tilde{\pi}_p(y)\| \leq \|p - \pi(x)\| \leq 2C_2\varepsilon\text{fs}(p).$$

This implies $t < 4C_2\varepsilon\text{fs}(p)$.

Using the fact that $\pi_p^*(x) = \tilde{\pi}_p^{-1}(\pi(x)) = y$, Eq. (4.8) and (4.9), and $t < 4C_2\varepsilon\text{fs}(p)$, we have

$$\begin{aligned} \|x - \pi_p^*(x)\| &\leq \|x - \pi(x)\| + \|\pi(x) - \pi_p^*(x)\| \\ &\leq \|x - \pi(x)\| + \|\pi(x) - y\| \\ &\leq (2C_2 + A)C_2\varepsilon^2\text{fs}(p) + 2t^2\text{fs}(p) \\ &< (2C_2 + A)C_2\varepsilon^2\text{fs}(p) + 32C_2^2\varepsilon^2\text{fs}(p) \\ &= (A + 34C_2)C_2\varepsilon^2\text{fs}(p) \end{aligned} \quad (4.10)$$

Combining Eq. (4.7), (4.8) and (4.10), we get our result. $\square$

Definition 4.3.5 (C^1 -embedding of simplices in $\mathcal{M}$) Let σ be an i -simplex, and let $f : \sigma \rightarrow \mathcal{M}$ be a C^1 -function. The simplex σ is C^1 -embedded by f in $\mathcal{M}$ if f is an injective mapping and for all $x \in \sigma$, the rank of the linear map $df(x) : \mathbb{R}^i \rightarrow T_{f(x)}\mathcal{M}$ is i , where $T_{f(x)}\mathcal{M}$ is the tangent space to $\mathcal{M}$ at $f(x)$.

We will use the following technical lemma bounding the fatness of a perturbed simplex in the proof of Lemma 4.3.8. The inequality in the lemma is similar to the one obtained in the proof of Lemma 14c in [Whi57a, Chapter 4].

Lemma 4.3.6 Let $\tau = [p_0, \dots, p_j]$ be a j -dimensional simplex in $\mathbb{R}^d$, and let for all $i \in \{0, \dots, j\}$, $\|p_i - p'_i\| \leq \rho\Delta_\tau$ with $\rho < 1/2$. Then,

$$\Theta_\sigma^{2j} \geq \frac{\Theta_\tau^{2j}}{(1 + 2\rho)^{2j}} - \frac{2j\rho}{1 + 2\rho}$$

where $\sigma = [p'_0, \dots, p'_j]$. This directly implies if $\rho \leq \frac{\Theta_\sigma^{2j}}{j2^{j+1}} \left(1 - \frac{1}{2j}\right)$ then $\Theta_\tau \geq \Theta_\sigma/2$

We will need the following technical result to prove Lemma 4.3.6.

Lemma 4.3.7 *Let $A = (a_{ij})_{1 \leq i, j \leq m}$, $B = (b_{ij})_{1 \leq i, j \leq m}$ be two $m \times m$ matrices satisfying the following inequalities:*

$$\begin{aligned} |a_{ij} - b_{ij}| &\leq \rho \\ |a_{ij}|, |b_{ij}| &\leq \eta \end{aligned}$$

Then

$$|\det A - \det B| \leq m^{\frac{m}{2}+1} \rho \eta^{m-1}.$$

Proof Let $B_1 = (c_{ij})_{1 \leq i, j \leq m}$ and $B'_1 = (c'_{ij})_{1 \leq i, j \leq m}$, where

$$c_{ij} = \begin{cases} a_{ij} & \text{if } i = 1 \\ b_{ij} & \text{otherwise} \end{cases}$$

and

$$c'_{ij} = \begin{cases} b_{ij} - a_{ij} & \text{if } i = 1 \\ b_{ij} & \text{otherwise} \end{cases}$$

Note that $|\det B'_1| \leq m^{\frac{m}{2}} \rho \eta^{m-1}$. Using expansion properties of determinants, we have

$$\begin{aligned} |\det A - \det B| &\leq |\det A - (\det B_1 + \det B'_1)| \\ &\leq |\det A - \det B_1| + m^{\frac{m}{2}} \rho \eta^{m-1} \end{aligned} \quad (4.11)$$

Again, let $B_2 = (d_{ij})_{1 \leq i, j \leq m}$ and $B'_2 = (d'_{ij})_{1 \leq i, j \leq m}$ where

$$d_{ij} = \begin{cases} a_{ij} & \text{if } i = 1, 2 \\ c_{ij} & \text{otherwise} \end{cases}$$

and

$$d'_{ij} = \begin{cases} a_{ij} & \text{if } i = 1 \\ b_{ij} - a_{ij} & \text{if } i = 2 \\ b_{ij} & \text{otherwise} \end{cases}$$

Note that $|\det B'_2| \leq m^{\frac{m}{2}} \rho \eta^{m-1}$. Using expansion property of determinants and Eq. (4.11) we get

$$\begin{aligned} |\det A - \det B| &\leq |\det A - (\det B_2 + \det B'_2)| + m^{\frac{m}{2}} \rho \eta^{m-1} \\ &\leq |\det A - \det B_2| + 2m^{\frac{m}{2}} \rho \eta^{m-1} \end{aligned} \quad (4.12)$$

Applying this procedure m times, we get

$$|\det A - \det B| \leq |\det A - (\det B_m + \det B'_m)| + (m-1)m^{\frac{m}{2}} \rho \eta^{m-1} \quad (4.13)$$

where $B_m = (e_{ij})_{1 \leq i, j \leq m}$ and $B'_m = (e'_{ij})_{1 \leq i, j \leq m}$,

$$e_{ij} = a_{ij}, \quad \forall i, j \in \{1, \dots, m\}$$

and

$$e'_{ij} = \begin{cases} a_{ij} & \text{if } i \neq m \\ b_{ij} - a_{ij} & \text{otherwise} \end{cases}$$

Note that $|\det B'_m| \leq m^{\frac{m}{2}} \rho \eta^{m-1}$. Using these facts and Eq. (4.13), we get

$$\begin{aligned} |\det A - \det B| &\leq |\det A - (\det B_m + \det B'_m)| + (m-1)m^{\frac{m}{2}} \rho \eta^{m-1} \\ &\leq |\det A - \det B_m| + m^{\frac{m}{2}+1} \rho \eta^{m-1} \\ &= m^{\frac{m}{2}+1} \rho \eta^{m-1} \quad \text{as } B_m = A \end{aligned}$$

□

Proof of Lemma 4.3.6 Let P_τ and P_σ denotes the two $d \times j$ matrices whose i^{th} -column are $(p_i - p_0)$ and $(p'_i - p'_0)$ respectively. The following inequalities follows directly from the fact that $\|p'_i - p_i\| \leq \rho \Delta_\tau$

$$\begin{aligned} |(p'_i - p'_0)^T (p'_j - p'_0) - (p_i - p_0)^T (p_j - p_0)| &\leq 2\rho(1 + 2\rho)\Delta_\tau^2 \\ |(p'_i - p'_0)^T (p'_j - p'_0)| &\leq (1 + 2\rho)^2 \Delta_\tau^2 \end{aligned}$$

From the above inequalities and the fact that

$$\text{vol}(\tau)^2, \text{vol}(\sigma)^2 = \frac{\det(P_\tau^T P_\tau)}{j!^2}, \frac{\det(P_\sigma^T P_\sigma)}{j!^2}$$

where A^T denotes the transpose of the matrix A , we get

$$\begin{aligned} |\text{vol}(\tau)^2 - \text{vol}(\sigma)^2| &= \frac{1}{j!^2} |\det(P_\tau^T P_\tau) - \det(P_\sigma^T P_\sigma)| \quad \text{from Lemma 4.3.7} \\ &\leq \frac{1}{j!^2} \times 2j^{j/2+1} \rho(1 + 2\rho)^{2j-1} \Delta_\tau^{2j} \\ &\leq 2j\rho(1 + 2\rho)^{2j-1} \Delta_\tau^{2j} \quad \text{as } j!^2 \geq j^{j/2} \end{aligned} \quad (4.14)$$

Therefore, using the fact that $\Delta_\sigma \leq (1 + 2\rho)\Delta_\tau$ and Eq. (4.14), we have

$$\Theta_\sigma^{2j} = \frac{\text{vol}(\sigma)}{\Delta_\sigma^{2j}} \geq \frac{\Theta_\tau^{2j}}{(1 + 2\rho)^{2j}} - \frac{2j\rho}{1 + 2\rho}$$

□

Lemma 4.3.8 (π C^1 -embeds $\tau \in \text{st}(p, \hat{\mathcal{M}})$) Assume condition **C1** and **C2** of Theorem 4.0.3. Let $\tau = [p, p_1, \dots, p_k]$ be a k -simplex in $\text{st}(p, \hat{\mathcal{M}})$ and let $\sigma = [p, q_1, \dots, q_k]$ be a k -simplex with $\|q_i - p_i\| \leq C\varepsilon^2 \text{lf}_s(p)$ for all $i \in \{1, \dots, k\}$. If ε is sufficiently small then the map π C^1 -embeds the simplex σ in $\mathcal{M}$.

Proof From condition **C1** and Lemma 4.2.1, if ε is sufficiently small then for all $p \in \mathcal{P}$, $\text{star}(p) = \text{st}(p, \hat{\mathcal{M}})$.

Using the fact that $\varepsilon < 1$ and Lemma 3.4.3 (1), $\sigma \subseteq B(p, (C + C_2\varepsilon)\varepsilon \text{lf}_s(p)) \subset B(p, (C + C_2)\varepsilon \text{lf}_s(p))$. Therefore, if $\varepsilon < \frac{1}{C+C_2}$, $\sigma \cap \mathcal{O} = \emptyset$. Recall that $\mathcal{O}$ denotes the medial axis

of $\mathcal{M}$. This implies, from Lemma 4.0.2 (1), that the restriction of the map π to σ is a C^1 -function. Rest of the proof is devoted to showing that π restricted to σ is injective and $d\pi$ is nonsingular.

We will assume $\varepsilon \leq \min \left\{ \frac{1}{4CC_3\eta_0}, \frac{\Theta_0^{2k}}{k2^{k+1}C\eta_0} \left(1 - \frac{1}{2^k}\right), \frac{1}{A_1+12(C+C_2)} \right\}$ where the term A_1 will be defined later in the proof.

Let

$$\rho \stackrel{\text{def}}{=} \max_{i \in \{1, \dots, k\}} \frac{\|q_i - p_i\|}{\Delta_\tau}.$$

Using the facts that $\Delta_\tau \geq \delta \text{lfs}(p)$, $\varepsilon/\delta \leq \eta_0$, and $\varepsilon \leq$

$$\rho \leq \frac{C\varepsilon^2 \text{lfs}(p)}{\Delta_\tau} \leq C\eta_0\varepsilon \leq \frac{\Theta_0^{2k}}{k2^{k+1}} \left(1 - \frac{1}{2^k}\right),$$

and therefore from the Lemma 4.3.6 and the fact that $\Theta_\tau \geq \Theta_0$, we have $\Theta_\sigma \geq \Theta_0/2$. From $\Gamma_\tau \leq C_3$, $\rho \leq C\eta_0\varepsilon$ and, we have

$$\begin{aligned} L_\sigma &\geq L_\tau - 2\rho\Delta_\tau \\ &= L_\tau(1 - 2\rho\Gamma_\tau) \\ &\geq C_3\delta \text{lfs}(p)(1 - 2CC_3\eta_0\varepsilon) \\ &\geq C_3\delta \text{lfs}(p)/2 \end{aligned}$$

Let $\eta = C\varepsilon^2 \text{lfs}(p)$, therefore from Corollary 2.3.3

$$\sin \angle(\text{aff}(\tau), \text{aff}(\sigma)) \leq \frac{2\eta}{\Theta_\sigma^k L_\sigma} \leq \left(\frac{2^{k+1}C\eta_0}{\Theta_0^k C_3} \right) \varepsilon$$

and from Lemma 2.3.3

$$\begin{aligned} \sin \angle(T_p\mathcal{M}, \text{aff}(\sigma)) &\leq \sin \angle(T_p\mathcal{M}, \text{aff}(\tau)) + \sin \angle(\text{aff}(\tau), \text{aff}(\sigma)) \\ &\leq \left(A + \frac{2^{k+1}C\eta_0}{\Theta_0^k C_3} \right) \varepsilon \stackrel{\text{def}}{=} A_1 \varepsilon \end{aligned} \tag{4.15}$$

For $\forall x \in \tau$ we have from Lemmas 3.4.3 (1) and 4.3.4, and the fact that $\varepsilon < 1$, $\|p - \pi(x)\| \leq (C_2 + C\varepsilon)\varepsilon \text{lfs}(p) < (C_2 + C)\varepsilon \text{lfs}(p)$. Therefore from Lemma 4.3.2 and Eq. (4.15),

$$\begin{aligned} \sin \angle(T_{\pi(x)}\mathcal{M}, \text{aff}(\sigma)) &\leq \sin \angle(T_{\pi(x)}\mathcal{M}, T_p\mathcal{M}) + \sin \angle(T_p\mathcal{M}, \text{aff}(\sigma)) \\ &< (12(C + C_2) + A_1)\varepsilon < 1 \end{aligned} \tag{4.16}$$

Eq. (4.16) implies that π restricted to σ is injective. Otherwise, if there exist $x_1, x_2 \in \sigma$ such that $\pi(x_1) = \pi(x_2)$ then the line segment in $\text{aff}(\sigma)$ joining the points x_1 and x_2 is parallel to $N_{\pi(x_1)}\mathcal{M}$. This contradicts Eq. (4.16).

Also from Eq. (4.16) and Lemma 4.0.2 (2) we have derivative $d\pi$ is non-singular when π is restricted to σ . This completes the proof. $\square$

Remark 4.3.9 (On orientation) In Lemmas 4.3.12 and 4.3.15, we will use the notion of orientation of Euclidean space/simplices, positively oriented simplices, and oriented PL manifold. For definitions and properties related to orientation, refer to any standard book on PL topology [Zee66, ?] or [Whi57a, Appendix II].

Remark 4.3.10 (Orienting $T_p\mathcal{M}$, $\text{st}(p, \text{Del}^{\xi_p}(P_p))$ and $\text{st}(p, \hat{\mathcal{M}})$) Assume condition **C1** of Theorem 4.0.3. Note that π_p restricted to $\text{star}(p)$ provides a simplicial isomorphism between $\text{star}(p)$ and $\text{st}(p, \text{Del}^{\xi_p}(P_p))$ from Lemma 2.4.2. And from Lemma 4.2.1, we have $\text{star}(p) = \text{st}(p, \hat{\mathcal{M}})$. First orient $T_p\mathcal{M}$ and then orient positively the k -dimensional simplices of $\text{st}(p, \text{Del}^{\xi_p}(P_p))$. Orient the k -dimensional simplices of $\text{st}(p, \hat{\mathcal{M}})$ isomorphically using the orientation of the corresponding k -dimensional simplices of $\text{st}(p, \text{Del}^{\xi_p}(P_p))$.

Since the $\text{lk}(p, \hat{\mathcal{M}})$ is a PL $(j-1)$ -sphere (in the proof of Lemma 4.2.4), we get $\text{st}(p, \hat{\mathcal{M}})$ is a PL k -ball as a direct consequence. Note that with this orientation $\text{st}(p, \hat{\mathcal{M}})$ becomes a oriented PL manifold.

Definition 4.3.11 (Simplexwise positive map) Let σ be a positively oriented i -simplex of $\mathbb{R}^i$, and let $f : \sigma \rightarrow \mathbb{R}^i$ be a C^1 -function. The map f is called simplexwise positive if $\det(J(f)) > 0$ for all $x \in \sigma$, where $J(f)$ and $\det(J(f))$ denote the Jacobian and the determinant of the Jacobian of the map f respectively.

The proof of the following lemma is similar to Lemma 23a in Chapter 4 from [Whi57a] and it will be used for proving that π_p^* restricted to $\text{st}(p, \hat{\mathcal{M}})$ is injective in Lemma 4.3.15. We give a proof for completeness.

Lemma 4.3.12 Assume conditions **C1** and **C2** of Theorem 4.0.3. For ε sufficiently small, π_p^* is a simplexwise positive mapping of $\text{st}(p, \hat{\mathcal{M}})$ into $T_p\mathcal{M}$.

Proof For any k -simplex $\sigma = [p, p_1, \dots, p_k] \in \text{st}(p, \hat{\mathcal{M}})$. For all points $q \in \sigma$, let $\pi_{p,t}(q) = (1-t)q + t\pi_p(q)$, and let $\sigma_t = \pi_{p,t}(\sigma)$. Since π_p is affine on each simplex, $\pi_{p,t}$ is also affine. Therefore σ_t is a simplex with vertices $p, p_{1t}, \dots, p_{kt}$ with $p_{jt} = \pi_{p,t}(p_j)$, $j \in \{1, \dots, k\}$.

It directly follows from Lemma 4.3.4 that

$$\forall j \in \{1, \dots, k\}, \|p_j - p_{jt}\| \leq C\varepsilon^2 \text{lfs}(p) \quad (4.17)$$

where C is the constant defined in Lemma 4.3.4. We deduce

Claim 4.3.13 For ε sufficiently small, π_p^* C^1 -embeds σ_t in $T_p\mathcal{M}$.




Figure 4.1: Refer to the proof of the Claim 4.3.13.

Proof For ε sufficiently small, we can show that $\pi(\sigma_t) \subset B(p, \frac{14\text{lf}_s(p)}{16^2})$. From the definition of π_p^* and the fact that $\pi(\sigma_t) \subset B(p, \frac{14\text{lf}_s(p)}{16^2})$, we have $\pi_p^* = \tilde{\pi}_p^{-1} \circ \pi$ as in part 3 of the proof of the Lemma 4.3.4. Refer to Fig. 4.1. That the map π_p^* C^1 -embeds σ_t onto $T_p\mathcal{M}$ follows from the facts that $\pi(\sigma_t) \subset B(p, \frac{14\text{lf}_s(p)}{16^2}) \cap \mathcal{M}$; $\tilde{\pi}_p^{-1}$ is a diffeomorphism when restricted to $B(p, \frac{14\text{lf}_s(p)}{16^2}) \cap \mathcal{M} \subseteq \tilde{\pi}_p(T_p^{1/16})$ (from Lemma 4.3.3); π C^1 -embeds σ_t into $\mathcal{M}$ (from Lemma 4.3.8); and $\pi_p^* = \tilde{\pi}_p^{-1} \circ \pi$. $\square$

Since the simplex σ_1 is in $T_p\mathcal{M}$, π_p^* is the identity in σ_1 . Therefore, $\det(J(\pi_p^*)) > 0$ in σ_1 . Since from Claim 4.3.13, $\det(J(\pi_p^*)) \neq 0$ in σ_t for all $t \in [0, 1]$, we also have $\det(J(\pi_p^*)) > 0$ in σ_0 . This concludes the proof of the lemma. $\square$

We will need the following standard lemma from convex geometry which bounds the distance between an interior point and a point on the boundary of a simplex. We have included the proof from [Whi57a, Lem. 14b of Chap. 4] for completeness.

Lemma 4.3.14 *Let $\sigma = [p_0, \dots, p_j]$ be a j -simplex and $p = \sum_{i=0}^j \mu_i p_i \in \sigma$ with $\sum_{i=0}^j \mu_i = 1$ and $\mu_i \geq 0$. Then,*

$$\text{dist}(p, \partial\sigma) \geq j! \Theta_\sigma^j \Delta_\sigma \times \min\{\mu_0, \dots, \mu_j\}$$

Proof Without loss of generality let's assume that the point p' closest to p lies on σ_{p_0} . This implies

$$\text{dist}(p, \partial\sigma) = \|p - p'\| = \mu_0 D_p(\sigma) \geq D_p(\sigma) \times \min\{\mu_0, \dots, \mu_j\}$$

Using the above equation and the fact that $D_p \geq j! \Theta_\sigma^j \Delta_\sigma$, we get the lemma. $\square$

Using Lemmas 4.3.12 and 4.3.14, we can now prove that π_p^* restricted to the open star $\text{st}(p, \hat{\mathcal{M}})$ of p is injective.

Lemma 4.3.15 (Injectivity of π_p^* restricted to $\text{st}(p, \hat{\mathcal{M}})$) *Assume conditions C1 and C2 of Theorem 4.0.3. Let ε be sufficiently small. For each point $p \in P$, the map π_p^* is injective on the open star $\text{st}(p, \hat{\mathcal{M}})$.*

Proof Note that since $\text{lk}(p, \hat{\mathcal{M}})$ is a PL $(k-1)$ -sphere from Lemmas 4.2.2 (2) and 4.2.4, we get $\text{st}(p, \hat{\mathcal{M}})$ to be a PL k -ball. This implies $\partial \text{st}(p, \hat{\mathcal{M}}) = \text{lk}(p, \hat{\mathcal{M}})$.

For convenience, rewrite $f = \pi_p^*$ and $S = \text{st}(p, \hat{\mathcal{M}})$. From Remark 4.3.10, we get S to be a oriented PL k -ball, and, by Lemma 4.3.12, $f = \pi_p^*$ is a simplexwise positive mapping of S into $T_p\mathcal{M}$. Let $f(S^{k-1})$ be the image by f of the $(k-1)$ -skeleton of S (i.e. the set of faces of S of dimension at most $k-1$) and let R be any connected open subset of $T_p\mathcal{M} - f(\partial S)$. From Lemma 2.5.21, any two points of $R \setminus f(S^{k-1})$ are covered the same number of times. If this number is 1, then f , restricted to the open subset $f^{-1}(R)$ of $\text{star}(p)$, is injective onto R .

To prove the lemma, it is therefore sufficient to prove that there exists a point g in $S \setminus S^{k-1}$ whose image $f(g)$ is not covered by any other point x of $S \setminus S^{k-1}$, i.e. $f(g) \neq f(x)$ for all $x \neq g$ and $x \in S \setminus S^{k-1}$. Let σ be a k -simplex $[q_0, \dots, q_k]$ of S and let

$$g = \frac{1}{k+1} \sum_{i=0}^k q_i.$$

Using Lemma 4.3.14 and the facts that $\Theta_\sigma \geq \Theta_0$ (since the simplices of $\hat{\mathcal{M}}$ are Θ_0 -fat), $\Delta_\sigma \geq \delta \text{lf}(p)$, and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 3.4.1), we have

$$\begin{aligned} \text{dist}(g, \partial\sigma) &\geq \frac{k! \Theta_\sigma^k \Delta_\sigma}{k+1} \geq \frac{k! \Theta_0^k}{k+1} \times \delta \text{lf}(p) \\ &\geq \frac{k! \Theta_0^k}{k+1} \times \frac{\varepsilon \text{lf}(p)}{\eta_0} = C' \varepsilon \text{lf}(p) \end{aligned} \quad (4.18)$$

where $C' = \frac{k! \Theta_0^k}{(k+1)\eta_0}$.

Also, from Lemma 14b of Chapter 4 of [Whi57a], Eq. (4.18) and using the fact that $\|\pi_p(x) - x\| \leq C \varepsilon^2 \text{lf}(p)$ for all $x \in \sigma$ (Lemma 4.3.4), we have

$$\text{dist}(\pi_p(g), \partial\pi_p(\sigma)) \geq C' \varepsilon \text{lf}(p) - 2C \varepsilon^2 \text{lf}(p), \quad (4.19)$$

Since π_p embeds S into $T_p \mathcal{M}$ (from Lemma 2.4.2) with $\pi_p(S) = \text{st}(p, \text{Del}^{\xi_p}(P_p))$ and $\text{Del}^{\xi_p}(P_p)$ is a trinagulation of a k -dimensional convex hull (see the discussion on the general position assumption at the beginning of this chapter), Eq. (4.19) implies that, for all $x \in S \setminus \text{int } \sigma$

$$\|\pi_p(g) - \pi_p(x)\| \geq C' \varepsilon \text{lf}(p) - 2C \varepsilon^2 \text{lf}(p). \quad (4.20)$$

From Eq. (4.20), we have for all $x \in S \setminus \text{int } \sigma$

$$\begin{aligned} \|\pi_p^*(x) - \pi_p^*(g)\| &\geq \|\pi_p(x) - \pi_p(g)\| - (\|\pi_p(x) - x\| + \|\pi_p^*(x) - x\|) \\ &\quad - (\|\pi_p(g) - g\| + \|\pi_p^*(g) - g\|) \\ &\geq C' \varepsilon \text{lf}(p) - 6C \varepsilon^2 \text{lf}(p) > 0 \end{aligned} \quad (4.21)$$

The last inequality holds for a sufficiently small ε .

Claim 4.3.13 together with Eq. (4.21) show that $\pi_p^*(x) \neq \pi_p^*(g)$ for all $x \in S \setminus \{g\}$. Hence, g is not covered by any other point of S , and the lemma follows. $\square$

We will now prove that π is also injective when restricted to $\text{st}(p, \hat{\mathcal{M}})$.

Lemma 4.3.16 (Injectivity of π restricted to $\text{st}(p, \hat{\mathcal{M}})$) *Assume conditions C1 and C2 of Theorem 4.0.3. Let ε be sufficiently small. For all p in P , the map π restricted to $\text{st}(p, \hat{\mathcal{M}})$ is injective.*

Proof To reach a contradiction, assume that there exist x_1, x_2 ($x_1 \neq x_2$) in $\text{star}(p) \setminus \partial \text{star}(p)$ such that $\pi(x_1) = \pi(x_2)$. Then $\pi_p^*(x_1) = \pi_p^*(x_2) = N_{\pi(x_1)} \mathcal{M} \cap T_p \mathcal{M}$. Which contradicts the fact that π_p^* is injective when restricted to $\text{st}(p, \hat{\mathcal{M}})$ from Lemma 4.3.15. $\square$

We will now show that for all $p \in P$, we have $\pi^{-1}(p) = \{p\}$ when π is restricted to $\hat{\mathcal{M}}$. The following lemma will be used to show that π restricted to $\hat{\mathcal{M}}$ is a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$ in Lemma 4.3.1.

Lemma 4.3.17 *Let ε be sufficiently small. For all p in P and restricting the map π to $\hat{\mathcal{M}}$, we have $\pi^{-1}(p) = \{p\}$.*

Proof To reach a contradiction, we assume that there exists a k -simplex $\tau = [q_0, \dots, q_k]$ in $\hat{\mathcal{M}}$ such that there exists $x \in \tau$ with $x \neq p$ and $\pi(x) = p$.

We will have to consider the following two cases.

Case 1. p is a vertex of τ . This implies that the unit vector $u \in \text{aff}(\tau)$ along the line joining the points p and x lies in $N_p\mathcal{M}$. But from Lemma 4.1.1 and ε sufficiently small, we have $\sin \angle(\text{aff}(\tau), T_p\mathcal{M}) \leq A\varepsilon < 1$.

Case 2. p is not a vertex of τ . The outline of the proof for case 2 is the following: we will divide the k -simplex τ into a union of $k+2$ convex cells and show that for each convex cell the distance of any point of the cell to p is $\Omega(\varepsilon \text{lfs}(q))$, where q is a vertex of τ . For ε sufficiently small, we will reach a contradiction from Lemma 4.3.4.

Since $\hat{\mathcal{M}}$ has no inconsistent configuration, τ is in $\text{star}(q_0)$. Let $m = \text{Vor}^\omega(\tau) \cap T_{q_0}\mathcal{M}$. From Lemma 3.4.2, we have

$$\begin{aligned} R &= \sqrt{\|m - q_0\|^2 - \omega(q_0)^2} \\ &\leq \|m - q_0\| \\ &\leq C_1 \varepsilon \text{lfs}(q_0). \end{aligned} \quad (4.22)$$

Using the facts that for all vertices q_i, q_j of τ , $\|q_i - q_j\| \leq C_2 \varepsilon \text{lfs}(q_i)$ (from Lemma 3.4.3 (1)), lfs is 1-Lipschitz and ε sufficiently small, we have

$$\begin{aligned} \text{lfs}(q_j) &\geq \text{lfs}(q_i) - \|q_i - q_j\| \\ &\geq (1 - C_2 \varepsilon) \text{lfs}(q_i) \\ &\geq \frac{\text{lfs}(q_i)}{2} \end{aligned} \quad (4.23)$$

Therefore using Eq. (4.22) and (4.23), we have for all vertices q_i of τ

$$R \leq C_1 \varepsilon \text{lfs}(q_0) \leq 2C_1 \varepsilon \text{lfs}(q_i). \quad (4.24)$$

Consider the edge $q_i q_j$ of τ and let c_{ij} be the projection of m onto the line segment $[q_i q_j]$. Observe that the ball of radius $r_{ij} = \sqrt{\|c_{ij} - q_i\|^2 - \omega(q_i)^2} \leq R$ centered at c_{ij} is orthogonal to the balls $B(q_i, \omega(q_i))$ and $B(q_j, \omega(q_j))$. Using the fact that $\mathcal{P}$ is a (ε, δ) -lfs sample of $\mathcal{M}$ and Lemma 2.4.3 (2), we have

$$\begin{aligned} r_{ij} &\geq \frac{\sqrt{1 - 4\omega_0^2} \|q_i - q_j\|}{2} \\ &\geq b\delta \max\{\text{lfs}(q_i), \text{lfs}(q_j)\} \end{aligned} \quad (4.25)$$

where $b = \frac{\sqrt{1 - 4\omega_0^2}}{2}$.

For all $q_i \in \{q_0, \dots, q_k\}$, let

$$\lambda(q_i) = \max \left\{ \omega(q_i), \frac{b}{4} \delta \text{lfs}(q_i) \right\}.$$

Observe that $\lambda(q_i) < \text{m}(q_i)/2$. For $\varepsilon \leq 1/2$ and Lemma 2.2.1 (1),

$$\begin{aligned} \lambda(q_i) &< \frac{\text{m}(q_i)}{2} \\ &\leq \frac{\varepsilon}{1 - \varepsilon} \text{lfs}(q_i) \\ &\leq 2\varepsilon \text{lfs}(q_i). \end{aligned} \quad (4.26)$$




Figure 4.2: Refer to the proof of Lemma 4.3.17, Case 2.

Let $x_{ij} = [c_{ij} q_i] \cap \partial B(c_{ij}, r_{ij})$ and $y_{ij} = [c_{ij} q_i] \cap \partial B(q_i, \lambda(q_i))$. Note that $x_{ij} = \partial B(c, R) \cap [c_{ij} q_i]$, see Figure 4.2.

Therefore

$$\begin{aligned}
 \|x_{ij} - y_{ij}\| &= r_{ij} + \lambda(q_i) - \sqrt{r_{ij}^2 + \omega(q_i)^2} \\
 &= \frac{2r_{ij}\lambda(q_i) + \lambda(q_i)^2 - \omega(q_i)^2}{r_{ij} + \lambda(q_i) + \sqrt{r_{ij}^2 + \omega(q_i)^2}} \\
 &\geq \frac{2r_{ij}\lambda(q_i)}{R + \lambda(q_i) + \sqrt{R^2 + \lambda(q_i)^2}} && \text{since } \lambda(q_i) \geq \omega(q_i) \text{ and } r_{ij} \leq R \\
 &\geq \frac{2r_{ij}\lambda(q_i)}{\left(2C_1\varepsilon + 2\varepsilon + \sqrt{4C_1^2\varepsilon^2 + 4\varepsilon^2}\right)\text{lfs}(q_i)} && \text{from Eq. (4.24) and (4.26)} \\
 &\geq \frac{b^2\delta^2\text{lfs}(q_i)}{4\varepsilon(C_1 + 1 + \sqrt{C_1^2 + 1})} && \text{using } \lambda(q_i) \geq \frac{b\delta\text{lfs}(q_i)}{4} \text{ and Eq. (4.25)} \\
 &\geq \frac{b^2\delta\text{lfs}(q_i)}{4\eta_0(C_1 + 1 + \sqrt{C_1^2 + 1})} && \text{since } \varepsilon/\delta \leq \eta_0 \text{ from Hypothesis 3.4.1}
 \end{aligned}$$

Using the fact that $r_{ij} \geq \|x_{ij} - y_{ij}\|$, we have

$$\begin{aligned}
 A_{ij}^2 &\stackrel{\text{def}}{=} (2r_{ij} - \|x_{ij} - y_{ij}\|) \times \|x_{ij} - y_{ij}\| \\
 &\geq r_{ij} \times \|x_{ij} - y_{ij}\| \\
 &\geq C_6^2 \delta^2 \text{lfs}(q_i)^2
 \end{aligned}$$

where C_6 is a constant that depends on ω_0 and η_0 .




Figure 4.3: Intersecting Chords Theorem.

Let z_{ij} denote the point closest to y_{ij} on $\partial B(c, R)$.

From the Intersecting Chords Theorem (see Figure 4.3) for circles, we have :

$$(2R - \|z_{ij} - y_{ij}\|) \times \|z_{ij} - y_{ij}\| = A_{ij}^2.$$

By the definition of z_{ij} , the solution to the above quadratic equation in $\|z_{ij} - y_{ij}\|$ is the smaller root:

$$\begin{aligned} \|z_{ij} - y_{ij}\| &= R - \sqrt{R^2 - A_{ij}^2} \\ &= \frac{A_{ij}^2}{R + \sqrt{R^2 - A_{ij}^2}} \geq \frac{A_{ij}^2}{2R} \geq \frac{C_6^2}{4C_1\eta_0^2} \varepsilon \text{lfs}(q_i) \stackrel{\text{def}}{=} 2C_7\varepsilon \text{lfs}(q_i) \end{aligned}$$

The last inequality follows from the facts that $R \leq 2C_1\varepsilon \text{lfs}(q_i)$ (from Eq. (4.24)) and $\varepsilon/\delta \leq \eta_0$.

Using the fact that ε is sufficiently small and Eq. (4.23), we have for all vertices q of τ

$$\|z_{ij} - y_{ij}\| \geq 2C_7\varepsilon \text{lfs}(q_i) \geq C_7\varepsilon \text{lfs}(q)$$

Let $\text{conv}(S)$ denote the convex hull of the points y_{ij} :

$$S = \{y_{ij} \mid i, j (\neq i) \in \{0, \dots, k\}\}.$$

Using convexity, we have

$$\text{dist}(\text{conv}(S), \partial B(c, R)) = \min_{i, j (\neq i) \in \{0, \dots, k\}} \|z_{ij} - y_{ij}\| \geq C_7\varepsilon \text{lfs}(q) \quad (4.27)$$

for all vertices $q \in \{q_0, \dots, q_k\}$.

Recall that we have assumed in the beginning of the proof that $x \in \tau$ and $\pi(x) = p$. Eq. (4.27) implies that $x \notin \text{conv}(S)$. Indeed, if $x \in \text{conv}(S)$, then from Eq. (4.27) and the fact that the ball $B(c, R)$ is empty, we will have

$$\|x - p\| \geq C_7\varepsilon \text{lfs}(q)$$

for all vertices q of τ . But from Lemma 4.3.4, we have

$$\|x - p\| = \|x - \pi(x)\| \leq C\varepsilon^2 \text{fs}(q)$$

for all vertices q of τ . We have reached a contradiction for ε sufficiently small. Hence $x \notin \text{conv}(S)$.

Let $S_i = \{y_{ij} \mid j \in \{0, \dots, k\} \setminus \{i\}\} \cup \{q_i\}$. The convex hulls of S_i , $\text{conv}(S_i)$, satisfies the following properties :

$$\begin{aligned} \text{conv}(S_i) &\subset B(q_i, \lambda(q_i)) \\ \tau &= \text{conv}(S) \cup \left(\bigcup_{i \in \{0, \dots, k\}} \text{conv}(S_i) \right) \end{aligned} \quad (4.28)$$

If $x \in \text{conv}(S_i)$, then from Lemma 4.3.4, we have for ε sufficiently small

$$\begin{aligned} \|q_i - p\| &\leq \|q_i - x\| + C\varepsilon^2 \text{fs}(q_i) \\ &\leq \lambda(q_i) + C\varepsilon^2 \text{fs}(q_i) \\ &< \frac{\text{nn}(q_i)}{2} + C\varepsilon^2 \text{fs}(q_i) && \text{since } \lambda(q_i) < \frac{\text{nn}(q_i)}{2} \\ &< \text{nn}(q_i) \end{aligned}$$

The last inequality holds for $\varepsilon < \frac{1}{2C\eta_0}$. We have reached a contradiction.

Since $x \notin \text{conv}(S_i)$ and $x \notin \text{conv}(S)$, it follows that $x \notin \tau$. We have therefore reached a contradiction, i.e. there does not exist $x \in \tau$ such that $\pi(x) = p$. $\square$

Definition 4.3.18 Let $\hat{\mathcal{M}} \stackrel{\text{def}}{=} \bigcup_i \hat{\mathcal{M}}_i$ and $\mathcal{M} \stackrel{\text{def}}{=} \bigcup_j \mathcal{M}_j$ where $\hat{\mathcal{M}}_i$ and $\mathcal{M}_j$ denotes the connected components, see Definition 2.5.12 in Section 2.5, of $\hat{\mathcal{M}}$ and $\mathcal{M}$ respectively. Also, let $\mathcal{Q}_i \stackrel{\text{def}}{=} \mathcal{P} \cap \hat{\mathcal{M}}_i$.

We will need the following technical lemma.

Lemma 4.3.19 1. $\hat{\mathcal{M}} = \bigcup_{p \in \mathcal{P}} \text{st}(p, \hat{\mathcal{M}})$.

2. $\hat{\mathcal{M}}_i = \bigcup_{p \in \mathcal{P}_i} \text{st}(p, \hat{\mathcal{M}}) = \bigcup_{p \in \mathcal{P}_i} \text{st}(p, \hat{\mathcal{M}})$. This implies $\hat{\mathcal{M}}_i$ is a compact and connected topological k -manifold without boundary.

Proof 1. Let $x \in \hat{\mathcal{M}}$, and $\sigma \in \hat{\mathcal{M}}$ be the minimal simplex¹ containing x . For all vertices p of σ , $x \notin \text{lk}(p, \hat{\mathcal{M}})$ and this implies $x \in \text{st}(p, \hat{\mathcal{M}})$. This in turn implies

$$\hat{\mathcal{M}} = \bigcup_{p \in \mathcal{P}} \text{st}(p, \hat{\mathcal{M}}).$$

2. To reach a contradiction lets assume that there exists $p \in \mathcal{Q}_i$ such that $\text{st}(p, \hat{\mathcal{M}}) \not\subset \hat{\mathcal{M}}_i$. Using the fact that $\text{st}(p, \hat{\mathcal{M}})$ is connected² and Lemma XXX we have $\text{st}(p, \hat{\mathcal{M}}) \cup \hat{\mathcal{M}}_i \subsetneq \hat{\mathcal{M}}_i$ from the initial assumption) is connected. We have reached a contradiction.

The proof that $\hat{\mathcal{M}}_i = \bigcup_{p \in \mathcal{Q}_i} \text{st}(p, \hat{\mathcal{M}})$ can be shown using the same arguments as the ones used in the proof of part 1. $\square$

¹ σ being a minimal simplex containing x means $x \in \sigma$ and for all proper subsimplices $\tau \subsetneq \sigma$, $x \notin \tau$.

² As $\text{st}(p, \hat{\mathcal{M}})$ is a union of simplices incident on p and simplices are connected, therefore, from Lemma lemma-union-connected-disjoint-connected, we get $\text{st}(p, \hat{\mathcal{M}})$ is connected.

Following lemma (and the proof) is similar to [GW04a, Lem. 3].

Lemma 4.3.20 *Let P be a (ε, δ) -lfs sample of $\mathcal{M}$ with $\varepsilon < 1$. For all connected components $\mathcal{M}_j$ of $\mathcal{M}$, we have $P \cap \mathcal{M}_j \neq \emptyset$.*

Proof We will use the following technical result on local feature size.

Claim 4.3.21 *Let $\mathcal{M}_i$ and $\mathcal{M}_j$ ($\neq \mathcal{M}_i$) be two connected components of $\mathcal{M}$. For all $p \in \mathcal{M}_i$, if $r < \text{lfs}(p)$ then*

$$\bar{B}(p, r) \cap \mathcal{M}_j \neq \emptyset.$$

Proof To reach a contradiction let's assume that $B(p, \varepsilon \text{lfs}(p)) \cap \mathcal{M}_j \neq \emptyset$. Let q be a point in $B(p, \varepsilon \text{lfs}(p)) \cap \mathcal{M}_j \neq \emptyset$ and $\tilde{\mathcal{M}} = \mathcal{M} \setminus \mathcal{M}_i$. Consider the following function:

$$f : [p, q] \rightarrow [0, \infty), \quad f(x) = \text{dist}(x, \mathcal{M}_i) - \text{dist}(x, \tilde{\mathcal{M}})$$

It is easy to see that f is continuous, and $f(p) < 0$ and $f(q) > 0$. From continuity of f it follows that there exists $x \in [p, q]$ such that $\text{dist}(x, \mathcal{M}_i) = \text{dist}(x, \tilde{\mathcal{M}})$. This implies x belongs to the medial axis $\mathcal{O}$ of $\mathcal{M}$. Which in turn implies, from the definition of lfs,

$$\text{lfs}(p) = \text{dist}(p, \mathcal{O}) \leq \|p - x\| \leq r < \text{lfs}(p).$$

We have reached a contradiction. □

Let p be a point in the connected component $\mathcal{M}_i$ of $\mathcal{M}$. Since P is an (ε, δ) -lfs sample of $\mathcal{M}$, $\bar{B}(p, \varepsilon \text{lfs}(p)) \cap P \neq \emptyset$. But from the above claim, we have $\bar{B}(p, \varepsilon \text{lfs}(p)) \cap \mathcal{M}_j = \emptyset$ for all connected component $\mathcal{M}_j \neq \mathcal{M}_i$ of $\mathcal{M}$. This implies $\mathcal{M}_i \cap P \neq \emptyset$. □

The following lemma is a direct consequence of Lemmas 4.3.16, 4.3.19 (1), and Theorem 2.5.16.

Lemma 4.3.22 (Open map) *Assume condition **C1** and **C2** of Theorem 4.0.3. If ε is sufficiently small, then π restricted to $\hat{\mathcal{M}}$ (or $\hat{\mathcal{M}}_i$) is an open map to $\mathcal{M}$.*

Proof Using the fact that $\text{st}(p, \hat{\mathcal{M}})$ is an open set of $\hat{\mathcal{M}}$, Lemmas 4.3.16 and 4.3.19 (1), and Theorem 2.5.16, we have π restricted to $\hat{\mathcal{M}}$ is an open map. Similarly, we can show that π is an open map when restricted to $\hat{\mathcal{M}}_i$. □

We can show that π restricted to $\hat{\mathcal{M}}_i$, a connected component of $\hat{\mathcal{M}}$, is a covering map for some connected component $\mathcal{M}_j$ of $\mathcal{M}$, i.e., the following lemma.

Lemma 4.3.23 (Covering) *Assume condition **C1** and **C2** of Theorem 4.0.3 and ε is sufficiently small. For all $\hat{\mathcal{M}}_i$ there exists $\mathcal{M}_j$ such that*

1. $\pi(\hat{\mathcal{M}}_i) = \mathcal{M}_j$, and
2. $\pi : \hat{\mathcal{M}}_i \rightarrow \mathcal{M}_j$ is a covering map

Proof 1. Since $\hat{\mathcal{M}}_i$ is a compact (from Lemma 4.3.19 (2)) and π is continuous, therefore from Lemma 2.5.6, $\pi(\hat{\mathcal{M}}_i)$ is compact. Using the facts that $\pi(\hat{\mathcal{M}}_i)$ is compact, $\mathcal{M}$ is a Hausdorff space and Lemma 2.5.5 we get $\pi(\hat{\mathcal{M}}_i)$ is closed. This implies there exists an open set $U_1 \subseteq \mathcal{M}$ such that $\mathcal{M} - \pi(\hat{\mathcal{M}}_i) = U_1$. Let $U_1 = U \cap \mathcal{M}_j$.

As $\hat{\mathcal{M}}_i$ is connected and π is continuous, we have from Lemma 2.5.9, $\pi(\hat{\mathcal{M}}_1)$ is connected. This implies, together with the Lemma 2.5.13, that there exists $\mathcal{M}_j$ such that $\pi(\hat{\mathcal{M}}_i) \subseteq \mathcal{M}_j$. Using the facts that π restricted $\mathcal{M}$ is an open map (from Lemma 4.3.22) and $\text{st}(p, \hat{\mathcal{M}})$ is an open set of $\hat{\mathcal{M}}$, we get $\pi(\hat{\mathcal{M}}_i) = \bigcup_{p \in P_i} \pi(\text{st}(p, \hat{\mathcal{M}}))$ is an open set of $\mathcal{M}$.

Let $U_1 = U \cap \mathcal{M}_j$ and $U_2 = \pi(\hat{\mathcal{M}}_i)$. Note that $U_2 \subseteq \mathcal{M}_j$.

Since $\mathcal{M}_j = U_1 \cup U_2$ with $U_1 \cap U_2 = \emptyset$, $\mathcal{M}_j$ connected, $\pi(|\hat{\mathcal{M}}_i|) \subseteq \mathcal{M}_j$ and $\pi(|\hat{\mathcal{M}}_i|) \neq \emptyset$, this implies $U_1 = \emptyset$ which in turn implies $\pi(|\hat{\mathcal{M}}_i|) = \mathcal{M}_j$.

2. Let $x \in \mathcal{M}_j$, and let $S = \pi^{-1}(x) \cap \hat{\mathcal{M}}_i$. For all $p \in P$, π restricted to $\text{st}(p, \hat{\mathcal{M}})$ is injective (Lemma 4.3.16), and for all $\hat{\mathcal{M}}_i$, $\hat{\mathcal{M}}_i = \bigcup_{p \in Q_i} \text{st}(p, \hat{\mathcal{M}})$ (Lemma 4.3.19). Therefore, $|S| \leq |Q_i|$. Without loss of generality, let's assume that $S = \{x_1, \dots, x_l\}$.

Since $\pi : \hat{\mathcal{M}} \rightarrow \mathcal{M}$ is an open map (from Lemma 4.3.22), $\pi(\text{st}(p, \hat{\mathcal{M}}))$ is open in $\mathcal{M}$. Since $\hat{\mathcal{M}}$ is a Hausdorff space, for all $r \in \{1, \dots, l\}$ there exists open set $V_r \subset \mathcal{M}$ such that (a) $x_r \in V_r$, (b) $V_r \subset \text{st}(p_r, \hat{\mathcal{M}})$, and (c) for all $r_1, r_2 (\neq r_1) \in \{1, \dots, l\}$, $V_{r_1} \cap V_{r_2} = \emptyset$. Also note that $\pi(V_r) \subset \mathcal{M}_j$. Since $\pi|_{\hat{\mathcal{M}}}$ is an open map, $V = \bigcup_r V_r \subset \mathcal{M}_j$ is an open set of $\mathcal{M}_j$. The open sets $U_r = \pi^{-1}(V) \cap \hat{\mathcal{M}}_i$ satisfy the following conditions:

- $\pi|_{\hat{\mathcal{M}}_i}^{-1}(V) = \bigcup_r U_r$.
- $\pi : U_r \rightarrow V$ is a continuous bijective map. This follows from the fact that $\pi(U_r) = V$ and π restricted to $\text{st}(p_r, \hat{\mathcal{M}}) (\supseteq U_r)$ is injective (from Lemma 4.3.16).
- $\forall r_1, r_2 (\neq r_1) \in \{1, \dots, l\}$, $U_{r_1} \cap U_{r_2} = \emptyset$.
- Since π restricted to $\mathcal{M} (\supset U_r)$ is an open map (from Lemma 4.3.22), π restricted to U_r is a continuous bijective map, we get, from Lemma 2.5.15, $\pi : U_r \rightarrow V$ is a homeomorphism.

Therefore $\pi : \hat{\mathcal{M}}_i \rightarrow \mathcal{M}_j$ is a covering map. See, Definition 2.5.19 in Chapter 2. $\square$

Proof of homeomorphism, i.e., Lemma 4.3.1 1. For all connected components $\mathcal{M}_j$ of $\mathcal{M}$, there exists $p \in \mathcal{M}_j \cap P$ from Lemma 4.3.20. Without loss of generality let $p \in \hat{\mathcal{M}}_1$. This implies $\pi(p) = p \in \mathcal{M}_j \cap \pi(\hat{\mathcal{M}}_1)$. From Lemma 4.3.23 (1) and the fact that $\mathcal{M}_j \cap \pi(\hat{\mathcal{M}}_1) \neq \emptyset$, we get $\pi(\hat{\mathcal{M}}_1) = \mathcal{M}_j$.

As this is true for all the components of $\mathcal{M}$, therefore we have proved that $\pi(\hat{\mathcal{M}}) = \mathcal{M}$.

2. We will first show that for all $\mathcal{M}_j$ there exists $\hat{\mathcal{M}}_i$ such that

$$\pi|_{\hat{\mathcal{M}}}^{-1}(\mathcal{M}_j) = \hat{\mathcal{M}}_i.$$

To reach a contradiction let's assume that $\pi(\hat{\mathcal{M}}_1) = \pi(\hat{\mathcal{M}}_2) = \mathcal{M}_j$. Since $\hat{\mathcal{M}}_1$ and $\hat{\mathcal{M}}_2$ are components of $\hat{\mathcal{M}}$, $\hat{\mathcal{M}}_1 \cap \hat{\mathcal{M}}_2 = \emptyset$ (from Lemma 2.5.13 (3)). Together with the facts that $\pi(\hat{\mathcal{M}}_1) = \pi(\hat{\mathcal{M}}_2) = \mathcal{M}_j$ and $\hat{\mathcal{M}}_1 \cap \hat{\mathcal{M}}_2 = \emptyset$, we have for all $x \in \mathcal{M}_j$, the size of the set $\pi^{-1} \cap \hat{\mathcal{M}}_1 \cap \hat{\mathcal{M}}_2$ is not less than 2. But from Lemma 4.3.17, we know that for all $p \in P_1 \cup P_2$, $\pi^{-1}(p) = \{p\}$. And we have reached a contradiction.

We will now show that $\pi : \hat{\mathcal{M}}_i \rightarrow \mathcal{M}_j$ where $\mathcal{M}_j = \pi(\hat{\mathcal{M}}_i)$ and a component of $\mathcal{M}$. From Lemma 4.3.23 (2), we know that $\pi : \hat{\mathcal{M}}_i \rightarrow \mathcal{M}_j$ is a covering map. This implies, using the facts that $\forall p \in P_i \pi^{-1}(p) = \{p\}$, $\mathcal{M}_j$ is connected (see Lemma 2.5.13 (1)) and Lemma 2.5.20 we get π restricted to $\hat{\mathcal{M}}_i$ is injective.

3. We have proved that the map $\pi : \hat{\mathcal{M}} \rightarrow \mathcal{M}$ is both surjective and injective. Using the facts that π restricted to $\hat{\mathcal{M}}$ is a continuous bijective map, $\hat{\mathcal{M}}$ is compact and $\mathcal{M}$ is a Hausdorff space, and Theorem 2.5.7, we get $\pi : \hat{\mathcal{M}} \rightarrow \mathcal{M}$ is a homeomorphism. $\square$

4.4 Pointwise approximation

Following lemma is a direct consequence is Lemmas 4.3.4 and 4.3.1, and the fact that lfs is 1-Lipschitz function.

Lemma 4.4.1 (Pointwise approximation) *Under conditions C1 and C2 of Theorem 4.0.3, and ε sufficiently, we have $\text{dist}(x, \pi|_{\hat{\mathcal{M}}}^{-1}(x)) = 2C\varepsilon^2 \text{lfs}(x)$. The constant C is defined in Lemma 4.3.4 and depends on k, ω_0, η_0 and Θ_0 .*

Proof From Lemma 4.3.1, π restricted to $\hat{\mathcal{M}}$ is a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$. In this proof we will only consider π restricted to $\hat{\mathcal{M}}$.

For $x \in \mathcal{M}$, let $x' = \pi^{-1}(x) \in \hat{\mathcal{M}}$. Let p be a point in P such that $x' \in \text{st}(p, \hat{\mathcal{M}})$. From Lemma 4.3.4, we have

$$\|x - x'\| \leq C\varepsilon^2 \text{lfs}(p). \quad (4.29)$$

Using the fact that $\|p - x'\| \leq C_2\varepsilon \text{lfs}(p)$ (from Lemma 3.4.3 (1)) and Eq. (4.29),

$$\|p - x\| \leq \|p - x'\| + \|x - x'\| \leq (C_2\varepsilon + C\varepsilon^2) \text{lfs}(p)$$

Using the fact that lfs is 1-Lipschitz and ε sufficiently small, we have

$$\text{lfs}(x) \geq \text{lfs}(p) - \|p - x\| \geq (1 - C_2\varepsilon - C\varepsilon^2) \text{lfs}(p) \geq \text{lfs}(p)/2 \quad (4.30)$$

Using Eq. (4.29) and (4.30), $\|x - x'\| \leq 2C\varepsilon^2 \text{lfs}(x)$. $\square$

4.5 Isotopy

We will use the following structural result from [GW04a].

Lemma 4.5.1 *For $p \in \mathcal{M}$ and $u \in N_p \mathcal{M}$, $\|u\| = 1$. Then the ball B of radius $\text{lfs}(p)$ centered at $p + \text{lfs}(p)u$ does not contain any points from $\mathcal{M}$ in its interior.*

Using Lemma 4.5.1, we get the following result:

Lemma 4.5.2 *Assume conditions C1 and C2 of Theorem 4.0.3. If ε is sufficiently small, for all points $x, y (\neq p)$ in $\hat{\mathcal{M}}$,*

$$[x \pi(x)] \cap [y \pi(y)] = \emptyset$$

Proof From Lemma 4.3.1, if ε is sufficiently small, then the restriction of π to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$. Moreover, from Lemma 4.4.1 and for ε sufficiently small, we have for all $x \in \hat{\mathcal{M}}$, $\|x - \pi(x)\| \leq 2C\varepsilon^2 \text{lfs}(\pi(x)) < \text{lfs}(\pi(x))$.

To reach a contradiction let's assume there exist $x, y (\neq x) \in \hat{\mathcal{M}}$ such that $[x\pi(x)] \cap [y\pi(y)]$, and let $z = [x\pi(x)] \cap [y\pi(y)]$. As π restricted to $\hat{\mathcal{M}}$ is a homeomorphism, $\pi(x) \neq \pi(y)$. Since π is a projection onto the manifold $\mathcal{M}$, z belongs to the affine spaces $N_{\pi(x)}\mathcal{M}$ and $N_{\pi(y)}\mathcal{M}$. Without loss of generality, we can assume that $\|\pi(y) - z\| \leq \|\pi(x) - z\|$. We observe that $\pi(x) \neq z$ as

$$\pi(x) = z \implies \pi(x) = \pi(y) \implies x = y$$

The 2nd last inequality follows from the definition of π and the last inequality follows from the fact that π restricted to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$.

Let u be a unit vector in $N_{\pi(x)}\mathcal{M}$ oriented from x to $\pi(x)$. Note that

$$\pi(x), \pi(y) \in \bar{B}(z, r) \subset \bar{B}(z_1, r_1)$$

where $r = \|x - \pi(x)\|$, $r_1 = \text{lfs}(\pi(x))$ and $z_1 = \pi(x) + \text{lfs}(\pi(x))u$. As $\partial B(z, r) \cap \partial B(z_1, r_1) = \pi(x)$, we have $\pi(y) \in B(z_1, \text{lfs}(\pi(x)))$. We have reached a contradiction from Lemma 4.5.1. $\square$

Consider the following map:

$$F : \hat{\mathcal{M}} \times [0, 1] \rightarrow \mathbb{R}^d, \text{ with } F(x, t) = x + t(\pi(x) - x).$$

We will denote by $F_t = F(\cdot, t)$ and $M_t = F_t(\hat{\mathcal{M}})$ for all $t \in [0, 1]$. It is easy to see that $F(x, t)$ (and F_t) is continuous, $M_0 = \hat{\mathcal{M}}$, and since π restricted to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and $\mathcal{M}$, $M_1 = \mathcal{M}$.

Using the facts that $\hat{\mathcal{M}}$ is compact, M_t is a Hausdorff space (as M_t is a subspace with the subspace topology of $\mathbb{R}^d$), and F_t is continuous, we get, using Theorem 2.5.7, the following result:

Lemma 4.5.3 *Assume conditions C1 and C2 of Theorem 4.0.3. If ε is sufficiently small, then for all $t \in [0, 1]$, $F_t : \hat{\mathcal{M}} \rightarrow M_t$ is a homeomorphism.*

This implies:

Lemma 4.5.4 (Isotopy) *Assume conditions C1 and C2 of Theorem 4.0.3. If ε is sufficiently small then*

$$F : \hat{\mathcal{M}} \times [0, 1] \longrightarrow \mathbb{R}^d$$

is an isotopy.

Chapter 5

Sampling and meshing of submanifolds

We propose an algorithm to sample and mesh a k -submanifold $\mathcal{M}$ of positive reach embedded in $\mathbb{R}^d$. The algorithm first constructs a crude sample of $\mathcal{M}$. It then refines the sample according to a prescribed parameter ε , and builds a mesh that approximates $\mathcal{M}$. Differently from most algorithms that have been developed for meshing surfaces of $\mathbb{R}^3$, the refinement phase does not rely on a subdivision of $\mathbb{R}^d$ (such as a grid or a triangulation of the sample points) since the size of such scaffoldings depends exponentially on the ambient dimension d . Instead, we only compute local stars consisting of k -dimensional simplices around each sample point. By refining the sample, we can ensure that all stars become coherent leading to a k -dimensional triangulated manifold $\hat{\mathcal{M}}$. The algorithm uses only simple numerical operations. We show that the size of the sample is $O(\varepsilon^{-k})$ and that $\hat{\mathcal{M}}$ is a good triangulation of $\mathcal{M}$. More specifically, we show that $\mathcal{M}$ and $\hat{\mathcal{M}}$ are isotopic, that their Hausdorff distance is $O(\varepsilon^2)$ and that the maximum angle between their tangent bundles is $O(\varepsilon)$. The asymptotic complexity of the algorithm is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (for fixed $\mathcal{M}$, d and k).

5.1 Introduction

We intend to sample and mesh a k -manifold $\mathcal{M}$ of positive reach embedded in $\mathbb{R}^d$. Manifolds of positive reach have been introduced by Federer [Fed59, Fed69] and include in particular C^2 manifolds. By mesh, we mean an embedded polyhedral approximation of $\mathcal{M}$ made up of simplices. We are especially interested in the case where the dimension k of $\mathcal{M}$ is much smaller than d , and intend to design an algorithm whose complexity depends on k rather than on d . Applications can be found in scientific computing for solving partial differential equations where the domain of interest has the structure of a manifold, in dynamical systems for computing the topology of space attractors, and in statistics and machine learning to approximate statistical manifolds.

Related work. The problem of triangulating manifolds has a long history in the mathematical literature. In differential topology, seminal contributions are due to Whitney [Whi57a], Cairns [Cai61], Munkres [Mun66], Whitehead [Whi40] to name a few. Although these papers are not of an algorithmic nature, they introduce and study several interesting concepts that have been extensively used in Computational Geometry recently such as Voronoi diagrams restricted to a manifold, ε -rch sample of a manifold, fat (or

thick) triangulations. However, these papers do not discuss the geometric quality of the approximation nor the size of the sample. The optimal sampling and approximation of convex bodies is also a long standing problem in convex optimization with major contributions by Gruber [Gru93, Gru04] and Dudley [Dud74]. Recently, Clarkson [Cla06] extended this line of work to non-convex smooth manifolds of arbitrary dimensions. However, his algorithm follows an intrinsic point of view which makes it difficult to use in practice since it requires to compute geodesic distances on the manifold which may be quite complicated in practice [PC05]. Other, more practical algorithms for approximating convex bodies, including the well-known sandwich algorithm, have been analyzed by Kamenev [Kam08]. We are not aware of similar studies for non convex manifolds except for the case of surfaces embedded in $\mathbb{R}^3$ which has been extensively studied in the Computational Geometry literature. See [CG06] for a recent survey. These methods start by computing some subdivision of the embedding space (such as a grid or a triangulation of the sample points) and their direct extension to higher dimensions would face an exponential dependence on d . A step in this direction is the extension of the celebrated Marching Cube algorithm to manifolds of higher dimensions [Min03, BWC02]. Continuation methods do not use any subdivision of the ambient space and are close in spirit to our approach. They construct a triangulated approximation of a k -dimensional submanifold in a greedy way and extend the current k -dimensional triangulated domain by adding a neighborhood of a boundary point. Some experimental results can be found in [Hen02] but no theoretical analysis of continuation methods is available.

Our approach. In this chapter, we will follow the extrinsic approach but show that we can avoid using any d -dimensional data structure (except in the initialization step). We will extend a technique developed for anisotropic mesh generation [BWY08] and build on results from Chapter 3 investigating the related problem of manifold reconstruction [BG11].

We assume that the manifold $\mathcal{M}$ to be meshed has a positive reach and that we know a lower bound on the reach. In addition we assume that we can compute, for each point $p \in \mathcal{M}$, the k -dimensional tangent space $T_p\mathcal{M}$ of $\mathcal{M}$ at p .

The algorithm starts with a sufficiently dense sample of $\mathcal{M}$ and then refines the sample and builds a mesh that approximates $\mathcal{M}$ so as to satisfy a prescribed sampling rate ε . The size of the initial sample does not depend on ε but only on $\mathcal{M}$. For each sample point $p \in P$, we compute its k -dimensional *star*, like in Chapter 3, in the restriction of the d -dimensional Delaunay triangulation of the sample P to the tangent space $T_p\mathcal{M}$ at p . Such a star can be computed in the k -dimensional flat $T_p\mathcal{M}$ once we have projected P onto $T_p\mathcal{M}$.

As we have seen in Chapter 3, in general, the stars do not glue coherently and it may well happen that q is a vertex in the star of p while p is not a vertex in the star of q . The crucial observation is that by refining the sample, we can ensure that all the stars become coherent leading to a k -dimensional mesh $\hat{\mathcal{M}}$. For ε small enough, we show that the size of the sample is $O(\varepsilon^{-k})$ and that $\hat{\mathcal{M}}$ is a good approximation of $\mathcal{M}$. Specifically, we show that $\mathcal{M}$ and $\hat{\mathcal{M}}$ are isotopic, that their Hausdorff distance is $O(\varepsilon^2)$ and that the maximum angle between their tangent spaces is $O(\varepsilon)$. Our bound on the Hausdorff distance matches the lower bound of Clarkson [Cla06] (up to a multiplicative constant that depends on $\mathcal{M}$). The bound on the angle between the tangent spaces seems to be new.

To refine the mesh according to a sampling parameter ε , we need an *oracle* to query the manifold and to compute new points on $\mathcal{M}$. This is a critical issue with respect to practical efficiency. In our algorithm, we only need to compute a point in the (0-dimensional) intersection of $\mathcal{M}$ with a $(d - k)$ -flat. Except from the oracle and the projection of points onto k -dimensional flats (the tangent spaces at the points of P), all computations are performed in those k -flats. As a consequence, the asymptotic complexity of the algorithm is $O(\varepsilon^{-k^2-k})$ for fixed k , d , and $\mathcal{M}$. Hence, while our approach is extrinsic, the ambient dimension appears only as an additive constant hidden in the big- O .

This work combines four main ideas that have been introduced separately before : the general mechanism of Delaunay refinement [Che97, Rup95, BO05], the concept and properties of Delaunay triangulations restricted to a manifold [ACDL02a, CDR05a, Ede01], the notion of tangential Delaunay complex [BF04, BG11], and a perturbation technique due to Li to remove flat simplices [Li03a]. Several of the structural results we need are borrowed from [BG11]. However, the algorithm in [BG11] takes as input a given set of points while here the algorithm has to construct the sample as well as the mesh, which makes the algorithm different and its analysis more delicate. This work also aims at clarifying the basic operations that are required to triangulate a manifold.

Organization of the chapter. We recall in Section 5.2 the definition of tangential Delaunay complex in the unweighted case. More general definition is given in Chapter 3. This complex is embedded in $\mathbb{R}^d$ but is not in general a k -dimensional triangulation due to the presence of so-called inconsistent configurations to be studied in Section 5.2. To remove inconsistent configurations, we propose an algorithm that refines the complex.

The algorithm is described in Section 5.3 and analyzed in Section 5.4.

Lastly, in Section 5.5, we show that the output of the algorithm is a good approximation of $\mathcal{M}$ using results from Chapter 4.

5.2 Definitions and preliminaries

This section recalls some definitions and results of tangential Delaunay complexes, borrowed from Chapter 3.

Note that in this chapter we will only use the notation of unweighted tangential Delaunay complexes. For completeness, proof of the results for this special case will be given in the appendix.

5.2.1 Revisiting tangential Delaunay, unweighted

Let P be a finite set of points on $\mathcal{M}$ and $\text{Del}(P)$ be the d -dimensional Delaunay triangulation of P , i.e. the collection of all the simplices with vertices in P that admit an empty circumscribing d -dimensional ball. A ball (or more generally any domain of $\mathbb{R}^d$) is called *empty* if its interior contains no point of P . Let in addition $\text{Del}_{p_i}(P)$ be the Delaunay triangulation of P restricted to the tangent space $T_{p_i}\mathcal{M}$, i.e. the collection of all the simplices with vertices in P that admit an empty circumscribing d -dimensional ball centered on $T_{p_i}\mathcal{M}$. Equivalently, the simplices of $\text{Del}_{p_i}(P)$ are the simplices of $\text{Del}(P)$ whose Voronoi dual face intersect $T_{p_i}\mathcal{M}$. Observe that $\text{Del}_{p_i}(P)$ is in general a k -dimensional complex and can always be made k -dimensional by applying some infinitesimal perturbation on

P. We will assume that the points of P are in *general position* in the rest of the chapter, meaning that all $\text{Del}_{p_i}(P)$ are k -dimensional triangulations. Finally, write $\text{star}(p_i)$ for the *star* of p_i in $\text{Del}_{p_i}(P)$, i.e. the set of simplices that are incident to p_i in $\text{Del}_{p_i}(P)$.

We recall the definition of tangential Delaunay complex and some known results from [BF04, BG11].

Definition 5.2.1 (Tangential Delaunay complex) *We call tangential Delaunay complex the simplicial complex $\text{Del}_{T\mathcal{M}}(P) = \{\tau, \tau \in \text{star}(p), p \in P\}$.*

Plainly, $\text{Del}_{T\mathcal{M}}(P)$ is a subcomplex of $\text{Del}(P)$. The following easy lemma is crucial since it shows that computing the tangential Delaunay complex reduces to computing n weighted Delaunay triangulations in k -dimensional flats if n denotes the cardinality of P.

We denote by $\pi_i : P \rightarrow T_{p_i}\mathcal{M}$ the orthogonal projection of P onto $T_{p_i}\mathcal{M}$ and by $F_i : P \rightarrow T_{p_i}\mathcal{M} \times \mathbb{R}$ the 1-1 mapping that associates to a point p the weighted point defined by $F_i(p) = (\pi_i(p), \sqrt{\beta^2 - \|\pi_i(p) - p\|^2})$ where

$$\beta = \max_{x \in P} \|\pi_i(x) - x\|.$$

Following lemma is a direct consequence of Lemma 2.4.2.

Lemma 5.2.2 *$\text{Del}_{p_i}(P)$ is the pullback by F_i of the k -dimensional weighted Delaunay triangulation of $F_i(P)$.*

Let $\tau = [p_0, \dots, p_k]$ be a k -simplex with vertices in $\mathcal{M}$ and let denote by $B_{p_i}(\tau)$ the d -dimensional ball circumscribing τ that is centered on $T_{p_i}\mathcal{M}$. The corresponding center and radii are denoted by $c_\tau(p_i)$ and $R_\tau(p_i)$. The following lemma, which is a variant of Lemma 10 in [BG11], bounds the size of the simplices of $\text{Del}_{T\mathcal{M}}(P)$ as a function of the sampling density. See Appendix B.1 for a proof.

Lemma 5.2.3 *Let P be an ε -rch sample of a manifold $\mathcal{M}$ with $\varepsilon \leq 1/8$. Then we have:*

1. $\text{Vor}(p) \cap T_p\mathcal{M} \subseteq B(p, 4\varepsilon \text{rch}(\mathcal{M}))$.
2. *for any k -simplex $\tau \in \text{star}(p)$, $R_p(\tau) \leq 4\varepsilon \text{rch}(\mathcal{M})$.*
3. *for all edges $pq \in \text{Del}_{T\mathcal{M}}(P)$, $\|p - q\| \leq 8\varepsilon \text{rch}(\mathcal{M})$.*

In general, the tangential Delaunay complex is *not* a triangulated k -manifold. This is due to the presence of so-called inconsistent simplices. Let τ be a k -simplex in the star of p_i which is not in the star of p_j . Equivalently, the Voronoi $(d - k)$ -dimensional face $\text{Vor}(\tau)$ dual to τ intersects $T_{p_i}\mathcal{M}$ (at a point $c_{p_i}(\tau)$) but does not intersect $T_{p_j}\mathcal{M}$. 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 $\text{aff}(\text{Vor}(\tau))$ with $T_{p_j}\mathcal{M}$. 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. Accordingly, the line segment $[c_{p_i}(\tau)c_{p_j}(\tau)]$ intersects the interior of some Voronoi cells (in particular, 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 configuration*.

Definition 5.2.4 (Inconsistent configuration) Let $\phi = [p_1, p_2, \dots, p_{k+2}]$ and let $p_i, p_k, p_l \in \phi$. We say that ϕ is an inconsistent configuration of $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ witnessed by p_i, p_j, p_l if

1. The k -simplex $\tau = \phi \setminus \{p_l\}$ is in $\text{star}(p_i)$ but not in $\text{star}(p_j)$, i.e. $T_{p_i}\mathcal{M} \cap \text{Vor}(\tau) \neq \emptyset$ and $T_{p_j}\mathcal{M} \cap \text{Vor}(\tau) = \emptyset$.
2. $\text{Vor}(p_l)$ is one of the first cell of $\text{Vor}(\mathcal{P})$ whose interior is intersected by the line segment $[c_{p_i}(\tau)c_{p_j}(\tau)]$, where $c_{p_i}(\tau) = T_{p_i}\mathcal{M} \cap \text{Vor}(\tau)$ and $c_{p_j}(\tau) = T_{p_j}\mathcal{M} \cap \text{aff}(\text{Vor}(\tau))$, and $[c_{p_i}(\tau)c_{p_j}(\tau)]$ is oriented from $c_{p_i}(\tau)$ to $c_{p_j}(\tau)$.

Write i_ϕ for the first point of $\text{Vor}(p_l)$ hit by the oriented segment $[c_{p_i}(\tau)c_{p_j}(\tau)]$. i_ϕ is the center of an empty d -dimensional ball that circumscribes ϕ . Hence ϕ is a Delaunay $(k+1)$ -simplex and i_ϕ is the point on $[c_{p_i}(\tau)c_{p_j}(\tau)]$ that belongs to $\text{Vor}(\phi)$, the Voronoi face dual to ϕ . Since we assumed that the points are in general position, an inconsistent configuration cannot belong to the tangential Delaunay complex (which does not contain faces of dimension greater than k). Observe also that some of the subfaces of an inconsistent configuration may not belong to the tangential Delaunay complex.

We will use the same notations for inconsistent configurations as for simplices, e.g. r_ϕ and c_ϕ for the circumradius and the circumcenter of ϕ , ρ_ϕ and Θ_ϕ for its radius-edge ratio and fatness respectively. We also write $\widetilde{R}_\phi = \|i_\phi - p_i\|$, where p_i is a vertex of ϕ . Note that $\widetilde{R}_\phi = \|i_\phi - p_i\| \geq \|c_\phi - p_i\| = R_\phi$.

The following important lemma bounds the radius and fatness of an inconsistent configuration.

Lemma 5.2.5 Let ϕ be an inconsistent configuration witnessed by p, q and r , and let $\tau = \phi \setminus \{r\}$. Assume that $R_\phi < \text{rch}(\mathcal{M})/4$, and write $\theta = \max \theta_x$ where $\theta_x = \angle(\text{aff}(\tau), T_x\mathcal{M})$ and x is a vertex of τ ($\sin \theta \leq \frac{4\rho_\tau R_\tau}{\Theta_\tau \text{rch}(\mathcal{M})}$ by Corollary 2.3.3). We have

1. $\widetilde{R}_\phi \leq \frac{R_\tau}{\cos \theta}$.
2. $\Theta_\phi \leq \frac{R_\tau}{\text{rch}(\mathcal{M}) \cos \theta} \left(1 + \frac{4\rho_\tau}{\Theta_\tau}\right)$.

5.3 Algorithm

We now describe our meshing algorithm. The algorithm assumes that we know the dimension k of $\mathcal{M}$ and that we can get the tangent space $T_p\mathcal{M}$ at any point $p \in \mathcal{M}$. In addition, we assume to know a positive lower bound on the reach of the manifold $\mathcal{M}$. We write it also $\text{rch}(\mathcal{M})$ for simplicity.

The algorithm takes as input parameters $\varepsilon, \rho_0 \geq 1/2, \Theta_0 < 1/2$. The sampling parameter ε will be used in Section 5.5 to bound the size of the sample and the approximation error. The two constants ρ_0 and Θ_0 are used below to define (Θ_0, ρ_0) -good simplices and (Θ_0, ρ_0) -slivers. These are variants of definitions introduced in Section 2.3 and will be useful for analyzing the refinement algorithm of this chapter. Additional parameters and conditions will be specified in Section 5.4.

The algorithm first constructs an initial sample P_0 of $\mathcal{M}$ of constant size. Then, it upsamples P_0 by inserting new points on $\mathcal{M}$ in a greedy way so as to satisfy a sampling

condition expressed in terms of parameter ε , and making sure that all the stars are consistent.

We now detail the main features of the algorithm.

5.3.1 Primitive operations

We assume that the manifold is generic in the sense that the intersection of any $(d - k)$ -flat with the manifold is a bounded set of points. The only primitive operation of our algorithm that involves $\mathcal{M}$, namely $\mathbf{ints}(\mathcal{M}, F)$, computes the intersection of $\mathcal{M}$ with a $(d - k)$ -flat F . This primitive operation can be implemented for various representations of $\mathcal{M}$: e.g. when $\mathcal{M}$ is given implicitly as a system of $d - k$ algebraic equations, computing $\mathbf{ints}(\mathcal{M}, F)$ reduces to solving a 0-dimensional system of d -variate algebraic equations.

We will also need to pick random points in Euclidean balls of $\mathbb{R}^k$.

5.3.2 Computing the initial sample P_0

The construction of the initial sample P_0 can be done in various ways. We can use the continuation method of [Hen02] or use a simpler grid. We sketch the grid method which is easy to implement although the construction requires $2^{O(d \log d)}$ time. Take a uniform d -dimensional grid with cells of diameter $\text{rch}(\mathcal{M})/16$ and pick the intersection points between the manifold and the $(d - k)$ -faces of the grid to build a set $S \subset \mathcal{M}$ which is an $\frac{1}{32}$ -rch sample of $\mathcal{M}$. To make the sample sparse, we do the following:

1. Set $P_0 = \emptyset$ and $\bar{S} = S$;
2. Take a point p from $\bar{S}$, insert p in P_0 , and remove from $\bar{S}$ the points that belong to $B(p, \text{rch}(\mathcal{M})/32)$.
3. Repeat Step 2 until $\bar{S} = \emptyset$.

It is easy to see that the subsample $P_0 \subseteq S \subset \mathcal{M}$ will be a $(\frac{1}{16}, \frac{1}{32})$ -rch sample of $\mathcal{M}$.

5.3.3 Good simplices and slivers

We adapt the following definitions from [Li03a].

Definition 5.3.1 (((Θ_0, ρ_0) -good simplex) A simplex τ is a (Θ_0, ρ_0) -good simplex if $\rho_\tau \leq \rho_0$ and

$$\min_{\substack{\sigma \subseteq \tau, \\ \dim(\sigma) > 0}} \Theta_\sigma^{\frac{1}{\dim(\sigma)}} \geq \Theta_0,$$

where $\dim(\sigma)$ denotes the dimension of the simplex σ .

Definition 5.3.2 (((Θ_0, ρ_0) -sliver) A j -simplex τ is called a (Θ_0, ρ_0) -sliver if $j > 1$, $\rho_\tau \leq \rho_0$, $\Theta_\tau < \Theta_0^j$, and all of its proper subfaces are (Θ_0, ρ_0) -good simplices.

Remark 5.3.3 For the rest of this chapter when we talk about good simplex and slivers we will mean (Θ_0, ρ_0) -good simplex and (Θ_0, ρ_0) -sliver respectively.

The next lemma follows from Lemma 5.2.5 (2). It relates inconsistent configurations and slivers.

Lemma 5.3.4 *Let ϕ be an inconsistent configuration witnessed by p, q and r , and let τ be the k -dimensional simplex $\phi \setminus \{r\}$. Assume that $\rho_\phi \leq \rho_0$, $R_\tau \leq \varepsilon \text{rch}(\mathcal{M})$ and that the subfaces of ϕ are good simplices. Then, if*

$$\varepsilon \leq \frac{\Theta_0^{k+1}}{\sqrt{\left(1 + \frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$$

then ϕ is a sliver.

Proof From Lemma 5.2.5 (2), we have $\Theta_\phi \leq \frac{R_\tau}{\text{rch}(\mathcal{M})\cos\theta} \left(1 + \frac{4\rho_\tau}{\Theta_\tau}\right)$ where $\sin\theta \leq \frac{4\rho_\tau R_\tau}{\Theta_\tau \text{rch}(\mathcal{M})}$. Using the fact that $R_\tau \leq R_p(\tau) \leq \varepsilon \text{rch}(\mathcal{M})$, $\rho_\tau \leq \rho_0$, and $\Theta_\tau \geq \Theta_0^k$, we deduce that, if

$$\varepsilon \leq \frac{\Theta_0^{k+1}}{\sqrt{\left(1 + \frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$$

then $\Theta_\phi < \Theta_0^{k+1}$. The lemma follows. $\square$

Hence, if ε is small enough, removing all slivers of dimensions at most $k+1$ will result in removing inconsistencies from $\text{Del}_{T\mathcal{M}}(\mathcal{P})$.

This remark motivates the following definition. The *augmented complex*, defined below, will be maintained by the algorithm and slivers will be removed from this complex.

Definition 5.3.5 (Augmented complex) *The augmented complex $C(\mathcal{P})$ consists of the following simplices and their subfaces:*

- (i) *The k -simplices of $\text{Del}_{T\mathcal{M}}(\mathcal{P})$.*
- (ii) *The inconsistent configurations ϕ witnessed by p, q and r , such that 1. $R_p(\tau) \leq \varepsilon \text{rch}(\mathcal{M})$ and 2. $\tau = \phi \setminus \{p_l\}$ is a good simplex.*

5.3.4 Picking region and good points

A new point to be inserted is chosen so as to remove a bad simplex σ of $C(\mathcal{P})$. It will be taken from the so-called *picking region* of σ which we define now. We introduce two new parameters, $\beta > 1$ and $\alpha \in [0, 1)$.

Definition 5.3.6 (Picking region $\Pi(\sigma, \alpha)$) *We consider the following two cases:*

1. *If $\sigma = \tau$ is a k -dimensional simplex in $\text{star}(p)$, then the picking region of τ is defined as $\Pi(\tau, \alpha) = B(c_p(\tau), \alpha R_p(\tau)) \cap \mathcal{M}$.*
2. *If $\sigma = \phi$ is an inconsistent configuration, then the picking region of ϕ is defined as $\Pi(\phi, \alpha) = B(i_\phi, \alpha \bar{R}_\phi) \cap \mathcal{M}$.*

Definition 5.3.7 (Tiny sliver) A simplex τ is called a *tiny sliver* with respect to a simplex σ if τ is a sliver and $R_\tau \leq \beta R_\sigma$.

Definition 5.3.8 (Good point) A point x in a picking region $\Pi(\sigma, \alpha)$ is called a *good point* if inserting x does not create any j -dimensional sliver that is both incident to x and tiny with respect to σ , $j \leq k + 1$.

The algorithm makes use of the following two functions:

1. **pick**(x, p) : The function **pick**(x, p) takes as input two points $x \in \mathbb{R}^d$ and $p \in \mathcal{M}$. The function returns a point closest to x from the set $F \cap \mathcal{M}$, where F is the $(d - k)$ -dimensional flat passing through x and parallel to $N_p \mathcal{M}$.
2. **good-pick**(σ, α) : This function takes as input a simplex σ and $\alpha \in [0, 1)$. It returns a good point x in $\Pi(\sigma, \alpha)$. (Here σ can be a k -dimensional simplex of $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ or a $(k + 1)$ -dimensional inconsistent configuration of $\mathcal{C}(\mathcal{P})$.)

To implement **pick**(x, p), we use the primitive **ints**($\mathcal{M}, F$) to get the set of intersection points (generically finite) and then return the intersection point closest to x .

We implement **good-pick**(σ, α) as follows. If σ is a k -simplex τ in $\text{star}(p)$, we apply the following procedure:

- S1.** Pick a random point $y \in B(c_p(\tau), \alpha R_p(\tau)) \cap T_p \mathcal{M}$ and calculate $x = \text{pick}(y, p)$.
- S2.** If $x \in B(c_p(\tau), \alpha R_p(\tau))$ then go to **S3** else go back to **S1** and start over.
- S3.** We check if x forms a j -dimensional sliver τ_1 ($2 \leq j \leq k + 1$) with other sample points contained in the ball $B(c_p(\tau), \alpha R_p(\tau) + 2\beta R_\tau)$. If not, x is a *good point* and we return x . Otherwise, we go back to **S1** and start over.

Observe that **S3** prevents to create simplices incident to x that are tiny with respect to τ .

If σ is an inconsistent configuration ϕ , we proceed as follows. Let ϕ be witnessed by p, q and r . According to the definition of an inconsistent configuration, the k -dimensional simplex $\tau = \phi \setminus \{r\}$ belongs to $\text{star}(p)$ and not to $\text{star}(q)$. We implement **good-pick**(ϕ, α) as in Case 1 except that we pick random points from the k -dimensional ball $B(c_p(\tau), r) \cap T_p \mathcal{M}$ where $r = \alpha \widetilde{R}_\phi + \|i_\phi - c_p(\tau)\|$.

In Section 5.4, we will prove the existence of good points in $\Pi(\sigma, \alpha)$.

5.3.5 Refinement Algorithm

We can now give the details of the algorithm.

input a finite $(\frac{1}{16}, \frac{1}{32})$ -rch sample $\mathcal{P}_0$ of $\mathcal{M}$, and parameters $\varepsilon, \rho_0, \Theta_0, \beta$ and α (the parameters should be chosen so that they satisfy the conditions given in Theorem 5.4.9 in Section 5.4)

output a sample $\mathcal{P}$ of $\mathcal{M}$ and $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}(\mathcal{P})$

The refinement algorithm consists of applying the following rules. Rule (i) is only applied if Rule (j) with $j < i$ cannot be applied. Each rule kills a simplex σ (i.e. removes σ from a star) by inserting a new point in its picking region. To insert (or remove) a point means here to update P , as well as the augmented complex $C(P)$. We call *new simplex* a simplex of $C(P)$ that is created when inserting a new point.

Notice that all the new k -simplices in $\text{Del}_{T\mathcal{M}}(P)$ and all the new $(k+1)$ -simplices in $C(P)$ will be incident to the newly inserted point p . Observe however that a simplex (possibly a sliver) that existed in the Delaunay triangulation $\text{Del}(P)$ but not in $C(P)$ before the insertion of p may become a (subface of a new) simplex of $C(P)$ after the insertion of p .

Rule 1 *Big simplices* : if there exists a k -simplex τ in $\text{star}(p)$ s.t. $R_p(\tau) > \varepsilon \text{rch}(\mathcal{M})$, insert $x = \text{pick}(c_p(\tau), p)$.

Rule 2 *Bad radius-edge ratio* :

- a If there exists a k -simplex τ in $\text{star}(p)$ such that $R_\tau > \rho_0 L_\tau$, insert $x = \text{pick}(c_p(\tau), p)$.
- b Similarly, if ϕ is an inconsistent configuration witnessed by p, q and l , such that $R_\phi > \rho_0 L_\phi$, insert $x = \text{pick}(i_\phi, p)$.

Rule 3 *Type-1 sliver* : If there exists a k -simplex τ of $\text{Del}_{T\mathcal{M}}(P)$ that is a sliver or has a subsimplex that is a sliver, insert $x = \text{good-pick}(\tau, \alpha)$.

Rule 4 *Type-2 sliver* : If an inconsistent configuration $\phi \in C(P)$ is a sliver or has a subsimplex that is a sliver, insert $x = \text{good-pick}(\phi, \alpha)$.

Once the algorithm terminates, all slivers and inconsistent configurations have been killed. Hence, all stars are consistent and a simple sweep allows to merge all the stars into the final mesh $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}(P)$.

5.4 Analysis of the algorithm

To prove that the algorithm terminates, we first bound the volume of the so-called forbidden regions. This will be helpful in proving that there exist good points in the picking regions. Termination of the algorithm is then proved by showing that the interpoint distance remains bounded from below. Lastly, we analyze the time complexity of the algorithm.

The following lemma is proved in Appendix B.2.

Lemma 5.4.1 *Let p be a point on $\mathcal{M}$. There exist constants ξ and A that depend only on k such that, for all $t \leq \xi$ and $r = t \text{rch}(\mathcal{M})$, we have*

$$0 < 1 - A t \leq \frac{\text{vol}(B(p, r) \cap \mathcal{M})}{\phi_k r^k} \leq 1 + A t$$

where ϕ_k is the volume of the k -dimensional unit Euclidean ball.

5.4.1 Forbidden regions

For a given j -simplex ($1 \leq j \leq k$) with vertices on $\mathcal{M}$, the *forbidden region* F_μ of μ is defined as

$$F_\mu = \{x \in \mathcal{M} : \mu \cup \{x\} \text{ forms a } (j+1)\text{-dimensional sliver}\}.$$

Remember that μ must be a good simplex by definition of a sliver. We will now bound the volume of F_μ .

Lemma 5.4.2 *Let μ be a good j -dimensional simplex with $2 \leq j \leq k$ with vertices on $\mathcal{M}$, $R_\mu \leq t \operatorname{rch}(\mathcal{M})$. If (i) $t \leq \xi/2$ (ξ is defined in Lemma 5.4.1), (ii) $\left(\frac{4\rho_0}{\Theta_0^k} + 2\right)t < \Theta_0$ and (iii) $(B+1)\Theta_0 \leq 1$ for some B that depends on k and ρ_0 , then*

$$\operatorname{vol}(F_\mu) \leq D \Theta_0 R_\mu^k,$$

where D depends also on k and ρ_0 .

Proof Let $x \in F_\mu$ and x^* be the point closest to x on $\partial B(c_\mu, R_\mu) \cap \operatorname{aff}(\mu)$. We denote by τ the $(j+1)$ -dimensional simplex $\tau = \mu \cup \{x\}$, and τ is a $(j+1)$ -dimensional sliver since $x \in F_\mu$. From Lemma 2.3.1 (2) and (3), the facts that $\Delta_\tau \leq 2R_\tau \leq 2\rho_\tau L_\tau \leq 2\rho_\tau L_\mu \leq 4\rho_\tau R_\mu$, $\rho_\tau \leq \rho_0$ and $\frac{\Theta_\tau}{\Theta_\mu} < \Theta_0$ (since τ is a $(j+1)$ -dimensional sliver), we get

$$\begin{aligned} \|x - x^*\| &\leq b(\rho_\tau) \operatorname{dist}(x, \operatorname{aff}(\mu)) \\ &\leq (j+1)2^j \rho_\tau^j b(\rho_\tau) \frac{\Theta_\tau}{\Theta_\mu} \Delta_\tau \\ &\leq (j+1)2^{j+2} \rho_0^{j+1} b(\rho_0) \Theta_0 R_\mu \\ &\leq (k+1)2^{k+2} \max(1, \rho_0)^{k+1} b(\rho_0) \Theta_0 R_\mu \stackrel{\text{def}}{=} B \Theta_0 R_\mu \end{aligned} \quad (5.1)$$




Figure 5.1: For the proof of Lemma 5.4.2.

Let p be a vertex of μ . Let c be the point closest to c_μ on $T_p\mathcal{M}$ and c^* be the point closest to c on $\mathcal{M}$ (see Figure 5.1). From Corollary 2.3.3, we have

$$\begin{aligned}
 \|c - c_\mu\| &\leq \sin \angle(T_p\mathcal{M}, \text{aff}(\mu)) \times R_\mu \\
 &\leq \frac{4\rho_0 t}{\Theta_0^j} \times R_\mu \\
 &\leq \frac{4\rho_0 t}{\Theta_0^k} \times R_\mu \quad \text{since } \Theta_0 < 1 \\
 &\stackrel{\text{def}}{=} C t R_\mu
 \end{aligned} \tag{5.2}$$

From Lemma 2.2.2 (2) we have

$$\|c - c^*\| \leq \frac{2\|c - p\|^2}{\text{rch}(\mathcal{M})} \leq \frac{2R_\mu^2}{\text{rch}(\mathcal{M})} \leq 2t R_\mu. \tag{5.3}$$

Using the fact that $\|c_\mu - x^*\| = R_\mu$ and Eq. (5.1), (5.2) and (5.3), we get

$$\begin{aligned}
 \|c^* - x\| &\leq \|c^* - c\| + \|c - c_\mu\| + \|c_\mu - x^*\| + \|x^* - x\| \\
 &\leq R_\mu(1 + (B\Theta_0 + (C + 2)t)) \\
 &< R_\mu(1 + (B + 1)\Theta_0),
 \end{aligned}$$

the last inequality follows from hypothesis (ii), which implies $(C + 2)t \leq \Theta_0$. We can similarly prove that $\|c - x\| \geq r(1 - (B + 1)\Theta_0)$.

Writing $\delta = (B + 1)\Theta_0 \leq 1$ (from hypothesis (iii)), we deduce from the inequalities above that $\|c^* - x\| \in [R_\mu(1 - \delta), R_\mu(1 + \delta)]$. Therefore, the forbidden region F_μ is included in $B(c^*, R_\mu(1 + \delta)) \cap \mathcal{M} \setminus B(c^*, R_\mu(1 - \delta)) \cap \mathcal{M}$. We now use Lemma 5.4.1 to bound the volume of F_μ . Observe that Lemma 5.4.1 can be applied since $R_\mu(1 + \delta) \leq 2R_\mu \leq 2t \text{rch}(\mathcal{M}) \leq \xi \text{rch}(\mathcal{M})$ (as $t \leq \xi/2$ and $\delta \leq 1$). We have

$$\begin{aligned}
 \frac{\text{vol}(F_\mu)}{\phi_k} &\leq \frac{\text{vol}(B(c^*, R_\mu(1 + \delta)) \cap \mathcal{M} \setminus B(c^*, R_\mu(1 - \delta)) \cap \mathcal{M})}{\phi_k} \\
 &\leq (1 + A(1 + \delta)t) R_\mu^k (1 + \delta)^k - (1 - A(1 - \delta)t) R_\mu^k (1 - \delta)^k \\
 &= R_\mu^k ((1 + \delta)^k - (1 - \delta)^k) + A t R_\mu^k ((1 + \delta)^{k+1} + (1 - \delta)^{k+1}) \\
 &\leq 2^k \delta R_\mu^k + A(2^{k+1} + 1) t R_\mu^k
 \end{aligned} \tag{5.4}$$

the last inequality follows from the fact that $(1 + x)^k - (1 - x)^k \leq 2^k x$ for $x \in [0, 1]$.

From hypothesis (ii) and the fact that $\Theta_0 < 1$, we have

$$t < \frac{\Theta_0}{\frac{4\rho_0}{\Theta_0^k} + 2} \leq \frac{\Theta_0}{4\rho_0 + 2}.$$

Using this inequality and Eq. (5.4), yields the result. $\square$

5.4.2 Proof of termination

To prove that the refinement algorithm terminates, we prove that the distance between any two points inserted by the algorithm is bounded away from 0, which is sufficient since we assumed that $\mathcal{M}$ is compact.

Remember that there are two types of simplices that are refined by the algorithm. Let σ denote a k -simplex in $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ or a $(k+1)$ -dimensional simplex in $C(\mathcal{P})$. A new point that is inserted in the picking region of σ is said to refine σ . We denote by $e(\sigma)$ the minimal distance between such a new point and the current sample.

We assume without loss of generality that $\text{rch}(\mathcal{M}) = 1$ for the rest of this section. We here give the hypotheses that will be used in this section.

$$\text{H1. } \beta \geq \frac{2}{1-\alpha}$$

$$\text{H2. } \rho_0 \geq \frac{4}{1-\alpha}$$

$$\text{H3. } \Theta_0 < \left\{ \frac{E\alpha^k}{N^{k+1}\beta^k D}, \frac{1}{B+1} \right\}$$

$$\text{H4. } \varepsilon < \min \left\{ \frac{\xi}{4\beta}, \frac{8\xi}{1+31\alpha+32\beta}, \frac{\alpha}{8(C+1)} \right\}$$

$$\text{H5. } \varepsilon \leq \frac{\Theta_0}{2\beta(C+2)}$$

In the hypotheses ξ is the constant defined in Lemma 5.4.1, B is defined in Lemma 5.4.2, $C = \frac{4\rho_0}{\Theta_0^k}$, D is defined in Lemma 5.4.2, E will be defined in Lemma 5.4.6, N will be defined in the proof of Lemma 5.4.8, and $\varepsilon, \alpha, \beta, \rho_0$ and Θ_0 are parameters of the algorithm.

Observe that once α is fixed in $[0, 1)$, β and ρ_0 can be fixed so as to satisfy H1 and H2. Then, we can fix Θ_0 so that H3 is satisfied, and lastly we can fix ε . H5 provides a trade-off between improving the quality of the simplices (by fixing a high Θ_0) and minimizing the size of the sample.

Lemma 5.4.3 *Let p be a point on $\mathcal{M}$ and q be a point on $T_p\mathcal{M}$ such that $\|p - q\| \leq 1/4$. Then $\|q - \text{pick}(q, p)\| \leq 2\|p - q\|^2$.*

Proof Let $A = B(p, r) \cap \mathcal{M}$ where $r = 2\|p - q\|$. Let $f : A \rightarrow T_p\mathcal{M}$ be the orthogonal projection map of A to $T_p\mathcal{M}$. It is proved in [NSW08b] (Lemma 5.3) that $B(p, r \cos \theta) \cap T_p\mathcal{M} \subseteq f(A)$ where $\sin \theta = \|p - q\| \leq 1/4$.

$\text{pick}(q, p)$ returns the point x closest to q in $\mathcal{M} \cap F$ where F is a $(d-k)$ -dimensional flat passing through q and parallel to $N_p\mathcal{M}$. Since $q \in B(p, r \cos \theta) \cap T_p\mathcal{M}$ and $B(p, r \cos \theta) \cap T_p\mathcal{M} \subseteq f(A)$, $x \in A$. Therefore, from Lemma 2.2.2 (1) and the fact that $\|p - x\| \leq 2\|p - q\|$, we have

$$\|q - x\| = \|p - x\| \sin(T_p\mathcal{M}, px) \leq 2\|p - q\|^2.$$

□

The following Lemmas 5.4.4, 5.4.5 and 5.4.8 will bound the minimum interpoint distance.

Lemma 5.4.4 (Rule 1) *If τ is a k -simplex of $\text{star}(p)$ for which Rule 1 is applied, i.e. $R_p(\tau) > \varepsilon$, then $e(\tau) \geq R_p(\tau)/2 > \varepsilon/2$.*

Proof Let $x = \mathbf{pick}(c_p(\tau), p)$ be the point inserted by Rule 1 to refine τ . Since P is a $1/16$ -sample of $\mathcal{M}$, it follows from Lemma 5.2.3 that $R_p(\tau) \leq 1/4$.

Using $\|c_p(\tau) - p\| = R_p(\tau)$, $R_p(\tau) \leq 1/4$, and Lemma 5.4.3, we get

$$\|c_p(\tau) - x\| \leq 2R_p(\tau)^2 \leq R_p(\tau)/2.$$

Therefore $x \in B(c_p(\tau), R_p(\tau)/2)$. For any vertex v inserted before x , we have

$$\|v - x\| \geq R_p(\tau) - \|c_p(\tau) - x\| \geq R_p(\tau)/2 > \varepsilon/2.$$

□

Lemma 5.4.5 (Rule 2) *Under Hypotheses H2 and H4, for a simplex σ being refined by Rule 2, i.e. $\sigma \in C(P)$ with $\rho_\sigma > \rho_0$, we have $e(\sigma) \geq R_\sigma/2 > \rho_0 L_\sigma/2 > 2L_\sigma$.*

Proof 1. Consider first the case where $\sigma = \tau$ is a k -simplex of $\text{star}(p)$ for some p . Let $x = \mathbf{pick}(c_p(\tau), p)$ be the point inserted for refining τ . Using the fact that P is a $1/16$ -sample of $\mathcal{M}$, and arguments similar to the ones used in the proof of Lemma 5.4.4, for any vertex v inserted before x , we have

$$\|v - x\| \geq R_p(\tau)/2 \geq R_\tau/2 > \rho_0 L_\tau/2 \geq 2L_\tau.$$

The last inequality follows from the fact $\rho_0 \geq 4$ (Hypothesis H2).

2. Consider now the case where $\sigma = \phi$ is an inconsistent configuration in $C(P)$ witnessed by p, q and r , and let $\tau = \phi \setminus \{r\}$ be a k -dimensional simplex. By definition of an inconsistent configuration, τ belongs to $\text{star}(p)$. Since ϕ belongs to $C(P)$, we have $R_p(\tau) \leq \varepsilon$ (by the definition of $C(P)$) and from Corollary 2.3.3, $\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq 4\rho_0 \varepsilon / \Theta_0^k = C\varepsilon$, as τ is a good simplex.

Let $x = \mathbf{pick}(i_\phi, p)$ be the point inserted by Rule 2 to refine ϕ . Let i' denote the projection of i_ϕ onto $T_p \mathcal{M}$ and $i'' = \mathbf{pick}(i', p)$.

From Hypothesis H4, we have $\varepsilon \leq \frac{\alpha}{8(1+C)}$ which implies $C\varepsilon < 1/2$ (a crude bound for simplicity). Using the same arguments as in the proof of Lemma 5.2.5 and $\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq C\varepsilon$, we have

$$\|c_\tau - c_p(\tau)\|, \|c_\tau - i_\phi\| \leq R_\tau \tan \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq \frac{C\varepsilon R_\tau}{\sqrt{1 - C^2 \varepsilon^2}} \leq 2C\varepsilon R_\tau \quad (5.5)$$

and

$$r_\phi \leq \widetilde{R}_\phi \leq r_\tau + \|c_\tau - i_\phi\| \leq (1 + 2C\varepsilon) R_\tau \leq 2R_\tau. \quad (5.6)$$

Using the facts that $\|p - i'\| \leq \widetilde{R}_\phi \leq 2R_\tau$, $R_\tau \leq R_p(\tau) \leq \varepsilon$, and Lemma 5.4.3, we have

$$\|i' - i''\| \leq 2\|p - i'\|^2 \leq 4\varepsilon \widetilde{R}_\phi. \quad (5.7)$$

Since i' is the projection of i_ϕ onto $T_p \mathcal{M}$, hence

$$\|i_\phi - i'\| \leq \|i_\phi - c_p(\tau)\| \leq \|i_\phi - c_\tau\| + \|c_\tau - c_p(\tau)\| \leq 4C\varepsilon R_\tau, \quad (5.8)$$

the last inequality follows from Eq. (5.5).

Since the line segments $i_\phi i'$, $i' i''$ are parallel to $N_p \mathcal{M}$, hence the line segment $i_\phi i''$ is parallel to $N_p \mathcal{M}$. From the definition of $x = \text{pick}(i_\phi, p)$, Eq. (5.7) and (5.8) and $\varepsilon \leq \frac{\alpha}{8(1+C)}$, we have

$$\begin{aligned} \|i_\phi - x\| &\leq \|i_\phi - i''\| \leq \|i_\phi - i'\| + \|i' - i''\| \\ &\leq 4\varepsilon \widetilde{R}_\phi + 4C\varepsilon R_\tau \leq 4\varepsilon(1+C)\widetilde{R}_\phi \leq \widetilde{R}_\phi/2. \end{aligned}$$

Let v be a vertex that has been inserted before x . Since $B(i_\phi, \widetilde{R}_\phi)$ is empty, $\|v - i_\phi\| \geq \widetilde{R}_\phi$ and therefore we have

$$\|v - x\| \geq \widetilde{R}_\phi/2 \geq R_\phi/2 > \rho_0 L_\phi/2 \geq 2L_\phi.$$

The last inequality again follows from the fact that $\rho_0 \geq 4$. $\square$

It follows that the shortest interpoint distance is not decreased when Rule 2 is applied.

To prove a similar result for Rule 3 and 4 we use a volume argument. The next lemma provides a lower bound on the volume of the picking regions.

Lemma 5.4.6 (Volume of $\Pi(\sigma, \alpha)$) *Under Hypotheses H1 and H4, if σ is a simplex to be refined by either Rules 3 or 4, we have*

$$\text{vol}(\Pi(\sigma, \alpha)) \geq E \alpha^k R_\sigma^k$$

where E is a constant > 0 and depends only on k .

Proof 1. Consider first the case where $\sigma = \tau$ is a k -simplex of $\text{star}(p)$; then $\Pi(\tau, \alpha) = B(c_p(\tau), \alpha R_p(\tau)) \cap \mathcal{M}$. Let c be the point of $\mathcal{M}$ closest to $c_p(\tau)$. Since τ is being refined by Rule 3, hence $R_p(\tau) \leq \varepsilon$. Therefore, from Lemma 2.2.2 (2), we get

$$\|c_p(\tau) - c\| \leq 2\varepsilon R_p(\tau). \quad (5.9)$$

From Hypothesis H4, we have $\varepsilon \leq \frac{\alpha}{8(1+C)}$, and $\varepsilon \leq \frac{\alpha}{8(C+1)}$ implies $\varepsilon < \alpha/8$. From Eq. (5.9) and the fact that $\varepsilon \leq \alpha/8$, we get

$$\Pi(\tau, \alpha) \supseteq B(c, (\alpha - 2\varepsilon)R_p(\tau)) \cap \mathcal{M} \supseteq B(c, \frac{3\alpha R_p(\tau)}{4}) \cap \mathcal{M}.$$

From the above inequality and Lemma 5.4.1, we then have

$$\begin{aligned} \text{vol}(\Pi(\tau, \alpha)) &\geq \text{vol}(B(c, \frac{3\alpha R_p(\tau)}{4}) \cap \mathcal{M}) \\ &\geq \frac{3^k}{4^k} \left(1 - \frac{3A\alpha\varepsilon}{4}\right) \phi_k \alpha^k R_p(\tau)^k \geq \frac{3^k}{4^k} \left(1 - \frac{3A\xi}{32}\right) \phi_k \alpha^k R_\tau^k. \end{aligned}$$

The last inequality follows from the facts that $\alpha \leq 1$, $R_p(\tau) \geq R_\tau$ and $\varepsilon \leq \frac{\xi}{8}$ (since from Hypothesis H4 we have $\varepsilon \leq \frac{\xi}{4\beta}$ and from Hypothesis H1 we have $\beta > 2$).

2. Consider now the case where $\sigma = \phi$ is an inconsistent configuration witnessed by p, q and r . Then $\Pi(\phi, \alpha) = B(i_\phi, \alpha \widetilde{R}_\phi) \cap \mathcal{M}$. From the definition of inconsistent configurations, the k -dimensional simplex $\tau = \phi \setminus \{r\}$ is in $\text{star}(p)$. Since ϕ belongs to $C(P)$,

we have $R_p(\tau) \leq \varepsilon$ and, since τ is a good simplex, we have from Corollary 2.3.3 that $\sin \angle(\text{aff}(\tau), T_p \mathcal{M}) \leq 4\rho_0 \varepsilon / \Theta_0^k = C\varepsilon$.

As in the proof of Lemma 5.4.5 we use the crude bound $C\varepsilon < 1/2$ which follows from Hypothesis H4; therefore Eq. (5.5) and (5.6) follow. Also, using $C\varepsilon < 1/2$ and Eq. (5.5), we have

$$R_p(\tau) \leq R_\tau + \|c_p(\tau) - c_\tau\| \leq (1 + 2C\varepsilon)R_\tau \leq 2R_\tau. \quad (5.10)$$

Let c denote the point of $\mathcal{M}$ closest to $c_p(\tau)$. As in the first part of the proof, we have $\|c - c_p(\tau)\| \leq 2\varepsilon R_p(\tau)$. Using Eq. (5.5), (5.6) and (5.10), and the fact that $\varepsilon \leq \frac{\alpha}{8(1+C)}$, we get $B(c, r) \cap \mathcal{M} \subseteq \Pi(\phi, \alpha)$ where

$$\begin{aligned} r &= \alpha \widetilde{R}_\phi - \|i_\phi - c\| \\ &\geq \alpha R_\phi - \|c - c_p(\tau)\| - \|c_p(\tau) - c_\tau\| - \|c_\tau - i_\phi\| \\ &\geq \alpha R_\phi - 2\varepsilon R_p(\tau) - 4C\varepsilon R_\tau \\ &\geq \alpha R_\phi - 4(1 + C)\varepsilon R_\tau \\ &\geq \frac{\alpha R_\phi}{2}. \end{aligned}$$

Moreover, we deduce from Eq. (5.6)

$$\frac{\alpha R_\phi}{2} \leq \frac{\alpha \widetilde{R}_\phi}{2} \leq \alpha R_\tau \leq \alpha \varepsilon$$

We then deduce, using Lemma 5.4.1,

$$\begin{aligned} \text{vol}(\Pi(\phi, \alpha)) &\geq \text{vol}(B(c, \alpha R_\phi/2) \cap \mathcal{M}) \\ &\geq \frac{1}{2^k} (1 - A\alpha\varepsilon) \phi_k \alpha^k R_\phi^k \geq \frac{1}{2^k} \left(1 - \frac{A\xi}{8}\right) \phi_k \alpha^k R_\phi^k. \end{aligned}$$

The last inequality again follows from the facts that $\alpha < 1$ and $\varepsilon \leq \xi/8$. $\square$

Lemma 5.4.7 *Let B_p be a ball of radius R centered at a point $p \in \mathcal{M}$ and let V be a maximal set of points of $B_p \cap \mathcal{M}$ such that the smallest interdistance between the points is not less than $2r$. If $R + r \leq \xi$, $|V| = \frac{1+A\xi}{1-A\xi} \left(\frac{R}{r} + 1\right)^k$.*

Proof Denote by B_x the ball centered at $x \in V$ of radius r . Plainly, for any $x \in V$, $B_x \subset B_p^+ = B(p, R + r)$ and for any $x, y \in V$, $B_x \cap B_y = \emptyset$. By Lemma 5.4.1, $\text{vol}(B_p^+ \cap \mathcal{M}) \leq (1 + A(R + r))\phi_k (R + r)^k$ and $\text{vol}(B_x) \geq (1 - Ar)\phi_k R^k$. It follows that the number of points of V is at most

$$\frac{1 + A(R + r)}{1 - Ar} \left(\frac{R + r}{r}\right)^k \leq \frac{1 + A\xi}{1 - A\xi} \left(\frac{R + r}{r}\right)^k.$$

$\square$

Lemma 5.4.8 (Rules 3 & 4) *Under Hypotheses H1 to H5, application of Rule 3 or 4 on a simplex σ does not decrease the interpoint distance to less than $\frac{(1-\alpha)\varepsilon}{4}$ and does not create any tiny slivers.*

Proof The proof is by induction. Specifically, we prove that the algorithm maintains the following *invariants*

Invariant 1 When refining a simplex σ using Rules 3 or 4, no tiny slivers of dimension $\leq k+1$ with respect to σ are created in $\text{Del}(P)$.

Invariant 2 The interpoint distance remains greater than $\frac{(1-\alpha)\varepsilon}{4}$.

We first consider the case when $\sigma = \tau$ is a k -simplex in $\text{star}(p)$ to be refined by application of Rule 3. The case of an inconsistent configuration to be refined by Rule 4 is similar. Note that, since Rule 1 has not been applied, $R_p(\tau) \leq \varepsilon$.

Invariant 1 First observe that Invariant 1 is maintained if τ is refined by inserting a *good* point in $\Pi(\tau, \alpha)$.

We now prove the existence of good points in $\Pi(\tau, \alpha)$. We first show that the set of points of P that can form tiny slivers with respect to τ in $\text{Del}(P)$ upon insertion of a point $x \in \Pi(\tau, \alpha)$ are at distance at most $\alpha R_p(\tau) + 2\beta R_\tau < (\alpha + 2\beta)\varepsilon$ from $c_p(\tau)$. Indeed, recall that a tiny sliver with respect to τ has a circumradius less than βR_τ and that a point $x \in \Pi(\tau, \alpha)$ is a good point if x does not form a tiny sliver (of dimension $\leq k+1$ and with respect to τ) with the sample points. Hence, it is enough to consider the points of P that belong to $IB = B(c_p(\tau), \alpha R_p(\tau) + 2\beta R_\tau)$ since all the new simplices in $\text{Del}(P)$ upon insertion of x are incident to x . This proves the claim.

From Lemma 5.4.6, we have $\Pi(\tau, \alpha) \neq \emptyset$. Let $c \in \Pi(\tau, \alpha)$, then

$$\|c_p(\tau) - c\| \leq \alpha R_p(\tau).$$

We deduce

$$IB = B(c_p(\tau), \alpha R_p(\tau) + 2\beta R_\tau) \subseteq B(c, (2\alpha + 2\beta)\varepsilon) \stackrel{\text{def}}{=} IB^+. \quad (5.11)$$

We now apply Lemma 5.4.7 to bound the number n of sample points in IB^+ . Set $R = (2\alpha + 2\beta)\varepsilon$, $r = \frac{(1-\alpha)\varepsilon}{8}$ and observe that $R + r = \frac{(1+15\alpha+16\beta)\varepsilon}{8} \leq \xi$ by Hypothesis H4. We then get

$$n \leq N \stackrel{\text{def}}{=} \frac{1 + A\xi}{1 - A\xi} \left(\frac{32(\alpha + \beta)}{1 - \alpha} + 1 \right)^k$$

Let μ be a good simplex with vertices in IB with $R_\mu \leq \beta R_\tau \leq \beta\varepsilon$. From Hypothesis H3, H4 and H5, $\Theta_0 < \frac{1}{B+1}$, $\beta\varepsilon \leq \xi/2$ and $(C+2)\beta\varepsilon < \Theta_0$. We can apply Lemma 5.4.2, from which we deduce

$$F_\mu \leq D\Theta_0 R_\mu^k \leq D\Theta_0 \beta^k R_\tau^k.$$

Consider the j -simplices, $j \leq k+1$, that are included in $IB \subseteq IB^+$. The total number of such simplices is at most N^{k+1} . Hence, the total volume of the forbidden regions associated to all those simplices is at most

$$W = N^{k+1} \times D\Theta_0 \beta^k R_\tau^k. \quad (5.12)$$

On the other hand, from Lemma 5.4.6 and the fact that $\varepsilon \leq \frac{\alpha}{8(1+C)}$ (Hypothesis H4), we know that

$$\text{vol}(\Pi(\tau, \alpha)) \geq E\alpha^k R_\tau^k.$$




Figure 5.2: For the proof of Lemma 5.4.8.

By Hypothesis H3, the volume W of all the forbidden regions is less than $\text{vol}(\Pi(\tau, \alpha))$, the volume of the picking region of τ . This proves the existence of good points in the picking region $\Pi(\tau, \alpha)$ of τ .

Invariant 2 We will now show that Invariant 2 is also maintained.

Let $\tau' \subseteq \tau$ denote a simplex of τ that is a sliver. We denote by $p(\tau')$ the simplex whose killing gave birth to τ' . Let us now prove that the interpoint distance remains at least $\frac{(1-\alpha)\varepsilon}{4}$ after the insertion of x from the picking region $\Pi(\tau, \alpha)$ of τ . Let x^* denote the point whose insertion killed $p(\tau')$. Observe that x^* is a vertex of τ' , and also of τ as $\tau' \subseteq \tau$. See Figure 5.2. We distinguish the following cases.

Case 1 $p(\tau')$ is a big simplex killed by application of Rule (1). According to Lemma 5.4.4, the lengths of the edges incident to x^* in τ' are greater than $\varepsilon/2$. The distance between x and the other points is thus greater than

$$(1-\alpha)R_p(\tau) \geq \frac{(1-\alpha)\Delta_{\tau'}}{2} \geq \frac{(1-\alpha)\varepsilon}{4}.$$

The last inequality follows from the fact that $\Delta_{\tau'}$ will be greater than the lengths of the edges of τ' incident to x^* , which in turn are $> \varepsilon/2$ since the radius of $p(\tau')$ is greater than ε and we insert the new point in $\Pi(p(\tau'), 1/2)$.

Case 2 $p(\tau')$ is a simplex with a bad radius-edge ratio killed by application of Rule (2). From Lemma 5.4.5, we have $\Delta_{\tau'} \geq R_{p(\tau')}/2 > \rho_0 L_{p(\tau')}/2$ and the distance between x and the other points is greater than

$$(1-\alpha)R_p(\tau) \geq \frac{(1-\alpha)\Delta_{\tau'}}{2} > \frac{(1-\alpha)\rho_0 L_{p(\tau')}}{4} \geq L_{p(\tau')} \geq \frac{(1-\alpha)\varepsilon}{4}.$$

The last two inequalities follow from Hypothesis H2 and the induction hypothesis respectively.

Case 3 $p(\tau')$ has been killed by application of Rule (3) or (4). The radius $R_{\tau'}$ is bigger than $\beta R_{p(\tau')}$ since, by the induction hypothesis, no tiny slivers have been created until this point. If $R_{\tau'} > \beta R_{p(\tau')}$ then the distance between x and the other points is thus greater than

$$\begin{aligned} (1 - \alpha)R_p(\tau) &\geq (1 - \alpha)R_{\tau} \geq (1 - \alpha)R_{\tau'} \geq (1 - \alpha)\beta R_{p(\tau')} \\ &\geq \frac{(1 - \alpha)\beta L_{p(\tau')}}{2} \geq L_{p(\tau')} \geq \frac{(1 - \alpha)\varepsilon}{4}. \end{aligned}$$

The last two inequalities follow from H1 and induction hypothesis respectively.

In all cases, the invariants are maintained after refinement of τ . This completes the proof of the lemma. The case of an inconsistent simplex ϕ to be refined by Rule (4) is similar. $\square$

We sum up the results of the section in the following theorem.

Theorem 5.4.9 *Under Hypotheses H1 to H5, the algorithm terminates. If, in addition,*

$$\text{H6. } \varepsilon \leq \frac{\Theta_0^{k+1}}{\sqrt{\left(1 + \frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$$

the algorithm removes all inconsistent configurations from $\text{Del}_{T\mathcal{M}}(\mathcal{P})$.

Proof Termination of the algorithm is a consequence of Lemmas 5.4.4, 5.4.5 and 5.4.8. The additional Hypothesis H6 and Lemma 5.3.4 show that all inconsistent configurations have been removed since we removed all slivers from the augmented complex $C(\mathcal{P})$. $\square$

5.4.3 Combinatorial complexity analysis

We assume that Hypotheses H1 to H6 are satisfied. Hence the algorithm terminates and $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}(\mathcal{P})$ has no inconsistencies. Before we prove the results, we define the *normalized volume* of $\mathcal{M}$ as follows:

$$V(\mathcal{M}) = \frac{\text{vol}(\mathcal{M})}{\text{rch}(\mathcal{M})^k} \quad (5.13)$$

We also assume in this section that $\alpha \leq 1/2$.

Theorem 5.4.10 *The number of points inserted by the algorithm is at most*

$$|\mathcal{P}| = \frac{2^{O(k)} V(\mathcal{M})}{\varepsilon^k}, \quad (5.14)$$

where the constant of proportionality in the big-O is an absolute constant.

Proof Let L_P denote the smallest interpoint distance of the point set P . From Lemmas 5.4.4, 5.4.5 and 5.4.8 and $\alpha \leq 1/2$, the minimum interpoint distance in P satisfies

$$L_P \geq \frac{(1-\alpha)\varepsilon \text{rch}(\mathcal{M})}{4} \geq \frac{\varepsilon \text{rch}(\mathcal{M})}{8}.$$

Hence, for any p, q ($p \neq q$) in P , we have $B(p, r) \cap B(q, r) = \emptyset$ where $r = \frac{\varepsilon \text{rch}(\mathcal{M})}{16}$. Using the fact that $\varepsilon \leq \frac{\xi}{8}$ (from Hypotheses H1 and H4) and Lemma 5.4.1, we have $\text{vol}(B(p, r) \cap \mathcal{M}) \geq (1 - \frac{A\xi}{16})\phi_k r^k \geq (1 - \frac{A\xi}{128})\phi_k r^k$. By a packing argument, we get

$$|P| \leq \frac{16^k \text{vol}(\mathcal{M})}{(1 - \frac{A\xi}{128})\phi_k \varepsilon^k \text{rch}(\mathcal{M})^k} = \frac{2^{O(k)}V(\mathcal{M})}{\varepsilon^k}.$$

□

The following lemma, which is a direct application of Propositions 6.2 and 6.3 from [NSW08b], will be used in the proof of Lemma 5.4.12.

Lemma 5.4.11 *Let $p, q \in \mathcal{M}$ with $\|p - q\| \leq \frac{\text{rch}(\mathcal{M})}{2}$, then $\sin \angle(T_p \mathcal{M}, T_q \mathcal{M}) \leq \sqrt{\frac{2\|p - q\|}{\text{rch}(\mathcal{M})}}$.*

We will use Lemmas 5.4.12 and 5.4.14 to calculate the time and space complexity of the algorithm in Theorem 5.4.17.

Lemma 5.4.12 *Let $p \in \mathcal{M}$. Then, $|B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap P| \leq \frac{2^{O(k)}}{\varepsilon^k}$ where the constant in the big-O is an absolute constant.*

Proof We will first show that $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ can be covered by $2^{O(k)}$ balls (where the constant in the big-O is an absolute constant) of radius $\frac{\text{rch}(\mathcal{M})}{6}$ centered on $\mathcal{M}$. Then we will show that $|B(x, \frac{\text{rch}(\mathcal{M})}{6}) \cap P|$ is less than $\frac{2^{O(k)}}{\varepsilon^k}$ (the constant again in the big-O is an absolute constant) for any point x on $\mathcal{M}$. Combining the two results, we will get our lemma.

1. Let S_1 be the maximal set of points in $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ such that $\|x - y\| \geq \frac{\text{rch}(\mathcal{M})}{3}$ for all x, y ($x \neq y$) in S_1 . By definition for all $x \in S_1$, the balls $B_x = B(x, \frac{\text{rch}(\mathcal{M})}{6})$ are disjoint. Also, these balls are contained in $B = B(p, r_1)$, where $r_1 = \frac{\text{rch}(\mathcal{M})}{2} + \frac{\text{rch}(\mathcal{M})}{6} = \frac{2\text{rch}(\mathcal{M})}{3}$.

Let us consider the k -dimensional balls $\widetilde{B}_x = B_x \cap T_p \mathcal{M} = B(x, \frac{\text{rch}(\mathcal{M})}{6}) \cap T_p \mathcal{M}$ for all $x \in S_1$, and $\widetilde{B} = B(p, r) \cap T_p \mathcal{M}$. The balls $\widetilde{B}_x$ are disjoint since the balls B_x are disjoint. From Lemma 2.2.2 (1), the distance of $x \in S_1$ to $T_p \mathcal{M}$ is

$$\text{dist}(x, T_p \mathcal{M}) = \|p - x\| \times \sin \angle(p, T_p \mathcal{M}) \leq \frac{\text{rch}(\mathcal{M})}{8}. \quad (5.15)$$

Using the fact that the radius of the balls B_x ($x \in S_1$) is $\frac{\text{rch}(\mathcal{M})}{6}$, and the above Eq. (5.15), we get that the k -dimensional balls $\widetilde{B}_x$ has squared radius

$$\frac{\text{rch}(\mathcal{M})^2}{6^2} - \text{dist}(x, T_p \mathcal{M})^2 \geq \left(\frac{1}{6^2} - \frac{1}{8^2}\right) \text{rch}(\mathcal{M})^2 \stackrel{\text{def}}{=} r_2^2.$$

We will now bound $|S_1|$ using a packing argument. As the balls $\widetilde{B}_x$, $x \in S_1$, are disjoint and contained in $\widetilde{B}$, therefore

$$|S_1| \leq \frac{r_1^k}{r_2^k} \stackrel{\text{def}}{=} N_1 = 2^{O(k)},$$

where the constant in the big- O is an absolute constant.

Since S_1 is a maximal set of points such that $\|x - y\| \geq \frac{\text{rch}(\mathcal{M})}{3}$ for all $x, y (\neq x) \in S_1$, we claim that

$$B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M} \subseteq \bigcup_{x \in S_1} B(x, \frac{\text{rch}(\mathcal{M})}{3}). \quad (5.16)$$

Otherwise if there exist $\tilde{x} \in B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M} \setminus \bigcup_{x \in S_1} B(x, \frac{\text{rch}(\mathcal{M})}{3})$ then $\|\tilde{x} - x\| \geq \frac{\text{rch}(\mathcal{M})}{3}$ for all $x \in S_1$. We have reached a contradiction since we have assumed that S_1 is a maximal set of point in $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ such that $\|x - y\| \geq \frac{\text{rch}(\mathcal{M})}{3}$ for all $x, y (x \neq y) \in S_1$.

We have shown that $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ (equation (5.16)) can be covered by $2^{O(k)}$ balls centered in $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ with radius $\frac{\text{rch}(\mathcal{M})}{3}$. Following the same method, we can show that $B(x, \frac{\text{rch}(\mathcal{M})}{3})$ can be covered by $2^{O(k)}$ (the constant in the big- O is an absolute constant) balls centered in $B(x, \frac{\text{rch}(\mathcal{M})}{3})$ of radius $\frac{\text{rch}(\mathcal{M})}{6}$. Therefore from Eq. (5.16) and the above bound, we get that $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$ can be covered by $2^{O(k)}$ balls of radius $\frac{\text{rch}(\mathcal{M})}{6}$ centered in $B(p, \frac{\text{rch}(\mathcal{M})}{2}) \cap \mathcal{M}$.

2. We will now show that for all $q \in \mathcal{M}$, $|B(q, \frac{\text{rch}(\mathcal{M})}{6}) \cap \mathcal{P}| \leq \frac{2^{O(k)}}{\varepsilon^k}$. As in the proof of Lemma 5.4.10, we have from Lemmas 5.4.4, 5.4.5 and 5.4.8 and $\alpha \leq 1/2$, $B(x, \frac{\varepsilon \text{rch}(\mathcal{M})}{16}) \cap B(y, \frac{\varepsilon \text{rch}(\mathcal{M})}{16}) = \emptyset$ for all $x, y (\neq x) \in \mathcal{P}$.

Let $\hat{r} = (\frac{\text{rch}(\mathcal{M})}{6} + \frac{\varepsilon \text{rch}(\mathcal{M})}{16})$, and $f : B(q, \hat{r}) \cap \mathcal{M} \rightarrow T_q \mathcal{M}$ be the projection map of $B(q, \hat{r}) \cap \mathcal{M}$ to $T_q \mathcal{M}$.

We will bound the volume of $f(B(x, \frac{\varepsilon \text{rch}(\mathcal{M})}{16}) \cap \mathcal{M})$, where $x \in B(q, \frac{\text{rch}(\mathcal{M})}{6}) \cap \mathcal{P}$, by using the same arguments used in the proof of Lemma 5.3 in [NSW08b].

Claim 5.4.13 *The projection map f satisfies the following: (i) f is injective, and (ii) the derivative df is nonsingular for all $x \in B(q, \hat{r}) \cap \mathcal{M}$.*

Proof 1. Let θ be the angle made by the segment $[x_1, x_2]$ with $T_q \mathcal{M}$, where $x_1, x_2 \in B(q, \hat{r}) \cap \mathcal{M}$. Using Lemmas 2.2.2 and 5.4.11 and the fact that $\varepsilon < 1$, we have

$$\begin{aligned} \sin \theta &\leq \sin \angle(x_1 x_2, T_{x_1} \mathcal{M}) + \sin \angle(T_{x_1} \mathcal{M}, T_q \mathcal{M}) \\ &\leq \frac{\|x_1 - x_2\|}{2 \text{rch}(\mathcal{M})} + \sqrt{\frac{2\|x_1 - q\|}{\text{rch}(\mathcal{M})}} \leq \frac{1}{6} + \frac{\varepsilon}{16} + \sqrt{\frac{1}{3} + \frac{\varepsilon}{8}} < 1 \end{aligned} \quad (5.17)$$

This implies f is injective. Otherwise there will exist two points $x_1, x_2 \in B(q, \hat{r}) \cap \mathcal{M}$ such that the line segment $[x_1, x_2]$ is orthogonal to $T_q \mathcal{M}$, but this is not possible from Eq. (5.17).

2. If df is singular at some point $x \in B(q, \hat{r}) \cap \mathcal{M}$, then the line segment $[x, f(x)]$ lies in $T_x \mathcal{M}$. As f is the projection map onto $T_q \mathcal{M}$, therefore $[x, f(x)]$ is parallel to $N_q \mathcal{M}$. Since

the segment $[x, f(x)]$ is orthogonal to $T_q\mathcal{M}$ and lies on $T_x\mathcal{M}$, we have $\angle(T_q\mathcal{M}, T_x\mathcal{M}) = \pi/2$. But from Lemma 5.4.11 and $\varepsilon < 1$, we have

$$\sin \angle(T_x\mathcal{M}, T_q\mathcal{M}) \leq \sqrt{\frac{2\|q-x\|}{\text{rch}(\mathcal{M})}} \leq \sqrt{\frac{1}{3} + \frac{\varepsilon}{8}} < 1.$$

We have reached a contradiction. $\square$

We will now bound the $\text{vol}(f(B_x))$ where $B_x = B(x, \frac{\varepsilon \text{rch}(\mathcal{M})}{16}) \cap \mathcal{M}$, for all $x \in B_{\mathcal{M}}(q, \frac{\text{rch}(\mathcal{M})}{3}) \cap \mathcal{P}$. Let θ_x is the maximal angle made by any secant $s = [x, y]$ with $T_q\mathcal{M}$ where $y \in \bar{B}_x = \bar{B}_{\mathcal{M}}(x, \frac{\varepsilon \text{rch}(\mathcal{M})}{16})$. From Lemmas 2.2.2, 5.4.11, and $\varepsilon < 1$, we get

$$\begin{aligned} \sin(\theta_x) &\leq \max_{y \in \bar{B}_x} \sin \angle(xy, T_x\mathcal{M}) + \sin \angle(T_x\mathcal{M}, T_q\mathcal{M}) \\ &\leq \max_{y \in \bar{B}_x} \frac{\|x-y\|}{2\text{rch}(\mathcal{M})} + \sqrt{\frac{2\|q-x\|}{\text{rch}(\mathcal{M})}} \leq \frac{\varepsilon}{32} + \sqrt{\frac{1}{3}} < 0.80 \end{aligned} \quad (5.18)$$

Since f is nonsingular at x and therefore locally invertible, hence there exists a ball of radius r centered on x such that $f^{-1}(B(f(x), r) \cap T_q\mathcal{M}) \subseteq B_x$. Let r_x denote the maximal radius such that for all $r < r_x$, we have $f^{-1}(B(f(x), r_x) \cap T_q\mathcal{M}) \subseteq B_x$. By definition r_x is such that $f^{-1}(B(f(x), r_x) \cap T_q\mathcal{M}) \not\subseteq B_x$. This can happen only when there exist a point $y \in \bar{B}_x = \bar{B}_{\mathcal{M}}(x, \frac{\varepsilon}{16})$ such that either f is singular at y or else $y \notin B_x$. As we have shown in Claim 5.4.13 (ii) that f is nonsingular at all points in $B_{\mathcal{M}}(q, \hat{r}) \supset B_x$, hence $x \in \bar{B}_x \setminus B_x$. Which implies that $\|x-y\| = \frac{\varepsilon \text{rch}(\mathcal{M})}{16}$ and the angle made by the segment $[x, y]$ with $T_q\mathcal{M}$ is $\leq \theta_x$ (by definition of θ_x). Hence $r_x \geq \frac{\varepsilon \text{rch}(\mathcal{M})}{16} \cos \theta_x$. Therefore

$$\text{vol}(f(B_x)) \geq \text{vol}(B(f(x), r_x) \cap T_q\mathcal{M}) = \phi_k \frac{\varepsilon^k \text{rch}(\mathcal{M})^k}{16^k} \cos^k \theta_x. \quad (5.19)$$

Since the balls $B_x = B(x, \frac{\varepsilon \text{rch}(\mathcal{M})}{16}) \cap \mathcal{M}$ for all $x \in B(q, \frac{\text{rch}(\mathcal{M})}{3}) \cap \mathcal{M}$ are disjoint and f is injective, we get from Eq. (5.19) and (5.18)

$$\text{vol}(f(\cup_{x \in S} B_x)) = \sum_{x \in S} \text{vol}(f(B_x)) = |S| \frac{\varepsilon^k \text{rch}(\mathcal{M})^k}{2^{O(k)}}$$

where $S = B(x, \frac{\text{rch}(\mathcal{M})}{3}) \cap \mathcal{P}$. Using the fact that $f(\cup_{x \in S} B_x) \subseteq B(q, \hat{r}) \cap T_q\mathcal{M}$, we have

$$|S| \leq \frac{\text{vol}(B(q, \hat{r}) \cap T_q\mathcal{M})}{\frac{\varepsilon^k \text{rch}(\mathcal{M})^k}{2^{O(k)}}} = \frac{2^{O(k)}}{\varepsilon^k}.$$

$\square$

Lemma 5.4.14 *The expected number of times **pick()** is called within **good-pick**(σ, α) is $\frac{1}{T}$, where*

$$T = 1 - O\left(\frac{\Theta_0}{\alpha^k}\right).$$

The constant in the big-O depends on k and ρ_0 .

Proof In the algorithm, **good-pick()** is called either by Rule 3 to refine a k -simplex in $\text{Del}_{T\mathcal{M}}(P)$ or by Rule 4 to refine an inconsistent configuration in $C(P)$. We will consider the two cases separately.

1. We will first consider the case when $\sigma = \tau$ is a k -dimensional simplex in $\text{star}(p) \subset \text{Del}_{T\mathcal{M}}(P)$. Since τ is a k -dimensional simplex in $\text{star}(p)$, hence it is being refined by Rule 3 and $R_p(\tau) \leq \varepsilon \text{rch}(\mathcal{M})$. Let $B = B(c_p(\tau), \alpha R_p(\tau)) \cap \mathcal{M}$, and let

$$\tilde{f} : \mathcal{M} \rightarrow T_p\mathcal{M}$$

denote the orthogonal projection map of $\mathcal{M}$ onto $T_p\mathcal{M}$. As in the proof of Lemma 5.4.12, we can prove that the map $\tilde{f}$ restricted to the set B is injective and the derivative $d\tilde{f}$ is nonsingular for all points in B , which implies that $\tilde{f}$ is an open map when restricted to the set B .

Function **goodpick** (τ, α) picks a random point $x \in B_p = B(c_p(\tau), \alpha R_p(\tau)) \cap T_p\mathcal{M}$ and checks whether the two following conditions are satisfied: (C1) $x' = \mathbf{pick}(x, p)$ is in B , and (C2) x' does not form a j -dimension sliver ($2 \leq j \leq k+1$) with other sample points contained in the ball $B(c_p(\tau), \alpha R_p(\tau) + 2\beta R_\tau)$. If both conditions (C1) and (C2) are satisfied, then return x .

We will now bound the volume of the set

$$S_1 = \{x \in B_p \mid \mathbf{pick}(x, p) \notin B\},$$

i.e. the set of points in B_p that do not satisfy (C1).

Claim 5.4.15 $S_1 \subseteq B_p \setminus \tilde{f}(B)$

Proof Let $x \in S_1$. Since $x \in S_1$, implies that $\mathbf{pick}(x, p)$ is either empty, i.e. $(d-k)$ -flat, H_x , passing through x and parallel to $N_p\mathcal{M}$ does not intersect $\mathcal{M}$ or $x' = \mathbf{pick}(x, p) \notin B$. We claim that there does not exist a point in B whose image under the map $\tilde{f}$ is x . Otherwise if there exist a point $y \in B$ such that $\tilde{f}(y) = x$, then this would imply that the line segment $[x, y]$ lies in H_x . This would imply that $H_x \cap \mathcal{M}$ is not empty and

$$\begin{aligned} \|x - y\|^2 &= \|y - c_p(\tau)\|^2 - \|x - c_p(\tau)\|^2 && \text{(Pythagoras theorem)} \\ &< \alpha^2 R_p(\tau)^2 - \|x - c_p(\tau)\|^2 && \text{(since } y \in B) \\ &\leq \|x' - c_p(\tau)\|^2 - \|x - c_p(\tau)\|^2 && \text{(since } x' \notin B) \\ &= \|x - x'\|^2 && \text{(Pythagoras theorem)} \end{aligned} \quad (5.20)$$

We have reached a contradiction since $H_x \cap \mathcal{M} \neq \emptyset$, and $\|y - x\| < \|x' - x\|$ but by definition $x' = \mathbf{pick}(x, p)$ is the point closest to x in $H_x \cap \mathcal{M}$. This implies that $x \in B_p \setminus \tilde{f}(B)$ and the claim follows. $\square$

From the Claim 5.4.15, we have

$$\begin{aligned} \text{vol}(S_1) &\leq \text{vol}(B_p \setminus \tilde{f}(B)) = \text{vol}(B_p) - \text{vol}(\tilde{f}(B)) \\ &= \phi_k \alpha^k R_p(\tau)^k - \text{vol}(\tilde{f}(B)). \end{aligned} \quad (5.21)$$

We will upper bound $\text{vol}(S_1)$ by lower bounding $\text{vol}(\tilde{f}(B))$.

Let p' be the point closest to $c_p(\tau)$ on $\mathcal{M}$. From Lemma 2.2.2 (2) we have $\|p - p'\| \leq 2\varepsilon\|p - c_p(\tau)\| = 2\varepsilon R_p(\tau)$. Therefore, $B' = B(p', r) \subseteq B$ where $r = (\alpha - 2\varepsilon)R_p(\tau)$. As in the proof of Lemma 5.4.12, using the fact that $\tilde{f}$ is an open map when restricted to B , we can show that

$$\text{vol}(\tilde{f}(B')) \geq \phi_k r^k \cos^k \theta, \quad (5.22)$$

where θ is the maximal angle made by any secant $s = [p', x]$ with $T_p \mathcal{M}$ where $x \in \bar{B}' = \bar{B}(p', r)$. Using Lemmas 2.2.2 and 5.4.11, and $\varepsilon < 1$, we get

$$\sin \theta \leq \frac{(\alpha - 2\varepsilon)\varepsilon}{2} + \sqrt{2\varepsilon + 4\varepsilon^2} < 3\sqrt{\varepsilon}. \quad (5.23)$$

Therefore using Eq. (5.22) and (5.23), we get

$$\begin{aligned} \text{vol}(\tilde{f}(B')) &\geq \phi_k r^k \cos^k \theta \\ &\geq \phi_k \left(1 - \frac{2\varepsilon}{\alpha}\right)^k \alpha^k r_p^k(\tau) (1 - 9\varepsilon)^{\frac{k}{2}} \\ &\geq \phi_k \alpha^k R_p(\tau)^k \left(1 - \frac{2k\varepsilon}{\alpha}\right) \left(1 - \frac{9k\varepsilon}{2}\right) \\ &= \left(1 - O\left(\frac{\varepsilon}{\alpha}\right)\right) \phi_k \alpha^k R_p(\tau)^k, \end{aligned} \quad (5.24)$$

where the constant in the big-O depends on k . Then, from Eq. (5.21) and (5.24), we have

$$\text{vol}(S_1) = O\left(\frac{\varepsilon}{\alpha}\right) \times \phi_k \alpha^k R_p(\tau)^k. \quad (5.25)$$

We will now bound the volume of the set

$$S_2 = \left\{ x \in B_p \mid x' = \mathbf{pick}(x, p) \text{ forms a } j\text{-dimensional sliver} \right. \\ \left. (2 \leq j \leq k+1) \text{ with sample points in } B(c_p(\tau), \alpha R_p(\tau) + \beta R_\tau) \right\},$$

i.e. the set of points in B_p whose answer to (C2) is “yes”.

Let S be the set of sample points in $B(c_p(\tau), \alpha R_p(\tau) + \beta R_\tau)$. We have shown in the proof of Lemma 5.4.8, that $|S| \leq N$. Let μ be a good simplex with vertices in S with $R_\mu \leq \beta R_\tau$. Then, from Lemma 5.4.2, we have

$$\text{vol}(F_\mu) \leq D\Theta_0 R_\mu^k \leq D\Theta_0 \beta^k R_p(\tau)^k. \quad (5.26)$$

The number of simplices of dimension $\leq k$ that can be formed with vertices from S is less than N^{k+1} . Let the union of the forbidden regions of all the good simplices of dimension $\leq k$ with vertices in S be denoted by W . Then, using the fact that $|S| \leq N$ and Eq. (5.26), we have

$$\text{vol}(W) \leq N^{k+1} \times D\Theta_0 \beta^k R_\tau^k(p). \quad (5.27)$$

We need the following claim to upper bound the volume of S_2 .

Claim 5.4.16 $S_2 \subseteq \tilde{f}(W) \cap B_p$

Proof Let $x \in S_2$. Since $x' = \mathbf{pick}(x, p) \in H_x \cap \mathcal{M}$, where H_x is a $(d-k)$ -flat passing through x and parallel to $N_p \mathcal{M}$, then $\tilde{f}(x') = x$. As $x \in S_2$, we also have $x' \in W$. Combining the facts that $S_2 \subset B_p$, $\tilde{f}(x') \in W$ and $\tilde{f}(x') = x$, we get $x \in \tilde{f}(W) \cap B_p$. Therefore $S_2 \subseteq \tilde{f}(W) \cap B_p$. $\square$

From Claim 5.4.16 and the fact that $\tilde{f}$ is a projection map, we have

$$\text{vol}(S_2) \leq \text{vol}(\tilde{f}(W) \cap B_p) \leq \text{vol}(\tilde{f}(W)) \leq \text{vol}(W). \quad (5.28)$$

Combining Eq. (5.25), (5.27) and (5.28), we get that the probability of $x' = \mathbf{pick}(x, p)$ satisfying conditions (C1) and (C2) for any random point in $x \in B_p$ is greater than

$$\begin{aligned} T &\stackrel{\text{def}}{=} \frac{\text{vol}(B_p \setminus S_1 \cup S_2)}{\text{vol}(B_p)} \geq \frac{\text{vol}(B_p) - \text{vol}(S_1 \cup S_2)}{\text{vol}(B_p)} \\ &\geq \frac{\text{vol}(B_p) - \text{vol}(S_1) - \text{vol}(S_2)}{\text{vol}(B_p)} \\ &\geq \left(1 - O\left(\frac{\varepsilon}{\alpha}\right) - \frac{N^{k+1} \beta^k D \Theta_0}{\phi_k \alpha^k}\right) \\ &= 1 - O\left(\frac{\Theta_0}{\alpha^k}\right), \end{aligned} \quad (5.29)$$

the constant in the big- O depends only k , ρ_0 and β , since $\alpha \leq 1/2$, $\varepsilon \leq \Theta_0$ (from Hypothesis H5), D depends on k and ρ_0 (Lemma 5.4.2), and $N = 2^{O(k)}$ (N depends only on β and α , see Lemma 5.4.8). Therefore the expected number of times we have to pick random points $x \in B_p$ s.t $x' = \mathbf{pick}(x, p)$ satisfies both the conditions (C1) and (C2) is less than

$$\sum_{i=1}^{\infty} i(1-T)^{i-1} T = \frac{1}{T}.$$

2. We can similarly show that the result holds for the case when $\sigma = \phi \in C(P)$ is an inconsistent configuration. $\square$

Theorem 5.4.17 *Under Hypotheses H1 to H5, the time complexity for updating $C(P \cup \{p\})$ from $C(P)$ when a new point p is inserted to the current sample P by the algorithm is $O(\varepsilon^{-k^2})$. The expected time complexity of the algorithm is $O(\varepsilon^{-k^2-k})$ for fixed $\mathcal{M}$, d and k .*

Proof 1. **Initialization.** Assume without loss of generality that $\mathcal{M}$ is enclosed in a d -dimensional box of unit length. We partition the unit box into a grid with unit length $\frac{\text{rch}(\mathcal{M})}{32\sqrt{d}}$ and intersect it with the manifold $\mathcal{M}$ to obtain the initial 1/16-sample of $\mathcal{M}$. Since the complexity of the grid is $\frac{2^{O(d \log d)}}{\text{rch}^d(\mathcal{M})}$ hence the number of points in the initial point sample, denoted by P_0 , is $\frac{2^{O(d \log d)}}{\text{rch}^d(\mathcal{M})}$ where the constant in the big- O is an absolute constant. The time complexity to get a subsample of the initial sample which is 1/32-sparse 1/16-sample of $\mathcal{M}$ is $O(d|P_0|^2)$.

2. **Refinement.** When a new point p is inserted by the algorithm, $\text{Del}_{T\mathcal{M}}(P \cup \{p\})$ is updated by creating the star of p and modifying the stars of all the points in $B(p, \text{rch}(\mathcal{M})/2) \cap P$.

Inconsistent configurations are only considered when all big simplices in $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ have been removed by application of Rule 1. Hence, by Lemma 5.2.5, we only have to consider inconsistent configurations with diameter at most $2R_\phi \leq 4R_\tau \leq 4\epsilon \text{rch}(\mathcal{M})$ (Eq. (5.6)) and therefore, to update $C(\mathcal{P} \cup \{p\})$, it suffices to look at the stars of the points in $B(p, 4\epsilon \text{rch}(\mathcal{M})) \cap (\mathcal{P} \cup \{p\})$. As in the proof of Theorem 5.4.10, we can show that the smallest interpoint distance between the points of $\mathcal{P}$ is $L_P \geq \epsilon \text{rch}(\mathcal{M})/8$. Using Lemma 5.4.12, we have for all $x \in \mathcal{P}$,

$$|B(x, \text{rch}(\mathcal{M})/2) \cap (\mathcal{P} \cup \{p\})| \leq \frac{2^{O(k)}}{\epsilon^k}. \quad (5.30)$$

The star of point $x \in B(p, \text{rch}(\mathcal{M})/2) \cap (\mathcal{P} \cup \{p\})$ can be calculated by projecting all the points in $B(x, \text{rch}(\mathcal{M})/2) \cap (\mathcal{P} \cup \{p\})$ on $T_x\mathcal{M}$ and calculating the weighted Delaunay triangulations of these projected points (Lemma 5.2.2). The time complexity for modifying the stars of all points in $B(p, \text{rch}(\mathcal{M})/2)$ is this

$$\frac{d 2^{O(k)}}{\epsilon^k} + \frac{2^{O(k^2)}}{\epsilon^{k^2}}$$

Using the same arguments, we get that the time complexity for modifying the inconsistencies is

$$\frac{d 2^{O(k)} \text{Vol}(\mathcal{M})}{\epsilon^k} + \frac{d 2^{O(k)}}{\epsilon^k} + \frac{2^{O(k^2)}}{\epsilon^{k^2}},$$

where the first term is for calculating the points in $B(x, \text{rch}(\mathcal{M})/2) \cap (\mathcal{P} \cup \{p\})$, see equation (5.30).

By Theorem 5.4.10, the algorithm inserts $2^{O(k)} \text{vol}(\mathcal{M})/\epsilon^k$ many points. From Lemma 5.4.14, we get the expected number of times **pick()** is called within **good-pick()** is $\frac{1}{T}$, where

$$T = 1 - O\left(\frac{\Theta_0}{a^k}\right),$$

where the constant in the big-O depends on k and ρ_0 . Hence, the total time complexity of the refinement algorithm is

$$(\text{Vol}(\mathcal{M}) + 1) \frac{d 2^{O(k)} \text{Vol}(\mathcal{M})}{T \epsilon^{2k}} + \frac{2^{O(k^2)} \text{Vol}(\mathcal{M})}{T \epsilon^{k^2+k}}. \quad (5.31)$$

□

5.5 Topological and geometric guarantees

We assume that the conditions of Theorem 5.4.9 are satisfied. Therefore $\hat{\mathcal{M}}$ has no slivers and no inconsistencies. Let $\pi : \mathbb{R}^d \rightarrow \mathcal{M}$ map each point of $\mathbb{R}^d$ to its closest point of $\mathcal{M}$. The following result is a special case of Theorem 4.0.3 proved in Chapter 4 (except for item 5 which is a direct consequence of item 2).

Theorem 5.5.1 (Properties of $\hat{\mathcal{M}}$) *For ϵ sufficiently small, we have the following :*

1. $\hat{\mathcal{M}}$ is a piecewise-linear manifold without boundary.

2. Map π restricted to $\hat{\mathcal{M}}$ provides an isotopy between $\hat{\mathcal{M}}$ and $\mathcal{M}$.
3. $\forall x \in \mathcal{M}, \|x - \pi^{-1}(x)\| = O(\varepsilon^2 \text{rch}(\mathcal{M}))$, where the constant in the big-O depends on k, ρ_0 and Θ_0 .
4. $\forall x \in \mathcal{M}, \sin \angle(T_x \mathcal{M}, \text{aff}(\tau)) = \sin \angle(N_x \mathcal{M}, N_\tau) = O(\varepsilon)$, where τ is a k -simplex of $\hat{\mathcal{M}}$ containing the point $\pi^{-1}(x)$.
5. The output sample P is an $(\varepsilon + O(\varepsilon^2), \Omega(\varepsilon))$ -rch sample of $\mathcal{M}$, where the constant in the big-O depends on k, ρ_0 and Θ_0 , and the constant in the big- Ω depend on α .

5.6 Summary

We have shown how to sample and triangulate a k -dimensional submanifold of $\mathbb{R}^k$ up to a prescribed sampling rate ε using a variant of Delaunay refinement. The submanifold is assumed to be compact, closed and of positive reach, but not necessarily oriented. The requirement $\text{rch}(\mathcal{M}) > 0$ can be somehow relaxed and Lipschitz manifolds can be triangulated in very much the same way as manifolds of positive reach, as already shown for surfaces in [BO06].

We assumed to know the reach of $\mathcal{M}$ (or, at least, a positive lower bound) and to be able to compute the tangent space at any point $p \in \mathcal{M}$. If $\mathcal{M}$ is described by a set of equations, computing the reach of $\mathcal{M}$ reduces to solving a 0-dimensional system of equations [BO05]. Remarkably, our algorithm can be proved to tolerate some uncertainty in the estimation of the tangent spaces.

The algorithm is simple and relies only on simple computations performed in affine subspaces. In order to walk around the curse of dimensionality, we do not triangulate the ambient space and only maintain a k -dimensional data structure, the so-called tangential Delaunay complex. This leads to an algorithm that uses a restricted set of simple numerical operations and whose asymptotic complexity is $O(\varepsilon^{-k^2-k})$ for fixed $\mathcal{M}$, d and k .

We have shown that the size of the sample is $O(\varepsilon^{-k})$ and that the output mesh $\hat{\mathcal{M}}$ is a good approximation of $\mathcal{M}$ from a geometric and a topological points of view. Specifically, we showed that the Hausdorff distance between $\mathcal{M}$ and $\hat{\mathcal{M}}$ is $O(\varepsilon^2 \text{rch}(\mathcal{M}))$ and that the maximal angle between their normal bundles is $O(\varepsilon)$. The constant hidden in the big-O depends on the normalized volume of $\mathcal{M}$ (defined in Section 5.4.3). Up to the multiplicative constant that depends on $\mathcal{M}$, those bounds are optimal in view of Clarkson's results [Cla06] (note that Clarkson's bound is for the Hausdorff distance only).

If one knows at each point x of $\mathcal{M}$ the local feature size $\text{lfs}(x)$, it is easy to modify the algorithm so that the constants depend on $\int_{\mathcal{M}} \frac{dx}{\text{lfs}^k(x)} \leq V(\mathcal{M})$. This constant could even be improved if one combines the algorithm of this chapter with the related technique developed for anisotropic mesh generation [BWY08]. Provided that one can estimate the second fundamental form at any point of $\mathcal{M}$, such an extension would allow to construct *anisotropic meshes* that locally conform to the local metric of $\mathcal{M}$ and approximate $\mathcal{M}$ with a better convergence rate.

Part III

Stability of Delaunay Triangulation

Chapter 6

An obstruction to intrinsic Delaunay triangulations

Delaunay has shown that the Delaunay complex of a finite set of points $\mathcal{P}$ of Euclidean space $\mathbb{R}^m$ is a triangulation of $\mathcal{P}$, provided that $\mathcal{P}$ satisfies a mild genericity property. Voronoi diagrams and Delaunay complexes can be defined for arbitrary Riemannian manifolds. However, Delaunay's genericity assumption no longer guarantees that the Delaunay complex will be a triangulation; stronger assumptions on $\mathcal{P}$ are required. A natural one is to assume that $\mathcal{P}$ is sufficiently dense. Although partial results in this direction have been obtained (or claimed), we show in this chapter that, for manifolds of dimension greater than 2, sample density alone is insufficient to ensure that the Delaunay complex is a triangulation.

6.1 Delaunay complex and Delaunay triangulation

Let $\mathbb{M} = (\mathcal{M}, d_{\mathcal{M}})$ be a metric space, and let $\mathcal{P}$ be a finite set of points of $\mathcal{M}$. An *empty ball* is an open ball in the metric $d_{\mathcal{M}}$ that contains no point from $\mathcal{P}$. We say that an empty ball B is *maximal* if no other empty ball with the same centre properly contains B . A *Delaunay ball* is a maximal empty ball.

A simplex σ is a *Delaunay simplex* if there exists some Delaunay ball B that circumscribes σ , i.e. such that the vertices of σ belong to $\partial B \cap \mathcal{P}$. The *Delaunay complex* is the set of Delaunay simplices, and is denoted $\text{Del}_{\mathcal{M}}(\mathcal{P})$. It is an abstract simplicial complex.

The *Voronoi cell* associated with $p \in \mathcal{P}$ is given by

$$\text{Vor}_{\mathcal{M}}(p) = \{x \in \mathcal{M} : d_{\mathcal{M}}(x, p) \leq d_{\mathcal{M}}(x, q) \text{ for all } q \in \mathcal{P}\}.$$

More generally, a *Voronoi face* is the intersection of a set of Voronoi cells: given $\sigma = \{p_0, \dots, p_k\} \subset \mathcal{P}$, we define the associated Voronoi face as

$$\text{Vor}_{\mathcal{M}}(\sigma) = \bigcap_{i=0}^k \text{Vor}_{\mathcal{M}}(p_i).$$

It follows that σ is a Delaunay simplex if and only if $\text{Vor}_{\mathcal{M}}(\sigma) \neq \emptyset$. In this case, every point in $\text{Vor}_{\mathcal{M}}(\sigma)$ is the centre of a Delaunay ball for σ . Thus every Voronoi face corresponds to a Delaunay simplex. The Voronoi cells give a decomposition of $\mathcal{M}$, denoted $\text{Vor}_{\mathcal{M}}(\mathcal{P})$,

called the *Voronoi diagram*. The Delaunay complex of $\mathcal{P}$ is the nerve of the Voronoi diagram.

In the case of $\mathbb{R}^m$ equipped with the standard Euclidean metric, Delaunay [Del34] showed that, if $\mathcal{P}$ is *generic*, then the natural inclusion $\mathcal{P} \hookrightarrow \mathbb{R}^m$ induces an embedding of the Delaunay complex $\text{Del}_{\mathbb{R}^m}(\mathcal{P})$ of $\mathcal{P}$ in $\mathbb{R}^m$, called the Delaunay triangulation of $\mathcal{P}$. The point set $\mathcal{P}$ is generic if there is no Delaunay ball with more than $m + 1$ points of $\mathcal{P}$ on its boundary. Point sets that are not generic are often dismissed in theoretical work, because an arbitrarily small perturbation of the points can be made which will yield a generic point set. Thus in the sense of the standard measure in the configuration space $\mathbb{R}^{m \times |\mathcal{P}|}$, almost all point sets will yield a Delaunay triangulation.

A similar situation is known for certain standard non-Euclidean geometries, such as Laguerre (or Power diagrams) geometry [Aur87], or spaces equipped with a Bregman divergence [BNN10], or Riemannian manifolds of constant sectional curvature, e.g., hyperbolic spaces [DMT92].

Leibon and Letscher [LL00] announced sampling density conditions which would ensure that the Delaunay complex defined by the intrinsic metric of an arbitrary compact Riemannian manifold was a triangulation. In fact, as shown here, the stated result is incorrect: sampling density alone is insufficient to guarantee an intrinsic Delaunay triangulation (see Theorem 6.3.3). For any given sampling density, even if the point set is generic, topological defects can arise when the vertices lie too close to a degenerate configuration.

When triangulating submanifolds of dimension 3 and higher in Euclidean space using Delaunay techniques, it was since discovered that near degenerate “sliver” simplices pose problems which cannot be escaped simply by increasing the sampling density. In particular, developing an example on a 3-manifold presented by Cheng et al. [CDR05b], Boissonnat et al. [BGO09, Lemma 3.1] show that, using the metric of the ambient Euclidean space restricted to the submanifold, the resulting Delaunay complex (called the *restricted Delaunay complex*) need not be a triangulation of the original submanifold, even with dense well separated sampling.

In this note we develop this example from the perspective of the intrinsic metric of the manifold. It can be argued that this is an easier way to visualize the problem, since we confine our viewpoint to a three dimensional space and perturb the metric, without referring to deformations into a fourth ambient dimension. This viewpoint also provides an explicit counterexample to the results announced by Leibon and Letscher [LL00].

6.2 A qualitative argument

As we show in this section with a qualitative argument, the problem can be viewed as arising from the fact that when m is greater than two, the intersection of two metric spheres is not uniquely specified by m points. We demonstrate the issue in the context of Delaunay balls. The problem is developed quantitatively in terms of the Voronoi diagram in Section 6.3.

We work exclusively on a three dimensional domain, and we are not concerned with “boundary conditions”; we are looking at a coordinate patch on a densely sampled compact 3-manifold.

One core ingredient in Delaunay’s triangulation result [Del34] is that any triangle τ is




Figure 6.1: In three dimensions, three closed geodesic balls can all touch three points, u, v, p , on their boundary and yet no one of them is contained in the union of the other two.

the face of exactly two tetrahedra. This follows from the observation that a triangle has a unique circumcircle, and that any circumscribing sphere for τ must include this circle. The affine hull of τ cuts space into two components, and if $\tau \in \text{Del}_{\mathbb{R}^m}(\mathcal{P})$, then it will have an empty circumsphere centred at a point c on the line through the circumcentre and orthogonal to $\text{aff}(\tau)$. The point c is contained on an interval on this line which contains all the empty spheres for τ . The endpoints of the interval are the circumcentres of the two tetrahedra that share τ as a face.

The argument hinges on the assumption that the points are in general position, and the uniqueness of the circumcircle for τ . If there were a fourth vertex lying on that circumcircle, then there would be three tetrahedra that have τ as a face, but this configuration would violate the assumption of general position.

Now if we allow the metric to deviate from the Euclidean one, no matter how slightly, the guarantee of a well defined unique circumcircle for τ is lost. In particular, If three spheres S_1 , S_2 and S_3 all circumscribe τ , their pairwise intersections will be different in general. I.e.,

$$S_1 \cap S_3 \neq S_2 \cap S_3.$$

Although these intersections may be topological circles that are “arbitrarily close” assuming the deviation of the metric from the Euclidean one is small enough, “arbitrarily close” is not good enough when the only genericity assumption allows configurations that are arbitrarily bad.

An attempt to illustrate the problem is given in Figure 6.1, where $\tau = \{u, v, p\}$. Here, the sphere S_3 would be contained inside the spheres S_1 and S_2 if the metric were Euclidean, but any aberration in the metric may leave a part of S_3 exposed to the outside. This means that in principle another sample point w could lie on S_3 , while S_1 and S_2 remain empty. Thus there are three tetrahedra that share τ as a face.

The essential difference between dimension 2 and the higher dimensions can be observed by examining the topological intersection properties of spheres. Specifically,

two $(m-1)$ -spheres intersect transversely in an $(m-2)$ -sphere. For a non-Euclidean metric, even if this property holds for sufficiently small geodesic spheres, only in dimension two is the sphere of intersection of the Delaunay spheres of two adjacent m -simplices uniquely determined by the vertices of the shared $(m-1)$ -simplex. See Figure 6.1.

6.3 An obstruction to intrinsic Delaunay triangulations

We now explicitly show how density assumptions based upon the strong convexity radius, as proposed by Leibon and Letscher [LL00], cannot escape topological problems in the Delaunay complex. The configuration considered here may be recognised as essentially the same as that which was described qualitatively in Section 6.2, but here we consider the Voronoi diagram rather than Delaunay balls. As before we work exclusively in a local coordinate patch on a densely sampled compact 3-manifold.

6.3.1 Sampling density alone is insufficient

We say $\mathcal{P} \subset \mathcal{M}$ is ϵ -dense if $d_{\mathcal{M}}(x, \mathcal{P}) < \epsilon$ for any $x \in \mathcal{M}$. If $d_{\mathcal{M}}(p, q) \geq \bar{\epsilon}$ for all $p, q \in \mathcal{P}$, then $\mathcal{P}$ is $\bar{\epsilon}$ -separated. The set $\mathcal{P}$ is an ϵ -net if it is ϵ -dense and ϵ -separated.

Leibon and Letscher [LL00, p. 343] explicitly assume that the points are *generic* which they state as

Definition 6.3.1 *The set $\mathcal{P} \subset \mathcal{M}$, is generic if $\mathcal{M}$ is an m -manifold and $m+2$ points never lie on the boundary of a round ball.*

Here a round ball refers to a geodesic ball. This definition of genericity is natural, and corresponds to Delaunay's original definition [Del34], except Delaunay only imposed the constraint on empty balls. A question that Delaunay addressed explicitly, but which was not addressed by Leibon and Letscher, is whether or not such an assumption is a reasonable one to make. Delaunay showed that any (finite or periodic) point set in Euclidean space can be made generic through an arbitrarily small affine perturbation. That a similar construction of a perturbation can be made for points on a compact Riemannian manifold has not been explicitly demonstrated. However, in light of the construction we now present, it seems that the question is moot when $m > 2$, because an arbitrarily small perturbation from degeneracy will not be sufficient to ensure a triangulation.

Leibon and Letscher proposed adaptive density requirements based upon the *strong convexity radius*. These requirements are somewhat complicated, but they will be satisfied if a simple constant sampling density requirement is satisfied. Exploiting a theorem [Cha06, Thm. IX.6.1], that relates the strong convexity radius to the injectivity radius, $\text{inj}(\mathcal{M})$, and a positive bound on the sectional curvatures, they arrive at the following:

Claim 6.3.2 ([LL00, Lemma 3.3]) *Suppose $\mathcal{K}_0$ is a positive upper bound on the sectional curvatures of $\mathcal{M}$, and*

$$\eta(\mathcal{M}) = \min \left\{ \frac{\text{inj}(\mathcal{M})}{10}, \frac{\pi}{10\sqrt{\mathcal{K}_0}} \right\}. \quad (6.1)$$

If $\mathcal{P}$ is an $\eta(\mathcal{M})$ -sample set for $\mathcal{M}$ with respect to $d_{\mathcal{M}}$, then $|\text{Del}_{\mathcal{M}}(\mathcal{P})| \cong \mathcal{M}$.

In fact, we will show that no sampling conditions based on density alone will be sufficient to guarantee a homeomorphic Delaunay complex in general, even when a separation assumption is also demanded. We will show:

Theorem 6.3.3 *With $\eta(\mathcal{M})$ as defined in Eq. (6.1), for any $\epsilon > 0$, there exists a compact Riemannian manifold $\mathcal{M}$, and a finite set $\mathcal{P} \subset \mathcal{M}$, such that $\mathcal{P}$ is an $(\epsilon\eta(\mathcal{M}))$ -net for $\mathcal{M}$, with respect to the metric $d_{\mathcal{M}}$, but $\text{Del}_{\mathcal{M}}(\mathcal{P})$ is not homeomorphic to $\mathcal{M}$.*

6.3.1.1 A counter-example

We will construct the counter-example by considering a perturbation of a Euclidean metric. This is a local operation, and the global properties of the manifold are only relevant in so far as they affect $\eta(\mathcal{M})$ of Eq. (6.1). We may assume, for example, that the manifold $\mathcal{M}$ is a 3-dimensional torus: $\mathbb{M} = \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$, initially with a flat metric.

Thus assume there is some ϵ_0 such that any compact Riemannian manifold may be triangulated by the intrinsic Delaunay complex when $\mathcal{P}$ is an $\epsilon_0\eta(\mathcal{M})$ -net. For convenience, we choose a system of units so that $\epsilon_0\eta(\mathcal{M}) = 1$. We will first construct a point configuration and metric perturbation that leads to a problem, and then we will show that the sampling assumptions are indeed met.

We introduce a number of parameters which we will manipulate to produce the counter-example. We are exploiting the fact that the genericity assumption allows configurations that are arbitrarily close to being degenerate. The assumed ϵ_0 has been fixed.

We will work within a coordinate chart on $\mathcal{M}$, where the metric is Euclidean. We will perturb this metric by constructing a metric tensor $\tilde{g}$, and we will denote by $\tilde{\mathcal{M}}$ the manifold with this new metric.

Consider points u, v, w, p in the xz -plane arranged with u and v at $\pm a$ on the z axis, and w and p at $\pm(a + \xi)$ on the x axis, with $a = \frac{3}{4}$, and $0 < \xi < r_0\gamma$, where r_0 and γ will be specified below. The Voronoi diagram of these points in the xz -plane is shown in Figure 6.2. The main point here is that the Voronoi boundary between $\text{Vor}_{\mathcal{M}}(u)$ and $\text{Vor}_{\mathcal{M}}(v)$ may be arbitrarily small with respect to the distance between the sites, i.e., ξ will be very very small.

The three dimensional Voronoi diagram is the extension of this in the horizontal y -direction, so that every cross-section looks the same. Note that since the points are not co-circular, they do not represent a degeneracy by Delaunay's criteria [Del34], but this is irrelevant; we will also argue that the points will not represent a degenerate configuration with respect to the new metric.

We now introduce a small localized metric perturbation so as to change the Voronoi diagram near the origin. For example, we can demand that the matrix of the metric tensor




Figure 6.2: A vertical slice: the xz -plane of the initial Voronoi diagram, seen from the negative y axis.

in our coordinate system has the form

$$\tilde{g}(p) = \begin{pmatrix} 1 - f(|p|) & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

where $|p|$ is the parametric distance from p to the origin. The radial function f is non-negative, and it and its first two derivatives are bounded, e.g.,

$$f(r), |f'(r)|, |f''(r)| \leq \beta. \quad (6.2)$$

We also demand that there exists a positive $\gamma \leq \beta$ such that $f(r) \geq \gamma$ when $r \leq r_0$, and that $f(r) = 0$ if $r \geq 2r_0$. The parameter r_0 , defines the radius of the ball bounding the perturbed region. Now we have $\text{dist}wp < \text{dist}uv$ when $\xi < r_0\gamma$.




Figure 6.3: The $y = 0$ slice of the perturbed Voronoi diagram.

Since γ may be arbitrarily small compared to β , standard arguments supply a function f meeting these conditions. For example, the C^∞ construction described by Munkres [Mun68, p. 6] may be multiplied by a scalar sufficiently small to meet our needs.

The vertical $y = 0$ cross-section of the perturbed Voronoi diagram will look something like Figure 6.3: $\text{Vor}_{\tilde{\mathcal{M}}}(p)$ and $\text{Vor}_{\tilde{\mathcal{M}}}(w)$ now meet in the xz -plane, and $\text{Vor}_{\tilde{\mathcal{M}}}(u)$ and $\text{Vor}_{\tilde{\mathcal{M}}}(v)$ do not. However, since geodesics which do not intersect the ball $B_{\mathbb{R}^3}(0, 2r_0)$ will remain straight lines in the parameter space, the Voronoi diagram is unchanged outside of a neighbourhood of the origin. Thus looking from above at the slice of the Voronoi diagram in the xy -plane, we will see something like Figure 6.4(a). Figure 6.4(b) shows the yz -plane.

Two Voronoi vertices have been introduced, the red and blue points in Figure 6.4. These are the centres of distinct empty geodesic circumballs for $\{p, u, v, w\}$. Since they cannot lie in the region unaffected by the perturbation, a quick calculation shows that the parametric distance of these Voronoi vertices from the origin is bounded by $4r_0$, when $r_0 \leq \frac{1}{4}$, and it follows from another small calculation that the parametric distance from these Voronoi vertices to any of the four sample points is bounded by $a(1 + \frac{3\xi + 16r_0^2}{a^2})$. The distances between these Voronoi vertices and the sample points in the new metric will also be subjected to the same bound, since no distances increase. Also, The sparsity condition will not be affected by the perturbation. Thus, since we can make r_0 as small as we please, and ξ is chosen such that $\xi < r_0\gamma$, it follows that the radius of these balls may be made arbitrarily close to $a = \frac{3}{4} = \frac{3}{4}\epsilon_0\eta(\mathcal{M})$. We will argue next that we can make $|\eta(\mathcal{M}) - \eta(\tilde{\mathcal{M}})|$ as small as desired by reducing the size of β in Eq. (6.2). Then other sample points may be placed on the manifold so that the density criteria are met, and no degenerate configuration (violation of Definition 6.3.1) need be introduced.

This means that the Delaunay complex, defined as the nerve of the Voronoi diagram, will not be a triangulation of the manifold $\tilde{\mathcal{M}}$. As observed by Boissonnat et al. [BGO09], the triangle faces $\{p, w, u\}$ and $\{p, w, v\}$ will be adjacent to only a single tetrahedron, namely $\{p, u, v, w\}$. Thus $\text{Del}_{\tilde{\mathcal{M}}}(\mathcal{P})$ is not a manifold complex, i.e., there are vertices in $\text{Del}_{\tilde{\mathcal{M}}}(\mathcal{P})$

whose star is not isomorphic to the star of a vertex of a triangulation of $\mathbb{R}^m$. This is clearly a problem if the original manifold has no boundary.




Figure 6.4: Looking at cross-sections; the positive y -direction is to the right. The four points, p, u, v, w , admit two small circumballs with distinct centres (the red and blue points).

Although it is in some sense close to being degenerate, we emphasise that this configuration represents a problem that cannot be escaped by an arbitrarily small perturbation of the sample points. An argument based on the triangle inequality shows that in order to effect a change in the topology of the Voronoi diagram, a displacement of the points by a distance of $\Omega(r_0\gamma - \xi)$ is required.

More specifically, we observe that the configuration $\{p, u, v, w\}$ may be placed in an otherwise well behaved point set $\mathcal{P}$ such that within a small ball centred at the origin in our coordinate chart, all points will have $\{p, u, v, w\}$ as the four closest points in $\mathcal{P}$, and this would remain the case even if the point positions were perturbed a small amount. We may further assume that the other Delaunay simplices are well shaped, so that stability results, from Chapter 7, can be used to argue that they cannot be destroyed with an *arbitrarily* small perturbation. Then we argue that in order to obtain a triangulation by a perturbation $\mathcal{P} \rightarrow \mathcal{P}'$, we must ensure that the Voronoi cell $\text{Vor}_{\tilde{\mathcal{M}}}(\{p', w'\})$ must vanish: the edge $\{p', w'\}$ will never be incident to any tetrahedron other than $\{p', u', v', w'\}$. Then an argument based on the triangle inequality shows that for a ρ -perturbation with $\rho < \frac{r_0\gamma - \xi}{6}$, there will be a point in $\text{Vor}_{\tilde{\mathcal{M}}}(\{p', w'\})$ within a distance of 2ρ of the origin.

6.3.1.2 The sizing function under perturbation

We need to establish that the metric manipulation that we performed in order to construct the counter-example, does not have a dramatic effect on the sizing function $\eta(\tilde{\mathcal{M}})$. This follows from the fact that we have bounded $g - \tilde{g}$ together with its first and second derivatives.

Since the sectional curvature may be described as a continuous function of g and its first and second derivatives [dC92, pp. 56 & 93], the effect of our perturbation on the sectional curvatures can be made arbitrarily small by reducing β in Eq. (6.2).

Since we started with a flat metric anyway, the bound $\mathcal{K}_0$ can be made arbitrarily small, and so the second term in Eq. (6.1) will not be the smallest. We need to bound the change in the injectivity radius as well.

This follows from results in the literature [Ehr74, Sak83], which state that for a compact manifold, $\text{inj}(\mathcal{M})$ depends continuously on the metric and its first and second derivatives. Specifically,

Lemma 6.3.4 (Ehrlich) *Let $\mathfrak{M}$ be the space of C^3 Riemannian metric structures g on a compact manifold $\mathbb{M}$, and endow $\mathfrak{M}$ with the C^2 topology. The function $g \mapsto \text{inj}(\mathcal{M}_g)$ is continuous in this topology.*

This means that for any desired bound on $|\eta(\mathcal{M}) - \eta(\tilde{\mathcal{M}})|$, there will be a β that will satisfy the bound.

The construction of the counter-example is complete.

6.3.2 Discussion

We have shown that for constructing a Delaunay triangulation for an arbitrary Riemannian manifold, a sampling density requirement is not sufficient in general. The solution we propose here is to constrain the kind of sample sets that we consider. Another approach would be to constrain the kind of metrics that are assumed. However, even with a purely Euclidean metric, allowing configurations to be arbitrarily close to degeneracy means that arbitrarily poorly shaped simplices are to be expected. When the metric is no longer Euclidean, the “shape” of a simplex no longer has an obvious meaning, but the problems associated with point configurations near degeneracy will certainly be present.

Our analysis relied on the ability to make the support of the perturbation small. This is unlikely to be a necessary feature of the construction, but it facilitates our simplistic analysis.

Clarkson [Cla06] remarked that an implication of Leibon and Letscher’s claim [LL00] is that for four points close enough together, there is a unique circumsphere with small radius. Our counter-example shows that circumcentres need not be unique under these conditions. In fact the existence of unique circumcentres does not follow from the triangulation result. However, the argument sketched out by Leibon and Letscher claimed that the intrinsic Voronoi diagram is a cell complex (i.e., it satisfies the *closed ball property* [ES97]), and this would imply unique circumcentres for the top dimensional simplices.

It is worth emphasising that the problems discussed here only arise when the dimension is greater than 2. The same sampling criteria for two dimensional manifolds has been fully validated [Lei99, DZM08], however these works both assume genericity in the sample set, without demonstrating that it is a reasonable assumption.

Chapter 7

Stability of Delaunay triangulations

We introduce a parametrized notion of genericity for Delaunay triangulations which, in particular, implies that the Delaunay simplices of δ -generic point sets are thick. Equipped with this notion, we study the stability of Delaunay triangulations under perturbations of the metric and of the vertex positions. We quantify the magnitude of the perturbations under which the Delaunay triangulation remains unchanged.

Structural results from this chapter will play a crucial role in solving (partially) the obstruction, highlighted in Chapter 6, to getting intrinsic Delaunay triangulations of manifolds.

7.1 Introduction

One of the central properties of Delaunay complexes, which was demonstrated when they were introduced [Del34], is that under a very mild assumption they are embedded, i.e., they define a triangulation of Euclidean space. The required assumption is that there are not too many cospherical points; the points are “generic”. The assumption is not considered limiting because, as Delaunay showed, an arbitrarily small affine perturbation can transform any given point set into one that is generic.

Given the assumption of a generic point set, we are assured that the Delaunay complex defines a triangulation, but a couple of issues arise when working with these triangulations. One is that the Delaunay triangulation can be highly sensitive to the exact location of the points. For example, the Delaunay triangulation of a point set might be different if a coordinate transform is first performed using floating point arithmetic.

Another problem concerns the geometric quality of the simplices in the triangulation. We define the *thickness* of a simplex as a number proportional to the ratio of the smallest altitude to the longest edge length of the simplex, and we demonstrate why this is a useful measure of the geometric quality of the simplex. For points in the plane, if there is an upper bound on the ratio of the radius of a Delaunay ball to the length of the shortest edge of the corresponding triangle, then there is a lower bound on the thickness of any Delaunay triangle. However, when there are three or more spatial dimensions, the thickness of Delaunay simplices may become arbitrarily small in spite of any bound on the circumradius to shortest edge length.

Both of these issues are shown to be related to points being close to a degenerate (non-generic) configuration. We parameterize Delaunay’s original definition of genericity, saying that a point set $P \subset \mathbb{R}^m$ is δ -generic if every m -simplex in the Delaunay complex

has a Delaunay ball that is at a distance greater than δ to the remaining points in P . We show that a bound on δ leads to a bound on the thickness of the Delaunay simplices, and also that the Delaunay complex itself is stable with respect to perturbations of the points or of the metric, provided the perturbation is small enough with respect to δ in a way that we quantify.

The stability of Delaunay triangulations has not previously been studied in this way. Related work can be found in the context of kinetic data structures [AGG⁺10] or in the context of robust computation [BS04], and in particular, the concept of protection we introduce in Section 7.3 is embodied in the guarded insphere predicate which has been employed in a controlled perturbation algorithm for 2D Delaunay triangulation [FKMS05].

Our interest in the problem of near-degeneracy in Delaunay complexes stems from work on triangulating Riemannian manifolds. As we have seen in this previous chapters, an established technique is to compute the triangulation locally at each point in an approximating Euclidean metric, and then perform manipulations to ensure that the local triangulations fit together consistently. See, Chapters 3 and 5. The reason the manipulations are necessary is exactly the problem of the instability of the Delaunay triangulation, and sometimes this is most conveniently described as an instability with respect to a perturbation of the local Euclidean metric.

Although we make no explicit reference to Voronoi diagrams, the Delaunay complexes we study can be equivalently defined as the nerve of the Voronoi diagram associated with the metric under consideration. We provide criteria for ensuring that the Delaunay complex is a triangulation without explicit requirements on the properties of the Voronoi diagram [ES97], in contrast to a common practice in related work [LL00, LS03, CDR05b, DZM08, CG12].

Organization of the chapter. After presenting background material in Section 7.2, we introduce the concept of δ -generic point sets for Euclidean Delaunay triangulations in Section 7.3. We show that Delaunay simplices of δ -generic point sets are thick; they satisfy a quality bound. Then in Section 7.4 we quantify how δ -genericity leads to robustness in the Delaunay triangulation when either the points or the metric are perturbed. The primary challenge is bounding the displacement of simplex circumcentres. We conclude with some remarks on the construction and application of δ -generic point sets.

7.2 Background

To simplify the presentation for the readers of this and the following chapter, we will redefine (recapitulate) some of the terms from Chapter 2.

Within the context of the standard m -dimensional Euclidean space $\mathbb{R}^m$, when distances are determined by the standard norm, $\|\cdot\|$, we use the following conventions. The distance between a point p and a set $X \subset \mathbb{R}^m$, is the infimum of the distances between p and the points of X , and is denoted $d_{\mathbb{R}^m}(p, X)$. We refer to the distance between two points a and b as $\|b - a\|$ or $d_{\mathbb{R}^m}(a, b)$ as convenient. A ball $B_{\mathbb{R}^m}(c, r) = \{x \mid \|x - c\| < r\}$ is **open**, and $\overline{B}_{\mathbb{R}^m}(c, r)$ is its topological closure. We will consider other metrics besides the Euclidean one. A generic metric is denoted d , and the associated open and closed balls are $B_d(c, r)$, and $\overline{B}(c, r)$. Generally, we denote the topological closure of a set X by $\overline{X}$, the interior by

int X , and the boundary by ∂X . The convex hull is denoted $\text{conv}(X)$, and the affine hull is $\text{aff}(X)$.

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

$$\sin \angle(U, V) = \sup_{u \in U, \|u\|=1} \|u - \pi_V(u)\|, \quad (7.1)$$

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

7.2.1 Sampling parameters and perturbations

The structures of interest will be built from a finite set $P \subset \mathbb{R}^m$, which we consider to be a set of *sample points*. If $D \subset \mathbb{R}^m$ is a bounded set, then P is an ϵ -*sample set* for D if $d_{\mathbb{R}^m}(x, P) < \epsilon$ for all $x \in \overline{D}$. We say that ϵ is a *sampling radius* for D satisfied by P . If no domain D is specified, we say P is an ϵ -sample set if $d_{\mathbb{R}^m}(x, P \cup \partial \text{conv}(P)) < \epsilon$ for all $x \in \text{conv}(P)$. Equivalently, P is an ϵ -sample set if it satisfies a sampling radius ϵ for

$$D_\epsilon(P) = \{x \in \text{conv}(P) \mid d_{\mathbb{R}^m}(x, \partial \text{conv}(P)) \geq \epsilon\}.$$

The set P is λ -*separated* if $d_{\mathbb{R}^m}(p, q) > \lambda$ for all $p, q \in P$. We usually assume that the sparsity of a ϵ -sample set is proportional to ϵ , thus: $\lambda = \mu_0 \epsilon$.

We consider a perturbation of the points $P \subset \mathbb{R}^m$ given by a function $\zeta : P \rightarrow \mathbb{R}^m$. If ζ is such that $d_{\mathbb{R}^m}(p, \zeta(p)) \leq \rho$, we say that ζ is a ρ -*perturbation*. As a notational convenience, we frequently define $\tilde{P} = \zeta(P)$, and let $\tilde{p}$ represent $\zeta(p) \in \tilde{P}$. We will only be considering ρ -perturbations where ρ is less than half the sparsity of P , so $\zeta : P \rightarrow \tilde{P}$ is a bijection.

Points in P which are not on the boundary of $\text{conv}(P)$ are *interior points* of P .

7.2.2 Simplices

Given a set of $j+1$ points $\{p_0, \dots, p_j\} \subset P \subset \mathbb{R}^m$, a (geometric) j -*simplex* $\sigma = [p_0, \dots, p_j]$ is defined by the convex hull: $\sigma = \text{conv}(\{p_0, \dots, p_j\})$. The points p_i are the *vertices* of σ . Any subset $\{p_{i_0}, \dots, p_{i_k}\}$ of $\{p_0, \dots, p_j\}$ defines a k -simplex τ which we call a *face* of σ . We write $\tau \leq \sigma$ if τ is a face of σ , and $\tau < \sigma$ if τ is a *proper face* of σ , i.e., if the vertices of τ are a proper subset of the vertices of σ .

The *boundary* of σ , is the union of its proper faces: $\partial \sigma = \bigcup_{\tau < \sigma} \tau$. In general this is distinct from the topological boundary defined above, but we denote it with the same symbol. The *interior* of σ is $\text{int } \sigma = \sigma \setminus \partial \sigma$. Again this is generally different from the topological interior. In particular, a 0-simplex p is equal to its interior: it has no boundary. Other geometric properties of σ include its diameter (its longest edge), Δ_σ , and its shortest edge, L_σ .

For any vertex $p \in \sigma$, the *face opposite* p is the face determined by the other vertices of σ , and is denoted σ_p . If τ is a j -simplex, and p is not a vertex of τ , we may construct a $(j+1)$ -simplex $\sigma = p * \tau$, called the *join* of p and τ . It is the simplex defined by p and the vertices of τ , i.e., $\tau = \sigma_p$.

Our definition of a simplex has made an important departure from standard convention: we do not demand that the vertices of a simplex be affinely independent. A j -simplex σ is a *degenerate simplex* if $\dim \text{aff}(\sigma) < j$. If we wish to emphasise that a simplex

is a j -simplex, we write j as a superscript: σ^j ; but this always refers to the *combinatorial* dimension of the simplex.

A *circumscribing ball* for a simplex σ is any m -dimensional ball that contains the vertices of σ on its boundary. If σ admits a circumscribing ball, then it has a *circumcentre*, c_σ , which is the centre of the smallest circumscribing ball for σ . The radius of this ball is the *circumradius* of σ , denoted R_σ . In general a degenerate simplex may not have a circumcentre and circumradius, but in the context of the Euclidean Delaunay complexes we will work with, the degenerate simplices we may encounter do have these properties. We will make use of the affine space N_σ composed of the centres of the balls that circumscribe σ . This space is orthogonal to $\text{aff}(\sigma)$ and intersects it at the circumcentre of σ . Its dimension is $m - \dim \text{aff}(\sigma)$.

The *altitude* of a vertex p in σ is $D_\sigma(p) = d_{\mathbb{R}^m}(p, \text{aff}(\sigma_p))$. A poorly-shaped simplex can be characterized by the existence of a relatively small altitude. The *thickness* of a j -simplex σ is the dimensionless quantity

$$\Upsilon_\sigma = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p \in \sigma} \frac{D_\sigma(p)}{j \Delta_\sigma} & \text{otherwise.} \end{cases}$$

We say that σ is Υ_0 -thick, if $\Upsilon_\sigma \geq \Upsilon_0$. If σ is Υ_0 -thick, then so are all of its faces. Indeed if $\tau \leq \sigma$, then the smallest altitude in τ cannot be smaller than that of σ , and also $\Delta_\tau \leq \Delta_\sigma$.

Our definition of thickness is essentially the same as that employed by Munkres [Mun68]. Munkres defined the thickness of σ^j as $\frac{r(\sigma^j)}{\Delta_{\sigma^j}}$, where $r(\sigma^j)$ is the radius of the largest contained ball centred at the barycentre. This definition of thickness turns out to be equal to $\frac{j}{j+1} \Upsilon_{\sigma^j}$.

In the previous chapters we employed a volume-based measure of simplex quality, and variations on this, typically referred to as *fatness*, have been popular in works on higher dimensional Delaunay-based meshing [CDE⁺00b, Li03b, BWY11]. We find a direct bound on the altitudes to be more convenient, because it yields a cleaner and tighter connection between the geometry and the linear algebra of simplices. Typically, a bound on some geometric displacement related to a simplex is obtained by bounding the inverse of a matrix associated with the simplex, and thickness is well suited for this task.

As a motivating example, consider the problem of bounding the angle between the affine hull of a simplex and an affine space that lies close to all the vertices of the simplex. Such a bound is relevant when meshing submanifolds of Euclidean space, for example, where it is desired that the affine hulls of the simplices are in agreement with the nearby tangent spaces of the manifold.

Whitney [Whi57b, p. 127], see Lemma 2.3.2, obtained such a bound, which manifestly depends on the quality of the simplex. Using thickness as a quality measure we obtain a sharper result:

Lemma 7.2.1 (Revisiting Whitney's angle bound) *Suppose σ is a j -simplex whose vertices all lie within a distance η from a k -dimensional affine space, $H \subset \mathbb{R}^m$, with $k \geq j$. Then*

$$\sin \angle(\text{aff}(\sigma), H) \leq \frac{2\eta}{\Upsilon_\sigma \Delta_\sigma}.$$

The idea of the proof is to express the unit vector u in Eq. (7.1) in terms of a basis for $\text{aff}(\sigma)$ given by the edges of σ that emanate from some arbitrarily chosen vertex. The

projection $\tilde{u}$ of u into H can then be expressed in terms of the projected basis vectors, using the same vector of coefficients. Since the vertices of σ all lie close to H , the projected basis vectors do not differ significantly from the originals, so bounding the magnitude of the difference between u and $\tilde{u}$ comes down to bounding the magnitude of the vector of coefficients of the unit vector u . This bound depends on how well-conditioned the basis is, and this is closely related to the thickness of σ .

These observations can be conveniently expressed and made concrete in terms of the singular values of a matrix. An excellent introduction to singular values can be found in the book by Trefethen and Bau [TB97, Chap. 4 & 5], but for our purposes we are primarily concerned with the largest and the smallest singular values, which we now describe.

We denote the i^{th} singular value of a matrix A by $s_i(A)$. The singular values are non-negative and ordered by magnitude. The largest singular value can be defined as

$$s_1(A) = \sup_{\|x\|=1} \|Ax\|;$$

it is the magnitude of the largest vector in the range of the unit sphere. The first singular value also defines the operator norm: $\|A\| = s_1(A)$. The standard observation that a bound on the norms of the columns of A yields a bound on $\|A\|$ is obtained by a short calculation.

Lemma 7.2.2 *If $\eta > 0$ is the least upper bound on the norms of the columns of an $m \times j$ matrix A , then*

$$\eta \leq \|A\| \leq \sqrt{j}\eta.$$

We will also be interested in obtaining a lower bound on the smallest singular value which, for an $m \times j$ matrix A with $j \leq m$, may be defined as

$$s_j(A) = \inf_{\|x\|=1} \|Ax\|.$$

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

Lemma 7.2.3 *If A is an $m \times j$ matrix of rank $j \leq m$, then $s_i(A^\dagger) = s_{j-i+1}(A)^{-1}$.*

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

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




Figure 7.1: Choosing p_0 as the origin, the edges emanating from p_0 in $\sigma = [p_0, \dots, p_j]$ form a basis for $\text{aff}(\sigma)$. The proof of Lemma 7.2.4 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.

Lemma 7.2.4 (Thickness and singular value) *Let $\sigma = [p_0, \dots, p_j]$ be a non-degenerate j -simplex in $\mathbb{R}^m$, with $j > 0$, and let P be the $m \times j$ matrix whose i^{th} column is $p_i - p_0$. Then the i^{th} row of $P^\dagger$ is given by $w_i^\top$, where w_i is orthogonal to $\text{aff}(\sigma_{p_i})$, and*

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

We have the following bounds on the smallest singular value of P :

$$s_j(P) \geq \sqrt{j} \Upsilon_\sigma \Delta_\sigma.$$

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

$$\max_{1 \leq i \leq j} D_\sigma(p_i)^{-1} = \left(\min_{1 \leq i \leq j} D_\sigma(p_i) \right)^{-1} = (j \Upsilon_\sigma \Delta_\sigma)^{-1},$$

Lemma 7.2.2, yields

$$s_1(P^\dagger) \leq (\sqrt{j} \Upsilon_\sigma \Delta_\sigma)^{-1},$$

because $s_i(A^\top) = s_i(A)$ for any matrix A . The stated bounds on $s_j(P)$ follow from Lemma 7.2.3. □

The proof of Lemma 7.2.4 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 $s_j(P)$. By Lemma 7.2.2, $s_1(P) \leq \sqrt{j} \Delta_\sigma$, and thus $\Upsilon_\sigma \leq \frac{s_j(P)}{s_1(P)}$. In other words,

Υ_σ^{-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. For example, we will use Lemma 7.2.4 in the demonstration of Lemma 7.4.1, which is the technical lemma related to the stability of the space of circumcentres of a simplex. Lemma 7.2.4 also facilitates a concise demonstration of Whitney's angle bound:

Proof of Lemma 7.2.1 Suppose $\sigma = [p_0, \dots, p_j]$. Choose p_0 as the origin of $\mathbb{R}^m$, and let $U \subset \mathbb{R}^m$ be the vector subspace defined by $\text{aff}(\sigma)$. Let W be the k -dimensional subspace parallel to H , and let $\pi : \mathbb{R}^m \rightarrow W$ be the orthogonal projection onto W .

Let $u \in U$ be a unit vector. Since the vectors $v_i = (p_i - p_0)$, $i \in \{1, \dots, j\}$ form a basis for U , we may write $u = Pa$, where P is the $m \times j$ matrix whose i^{th} column is v_i , and $a \in \mathbb{R}^j$ is the vector of coefficients. Then, defining $X = P - \pi P$, we get

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

W is at a distance less than η from H , because $p_0 \in W$ and $d_{\mathbb{R}^m}(p_i, H) \leq \eta$ for all $0 \leq i \leq j$. It follows that $\|v_i - \pi v_i\| \leq 2\eta$, and Lemma 7.2.2 yields

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

Observing that $1 = \|u\| = \|Pa\| \geq s_j(P) \|a\|$, we find

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

and the result follows from Lemma 7.2.4 and Eq. 7.1. $\square$

7.2.3 Complexes

Given a finite set P , an *abstract simplicial complex* is a set of subsets $K \subset 2^P$ such that if $\sigma \in K$, then every subset of σ is also in K . The Delaunay complexes we study are abstract simplicial complexes, but their simplices carry a canonical geometry induced from the inclusion map $\iota : P \hookrightarrow \mathbb{R}^m$. (We assume ι is injective on P , and so do not distinguish between P and $\iota(P)$.) For each abstract simplex $\sigma \in K$, we have an associated geometric simplex $\text{conv}(\iota(\sigma))$, and normally when we write $\sigma \in K$, we are referring to this geometric object. Occasionally, when it is convenient to emphasise a distinction, we will write $\iota(\sigma)$ instead of σ .

Thus we view such a K as a set of simplices in $\mathbb{R}^m$, and we refer to it as a *complex*, but it is not generally a (geometric) simplicial complex. A *geometric simplicial complex* is a finite collection G of simplices in $\mathbb{R}^N$ such that if $\sigma \in G$, then all of the faces of σ also belong to G , and if $\sigma, \tilde{\sigma} \in G$ and $\sigma \cap \tilde{\sigma} \neq \emptyset$, then $\tau = \sigma \cap \tilde{\sigma}$ is a simplex and $\tau \leq \sigma$ and $\tau \leq \tilde{\sigma}$. Observe that the simplices in a geometric simplicial complex are necessarily non-degenerate. An abstract simplicial complex is defined from a geometric simplicial complex in an obvious way. A *geometric realization* of an abstract simplicial complex K is a geometric simplicial complex whose associated abstract simplicial complex may be identified with K . A geometric realization always exists for any complex. Details can be found in algebraic topology textbooks; the book by Munkres [Mun84] for example.

The *dimension of a complex* K is the largest dimension of the simplices in K . We say that K is an m -complex, to mean that it is of dimension m . The complex K is a *pure m -complex* if it is an m -complex, and every simplex in K is the face of an m -simplex.

The *carrier* of an abstract complex K is the underlying topological space $|K|$, associated with a geometric realization of K . Thus if G is a geometric realization of K , then $|K| = \bigcup_{\sigma \in G} \sigma$. For our complexes, the inclusion map ι induces a continuous map $\iota : |K| \rightarrow \mathbb{R}^m$, defined by barycentric interpolation on each simplex. If this map is injective, we say that K is *embedded*. In this case ι also defines a geometric realization of K , and we may identify the carrier of K with the image of ι .

A subset $K' \subset K$ is a *subcomplex* of K if it is also a complex. The *star* of a subcomplex $K' \subseteq K$ is the subcomplex generated by the simplices incident to K' . I.e., it is all the simplices that share a face with a simplex of K' , plus all the faces of such simplices. This is a departure from a common usage of this same term in the topology literature. The star of K' is denoted $\text{star}(K')$ when there is no risk of ambiguity, otherwise we also specify the parent complex, as in $\text{star}(K'; K)$. A simple example of the star of a complex is depicted in Figure 7.2.




Figure 7.2: The star of a subcomplex $K' \subset K$ is the subcomplex $\text{star}(K') \subset K$ that consists all the simplices that share a face with K' (this includes all of K' itself), and all the faces of these simplices. Here we show an embedded 2-complex, with all 2-simplices shaded. The subcomplex K' consists of the two indicated triangles, and their faces (blue). The simplices of $\text{star}(K')$ are shown in bold (red and blue). The other simplices do not belong to $\text{star}(K')$ (black).

A *triangulation* of $P \subset \mathbb{R}^m$ is an embedded complex K with vertices P such that $|K| = \text{conv}(P)$. A complex K is a *j -manifold complex* if the star of every vertex is isomorphic to the star of a triangulation of $\mathbb{R}^j$. In order to exploit the local nature of the definition of a manifold complex, it is convenient to have a local notion of triangulation for the star of a vertex in K , even if the whole of K is not a triangulation of its vertices:

Definition 7.2.5 (Triangulation at a point) A complex K is a triangulation at $p \in \mathbb{R}^m$ if:

1. p is a vertex of K .
2. $\text{star}(p)$ is embedded.
3. p lies in $\text{int}|\text{star}(p)|$.
4. For all $\tau \in K$, and $\sigma \in \text{star}(p)$, if $(\text{int } \tau) \cap \sigma \neq \emptyset$, then $\tau \in \text{star}(p)$.

In a general complex Condition 4 above is not a local property, however in the case of Delaunay complexes that interests us here, local conditions are sufficient to verify the condition, as we will show in Section 7.3.2.1. Observe also that Condition 4 also precludes intersections with degenerate simplices, since such a simplex would have a face that violates the condition.

If σ is a simplex with vertices in P , then any map $\zeta : P \rightarrow \tilde{P} \subset \mathbb{R}^m$ defines a simplex $\zeta(\sigma)$ whose vertices in $\tilde{P}$ are the images of vertices of σ . If K is a complex on P , and $\tilde{K}$ is a complex on $\tilde{P}$, then ζ induces a *simplicial map* $K \rightarrow \tilde{K}$ if $\zeta(\sigma) \in \tilde{K}$ for every $\sigma \in K$. We denote this map by the same symbol, ζ . We are interested in the case when ζ is an *isomorphism*, which means it establishes a bijection between K and $\tilde{K}$. We then say that K and $\tilde{K}$ are *isomorphic*, and write $K \cong \tilde{K}$, or $K \stackrel{\zeta}{\cong} \tilde{K}$ if we wish to emphasise that the correspondence is given by ζ .

A simplicial map $\zeta : K \rightarrow \tilde{K}$ defines a continuous map $\zeta : |K| \rightarrow |\tilde{K}|$, by barycentric interpolation on each simplex $\sigma \in K$. We observe the following consequence of Brouwer's invariance of domain:

Lemma 7.2.6 *Suppose K is a complex with vertices $P \subset \mathbb{R}^m$, and $\tilde{K}$ a complex with vertices $\tilde{P} \subset \mathbb{R}^m$. Suppose also that K is a triangulation at $p \in P$, and that $\zeta : P \rightarrow \tilde{P}$ induces an injective simplicial map $\text{star}(p) \rightarrow \text{star}(\zeta(p))$. If $\text{star}(\zeta(p))$ is embedded, then*

$$\zeta(\text{star}(p)) = \text{star}(\zeta(p)),$$

and $\zeta(p)$ is an interior point of $\tilde{P}$.

Proof We need to show that $\text{star}(\zeta(p)) \subseteq \zeta(\text{star}(p))$. Since $\text{star}(p)$ is embedded, ζ defines a continuous map $\zeta : |\text{star}(p)| \rightarrow |\text{star}(\zeta(p))|$ that is injective on each simplex. Since $\text{star}(\zeta(p))$ is also embedded, this continuous map is injective on $|\text{star}(p)|$. Since K is a triangulation at p , there is an open ball B centred at p such that $B \subset \text{int}|\text{star}(p)|$. Then $\zeta|_B : \mathbb{R}^m \supset B \rightarrow \zeta(B) \subset \mathbb{R}^m$ is a homeomorphism by Brouwer's invariance of domain [Dug66, Ch. XVII]. It follows that $\zeta(p)$ is an interior point of $\tilde{P}$.

Suppose $\sigma \in \text{star}(\zeta(p))$ and $\zeta(p)$ is a vertex of σ . Then, since σ is not degenerate, there is a point $x \in \zeta(B) \cap \text{int } \sigma$, and from the above argument, x also lies in the interior of some simplex $\tilde{\tau} \in \zeta(\text{star}(p)) \subseteq \text{star}(\zeta(p))$. Since $\text{star}(\zeta(p))$ is embedded, $\tilde{\tau} \cap \sigma$ is a face of σ and of $\tilde{\tau}$, but since x is in the interior of both simplices, it must be that $\tilde{\tau} = \sigma$. Thus $\sigma \in \zeta(\text{star}(p))$.

If $\sigma \in \text{star}(\zeta(p))$, then there is some $\tau \in \text{star}(\zeta(p))$ such that $\zeta(p)$ is a vertex of τ and $\sigma \leq \tau$. Since $\tau \in \zeta(\text{star}(p))$, we also have $\sigma \in \zeta(\text{star}(p))$, by the definition of a simplicial map. $\square$

7.3 Parameterized genericity

In this section we examine the Delaunay complex of $P \subset \mathbb{R}^m$, taking the view that poorly-shaped simplices arise from almost degenerate configurations of points. We introduce the concept of a protected Delaunay ball, which leads to a parameterized definition of genericity. We then show that a lower bound on the protection of the maximal simplices yields a lower bound on their thickness.

7.3.1 The Delaunay complex

An *empty ball* is one that contains no point from P .

Definition 7.3.1 (Delaunay complex) A Delaunay ball is a maximal empty ball. Specifically, $B = B_{\mathbb{R}^m}(x, r)$ is a Delaunay ball if any empty ball centred at x is contained in B . A simplex σ is a Delaunay simplex if there exists some Delaunay ball B such that the vertices of σ belong to $\partial B \cap P$. The Delaunay complex is the set of Delaunay simplices, and is denoted $\text{Del}(P)$.

The Delaunay complex has the combinatorial structure of an abstract simplicial complex, but $\text{Del}(P)$ is embedded only when P satisfies appropriate genericity requirements, as discussed in Section 7.3.2. Otherwise, $\text{Del}(P)$ contains degenerate simplices. We make here some observations that are not dependent on assumptions of genericity.




Figure 7.3: Lemma 7.3.2: If the affine hull of σ is not full dimensional, then a Delaunay ball has freedom to expand, and σ must be the face of a higher dimensional Delaunay simplex.

The union of the Delaunay simplices is $\text{conv}(P)$. A simplex $\sigma \in \text{Del}(P)$ is a *boundary simplex* if all its vertices lie on $\partial \text{conv}(P)$. We observe

Lemma 7.3.2 (Maximal simplices) If $\text{aff}(P) = \mathbb{R}^m$, then every Delaunay j -simplex, σ , is a face of a Delaunay simplex σ' with $\dim \text{aff}(\sigma') = m$. In particular, if $j \leq m$, then σ is a face of a Delaunay m -simplex. If σ is not a boundary simplex, and $\dim \text{aff}(\sigma) < m$, then there are at least two Delaunay $(j+1)$ -simplices that have σ as a face.

Proof Suppose $\dim \text{aff}(\sigma) < m$. Let $B = B_{\mathbb{R}^m}(c, r)$ be a Delaunay ball for σ . Let ℓ be the line through c and c_σ . If $c = c_\sigma$, let ℓ be any line through c and orthogonal to $\text{aff}(\sigma)$. There must be a point $\hat{c} \in \ell$ such that the circumscribing ball for σ centred at $\hat{c}$ is not empty. If this were not the case, we would have $\text{aff}(\sigma) = \text{aff}(P)$, and thus $\dim \text{aff}(P) < m$. It follows then (from the continuity of the radius of the circumballs parameterized by ℓ), that there is a point $c' \in [c, \hat{c}]$ that is the centre of a Delaunay ball for a simplex σ' that has σ as a proper face. The first assertion follows.

The second assertion follows from the same argument, and the observation that if σ is not on the boundary of $\text{conv}(P)$, then there must be non-empty balls centred on ℓ at either side of c . If $p \in P \setminus \text{aff}(\sigma)$ is on the boundary of an empty ball centred at one side of c , by the intersection properties of spheres, it cannot be on the boundary of an empty ball centred on the other side of c . Thus there must be at least two distinct Delaunay $(k+1)$ -simplices that share σ as a face. $\square$

Lemma 7.3.2 gives rise to the following observation, which plays an important role in Section 7.3.3, where we argue that protecting the Delaunay m -simplices yields a thickness bound on the simplices.

Lemma 7.3.3 (Separation) *If $\tau \in \text{Del}(P)$ is a j -simplex that is not a boundary simplex, and $q \in P \setminus \tau$, then there is a Delaunay m -simplex σ^m which has τ as a face, but does not include q .*

Proof Assume $j < m$, for otherwise there is nothing to prove. If $\sigma = q * \tau$ is not Delaunay, the assertion follows from the first part of Lemma 7.3.2. Assume σ is Delaunay and let $\tilde{\sigma}^m$ be a Delaunay m -simplex that has σ as a face. Thus $\tilde{\sigma}^m = q * \sigma^{m-1}$ for some Delaunay $(m-1)$ -simplex, σ^{m-1} . Since $\tau \leq \sigma^{m-1}$ does not belong to the boundary of $\text{conv}(P)$, neither does σ^{m-1} , so by the second part of Lemma 7.3.2, there is another Delaunay m -simplex σ^m that has σ^{m-1} (and therefore τ) as a face. Since σ^m is distinct from $\tilde{\sigma}^m$, it does not have q as a vertex. $\square$

7.3.1.1 The Delaunay complex in other metrics

We will also consider the Delaunay complex defined with respect to a metric d on $\mathbb{R}^m$ which differs from the Euclidean one. Specifically, if $P \subset U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is a metric, then we define the Delaunay complex $\text{Del}_d(P)$ with respect to the metric d .

The definitions are exactly analogous to the Euclidean case: A Delaunay ball is a maximal empty ball $B_d(x, r)$ in the metric d . The resulting Delaunay complex $\text{Del}_d(P)$ consists of all the simplices which are circumscribed by a Delaunay ball with respect to the metric d . The simplices of $\text{Del}_d(P)$ are, possibly degenerate, geometric simplices in $\mathbb{R}^m$. As for $\text{Del}(P)$, $\text{Del}_d(P)$ has the combinatorial structure of an abstract simplicial complex, but unlike $\text{Del}(P)$, $\text{Del}_d(P)$ may fail to be embedded even when there are no degenerate simplices.

7.3.2 Protection

A Delaunay simplex σ is δ -protected if it has a Delaunay ball B such that $d_{\mathbb{R}^m}(q, \partial B) > \delta$ for all $q \in P \setminus \sigma$. We say that B is a δ -protected Delaunay ball for σ . If $\tau < \sigma$, then B is also a Delaunay ball for τ , but it cannot be a δ -protected Delaunay ball for τ . We say that σ is *protected* to mean that it is δ -protected for some unspecified $\delta \geq 0$.




Figure 7.4: A Delaunay simplex σ is δ -protected if it has a Delaunay ball $B_{\mathbb{R}^m}(c, r)$ such that $\bar{B}_{\mathbb{R}^m}(c, r + \delta) \cap (P \setminus \sigma) = \emptyset$.

Definition 7.3.4 (δ -generic) A finite set of points $P \subset \mathbb{R}^m$ is δ -generic if $\text{aff}(P) = \mathbb{R}^m$, and all the Delaunay m -simplices are δ -protected. The set P is simply generic if it is δ -generic for some unspecified $\delta \geq 0$.

Observe that we have employed a strict inequality in the definition of δ -protection. In particular, a δ -generic point set is generic even when $\delta = 0$. In order for the quantity δ to be meaningful, it should be considered with respect to a sampling radius ϵ for P . We will always assume that $\delta \leq \epsilon$.

In his seminal work, Delaunay [Del34] demonstrated that if there is no empty ball with $m + 2$ points from P on its boundary, then $\text{Del}(P)$ is realized as a simplicial complex in $\mathbb{R}^m$. In other words $\text{Del}(P)$ is an embedded complex, and in fact it is a triangulation of P , the *Delaunay triangulation*. If P is generic according to Definition 8.2.8, then Delaunay's criterion will be met. This is obvious if there are no degenerate m -simplices, and Definition 8.2.8 ensures that a degenerate m -simplex cannot exist in $\text{Del}(P)$, as shown by Lemma 7.3.5 below.

In particular, if P is generic if and only if there are no Delaunay simplices with dimension higher than m . We can say more. There are no degenerate Delaunay simplices. This can be inferred directly from Delaunay's result [Del34], but is also easily established from Lemma 7.3.2. In Section 7.3.3 we will quantify this observation with a bound on the thickness of the Delaunay simplices.

The δ -generic assumption means that all the Delaunay m -simplices are δ -protected, but the lower dimensional Delaunay do not necessarily enjoy this level of protection. The fact that there are no degenerate Delaunay simplices implies that all the simplices of all dimensions are $\tilde{\delta}$ -protected for some $\tilde{\delta} > 0$.

7.3.2.1 Local Delaunay triangulation

Delaunay avoided boundary complications by assuming a periodic point set, but we are particularly interested in the case where the point sets come from local patches of a well-sampled compact manifold without boundary. Periodic boundary conditions are not

appropriate in this setting, but this is not a problem because, as we show here, Delaunay's argument applies locally.

Delaunay's proof that the Delaunay complex of a generic periodic point set is a triangulation of $\mathbb{R}^m$ consists of two observations. First it is observed that if two Delaunay simplices intersect, then they intersect in a common face. This shows that $\text{Del}(P)$ is embedded. The argument is not complicated by the presence of boundary points:

Lemma 7.3.5 (Embedded star) *Suppose $\text{aff}(P) = \mathbb{R}^m$ and $p \in P$. If all the m -simplices in $\text{star}(p; \text{Del}(P))$ are protected, then $\text{star}(p; \text{Del}(P))$ is embedded, and it is a pure m -complex.*

Proof We first observe that the m -simplices in $\text{star}(p)$ are not degenerate. If σ^m is degenerate, then by Lemma 7.3.2, there is a simplex τ with $\text{aff}(\tau) = \mathbb{R}^m$, and $\sigma^m < \tau$. We have $\tau \in \text{star}(p)$, since $p \in \tau$. An affinely independent set of $m+1$ vertices from τ defines a non-degenerate m -simplex $\tilde{\sigma}^m < \tau$, and since its unique circumball is also a Delaunay ball for τ , it cannot be protected, a contradiction.

Now suppose that $\sigma, \tau \in \text{star}(p)$ and $\sigma \cap \tau \neq \emptyset$. We need to show that they intersect in a common face. By Lemma 7.3.2, we may assume that σ and τ are m -simplices. Assume $\sigma \neq \tau$, and let B_1 and B_2 be the Delaunay balls for σ and τ . Then $\text{aff}(\partial B_1 \cap \partial B_2)$ defines an $(m-1)$ -flat, H . Since B_1 and B_2 are empty balls, H separates the interiors of σ and τ , and thus they must intersect in H , i.e., at the common face defined by the vertices in $\partial B_1 \cap \partial B_2$. $\square$

The second observation Delaunay made is that, in the case of a periodic (infinite) point set, every $(m-1)$ -simplex is the face of two m -simplices (Lemma 7.3.2). The implication here is that $\text{Del}(P)$ cannot have a boundary, and therefore must cover $\mathbb{R}^m$. Here we flesh out the argument for our purposes: If an embedded finite complex contains m -simplices then its topological boundary must contain $(m-1)$ -simplices. We first observe that the topological boundary of an embedded complex is defined by a subcomplex:

Lemma 7.3.6 (Boundary complex) *If K is an embedded (finite) complex in $\mathbb{R}^m$, then the topological boundary of $|K| \subset \mathbb{R}^m$ is defined by a subcomplex: $\partial|K| = |\text{bd}(K)|$, where the subcomplex $\text{bd}(K) \subset K$ is called the boundary complex of K .*




Figure 7.5: Diagram for the proof of Lemma 7.3.6.

Proof Since K is finite, $\partial|K|$ is contained in $|K|$. Suppose $x \in \partial|K|$. Then $x \in \text{int } \sigma^j$ for some $\sigma^j \in K$. We wish to show that $\sigma^j \subset \partial|K|$. Suppose to the contrary that $y \in \text{int } \sigma^j$, but y does not belong to $\partial|K|$. This means that $y \in \text{int } |K|$.

Consider the segment $\ell = [x, y] \subset \text{int } \sigma^j$. Let $Z \subset K$ be the subcomplex consisting of those simplices that do not contain σ^j . Let

$$r_1 = \min_{\sigma \in Z} d_{\mathbb{R}^m}(\ell, \sigma).$$

Choosing $r \leq r_1$, and $x' \in B_{\mathbb{R}^m}(x, r) \setminus |K|$, let $y' = y + (x' - x)$. Since $y \in \text{int } |K|$, we may assume that r is small enough so that $y' \in \text{int } |K|$.

Consider the segment $\ell' = [x', y']$. By construction, $\ell' \cap |Z| = \emptyset$. However, consider the point $u \in \text{int } \ell'$ that is the point in $\ell' \cap |K|$ that is closest to x' . The point u lies in the interior of some simplex $\sigma' \in K$, but we cannot have $\sigma^j \leq \sigma'$. Indeed if this were the case, x' would lie in $\text{aff}(\sigma')$, and so $u \in \partial \sigma'$, contradicting the assumption that $u \in \text{int } \sigma'$.

But this means that $\sigma' \in Z$, which contradicts the fact that $\ell' \cap |Z| = \emptyset$. Therefore we must have $y \in \partial |K|$ for all $y \in \text{int } \sigma^j$.

Finally, observe that if $\tau < \sigma^j$, then $\tau \subset \partial |K|$, since $\partial |K|$ is closed. $\square$

Lemma 7.3.7 (Pure boundary complex) *If K is a (finite) pure m -complex embedded in $\mathbb{R}^m$, then its boundary complex is a pure $(m-1)$ -complex.*




Figure 7.6: Diagram for the proof of Lemma 7.3.7.

Proof Since K is finite, $\text{bd}(K)$ is nonempty; it contains at least the vertices in $\partial \text{conv}(|K|)$. We will show that if $\sigma^j \in \text{bd}(K)$, is a j -simplex, with $0 \leq j < m-1$, then there is a $\sigma^k \in \text{bd}(K)$ with $\sigma^j < \sigma^k$. The result then follows, since $\text{bd}(K)$ cannot contain m -simplices, because K is embedded.

Suppose $\sigma^j \in \text{bd}(K)$, and $x \in \text{int } \sigma^j$. Let $Z \subset K$ be the subcomplex consisting of simplices that do not contain σ^j , and let

$$r = \min_{\sigma \in Z} d_{\mathbb{R}^m}(x, \sigma).$$

Let $B = B_{\mathbb{R}^m}(x, r)$, and choose $y \in B \setminus |K|$. Let F be the $(m-j)$ -dimensional affine space orthogonal to $\text{aff}(\sigma^j)$ and containing y , and let $S^{m-j-1} = F \cap \partial B_{\mathbb{R}^m}(x, r')$, where $r' = \|x - y\|$. See Figure 7.6.

Since K is pure, there is an m -simplex σ^m with $\sigma^j < \sigma^m$. We have $\sigma^m \cap S^{m-j-1} \neq \emptyset$. Indeed, choose $w \in \text{int } \sigma^m$, and $u \in \sigma^j$ different from x , and observe that the plane Q defined by x, w, u intersects $B \cap \sigma^m$ in a semi-disk, by construction of B . By the construction of S^{m-j-1} , it must intersect this semidisk.

Let $z \in S^{m-j-1}$ be a point that minimises the geodesic distance in S^{m-j-1} to y . Then $z \in \partial|K|$. Thus $z \in \text{int } \sigma^k$ for some $\sigma^k \in \text{bd}(K)$, and since $z \in B$, σ^k cannot belong to Z . Thus $\sigma^j \leq \sigma^k$, but since $S^{m-j-1} \cap \sigma^j = \emptyset$, we have $\sigma^j < \sigma^k$. $\square$

From Lemma 7.3.2, Lemma 7.3.5, and Lemma 7.3.7, one can verify that if P is generic then $\partial\text{conv}(P) = \partial|\text{Del}(P)|$, and thus obtain the standard result that $\text{Del}(P)$ is a triangulation of P . However, we are interested localizing the result, without the assumption that the entire point set is generic. We have the following local version of Delaunay's triangulation result:

Lemma 7.3.8 (Local Delaunay triangulation) *If $p \in P$ is an interior point, and the Delaunay m -simplices incident to p are protected, then $\text{Del}(P)$ is a triangulation at p .*

Proof By Lemma 7.3.5, $\text{star}(p)$ is a pure m -complex, and it is embedded. It follows then from Lemma 7.3.7, that the boundary complex $\text{bd}(\text{star}(p))$ is a pure $(m-1)$ -complex. Thus p cannot belong to $\text{bd}(\text{star}(p))$. Indeed, it follows from Lemma 7.3.2 that any $(m-1)$ -simplex $\sigma \in \text{star}(p)$ is the face of at least two m -simplices in $\text{Del}(P)$, and if $p \in \sigma$, then both of these m -simplices belong to $\text{star}(p)$, and are embedded, with intersection σ . Thus p cannot belong to an $(m-1)$ -simplex in $\text{bd}(\text{star}(p))$, and therefore $p \in \text{int}|\text{star}(p)|$.

It remains to verify Condition 4 of Definition 7.2.5. The argument is similar to the proof of Lemma 7.3.5: Suppose $x \in (\text{int } \tau) \cap \sigma$ for $\sigma \in \text{star}(p)$. We may assume that σ is an m -simplex. Then consider the Delaunay balls B_1 for σ and B_2 for τ . If $B_1 = B_2$, then, since σ is protected, τ must be a face of σ , and so belong to $\text{star}(p)$. Assume then that $B_1 \neq B_2$, and let H be the $(m-1)$ -flat defined by $\text{aff}(\partial B_1 \cap \partial B_2)$. Since B_1 is empty, $x \in \text{int } \tau$ cannot lie in the open half-space defined by H and containing σ . Since $x \in \sigma$ also, it must lie in H , and therefore all vertices of τ lie in $H \cap \partial B_2 = H \cap \partial B_1$, and so τ is a face of σ . $\square$

7.3.2.2 Safe interior simplices

We wish to consider the properties of Delaunay triangulations in regions which are comfortably in the interior of $\text{conv}(P)$, and avoid the complications that arise as we approach the boundary of the point set. We introduce some terminology to facilitate this.

If none of the vertices of σ lie on $\partial\text{conv}(P)$, then it is an *interior simplex*. We wish to identify a subcomplex of the interior simplices of $\text{Del}(P)$ consisting of those simplices whose neighbour simplices are also all interior simplices with small circumradius. An interior simplex near the boundary of $\text{conv}(P)$ does not necessarily have its circumradius constrained by the sampling radius. However, we have the following:

Lemma 7.3.9 *If P is an ϵ -sample set, and $\sigma \in \text{Del}(P)$ has a vertex p such that $d_{\mathbb{R}^m}(p, \partial\text{conv}(P)) \geq 2\epsilon$, then $R_\sigma < \epsilon$ and σ is an interior simplex.*

Proof Let $B_{\mathbb{R}^m}(c, r)$ be a Delaunay ball for σ . We will show $r < \epsilon$. Suppose to the contrary. Let x be the point on $[c, p]$ such that $d_{\mathbb{R}^m}(p, x) = \epsilon$. Then p is the closest point in P to x , and so the sampling criteria imply that $d_{\mathbb{R}^m}(x, \partial\text{conv}(P)) < \epsilon$. But then $d_{\mathbb{R}^m}(p, \partial\text{conv}(P)) \leq d_{\mathbb{R}^m}(p, x) + d_{\mathbb{R}^m}(x, \partial\text{conv}(P)) < 2\epsilon$, contradicting the hypothesis on p .

Thus $r < \epsilon$, and it follows that σ is an interior simplex because if $q \in \sigma$, then $d_{\mathbb{R}^m}(p, q) \leq 2r < d_{\mathbb{R}^m}(p, \partial \text{conv}(P))$. $\square$

This suggests the following:

Definition 7.3.10 (Deep interior points) Suppose $P \subset \mathbb{R}^m$ is an ϵ -sample set. The subset $P_I \subset P$ consisting of all $p \in P$ with $d_{\mathbb{R}^m}(p, \partial \text{conv}(P)) \geq 4\epsilon$ is the set of deep interior points.

By Lemma 7.3.9, all the simplices that include a deep interior point, as well as all the neighbours of such simplices, will have a small circumradius. For technical reasons it is inconvenient to demand that *all* the Delaunay m -simplices be δ -protected. We focus instead on a subset defined with respect to a set of deep interior points:




Figure 7.7: If P is δ -generic for P_J , then the safe interior simplices are the simplices in $\text{star}(P_J)$. Here P_J consists of the two large vertices (blue). They must be at least 4ϵ from $\partial \text{conv}(P)$ (which is not depicted in the figure). The safe interior simplices are shaded. All the simplices in $\text{star}(\text{star}(P_J))$ are δ -protected. These simplices have bold outlines (red), but are not necessarily shaded.

Definition 7.3.11 (δ -generic for P_J) The set $P \subset \mathbb{R}^m$ is δ -generic for P_J if $P_J \subseteq P_I$ and all the m -simplices in $\text{star}(\text{star}(P_J; \text{Del}(P)))$ are δ -protected. The safe interior simplices are the simplices in $\text{star}(P_J; \text{Del}(P))$.

Thus the safe interior simplices are determined by our choice of $P_J \subseteq P_I$, and our protection requirements ensure that all the m -simplices that share a face with a safe interior simplex are δ -protected and have a small circumradius. A schematic example is depicted in Figure 7.7.

7.3.3 Thickness from protection

Our goal here is to demonstrate that the safe interior simplices on a δ -generic point set are Υ_0 -thick. If $\delta = \nu_0\epsilon$, for some constant $\nu_0 \leq 1$, then we obtain a constant Υ_0 which depends only on ν_0 . The key observation is that together with Lemma 7.3.3, protection imposes constraints on all the Delaunay simplices; they cannot be too close to being

degenerate. In the particular case that $j = 0$, Lemma 7.3.3 immediately implies that the vertices of the safe interior simplices are δ -separated:

Lemma 7.3.12 (Separation from protection) *If P is δ -generic for P_J , then $L_\sigma > \delta$ for any safe interior simplex σ .*

Lemma 7.3.13 *Suppose that $B = B_{\mathbb{R}^m}(c, r)$ is a Delaunay ball for $\sigma = q * \tau$ with $r < \epsilon$ and that $L_\tau \geq \lambda$ for some $\lambda \leq \epsilon$. Suppose also that $\tau \leq \sigma'$ and that σ is not a face of σ' .*

If B' is a δ -protected Delaunay ball for σ' , and $H = \text{aff}(\partial B \cap \partial B')$, then

$$d_{\mathbb{R}^m}(q, H) > \frac{\sqrt{3}\delta}{4\epsilon}(\lambda + \delta).$$

It follows that, if P is δ -generic for P_J , with sampling radius ϵ , and τ is a safe interior simplex, then

$$D_\sigma(q) > \frac{\sqrt{3}\delta^2}{2\epsilon}.$$

Proof Let $B' = B_{\mathbb{R}^m}(c', r')$ be the δ -protected Delaunay ball for σ' . Our geometry will be performed in the plane, Q , defined by c , c' , and q . This plane is orthogonal to the $(m-1)$ -flat H , and it follows that the distance $d_{\mathbb{R}^m}(q, H)$ is realized by a segment in the plane Q : the projection, q^* , of q onto H lies in Q , and $d_{\mathbb{R}^m}(q, H) = d_{\mathbb{R}^m}(q, q^*)$.




Figure 7.8: Diagram for Lemma 7.3.13. (a) When H separates q and c then $d_{\mathbb{R}^d}(q, q^*) > \delta$. (b) Otherwise, a lower bound on the distance between q and its projection q^* on H is obtained by an upper bound on the angle $\angle qab$.

If H separates q from c , then $\partial B'$ separates q from q^* , and $d_{\mathbb{R}^m}(q, q^*) > d_{\mathbb{R}^m}(q, \partial B') > \delta$, since B' is δ -protected (Figure 7.8(a)). The lemma then follows since λ and δ are each no larger than ϵ . Thus assume that q and c lie on the same side of H , as shown in Figure 7.8(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 $d_{\mathbb{R}^m}(q, q^*)$ by finding an upper bound on the angle $\gamma = \angle qab$. This is the same as the standard calculation for upper-bounding the angles in a triangle with bounded circumradius to shortest edge ratio. 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 $d_{\mathbb{R}^m}(q, q^*) > \delta$ and the lemma is again trivially satisfied.

Since $d_{\mathbb{R}^m}(a, q) > \delta$, we have $d_{\mathbb{R}^m}(q, q^*) = d_{\mathbb{R}^m}(a, q) \sin \gamma > \delta \sin \gamma$. Also note that $d(a, b) \geq 2R_\tau \geq L_\tau \geq \lambda$. Let $\alpha = \angle qac$. Then $\cos \alpha = \frac{d_{\mathbb{R}^m}(a, q)}{2r} \geq \frac{\delta}{2\epsilon}$, which means that $\alpha \leq \arccos \frac{\delta}{2\epsilon} \leq \frac{\pi}{2}$. Similarly, with $\beta = \angle cab$, we have $\beta \leq \arccos \frac{\lambda}{2\epsilon} \leq \frac{\pi}{2}$. Thus $\frac{\pi}{2} \leq \gamma = \alpha + \beta \leq \gamma'$, where

$$\gamma' = \arccos \frac{\delta}{2\epsilon} + \arccos \frac{\lambda}{2\epsilon}.$$

Since $\sin \gamma \geq \sin \gamma'$, when $\frac{\pi}{2} \leq \gamma \leq \gamma' \leq \pi$, we have

$$\begin{aligned} d_{\mathbb{R}^m}(q, q^*) &> \delta \sin \gamma \geq \delta \sin \gamma' \\ &= \delta \sin \left(\arccos \frac{\delta}{2\epsilon} + \arccos \frac{\lambda}{2\epsilon} \right) \\ &\geq \delta \left(\frac{\lambda}{2\epsilon} \sin \left(\arccos \frac{\delta}{2\epsilon} \right) + \frac{\delta}{2\epsilon} \sin \left(\arccos \frac{\lambda}{2\epsilon} \right) \right) \\ &\geq \frac{\sqrt{3}\delta}{4\epsilon} (\lambda + \delta), \end{aligned}$$

where the last inequality follows from $\lambda \leq \epsilon$ and $\delta \leq \epsilon$.

Since $\text{aff}(\tau) \subset H$, it follows that $D_\sigma(q) \geq d_{\mathbb{R}^m}(q, H)$, and if P is δ -generic for P_J , then $\lambda \geq \delta$, and Lemma 7.3.3 ensures that there is a δ -protected σ' that contains τ but not q . $\square$

We thus obtain a bound on the thickness of the safe interior simplices when P is δ -generic for P_J . Since Lemma 7.3.13 yields a lower bound of $\frac{\sqrt{3}\delta^2}{2\epsilon}$ on the altitudes of the safe interior simplices, and since $\Delta_\sigma \leq 2\epsilon$, we have that $\Upsilon_\sigma \geq \frac{\sqrt{3}\delta^2}{4\epsilon^2}$ for all safe interior σ . If $\delta = \nu_0\epsilon$, we obtain a constant thickness bound.

Theorem 7.3.14 (Thickness from protection) *If $P \subset \mathbb{R}^m$ is δ -generic for P_J with $\delta = \nu_0\epsilon$, where ϵ is a sampling radius for P , then the safe interior simplices are Υ_0 -thick, with*

$$\Upsilon_0 = \frac{\sqrt{3}\nu_0^2}{4}.$$

7.4 Delaunay stability

We find upper bounds on the magnitude of a perturbation for which a protected Delaunay ball remains a Delaunay ball. We consider both perturbations of the sample points in Euclidean space, and perturbations of the metric itself. The primary technical challenge is bounding the effect of a perturbation on the circumcentre of an m -simplex. We then find the relationship between the perturbation parameter ρ and the protection parameter δ which ensures that a δ -protected Delaunay simplex will remain a Delaunay simplex.

7.4.1 Perturbations and circumcentres

As expected, a bound on the displacement of the circumcentre requires a bound on the thickness of the simplex.

7.4.1.1 Almost circumcentres

If σ is thick, a point whose distances to the vertices of σ are all almost the same, will lie close to N_σ .

Lemma 7.4.1 *If $\sigma = [p_0, \dots, p_k] \subset \mathbb{R}^m$ is a non-degenerate k -simplex, and $x \in \mathbb{R}^m$ is such that*

$$\left| \|p_i - x\|^2 - \|p_j - x\|^2 \right| \leq \check{\rho}^2 \quad \text{for all } i, j \in [0, \dots, k], \quad (7.2)$$

then there is a point $c \in N_\sigma$ such that $\|c - x\| \leq \eta$, where

$$\eta = \frac{\check{\rho}^2}{2\Upsilon_\sigma \Delta_\sigma}.$$

In particular, if σ is an m -simplex then $x \in \bar{B}_{\mathbb{R}^m}(c_\sigma, \eta)$.

If the inequalities in Equations (7.2) are made strict, then the conclusions may also be stated with strict inequalities.

Proof First suppose $k = m$. The circumcentre of σ is given by the linear equations $\|c_\sigma - p_i\|^2 = \|c_\sigma - p_0\|^2$, or

$$(p_i - p_0)^\top c_\sigma = \frac{1}{2}(\|p_i\|^2 - \|p_0\|^2).$$

Letting b be the vector whose i^{th} component is defined by the right hand side of the equation, and letting P be the $m \times m$ matrix, whose i^{th} column is $(p_i - p_0)$, we write the equations in matrix form:

$$P^\top c_\sigma = b. \quad (7.3)$$

Without loss of generality, assume p_0 is the vertex that minimizes the distance to x . Then, defining x_a to be the vector whose i^{th} component is $\frac{1}{2}(\|p_i - x\|^2 - \|p_0 - x\|^2)$, we have $\|p_i - x\|^2 = \|p_0 - x\|^2 + 2(x_a)_i$, and we find

$$P^\top x = b - x_a. \quad (7.4)$$

From Equations (7.3) and (7.4) we have

$$\|c_\sigma - x\| = \left\| (P^\top)^{-1} x_a \right\| \leq \|P^{-1}\| \|x_a\|.$$

From Eq. (7.2), it follows that $\|x_a\| \leq \frac{\sqrt{m}\check{\rho}^2}{2}$, and from Lemmas 7.2.3 and 7.2.4 we have $\|P^{-1}\| \leq (\sqrt{m}\Upsilon_\sigma \Delta_\sigma)^{-1}$, and thus the result holds for full dimensional simplices.

If σ is a k -simplex with $k \leq m$, then we consider $\hat{x}$, the orthogonal projection of x into $\text{aff}(\sigma)$. We observe that $\hat{x}$ also must satisfy Eq. (7.2), and we conclude from the above argument that $\|c_\sigma - \hat{x}\| \leq \eta$. Then letting $c = c_\sigma + (x - \hat{x})$ we have that $c \in N_\sigma$ and $\|c - x\| \leq \eta$. $\square$

It will be convenient to have a name for a point that is almost equidistant to the vertices of a simplex:

Definition 7.4.2 A $\tilde{\rho}$ -centre for a simplex $\sigma = [p_0, \dots, p_k] \subset \mathbb{R}^m$ is a point x that satisfies

$$\left| \|p_i - x\| - \|p_j - x\| \right| \leq \tilde{\rho} \quad \text{for all } i, j \leq k. \quad (7.5)$$

With a bound on the distance from x to the vertices of σ , Lemma 7.4.1 can be transformed into a bound on the distance from a $\tilde{\rho}$ -centre to the closest true centre in N_σ :

Lemma 7.4.3 If $\sigma = [p_0, \dots, p_k] \subset \mathbb{R}^m$ is non-degenerate, and for some $\tilde{\rho} > 0$ the point $x \in \mathbb{R}^m$ is a $\tilde{\rho}$ -centre such that

$$\|p_i - x\| < \tilde{\epsilon} \quad \text{for all } i, j \leq k,$$

then there exists a $c \in N_\sigma$ such that $\|x - c\| < \eta$, where

$$\eta = \frac{\tilde{\epsilon}\tilde{\rho}}{\Upsilon_\sigma \Delta_\sigma}.$$

In particular, if σ is an m -simplex, then $x \in B_{\mathbb{R}^m}(c_\sigma, \eta)$.

Proof Let $R = \max_i \|p_i - x\|$ and $r = \min_i \|p_i - x\|$. Then

$$R^2 - r^2 = (R + r)(R - r) < 2\tilde{\epsilon}(R - r) \leq 2\tilde{\epsilon}\tilde{\rho},$$

and the result then follows from Lemma 7.4.1. $\square$

7.4.1.2 Circumcentres and metric perturbations

We will show here that for an Υ_0 -thick m -simplex σ in $\mathbb{R}^m$, and a metric d that is close to $d_{\mathbb{R}^m}$, there will be a metric circumcentre c near c_σ . We require the metric d to be continuous in the topology defined by $d_{\mathbb{R}^m}$. Henceforth, whenever we refer to a *perturbation of the Euclidean metric*, this topological compatibility will always be assumed.

The proof is a topological argument based on considering a mapping into $\mathbb{R}^m$ of a small ball around the circumcentre of σ . The mapping is based on the metric and is such that circumcentres get mapped to the origin. In the mapping associated to the Euclidean metric, points that get mapped close to the origin are $\tilde{\rho}$ -centres, and since the $\tilde{\rho}$ -centres are in the interior of the ball, the boundary of the ball does not get mapped close to the origin. A small perturbation of the metric yields a small perturbation in the mapping, and so we can argue that there is a homotopy between the mapping associated with the Euclidean metric and the one associated to the perturbed metric, such that no point on the boundary of the ball ever gets mapped to the origin. The situation is depicted schematically in Figure 7.9. A consideration of the degree of the mapping allows us to conclude that the ball must contain a circumcentre for the perturbed metric.

We will demonstrate the following:

Lemma 7.4.4 (Circumcentres: metric perturbation) Let $U \subset \mathbb{R}^m$, and let $d : U \times U \rightarrow \mathbb{R}$ be a continuous metric with respect to the topology defined by $d_{\mathbb{R}^m}$, and such that for any $x, y \in U$ with $d_{\mathbb{R}^m}(x, y) < 2\epsilon$, we have $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$, with

$$\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{8}.$$




Figure 7.9: The maps f_e and f (described in the main text) map circumcentres to the origin. Since $f_e^{-1}(0)$ contains a unique point, and $f_e(\partial B)$ lies far from the origin, a consideration of the degree of the mappings, together with the fact that $f_e(\partial B)$ and $f(\partial B)$ are close, reveals that $f^{-1}(0)$ cannot be empty, and thus B must contain a circumcentre of σ^m with respect to the metric d .

If $\sigma = [p_0, \dots, p_m] \subset U$ is an Υ_0 -thick m -simplex with $R_\sigma < \epsilon$, and $L_\sigma \geq \mu_0 \epsilon$, and such that $d_{\mathbb{R}^m}(p_i, \partial U) \geq 2\epsilon$, then there is a point

$$c \in B = B_{\mathbb{R}^m}(c_\sigma, \eta) \quad \text{with } \eta = \frac{8\rho}{\Upsilon_0 \mu_0},$$

and such that $d(c, p_i) = d(c, p_j)$ for all $p_i, p_j \in \sigma$.

In order to prove Lemma 7.4.4, we will use a particular case of Lemma 7.4.3:

Lemma 7.4.5 Suppose σ is an Υ_0 -thick m -simplex such that $L_\sigma \geq \mu_0 \epsilon$. If x is a $\tilde{\rho}$ -centre for σ with $d_{\mathbb{R}^m}(x, p) < 2\epsilon$ for all $p \in \sigma$, then $x \in B_{\mathbb{R}^m}(c_\sigma, \eta)$, where $\eta = \frac{2\tilde{\rho}}{\Upsilon_0 \mu_0}$.

Let $B = B_{\mathbb{R}^m}(c_\sigma, \eta)$ be the open ball which contains the $\tilde{\rho}$ -centres for σ . We will show that if $\tilde{\rho} = 4\rho$, then a circumcentre c for σ with respect to d will also lie in B . However, we make no claim that c is unique. Note that $\bar{B} \subset U$.

Consider the function $f_e : \bar{B} \rightarrow \mathbb{R}^m$ given by

$$f_e(x) = (d_{\mathbb{R}^m}(x, p_1) - d_{\mathbb{R}^m}(x, p_0), \dots, d_{\mathbb{R}^m}(x, p_m) - d_{\mathbb{R}^m}(x, p_0))^T. \quad (7.6)$$

Observe that f_e maps the circumcentre of σ , and only the circumcentre, to the origin: $f_e^{-1}(0) = \{c_\sigma\}$.

We construct a similar function for the metric d ,

$$f(x) = (d(x, p_1) - d(x, p_0), \dots, d(x, p_m) - d(x, p_0))^T, \quad (7.7)$$

and we will show that there must be a $c \in f^{-1}(0) \subset B$. We first show that there is a homotopy between f and f_e such that the image of $\partial \bar{B}$ never touches the origin:

Lemma 7.4.6 Under the hypotheses of Lemma 7.4.4, if $\tilde{\rho} = 4\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{2}$, then there is a homotopy $F : \bar{B} \times [0, 1] \rightarrow \mathbb{R}^m$ between $f_e(x) = F(x, 0)$ and $f(x) = F(x, 1)$ with $F(x, t) \neq 0$ for all $x \in \partial \bar{B}$ and $t \in [0, 1]$.

Proof We define the homotopy $F : \bar{B} \times [0, 1] \rightarrow \mathbb{R}^m$ by

$$F(x, t) = (1 - t)f_e(x) + tf(x).$$

By the bounds on $\tilde{\rho}$ and R_σ , for every $x \in \bar{B}$, and $p \in \sigma$, we have

$$d_{\mathbb{R}^m}(x, p) \leq \frac{2\tilde{\rho}}{\Upsilon_0\mu_0} + R_\sigma < 2\epsilon.$$

Thus it follows from Lemma 7.4.5 that $x \in \partial\bar{B}$ cannot be a $\tilde{\rho}$ -centre.

It is convenient to consider the max norm on $\mathbb{R}^m$ defined by the largest magnitude of the components: $\|f_e(x)\|_\infty = \max_i |f_e(x)_i|$. (This gives us a better bound than working with the standard Euclidean norm.) If $\|f_e(y)\|_\infty \leq \frac{\tilde{\rho}}{2}$, then y must be a $\tilde{\rho}$ -centre. Indeed, for all i and j , we would have

$$\begin{aligned} \left| \|p_i - y\| - \|p_j - y\| \right| &\leq \left| \|p_i - y\| - \|p_0 - y\| \right| + \left| \|p_0 - y\| - \|p_j - y\| \right| \\ &\leq \frac{\tilde{\rho}}{2} + \frac{\tilde{\rho}}{2} = \tilde{\rho} \end{aligned}$$

Thus, since $x \in \partial\bar{B}$ is not a $\tilde{\rho}$ -centre, we must have $\|f_e(x)\|_\infty > \frac{\tilde{\rho}}{2}$.

Also, from the hypothesis on d , we have

$$\|f_e(x) - f(x)\|_\infty \leq 2\rho = \frac{\tilde{\rho}}{2},$$

for all $x \in \partial\bar{B}$. Therefore, for all $x \in \partial\bar{B}$, we get

$$\|F(x, t)\|_\infty \geq \|f_e(x)\|_\infty - t\|f(x) - f_e(x)\|_\infty > 0$$

□

We will need the following observation:

Lemma 7.4.7 *The origin is a regular value for the function f_e defined in Eq. (7.6).*

Proof Choose a coordinate system such that $c_\sigma \in B$ is the origin. We show by a direct calculation that $\det J(f_e)_0 \neq 0$, where $J(f_e)_0$ is the Jacobian matrix representing the derivative of f_e in our coordinate system.

Let $p_i = (p_{i0}, \dots, p_{im})^\top$ for all $p_i \in \{p_0, \dots, p_m\}$. For $x = (x_1, \dots, x_m)^\top \in \mathbb{R}^m$, let $f_e(x) = (f_0(x), \dots, f_m(x))^\top$, where

$$f_i(x) = \|p_i - x\| - \|p_0 - x\| = \sqrt{\sum_{k=1}^m (p_{ik} - x_k)^2} - \sqrt{\sum_{k=1}^m (p_{0k} - x_k)^2}.$$

We find

$$\left. \frac{\partial f_i}{\partial x_j} \right|_0 = \frac{p_{0j} - p_{ij}}{R_\sigma},$$

and thus

$$J(f_e)_0 = -\frac{1}{R_\sigma} P^\top, \tag{7.8}$$

where as usual P is the matrix whose columns are $p_i - p_0$. Since $\text{vol}(\sigma^m) = \frac{|\det(P)|}{m!}$, Eq. (7.8) implies

$$|\det J(f_e)_0| = \frac{m! \text{vol}(\sigma^m)}{R_\sigma^m}.$$

Thus since $f_e^{-1}(0) = \{0\}$, 0 is a regular value for f_e provided σ is non-degenerate. $\square$

Lemma 7.4.4 now follows from a consideration of the degree of the mappings f and f_e relative to zero. The *degree* of a smooth map $g : \bar{B} \rightarrow \mathbb{R}^m$ at a regular point $p \in g(B)$ is defined by

$$\deg(g, p, B) = \sum_{x \in g^{-1}(p)} \text{sign det } J(g)_x,$$

where $J(g)_x$ is the Jacobian matrix of g at x . The exposition by Dancer [Dan00] is a good reference for the degree of maps from manifolds with boundary. It is shown that the definition of $\deg(g, p, B)$ extends to continuous maps g that are not necessarily differentiable. If $h : \bar{B} \rightarrow \mathbb{R}^m$ is homotopic to g by a homotopy $H : \bar{B} \times [0, 1] \rightarrow \mathbb{R}^m$ such that $H(x, t) \neq p$ for all $t \in [0, 1]$, and $x \in \partial B$, then $\deg(g, p, B) = \deg(h, p, B)$.

Since $f_e^{-1}(0) = \{c_\sigma\}$, it follows from Lemma 7.4.7 that $\deg(f_e, 0, B) = \pm 1$. Then Lemma 7.4.6 implies $\deg(f, 0, B) = \deg(f_e, 0, B)$, and since this is nonzero, it must be that $f^{-1}(0) \neq \emptyset$. The demonstration of Lemma 7.4.4 is complete.

7.4.1.3 Circumcentres and point perturbations

The exact same argument as was used to demonstrate Lemma 7.4.4 can be used to show that an m -simplex $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_m]$ whose vertices lie close to a thick m -simplex σ , will also have a circumcentre that lies close to c_σ . We replace the function f defined in Eq. (7.7) by the function

$$\tilde{f}(x) = (d_{\mathbb{R}^m}(x, \tilde{p}_1) - d_{\mathbb{R}^m}(x, \tilde{p}_0), \dots, d_{\mathbb{R}^m}(x, \tilde{p}_m) - d_{\mathbb{R}^m}(x, \tilde{p}_0))^T,$$

and the same argument goes through. We obtain:

Lemma 7.4.8 (Circumcentres: point perturbation) *Suppose that $\sigma = [p_0, \dots, p_m]$ is an Υ_0 -thick m -simplex with $R_\sigma < \epsilon$ and $L_\sigma \geq \mu_0 \epsilon$. Suppose also that $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_m]$ is such that $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in [0, \dots, m]$. If*

$$\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{8}, \quad \text{then} \quad d_{\mathbb{R}^m}(c_{\tilde{\sigma}}, c_\sigma) < \frac{8\rho}{\Upsilon_0 \mu_0}.$$

7.4.2 Perturbations and protection

Suppose $\zeta : P \rightarrow \tilde{P}$ is a ρ -perturbation. If σ is a δ -protected m -simplex in $\text{Del}(P)$, then we want an upper bound on ρ that will ensure that $\tilde{\sigma} = \zeta(\sigma)$ is protected in $\text{Del}(\tilde{P})$. The following definition will be convenient:

Definition 7.4.9 (Secure simplex) *A simplex $\sigma \in \text{Del}(P)$ is secure if it is a δ -protected m -simplex that is Υ_0 -thick and satisfies $R_\sigma < \epsilon$ and $L_\sigma \geq \mu_0 \epsilon$.*

Our stability results apply to subcomplexes of secure simplices, the definition of which employs multiple parameters. For safe interior simplices Lemma 7.3.12 and Theorem 7.3.14 allow us to consolidate some of these parameters with the ratio δ/ϵ :

Lemma 7.4.10 (Safe interior simplices are secure) *If P satisfies a sampling radius ϵ and is δ -generic for P_I , with $\delta = \nu_0\epsilon$, then the safe interior m -simplices are secure, with $\mu_0 = \nu_0$, and $\Upsilon_0 = \frac{\sqrt{3}\nu_0^2}{4}$.*

Lemma 7.4.11 (Protection and point perturbation) *Suppose that $P \subset \mathbb{R}^m$ and $\sigma \in \text{Del}(P)$ is secure. If $\zeta : P \rightarrow \tilde{P}$ is a ρ -perturbation with*

$$\rho \leq \frac{\Upsilon_0\mu_0}{18}\delta,$$

then $\zeta(\sigma) = \tilde{\sigma} \in \text{Del}(\tilde{P})$ and has a $(\delta - \frac{18}{\Upsilon_0\mu_0}\rho)$ -protected Delaunay ball.

Proof Let $B = B_{\mathbb{R}^m}(c, r)$ be the δ -protected Delaunay ball for $\sigma \in \text{Del}(P)$, and let $\tilde{B} = B_{\mathbb{R}^m}(\tilde{c}, \tilde{r})$ be the circumball for the corresponding perturbed simplex $\tilde{\sigma}$. We wish to establish a bound on ρ that will ensure that $\tilde{B}$ is protected with respect to $\tilde{P}$.

Let $q \in P$ be a point not in σ . We need to ensure that the corresponding $\tilde{q}$ lies outside the closure of $\tilde{B}$, i.e., that $d_{\mathbb{R}^m}(\tilde{q}, \tilde{c}) > \tilde{r}$.

Since $\delta \leq \epsilon$, the hypothesis of Lemma 7.4.8 is satisfied by ρ , and we have $d_{\mathbb{R}^m}(\tilde{c}, c) < \eta\rho$, where $\eta = \frac{8}{\Upsilon_0\mu_0}$. Thus for $p \in \sigma$ and corresponding $\tilde{p} \in \tilde{\sigma}$ we have

$$\begin{aligned} \tilde{r} &\leq d_{\mathbb{R}^m}(c, p) + d_{\mathbb{R}^m}(c, \tilde{c}) + d_{\mathbb{R}^m}(p, \tilde{p}) \\ &< r + (\eta + 1)\rho. \end{aligned}$$

Also

$$\begin{aligned} d_{\mathbb{R}^m}(\tilde{q}, \tilde{c}) &\geq d_{\mathbb{R}^m}(q, c) - d_{\mathbb{R}^m}(\tilde{c}, c) - d_{\mathbb{R}^m}(\tilde{q}, q) \\ &> r + \delta - \rho(\eta + 1). \end{aligned}$$

Therefore $\tilde{q}$ will be outside of the closure of $\tilde{B}$ provided $r + \delta - \rho(\eta + 1) \geq r + (1 + \eta)\rho$, i.e., when $\delta \geq 2(\eta + 1)\rho$. The result follows from the definition of η and the observation that μ_0 and Υ_0 are each no larger than one. $\square$

A similar argument yields a bound on the metric perturbation that will ensure the Delaunay balls for the m -simplices remain protected:

Lemma 7.4.12 (Protection and metric perturbation) *Suppose $U \subset \mathbb{R}^m$ contains $\text{conv}(P)$ and $d : U \times U \rightarrow \mathbb{R}$ is a metric such that $|d_{\mathbb{R}^m}(x, y) - d(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\sigma \in \text{Del}(P)$ is secure. If*

$$\rho \leq \frac{\Upsilon_0\mu_0}{20}\delta,$$

and $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$, then σ also belongs to $\text{Del}_d(P)$, and has a $(\delta - \frac{20}{\Upsilon_0\mu_0}\rho)$ -protected Delaunay ball in the metric d .

Proof Let $B = B_{\mathbb{R}^m}(c, r)$ be the Euclidean δ -protected Delaunay ball for $\sigma \in \text{Del}(P)$, and let $\tilde{B} = B_{\mathbb{R}^m}(\tilde{c}, \tilde{r})$ be a circumball for σ in the metric d . We wish to establish a bound on ρ that will ensure that $\tilde{B}$ is protected with respect to d .

Let $q \in P$ be a point not in σ . We need to ensure that $d(q, \tilde{c}) > \tilde{r}$. Since $\delta \leq \epsilon$, the hypothesis ensures that $\rho \leq \frac{\Upsilon_0 \lambda}{8}$, and so Lemma 7.4.4 yields $d_{\mathbb{R}^m}(\tilde{c}, c) < \eta\rho$, where $\eta = \frac{8}{\Upsilon_0 \mu_0}$. Thus for $p \in \sigma$

$$\begin{aligned} \tilde{r} &\leq d(c, p) + d(c, \tilde{c}) \\ &< (r + \rho) + (\eta\rho + \rho) \\ &= r + (\eta + 2)\rho, \end{aligned}$$

and

$$\begin{aligned} d(q, \tilde{c}) &\geq d(q, c) - d(\tilde{c}, c) \\ &> r + \delta - (\eta + 2)\rho. \end{aligned}$$

Thus $\tilde{q}$ will be outside of the closure of $\tilde{B}$ provided $r + \delta - (\eta + 2)\rho \geq r + (\eta + 2)\rho$, i.e., when

$$\delta \geq 2(\eta + 2)\rho.$$

The result follows from the definition of η and the observation that μ_0 and Υ_0 are each no larger than one. $\square$

7.4.3 Perturbations and Delaunay stability

The results of Section 7.4.2 translate into stability results for Delaunay triangulations. In the case of point perturbations in Euclidean space, the connectivity of the Delaunay triangulation cannot change as long as the simplices corresponding to the initial m -simplices remain protected. This is a direct consequence of Delaunay's original result [Del34], but we explicitly lay out the argument.

In the case of metric perturbation, we can no longer take for granted that the Delaunay complex cannot change its connectivity if the m -simplices remain protected. This is because we are no longer guaranteed that the Delaunay complex will be a triangulation. Using the consequences of the point-perturbation result, we establish bounds that ensure that the Delaunay complex in the perturbed metric will be the same as the original Delaunay triangulation.

7.4.3.1 Point perturbations

A consequence of Delaunay's triangulation result is that if a perturbation does not destroy any m -simplices in the Delaunay complex of a generic point set, then no new simplices are created either, and the complex is unchanged. More precisely we have:

Lemma 7.4.13 *Suppose $P \subset \mathbb{R}^m$ is a generic sample set, and $Q \subseteq P$ is a subset of interior points. If $\zeta : P \rightarrow \tilde{P}$ is a perturbation such that $\zeta(\text{star}(Q; \text{Del}(P))) \subseteq \text{star}(\zeta(Q); \text{Del}(\tilde{P}))$, and every m -simplex $\tilde{\sigma}^m \in \zeta(\text{star}(Q))$ is protected in $\text{Del}(\tilde{P})$, then $\zeta(\text{star}(Q)) = \text{star}(\zeta(Q))$.*

Proof Let $p \in Q$. By Lemma 7.3.5, $\text{star}(\zeta(p))$ is embedded, and by Lemma 7.3.8, $\text{Del}(P)$ is a triangulation at p . Since $\zeta : P \rightarrow \tilde{P}$ is injective, it follows that the simplicial map induced by ζ must be injective, and the result follows from Lemma 7.2.6. $\square$

Lemma 7.4.11 establishes bounds on a ρ -perturbation $\zeta : P \rightarrow \tilde{P}$ which will guarantee that if $Q \subset P$, and the simplices in $\text{star}(Q)$ are secure, then $\zeta(\text{star}(Q)) \subseteq \text{Del}(\tilde{P})$. Lemma 7.4.11 also guarantees that, if ρ is small enough, the m -simplices in $\text{star}(\zeta(Q); \text{Del}(\tilde{P}))$ will be protected. Thus if Q consists only of interior points of P , Lemma 7.4.13 applies. We have the following stability theorem for protected Delaunay triangulations:

Theorem 7.4.14 (Stability under point perturbation) *Suppose $P \subset \mathbb{R}^m$ and $Q \subseteq P$ is a subset of interior points such that every m -simplex in $\text{star}(Q)$ is secure. If $\zeta : P \rightarrow \tilde{P}$ is a ρ -perturbation, with*

$$\rho \leq \frac{\Upsilon_0 \mu_0}{18} \delta$$

then

$$\text{star}(Q; \text{Del}(P)) \stackrel{\zeta}{\cong} \text{star}(\zeta(Q); \text{Del}(\tilde{P})).$$

The ρ -relaxed Delaunay complex for P was defined by de Silva [dS08] by the criterion that $\sigma \in \text{Del}^\rho(P)$ if and only if there is a ball $B = B_{\mathbb{R}^m}(c, r)$ such that $\sigma \subset \bar{B}$, and $d_{\mathbb{R}^m}(c, q) \geq r - \rho$ for all $q \in P$. Thus the simplices in $\text{Del}^\rho(P)$ all have “almost empty”, balls centred on a ρ -centre for σ . We have the following consequence of Theorem 7.4.14:

Corollary 7.4.15 (Stability under relaxation) *Suppose $P \subset \mathbb{R}^m$ and $Q \subseteq P$ is a set of interior points such that every m -simplex in $\text{star}(Q)$ is secure. If*

$$\rho \leq \frac{\Upsilon_0 \mu_0}{18} \delta,$$

then

$$\text{star}(Q; \text{Del}^\rho(P)) = \text{star}(Q; \text{Del}(P)).$$

Proof Suppose that $\sigma \in \text{star}(Q; \text{Del}^\rho(P))$. Then there is a ball B enclosing σ such that any point $q \in B$ is within a distance ρ from ∂B . Project all such points radially out to ∂B . Then we have a ρ -perturbation $\zeta : P \rightarrow \tilde{P}$, and σ has become $\tilde{\sigma} \in \text{star}(\zeta(Q); \text{Del}(\tilde{P}))$. By Theorem 7.4.14, $\text{star}(\zeta(Q); \text{Del}(\tilde{P})) \stackrel{\zeta}{\cong} \text{star}(Q; \text{Del}(P))$, and therefore $\sigma \in \text{star}(Q; \text{Del}(P))$. $\square$

7.4.3.2 Metric perturbation

For a perturbation of the metric, we can exploit the stability results obtained for perturbations of points in the Euclidean metric to ensure that no simplices can appear in $\text{star}(Q; \text{Del}_d(P))$ that do not already exist in $\text{star}(Q; \text{Del}(P))$.

Lemma 7.4.16 *Suppose $\text{conv}(P) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $Q \subseteq P$ is a set of interior points such that every m -simplex $\sigma \in \text{star}(Q)$ is secure and satisfies $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$. If*

$$\rho \leq \frac{\Upsilon_0 \mu_0}{36} \delta,$$

then

$$\text{star}(Q; \text{Del}_d(P)) \subseteq \text{star}(Q; \text{Del}(P)).$$

Proof Let $B_d(c, r)$ be a Delaunay ball for simplex $\sigma \in \text{star}(Q; \text{Del}_d(P))$. Then $d(c, p) \leq d(c, q)$ for all $p \in \sigma$, and $q \in P$. By the hypothesis on d , this implies that $d_{\mathbb{R}^m}(c, p) \leq d_{\mathbb{R}^m}(c, q) + 2\rho$ for all $p \in \sigma$ and $q \in P$, and therefore $\sigma \in \text{Del}^{2\rho}(P)$. The result now follows from Corollary 7.4.15. $\square$

The perturbation bounds required by Lemma 7.4.16, also satisfy the requirements of Lemma 7.4.12. This gives us the reverse inclusion, and thus we can quantify the stability under metric perturbation for subcomplexes of secure simplices in Delaunay triangulations:

Theorem 7.4.17 (Stability under metric perturbation) *Suppose $\text{conv}(P) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $Q \subseteq P$ is a set of interior points such that every m -simplex $\sigma \in \text{star}(Q)$ is secure and satisfies $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$. If*

$$\rho \leq \frac{\Upsilon_0 \mu_0}{36} \delta,$$

then

$$\text{star}(Q; \text{Del}_d(P)) = \text{star}(Q; \text{Del}(P)).$$

Using Lemma 7.4.10, and recognizing that the safe interior simplices also satisfy the distance from boundary requirement of Theorem 8.3.2, we can restate this metric perturbation stability result for Delaunay triangulations on δ -generic point sets:

Corollary 7.4.18 (Stability under metric perturbation) *Suppose P is δ -generic for P_J , with sampling radius ϵ and $\delta = \nu_0 \epsilon$. Suppose also that $\text{conv}(P) \subseteq U$, and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. If*

$$\rho \leq \frac{\nu_0^3}{84} \delta = \frac{\nu_0^4}{84} \epsilon,$$

then

$$\text{star}(P_J; \text{Del}_d(P)) = \text{star}(P_J; \text{Del}(P)).$$

7.5 Summary

We have quantified the close relationship between the genericity of a point set, the quality of the simplices in the Delaunay complex, and the stability of the Delaunay complex under perturbation. The problem of poorly shaped simplices in a higher dimensional Delaunay complex can be seen as a manifestation of point sets that are close to being degenerate. The introduction of thickness as a geometric quality measure for simplices facilitated the stability calculations, which develop around a consideration of the circumcentres of a simplex in the presence of a perturbation.

We considered a point set $P \subset \mathbb{R}^m$ meeting a sampling radius ϵ and showed a constant bound on the thickness of the Delaunay simplices provided P is δ -generic with $\delta = \nu_0 \epsilon$ for some constant ν_0 . The question then arises: What is the least upper bound on the feasible ν_0 as a function of the dimension m ?

An important aspect of the current work is that the triangulation results are localised. Since a manifold can be locally well approximated by Euclidean space, the objective is to

fit together local Euclidean Delaunay patches where the Euclidean metric varies slightly between patches. This is where the stability of the Delaunay patches is important. In this setting we can also accommodate small variations in the sampling radius between neighbouring patches. Thus the algorithm will be able to triangulate sample sets whose sampling radius is defined by a Lipschitz sizing function.

Using stability results from this chapter, we develop a refinement algorithm that produces an intrinsic Delaunay triangulations of submanifolds of Euclidean space in Chapter 7. This partially recovers Leibon and Letscher's [LL00] result for the case of submanifolds of Euclidean space. In future work we want to extend this to an algorithm for triangulating manifolds that will exploit only the local intrinsic metric properties of the manifold, with no requirement that it be embedded in an ambient space.

Chapter 8

Constructing intrinsic Delaunay triangulations

We describe an algorithm to construct an intrinsic Delaunay triangulation of a smooth closed submanifold of Euclidean space. Using structural results established in the previous chapter (Chapter 7) on the stability of Delaunay triangulations on δ -generic point sets, we establish sampling criteria which ensure that the intrinsic Delaunay complex coincides with the restricted Delaunay complex and also with the recently introduced tangential Delaunay complex. The algorithm generates a point set that meets the required criteria while the tangential complex is being constructed. In this way the computation of geodesic distances is avoided, the runtime is only linearly dependent on the ambient dimension, and the Delaunay complexes are guaranteed to be triangulations of the manifold.

8.1 Introduction

This chapter addresses the problem of constructing an intrinsic Delaunay triangulation of a smooth closed submanifold $\mathcal{M} \subset \mathbb{R}^N$. We present an algorithm which generates a point set $\mathcal{P} \subset \mathcal{M}$ and a simplicial complex on $\mathcal{P}$ that is homeomorphic to $\mathcal{M}$ and has a connectivity determined by the Delaunay triangulation of $\mathcal{P}$ with respect to the intrinsic metric of $\mathcal{M}$.

For a submanifold of Euclidean space, the restricted Delaunay complex [ES97], which is defined by the ambient metric restricted to the submanifold, was employed by Cheng et al. [CDR05b] as the basis for a triangulation. However, it was found that sampling density alone was insufficient to ensure a triangulation, and manipulations of the complex were employed.

In an earlier work, Leibon and Letscher [LL00] announced sampling density conditions which would ensure that the Delaunay complex defined by the intrinsic metric of the manifold was a triangulation. In fact, as shown in Chapter 6, the stated result is incorrect: sampling density alone is insufficient to guarantee an intrinsic Delaunay triangulation (see Theorem 6.3.3). *Topological defects* can arise when the vertices lie too close to a degenerate or “quasi-cospherical” configuration.

Our interest in the intrinsic Delaunay complex stems from its close relationship with other Delaunay-like structures that have been proposed in the context of non-homogeneous metrics. For example, anisotropic Voronoi diagrams [LS03] and anisotropic

Delaunay triangulations emerge as natural structures when we want to mesh a domain of $\mathbb{R}^m$ while respecting a given metric tensor field.

This chapter builds over preliminary results on anisotropic Delaunay meshes [BWY11], and manifold reconstruction (and meshing) using the tangential Delaunay complex (Chapters 3 and 5). The central idea in both cases is to define Euclidean Delaunay triangulations locally and to glue these local triangulations together by removing *inconsistencies* between them. We view the inconsistencies as arising from instability in the Delaunay triangulations, and exploit the structural results from Chapter 7 to define sampling conditions under which these inconsistencies cannot arise.

The algorithm is based on the tangential Delaunay complex [BG11], and is an adaptation of a Delaunay refinement algorithm given in Chapter 5 to avoid poorly shaped simplices called “flakes”. The tangential Delaunay complex is defined with respect to local Delaunay triangulations restricted to the tangent spaces at sample points. We demonstrate that the algorithm produces sampling conditions such that the tangential Delaunay complex coincides with the restricted Delaunay complex and the intrinsic Delaunay complex. Like the algorithm in Chapter 5, the refinement algorithm avoids the problem of slivers without the need to resort to a point weighting strategy [CDE⁺00b, CDR05b] (and Chapter 3), which alters the definition of the restricted Delaunay complex.

Organization of the chapter. We present background and foundational material in Section 8.2. Then, in Section 8.3, we exploit results established in Chapter 7 to demonstrate sampling conditions under which the intrinsic Delaunay complex, the restricted Delaunay complex, and the tangential Delaunay complex coincide and are manifold. The algorithm itself is presented in Section 8.4, and the analysis of the algorithm is presented in Section 8.5. The missing proof of Lemma 8.5.8 is given in the Appendix C.

8.2 Background

8.2.1 Notations from Chapter 7

Notations used in this chapter will be exactly same as the ones given in Section 7.2 of Chapter 7, specially ones related to

1. General notations given in the beginning of Section 7.2.
2. Sampling parameters and perturbations (Section 7.2.1).
3. Simplices (Section 7.2.2).
4. Complexes (Section 7.2.3).

8.2.2 Simplex perturbation

We will make use of two results displaying the robustness of thick simplices with respect to small perturbations of their vertices. The first observation bounds the change in thickness itself under small perturbations:

Lemma 8.2.1 (Thickness under perturbation) *Let $\sigma = [p_0, \dots, p_j]$ and $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_j]$ be j -simplices such that $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in \{0, \dots, j\}$. For any positive $\eta \leq 1$, if*

$$\rho \leq \frac{(1-\eta)\Upsilon_\sigma^2 L_\sigma}{14}, \quad (8.1)$$

then

$$D_{\tilde{\sigma}}(\tilde{p}_i) \geq \eta D_\sigma(p_i),$$

for all $i \in \{0, \dots, j\}$. It follows that

$$\Upsilon_{\tilde{\sigma}} \Delta_{\tilde{\sigma}} \geq \eta \Upsilon_\sigma \Delta_\sigma,$$

and

$$\Upsilon_{\tilde{\sigma}} \geq \left(1 - \frac{2\rho}{\Delta_\sigma}\right) \eta \Upsilon_\sigma \geq \frac{6}{7} \eta \Upsilon_\sigma.$$

Proof Let $p, q \in \sigma$ with $\tilde{p}, \tilde{q}$ the corresponding vertices of $\tilde{\sigma}$. Let $v = p - q$ and $\tilde{v} = \tilde{p} - \tilde{q}$. Define $\theta = \angle(v, \text{aff}(\sigma_p))$ and $\tilde{\theta} = \angle(\tilde{v}, \text{aff}(\tilde{\sigma}_{\tilde{p}}))$. Since $\Upsilon_\sigma \leq \Upsilon_{\sigma_p}$, Whitney's Lemma 7.2.1 lets us bound $\angle(\text{aff}(\sigma_p), \text{aff}(\tilde{\sigma}_{\tilde{p}}))$ by the angle α defined by

$$\sin \alpha = \frac{2\rho}{\Upsilon_\sigma \Delta_\sigma}.$$

Also, by an elementary geometric argument,

$$\sin \gamma = \frac{2\rho}{\|\tilde{v}\|}$$

defines γ as an upper bound on the angle between the lines generated by v and $\tilde{v}$.

Thus we have

$$D_{\tilde{\sigma}}(\tilde{p}) = \|\tilde{v}\| \sin \tilde{\theta} \geq (\|\tilde{v}\| - 2\rho) \sin(\theta - \alpha - \gamma).$$

Using the addition formula for sine together with the facts that for $x, y \in [0, \frac{\pi}{2}]$, $(1-x) \leq \cos x$; $2 \sin x \geq x$; and $\sin x + \sin y \geq \sin(x+y)$, we get

$$D_{\tilde{\sigma}}(\tilde{p}) \geq (\|\tilde{v}\| - 2\rho) \left[\left(1 - 2 \left(\frac{2\rho}{\Upsilon_\sigma \Delta_\sigma} + \frac{2\rho}{\|\tilde{v}\|} \right) \right) \frac{D_\sigma(p)}{\|\tilde{v}\|} - \left(\frac{2\rho}{\Upsilon_\sigma \Delta_\sigma} + \frac{2\rho}{\|\tilde{v}\|} \right) \right].$$

For convenience, define $\mu = \frac{2\rho}{L_\sigma} \geq \frac{2\rho}{\|\tilde{v}\|} \geq \frac{2\rho}{\Delta_\sigma}$. Then

$$\begin{aligned} D_{\tilde{\sigma}}(\tilde{p}) &\geq \|\tilde{v}\| (1 - \mu) \left[\left(1 - 2 \left(1 + \frac{1}{\Upsilon_\sigma}\right) \mu \right) \frac{D_\sigma(p)}{\|\tilde{v}\|} - \left(1 + \frac{1}{\Upsilon_\sigma}\right) \mu \right] \\ &\geq (1 - \mu) \left[\left(1 - \frac{4\mu}{\Upsilon_\sigma}\right) D_\sigma(p) - \frac{2\mu \|\tilde{v}\|}{\Upsilon_\sigma} \right] \\ &\geq (1 - \mu) \left[\left(1 - \frac{4\mu}{\Upsilon_\sigma}\right) D_\sigma(p) - \frac{2\mu \|\tilde{v}\|}{\Upsilon_\sigma^2 \Delta_\sigma} D_\sigma(p) \right] \\ &\geq (1 - \mu) \left(1 - \frac{4\mu}{\Upsilon_\sigma} - \frac{2\mu}{\Upsilon_\sigma^2}\right) D_\sigma(p) \\ &\geq \left(1 - \frac{7\mu}{\Upsilon_\sigma^2}\right) D_\sigma(p) \\ &\geq K D_\sigma(p) \quad \text{when } \mu \leq \frac{(1-K)\Upsilon_\sigma^2}{7}. \end{aligned}$$




Figure 8.1: Diagram for the proof of Lemma 8.2.2.

The condition on μ is satisfied when ρ satisfies Inequality (8.1).

The bound on $\Upsilon_{\tilde{\sigma}}\Delta_{\tilde{\sigma}}$ follows immediately from the bounds on the $D_{\tilde{\sigma}}(\tilde{p})$, and the bound on $\Upsilon_{\tilde{\sigma}}$ itself follows from the observation that

$$\frac{\Delta_{\sigma}}{\Delta_{\tilde{\sigma}}} \geq \frac{\Delta_{\sigma}}{\Delta_{\sigma} + 2\rho} \geq \left(1 - \frac{2\rho}{\Delta_{\sigma}}\right) \geq \left(1 - \frac{\Upsilon_{\sigma}^2}{7}\right) \geq \frac{6}{7},$$

when ρ satisfies Inequality (8.1). □

We will also make use of a bound relating circumscribing balls of a simplex that undergoes a perturbation:

Lemma 8.2.2 (Circumscribing balls under perturbation) *Let $\sigma = [p_0, \dots, p_j]$ and $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_j]$ be j -simplices such that $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in \{0, \dots, j\}$. Suppose $B = B_{\mathbb{R}^m}(c, r)$, with $r < \epsilon$, is a circumscribing ball for σ . If*

$$\rho \leq \frac{\Upsilon_{\sigma}^2 L_{\sigma}}{28},$$

then there is a circumscribing ball $\tilde{B} = B_{\mathbb{R}^d}(\tilde{c}, \tilde{r})$ for $\tilde{\sigma}$ with

$$\|\tilde{c} - c\| < \left(\frac{8\epsilon}{\Upsilon_{\sigma}\Delta_{\sigma}}\right)\rho \tag{8.2}$$

and

$$|\tilde{r} - r| < \left(\frac{9\epsilon}{\Upsilon_{\sigma}\Delta_{\sigma}}\right)\rho.$$

If, in addition, we have that $\tilde{p}_0 = p_0$, then $|\tilde{r} - r| \leq \|\tilde{c} - c\|$, and Eq. (8.2) serves also as a bound on $|\tilde{r} - r|$.

Proof By the perturbation bounds, the distances between c and the vertices of $\tilde{\sigma}$ differ by no more than 2ρ . Also, $\|c - p_i\| < \tilde{\epsilon} = \epsilon + \rho$, and so by Lemma 7.4.1 we have

$$d_{\mathbb{R}^m}(c, N_{\tilde{\sigma}}) < \frac{2\tilde{\epsilon}\rho}{\Upsilon_{\tilde{\sigma}}\Delta_{\tilde{\sigma}}}.$$

See Figure 8.2.2. The bound on ρ allows us to apply Lemma 8.2.1 with $K = \frac{1}{2}$, so $\Upsilon_{\tilde{\sigma}}\Delta_{\tilde{\sigma}} \geq \frac{1}{2}\Upsilon_{\sigma}\Delta_{\sigma}$, and we obtain the bound on $\|\tilde{c} - c\|$ with the observation that $\tilde{\epsilon} \leq 2\epsilon$. Indeed, $\rho \leq \epsilon$ because $L_{\sigma} \leq 2\epsilon$.

By the triangle inequality $|\tilde{r} - r| \leq \|\tilde{p}_0 - p_0\| + \|\tilde{c} - c\|$, and the stated bound on $|\tilde{r} - r|$ follows from the observation that $\frac{\epsilon}{\Upsilon_{\sigma}\Delta_{\sigma}} \geq 1$ if $j > 1$. Under the assumption that $\tilde{p}_0 = p_0$, the bound on $\|\tilde{c} - c\|$ also serves as a bound on $|\tilde{r} - r|$. $\square$

8.2.3 Flakes

For algorithmic reasons, it is convenient to have a more structured constraint on simplex geometry than that provided by a simple thickness bound. A simplex that is not thick has a relatively small altitude, but we wish to exploit a family of bad simplices for which *all* the altitudes are relatively small. As shown by Lemma 8.2.6 below, the Γ_0 -flakes form such a family. The flake parameter Γ_0 is a positive real number smaller than one.

Definition 8.2.3 (Γ_0 -good simplices and Γ_0 -flakes) *A simplex σ is Γ_0 -good if $\Upsilon_{\sigma^j} \geq \Gamma_0^j$ for all j -simplices $\sigma^j \leq \sigma$. A simplex is Γ_0 -bad if it is not Γ_0 -good. A Γ_0 -flake is a Γ_0 -bad simplex in which all the proper faces are Γ_0 -good.*

Observe that a flake must have dimension at least 2, since $\Upsilon_{\sigma^j} = 1$ for $j < 2$. Note that definition of Γ_0 -flake is very close in spirit with Θ_0 -sliver introduced in Chapter 2. We introduce this new definition as it will be easier to work rather than slivers.

Ensuring that all simplices are Γ_0 -good is the same as ensuring that there are no flakes. Indeed, if σ is Γ_0 -bad, then it has a j -face $\sigma^j \leq \sigma$ that is not Γ_0^j -thick. By considering such a face with minimal dimension we arrive at the following important observation:

Lemma 8.2.4 *A simplex is Γ_0 -bad if and only if it has a face that is a Γ_0 -flake.*

We obtain an upper bound on the altitudes of a Γ_0 -flake through a consideration of dihedral angles. In particular, we observe the following general relationship between simplex altitudes:

Lemma 8.2.5 *If σ is a j -simplex with $j \geq 2$, then for any two vertices $p, q \in \sigma$, the dihedral angle between σ_p and σ_q defines an equality between ratios of altitudes:*

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

Proof Let $\sigma_{pq} = \sigma_p \cap \sigma_q$, and let p_* be the projection of p into $\text{aff}(\sigma_{pq})$. Taking p_* as the origin, we see that $\frac{p-p_*}{D_{\sigma_q}(p)}$ has the maximal distance to $\text{aff}(\sigma_p)$ out of all the unit vectors in $\text{aff}(\sigma_q)$, and this distance is $\frac{D_{\sigma}(p)}{D_{\sigma_q}(p)}$. By definition this is the sine of the angle between $\text{aff}(\sigma_p)$ and $\text{aff}(\sigma_q)$. A symmetric argument is carried out with q to obtain the result. $\square$

We arrive at the following important observation about flake simplices:

Lemma 8.2.6 (Flakes have small altitudes) *If a k -simplex σ is a Γ_0 -flake, then for every vertex $p \in \sigma$, the altitude satisfies the bound*

$$D_{\sigma}(p) < \frac{k\Delta_{\sigma}^2\Gamma_0}{(k-1)L_{\sigma}}.$$

Proof Recalling Lemma 8.2.5 we have

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

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

$$D_\sigma(q) = k\Upsilon_\sigma \Delta_\sigma < k\Gamma_0^k \Delta_\sigma,$$

and

$$D_{\sigma_p}(q) \geq (k-1)\Upsilon_{\sigma_p} \Delta_{\sigma_p} \geq (k-1)\Gamma_0^{k-1} L_\sigma,$$

and

$$D_{\sigma_q}(p) \leq \Delta_{\sigma_q} \leq \Delta_\sigma,$$

and the bound is obtained. $\square$

From Lemma 2.3.1 (2), we can have similar bound on the altitudes for Θ_0 -slivers.

8.2.4 The Delaunay complex and protection

We will now recall the definition of Delaunay complex in terms of *empty balls*. An empty ball is one that contains no point from P .

Definition 8.2.7 (Delaunay complex) A Delaunay ball is a maximal empty ball. Specifically, $B = B_{\mathbb{R}^m}(x, r)$ is a Delaunay ball if any empty ball centred at x is contained in B . A simplex σ is a Delaunay simplex, if there exists some Delaunay ball B such that the vertices of σ belong to $\partial B \cap P$. The Delaunay complex is the set of Delaunay simplices, and is denoted $\text{Del}(P)$.

The Delaunay complex has the combinatorial structure of an abstract simplicial complex, but $\text{Del}(P)$ is embedded only when P satisfies appropriate genericity requirements, see Chapter 7.

A Delaunay simplex σ is δ -protected if it has a Delaunay ball B such that $d_{\mathbb{R}^m}(q, \partial B) > \delta$ for all $q \in P \setminus \sigma$. We say that B is a δ -protected Delaunay ball for σ . If $\tau < \sigma$, then B is also a Delaunay ball for τ , but it cannot be a δ -protected Delaunay ball for τ . We say that σ is *protected* to mean that it is δ -protected for some unspecified $\delta > 0$.

Definition 8.2.8 (δ -generic) A finite set of points $P \subset \mathbb{R}^m$ is δ -generic if all the Delaunay m -simplices are δ -protected. The set P is simply generic if it is δ -generic for some unspecified $\delta > 0$.

In Chapter 7 we have demonstrated that δ -generic point sets impart a *quantifiable stability* on the Delaunay complex. In Section 8.3 we review the main stability result and develop it to define the sampling conditions that will be met by the algorithm that we introduce in Section 8.4.

8.2.5 The Delaunay complex in other metrics

We will also consider the Delaunay complex defined with respect to a metric d on $\mathbb{R}^m$ which differs from the Euclidean one. Specifically, if $P \subset U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is a metric, then we define the Delaunay complex $\text{Del}_d(P)$ with respect to the metric d .

In this chapter we will deal with Delaunay complex defined with respect to metrics other than the standard Euclidean metric. For a given point set $P \subset U \subset \mathbb{R}^m$ and a metric $d : U \times U \rightarrow \mathbb{R}$ then we define Delaunay complex $\text{Del}_d(P)$ with respect to the metric d in the following way:

Definition 8.2.9 (Delaunay ball and Delaunay complex) *A Delaunay ball is a maximal empty ball $B_d(x, r)$ in the metric d . The resulting Delaunay complex $\text{Del}_d(P)$ consists of all the simplices which are circumscribed by a Delaunay ball with respect to the metric d .*

The simplices of $\text{Del}_d(P)$ are, possibly degenerate, geometric simplices in $\mathbb{R}^m$. As for $\text{Del}(P)$, $\text{Del}_d(P)$ has the combinatorial structure of an abstract simplicial complex, but unlike $\text{Del}(P)$, $\text{Del}_d(P)$ may fail to be embedded even when there are no degenerate simplices.

In this chapter we will use the following notations for standard Delaunay complexes.

1. For a point set $P \subset \mathbb{R}^k$, when we want to stress that the Delaunay complex we are considering is with respect to the standard Euclidean metric $d_{\mathbb{R}^k}$ in $\mathbb{R}^k$, we will use the notation $\text{Del}_{\mathbb{R}^k}(P)$. For a object X , Euclidean Delaunay complex restricted to X will be denoted by $\text{Del}_{X|\mathbb{R}^k}(P)$.
2. For a point set $P \subset \mathcal{M}$, $\text{Del}_{\mathcal{M}}(P)$ denotes the Delaunay complex with respect to the intrinsic metric $d_{\mathcal{M}}$ of $\mathcal{M}$.

8.2.6 The Voronoi diagram

We will occasionally make reference to the Voronoi diagram, which is a structure dual to the Delaunay complex. It offers an alternative way to interpret observations made with respect to the Delaunay complex.

The *Voronoi cell* associated with $p \in P$ with respect to the metric $d : U \times U \rightarrow \mathbb{R}$ is given by

$$\text{Vor}_d(p) = \{x \in U \mid d(x, p) \leq d(x, q) \text{ for all } q \in P\}.$$

More generally, a *Voronoi face* is the intersection of a set of Voronoi cells: given $\{p_0, \dots, p_k\} \subset P$, let σ denote the corresponding abstract simplex. We define the associated Voronoi face as

$$\text{Vor}_d(\sigma) = \bigcap_{i=0}^k \text{Vor}_d(p_i).$$

It follows that σ is a Delaunay simplex if and only if $\text{Vor}_d(\sigma) \neq \emptyset$. In this case, every point in $\text{Vor}_d(\sigma)$ is the centre of a Delaunay ball for σ . Thus every Voronoi face corresponds to a Delaunay simplex. The Voronoi cells give a decomposition of U , denoted $\text{Vor}_d(P)$, called the *Voronoi diagram*. Our definition of the Delaunay complex of P corresponds to the nerve of the Voronoi diagram.

In this chapter we will use the following notations for standard Voronoi diagrams.

1. For a point set $P \subset \mathbb{R}^k$, when we want to stress that the Voronoi diagram we are considering is with respect to the standard Euclidean metric $d_{\mathbb{R}^k}$ in $\mathbb{R}^k$, we will use the notation $\text{Vor}_{\mathbb{R}^k}(P)$. For an object X , Euclidean Voronoi diagram restricted to X will be denoted by $\text{Vor}_{\mathbb{R}^k|_X}(P)$.
2. For a point set $P \subset \mathcal{M}$, $\text{Vor}_{\mathcal{M}}(P)$ denotes the Voronoi diagram with respect to the intrinsic metric $d_{\mathcal{M}}$ of $\mathcal{M}$.

8.2.7 Background results for manifolds

The tangent space at $p \in \mathcal{M}$ is denoted $T_p\mathcal{M}$, and we identify it with an m -flat in the ambient space. The normal space, $N_p\mathcal{M}$, is the orthogonal complement of $T_p\mathcal{M}$ in $T_p\mathbb{R}^N$, and we likewise treat it as the affine subspace of dimension $m-k$ orthogonal to $T_p\mathcal{M} \subset \mathbb{R}^N$.

A ball $B = B_{\mathbb{R}^N}(c, r)$ is a *medial ball* at p if $B \cap \mathcal{M} = \emptyset$, it is tangent to $\mathcal{M}$ at p , and it is maximal in the sense that any ball which contains B either coincides with B or intersects $\mathcal{M}$. The *local reach* at p is the infimum of the radii of the medial balls at p , and the *reach* of $\mathcal{M}$, denoted $\text{rch}(\mathcal{M})$, is the infimum of the local reach over all points of $\mathcal{M}$. In order to approximate the geometry and topology with a simplicial complex, manifolds with small reach require a higher sampling density than those with a larger reach. As is typical, an upper bound on our sampling radius will be proportional to $\text{rch}(\mathcal{M})$. Since $\mathcal{M} \subset \mathbb{R}^N$ is a smooth, compact embedded submanifold, it has positive reach.

An estimate of how the tangent space locally deviates from the manifold is given by an observation of Federer [Fed59, Theorem 4.8(7)] (see also Giesen and Wagner [GW04b, Lemma 6]):

Lemma 8.2.10 (Distance to tangent space) *If $x, y \in \mathcal{M} \subset \mathbb{R}^N$ and $d_{\mathbb{R}^N}(x, y) \leq r < \text{rch}(\mathcal{M})$, then*

$$d_{\mathbb{R}^N}(y, T_x\mathcal{M}) \leq \frac{r^2}{2\text{rch}(\mathcal{M})},$$

and thus

$$\sin \alpha \leq \frac{r}{2\text{rch}(\mathcal{M})},$$

where α is the angle between $[x, y]$ and $T_x\mathcal{M}$.

A complementary result bounds the distance to the manifold from a point on a tangent space [BG10b, Lemma 4.3]:

Lemma 8.2.11 (Distance to manifold) *Suppose $v \in T_x\mathcal{M}$ with $\|v - x\| = r \leq \frac{\text{rch}(\mathcal{M})}{4}$. Let $y = \psi_x(v) \in \mathcal{M}$, where ψ_x is the inverse projection (8.7). Then,*

$$d_{\mathbb{R}^N}(v, y) \leq \frac{2r^2}{\text{rch}(\mathcal{M})}.$$

The previous two lemmas lead to a convenient bound on the angle between nearby tangent spaces. We prove here a variation on previous results [NSW08c, Prop. 6.2] [BG11, Lemma 5.5]:

Lemma 8.2.12 (Tangent space variation) *Let $x, y \in \mathcal{M}$ be such that $d_{\mathbb{R}^N}(x, y) = r \leq \frac{\text{rch}(\mathcal{M})}{4}$, and let α be the angle between $T_x\mathcal{M}$ and $T_y\mathcal{M}$. Then,*

$$\sin \alpha < \frac{6r}{\text{rch}(\mathcal{M})}.$$

Proof Let $v \in T_y\mathcal{M} \subset \mathbb{R}^N$ with $\|v - y\| = r$. We will bound the angle between $v - y$ and $T_x\mathcal{M}$. We have

$$\begin{aligned} \sin \alpha &\leq \frac{1}{\|v - y\|} (d_{\mathbb{R}^N}(y, T_x\mathcal{M}) + d_{\mathbb{R}^N}(v, T_x\mathcal{M})) \\ &\leq \frac{1}{\|v - y\|} (d_{\mathbb{R}^N}(y, T_x\mathcal{M}) + d_{\mathbb{R}^N}(v, \hat{v}) + d_{\mathbb{R}^N}(\hat{v}, T_x\mathcal{M})), \end{aligned} \quad (8.3)$$

where $\hat{v} \in \mathcal{M}$ is the closest point to v in $\mathcal{M}$.

By Lemma 8.2.10, we have $d_{\mathbb{R}^N}(y, T_x\mathcal{M}) \leq \frac{r^2}{2\text{rch}(\mathcal{M})}$, and by Lemma 8.2.11 we get $d_{\mathbb{R}^N}(v, \hat{v}) \leq \frac{2r^2}{\text{rch}(\mathcal{M})}$. For the third term in Eq. (8.3), we find

$$\begin{aligned} d_{\mathbb{R}^N}(x, \hat{v}) &\leq d_{\mathbb{R}^N}(x, y) + \|v - y\| + d_{\mathbb{R}^N}(v, \hat{v}) \\ &\leq 2r + \frac{2r^2}{\text{rch}(\mathcal{M})} \leq \frac{5r}{2} < \text{rch}(\mathcal{M}), \end{aligned}$$

and so we may apply Lemma 8.2.10 to obtain $d_{\mathbb{R}^N}(\hat{v}, T_x\mathcal{M}) \leq \frac{25r^2}{8\text{rch}(\mathcal{M})}$.

Putting these observations back into Eq. (8.3) we find

$$\begin{aligned} \sin \alpha &\leq \frac{1}{\|v - y\|} \left(\frac{r^2}{2\text{rch}(\mathcal{M})} + \frac{2r^2}{\text{rch}(\mathcal{M})} + \frac{25r^2}{8\text{rch}(\mathcal{M})} \right) \\ &= \frac{45r}{8\text{rch}(\mathcal{M})} < \frac{6r}{\text{rch}(\mathcal{M})}. \end{aligned}$$

□

The following observation is a direct consequence of results established by Niyogi et al. [NSW08c, Lemma 5.4]:

Lemma 8.2.13 *Let $W = B_{\mathbb{R}^N|_{\mathcal{M}}}(p, r)$, for some $p \in \mathcal{M}$ and $r < \text{rch}(\mathcal{M})/2$. When restricted to W , the orthogonal projection $\pi_p|_W : W \rightarrow T_p\mathcal{M}$ is a diffeomorphism onto its image.*

Proof Let $f = \pi_p|_W$. Niyogi et al. showed [NSW08c, Lemma 5.4] that the Jacobian of f is nonsingular on W , so that W is a covering space for $U = f(W) \subset T_p\mathcal{M}$. The Morse-theory argument of Boissonnat et al. [BFC01, Proposition 12] can be applied to demonstrate that W is a topological ball. It follows that U is connected, since any path in W projects to a path in U . Thus W must be a single-sheeted cover of U , since $f^{-1}(0) = \{p\}$. Indeed, if $q \in W$ with $q \neq p$ and $f(q) = 0$, then $[p, q]$ would be perpendicular to $T_p\mathcal{M}$, contradicting Lemma 8.2.10. Thus $f : W \rightarrow U$ is a diffeomorphism. □

Direct consequence of Lemma B.2.1, from Appendix B.2, is the following result:

Lemma 8.2.14 (Geodesic distance bound) *Let $x, y \in \mathcal{M}$ be such that $d_{\mathbb{R}^N}(x, y) \leq \frac{\text{rch}(\mathcal{M})}{2}$. Then*

$$d_{\mathcal{M}}(x, y) \leq d_{\mathbb{R}^N}(x, y) \left(1 + \frac{2d_{\mathbb{R}^N}(x, y)}{\text{rch}(\mathcal{M})} \right).$$

8.3 Equating Delaunay structures

We now turn to the task of triangulating $\mathcal{M}$, a smooth, compact m -manifold, without boundaries embedded in $\mathbb{R}^N$. In this section we demonstrate our main structural result, Theorem 8.3.6, which is stated at the end of Section 8.3.1. It says that the complex constructed by the algorithm we describe in Section 8.4 is in fact an intrinsic Delaunay triangulation of the manifold, which we introduce next.

8.3.1 Delaunay structures on manifolds

The *restricted Delaunay complex* is the Delaunay complex $\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$ obtained when distances on the manifold are measured with the metric $d_{\mathbb{R}^N|_{\mathcal{M}}}$. This is the Euclidean metric of the ambient space, restricted to the submanifold $\mathcal{M}$. In other words, $d_{\mathbb{R}^N|_{\mathcal{M}}}(x, y) = d_{\mathbb{R}^N}(x, y)$. We use this notation to avoid ambiguities in conjunction with the local Euclidean metrics discussed below. The Delaunay complex $\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$ is a substructure of $\text{Del}_{\mathbb{R}^N}(\mathcal{P})$.

Alternatively, distances on the manifold may be measured with $d_{\mathcal{M}}$, the *intrinsic metric* of the manifold. This metric defines the distance between x and y as the infimum of the lengths of the paths on $\mathcal{M}$ which connect x and y . Since the length of a path on $\mathcal{M}$ is defined as its length as a curve in $\mathbb{R}^N$, this metric is also induced from $d_{\mathbb{R}^N}$. The *intrinsic Delaunay complex* is the Delaunay structure $\text{Del}_{\mathcal{M}}(\mathcal{P})$ associated with this metric.

Although neither of these metrics are Euclidean, the idea is that locally, in a small neighbourhood of any point, these metrics may be well approximated by $d_{\mathbb{R}^m}$. Then, if the sampling satisfies appropriate δ -generic and ϵ -dense criteria in these local Euclidean metrics, the global Delaunay complex in the metric of the manifold will coincide locally with a Euclidean Delaunay triangulation, and we can thus guarantee a manifold complex.

8.3.1.1 Local Euclidean metrics




Figure 8.2: A *local coordinate chart* at a point p in $\mathcal{M}$

A *local coordinate chart* at a point $p \in \mathcal{M}$, is a pair (W, ϕ_p) , where $W \subset \mathcal{M}$ is an open neighbourhood of p , and $\phi_p : W \rightarrow U = \phi_p(W) \subset \mathbb{R}^m$ is a homeomorphism onto its image, with $\phi_p(p) = 0$. A local coordinate chart allows us to *pull back* the Euclidean metric to W . For all $x, y \in W$, the metric $d_{\phi_p}(x, y) = d_{\mathbb{R}^m}(\phi_p(x), \phi_p(y))$ is a *local Euclidean metric* for p on W . This metric depends upon the choice of ϕ_p ; there are different ways to impose a Euclidean metric on W .

It is convenient to take the reciprocal point of view, and work with a *local parameterization* at a point $p \in \mathcal{M}$. This is a pair (U, ψ_p) , such that $U \subset \mathbb{R}^m$, and (W, ψ_p^{-1}) is a local coordinate chart for p , where $W = \psi_p(U)$. We can then use ψ_p to pull back the metric of the manifold to U , and to simplify the notation we write $d_{\mathcal{M}}(x, y)$ for $x, y \in U$, where it is to be understood that this means $d_{\mathcal{M}}(\psi_p(x), \psi_p(y))$, and likewise for $d_{\mathbb{R}^m|_{\mathcal{M}}}(x, y)$. Indeed, once W and U have been coupled together by a homeomorphism, we can transfer the metrics between them and the distinction becomes only one of perspective; the standard metric $d_{\mathbb{R}^m}$ on U is a local Euclidean metric for p .

We wish to generate a sample set $\mathcal{P} \subset \mathcal{M}$ that will allow us to exploit the stability results for Delaunay triangulations from Chapter 7. We consider the stability of a Delaunay triangulation in a local Euclidean metric. The following definition is convenient when stating the stability results:

Definition 8.3.1 (Secure simplex) *A simplex $\sigma \in \text{Del}(\mathcal{P})$ is secure if it is a δ -protected m -simplex that is Υ_0 -thick and satisfies $R_\sigma < \epsilon$ and $L_\sigma \geq \mu_0 \epsilon$.*

We will make reference to the following result from the previous chapter:

Theorem 8.3.2 (Metric stability assuming thickness) *Suppose $\text{conv}(\mathcal{P}) \subseteq U \subset \mathbb{R}^m$ and the metric $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\mathcal{P}_J \subseteq \mathcal{P}$ is such that every m -simplex $\sigma \in \text{star}(\mathcal{P}_J; \text{Del}(\mathcal{P}))$ is secure and satisfies $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$. If*

$$\rho \leq \frac{\Upsilon_0 \mu_0}{36} \delta,$$

then

$$\text{star}(\mathcal{P}_J; \text{Del}_d(\mathcal{P})) = \text{star}(\mathcal{P}_J; \text{Del}(\mathcal{P})).$$

In our context the point set $\mathcal{P}$ used in Theorem 8.3.2 will come from a larger point set $\mathcal{P}$, such that $\mathcal{P} = W \cap \mathcal{P}$. We will write $\mathcal{P}_W$ in order to emphasise this dependence on W . We want to ensure that

$$\text{star}(\mathcal{P}_J; \text{Del}_d(\mathcal{P}_W)) = \text{star}(\mathcal{P}_J; \text{Del}_d(\mathcal{P})). \quad (8.4)$$

This requirement is attained by demanding that $\mathcal{P}$ satisfy a sampling radius of ϵ with respect to the metric $d_{\mathcal{M}}$. Since $d_{\mathbb{R}^m}(x, y) \leq d_{\mathcal{M}}(x, y)$ for all $x, y \in U \cong W$, by our particular choice of ψ_p , we will have that $\mathcal{P}_W$ is an ϵ -sample set with respect to the metric $d_{\mathbb{R}^m}$. We ensure that U is large enough so that $d_{\mathbb{R}^m}(p, \partial U) \geq 4\epsilon$ for all $p \in \mathcal{P}_J$. It then follows that $R_\sigma < \epsilon$ for any simplex $\sigma \in \text{star}(\mathcal{P}_J; \text{Del}(\mathcal{P}_W))$, because $\mathcal{P}_W$ is an ϵ -sample set, Lemma 7.3.9, and thus $d_{\mathbb{R}^m}(q, \partial U) \geq 2\epsilon$ for any $q \in \sigma$. It follows that $d_{\mathcal{M}}(q, \partial U) \geq 2\epsilon$ as well, and thus the sampling radius on $\mathcal{P}$ ensures that Eq. (8.4) is satisfied. For our purposes $\mathcal{P}_J$ will consist of a single point p , and the sampling radius ϵ is constrained by the requirement that U be small enough that the metric distortion introduced by ψ_p meets the requirements of Theorem 8.3.2.

8.3.1.2 The tangential Delaunay complex

The algorithm we describe in Section 8.4 is a variation of the algorithm described by Boissonnat and Ghosh [BG10b]. This algorithm builds the *tangential Delaunay complex*, which we denote by $\text{Del}_{\mathcal{T}\mathcal{M}}(\mathcal{P})$. This is not a Delaunay complex as we have defined them, since it cannot be defined by the Delaunay empty ball criteria with respect to

any single metric. However, it is a Delaunay-type structure, and as with $\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$, the tangential Delaunay complex is a substructure of $\text{Del}_{\mathbb{R}^N}\mathcal{P}$. We will demonstrate sampling conditions which ensure that $\text{Del}_{T\mathcal{M}}(\mathcal{P}) = \text{Del}_{\mathcal{M}}(\mathcal{P}) = \text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$.

Definition 8.3.3 (Tangential Delaunay complex) *The tangential Delaunay complex for $\mathcal{P} \subset \mathcal{M} \subset \mathbb{R}^N$ is defined by the criterion that $\sigma \in \text{Del}_{T\mathcal{M}}(\mathcal{P})$ if it has an empty circumscribing ball $B_{\mathbb{R}^N}(c, r)$ such that $c \in T_p\mathcal{M}$ for some vertex $p \in \sigma$.*

We define some local complexes to facilitate discussions of the tangential Delaunay complex. For all $p \in \mathcal{P}$, let

$$K(p) = \{\sigma \mid \text{Vor}_{\mathbb{R}^N}(\sigma) \cap T_p\mathcal{M} \neq \emptyset\},$$

and define

$$\text{star}(p) = \text{star}(p; K(p)). \quad (8.5)$$

Then the tangential Delaunay complex is the union of the complexes $\text{star}(p)$ for all $p \in \mathcal{P}$.

Remark 8.3.4 (Connection between $K(p)$ and weighted Del. complex) *In Lemma 2.4.2, from Chapter 2, we showed that $\text{Vor}_{\mathbb{R}^N}(\mathcal{P}) \cap T_p\mathcal{M}$ is equal to the m -dimensional weighted Voronoi diagram of $\mathcal{P}' \subset T_p\mathcal{M}$, where $\mathcal{P}'$ is the orthogonal projection of $\mathcal{P}$ onto $T_p\mathcal{M}$ and the squared weight of a point $p'_i \in \mathcal{P}'$ is*

$$-\|p_i - p'_i\|^2 + \max_{p_j \in \mathcal{P}} \|p_j - p'_j\|^2.$$

Therefore, $K(p)$ is isomorphic to a dual complex (the nerve) of the k -dimensional weighted Voronoi diagram of $\mathcal{P}'$.

8.3.1.3 Power protection

The algorithm introduced in Section 8.4.2 will ensure that for every simplex σ in the tangential Delaunay complex, and every vertex $p \in \sigma$, there is a Delaunay ball for σ that is centred on $T_p\mathcal{M}$ and is protected in the following sense:

Definition 8.3.5 (Power protection) *A simplex σ with Delaunay ball $B_{\mathbb{R}^N}(C, R)$ is δ^2 -power-protected if $d_{\mathbb{R}^N}(C, q)^2 - R^2 > \delta^2$ for all $q \in \mathcal{P} \setminus \sigma$.*

Observe that, if $C \notin \mathcal{M}$, the ball $B_{\mathbb{R}^N}(C, R)$ is not an object that can be described by the metric $d_{\mathbb{R}^N|_{\mathcal{M}}}$. In the context of the tangential Delaunay complex we use power-protection rather than the protection described in Section 7.3.2 because working with squared distances is convenient when we consider the Delaunay complex restricted to an affine subspace.

8.3.1.4 Main structural result

The rest of Section 8.3 is devoted to the proof of Theorem 8.3.6 below. It says that for the point set generated by our algorithm, the tangential Delaunay complex is isomorphic with the intrinsic Delaunay complex of $\mathcal{M}$. It then follows, from Theorem 4.0.3, that the intrinsic Delaunay complex is in fact homeomorphic to $\mathcal{M}$; it is an intrinsic Delaunay triangulation.

Thus we obtain a partial recovery of the kind of results attempted by Leibon and Letscher [LL00]. Our sampling conditions, and our algorithm (existence proof) rely on the embedding of $\mathcal{M}$ in $\mathbb{R}^N$; we leave purely intrinsic sampling conditions for future work.

Theorem 8.3.6 (Intrinsic Delaunay triangulation) *Suppose $\mathcal{P} \subset \mathcal{M}$ is $(\tilde{\mu}_0\epsilon)$ -sparse with respect to $d_{\mathbb{R}^N}$, and every m -simplex $\tilde{\sigma} \in \text{Del}_{T\mathcal{M}}(\mathcal{P})$ is $\tilde{\Upsilon}_0$ -thick, and has, for every vertex $p \in \tilde{\sigma}$, a $\tilde{\delta}^2$ -power-protected empty ball of radius less than ϵ centred on $T_p\mathcal{M}$, with $\tilde{\delta} \geq \delta_0\tilde{\mu}_0\epsilon$. If $\delta_0^2\tilde{\mu}_0^2 \leq \frac{1}{7}$, and*

$$\epsilon \leq \frac{\tilde{\Upsilon}_0^2 \tilde{\mu}_0^3 \delta_0^2 \text{rch}(\mathcal{M})}{1.5 \times 10^6},$$

then

$$\text{Del}_{T\mathcal{M}}(\mathcal{P}) = \text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P}) = \text{Del}_{\mathcal{M}}(\mathcal{P}),$$

and for ϵ sufficiently small, these will be homeomorphic to $\mathcal{M}$:

$$|\text{Del}_{\mathcal{M}}(\mathcal{P})| \cong \mathcal{M}.$$

8.3.2 Choice of local Euclidean metric




Figure 8.3: Local parameterization at $p \in \mathcal{M}$ using the orthogonal projection map $\pi_p : \mathbb{R}^N \rightarrow T_p\mathcal{M}$.

A local parameterization at $p \in \mathcal{M}$ will be constructed with the aid of the orthogonal projection

$$\pi_p : \mathbb{R}^N \rightarrow T_p\mathcal{M}, \quad (8.6)$$

restricted to $\mathcal{M}$. As shown in Lemma 8.2.13, Niyogi et al. [NSW08c, Lemma 5.4] demonstrated that if $r < \frac{\text{rch}(\mathcal{M})}{2}$, then π_p is a diffeomorphism from $W = B_{\mathbb{R}^N|_{\mathcal{M}}}(p, r)$ onto its image $U \subset T_p\mathcal{M}$. We will identify $T_p\mathcal{M}$ with $\mathbb{R}^m$, and define the homeomorphism

$$\psi_p = \pi_p|_W^{-1} : U \longrightarrow W. \quad (8.7)$$

Using ψ_p to pull back the metrics $d_{\mathcal{M}}$ and $d_{\mathbb{R}^N|_{\mathcal{M}}}$ to $\mathbb{R}^m$, we can view them as perturbations of $d_{\mathbb{R}^m}$. The magnitude of the perturbation is governed by the radius of the ball used to define W .

Definition 8.3.7 We call a neighbourhood W of $p \in \mathcal{M}$ admissible if $W \subseteq B_{\mathbb{R}^N|_{\mathcal{M}}}(p, r)$, with $r \leq \frac{\text{rch}(\mathcal{M})}{100}$.

In all that follows, any mention of a local Euclidean metric refers to the one defined by π_p restricted to an admissible neighbourhood. The requirement $r \leq \frac{\text{rch}(\mathcal{M})}{100}$ is simply a convenient bound that yields a small integer constant in the perturbation bound of the following lemma, and does not constrain subsequent results. The bound could be relaxed to $r \leq \frac{\text{rch}(\mathcal{M})}{4}$ at the expense of a weaker bound on the perturbation.

Lemma 8.3.8 (Metric distortion) Suppose (U, ψ_p) is a local parameterisation at $p \in W \subset \mathcal{M}$ with $W = \psi_p(U)$. If $W \subseteq B_{\mathbb{R}^N|_{\mathcal{M}}}(p, r)$, with $r \leq \frac{\text{rch}(\mathcal{M})}{100}$, then for all $x, y \in U$,

$$|d_{\mathbb{R}^N|_{\mathcal{M}}}(x, y) - d_{\mathbb{R}^m}(x, y)| \leq |d_{\mathcal{M}}(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \frac{23r^2}{\text{rch}(\mathcal{M})}.$$

Proof Let $u, v \in W \subset B_{\mathbb{R}^N|_{\mathcal{M}}}(p, r)$, and let θ be the angle between the line segments $[u, v]$ and $[\pi_p(u), \pi_p(v)]$, θ_1 the angle between $[u, v]$ and $T_u\mathcal{M}$, and θ_2 the angle between $T_p\mathcal{M}$ and $T_u\mathcal{M}$. Thus $\theta \leq \theta_1 + \theta_2$, and $d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) = d_{\mathbb{R}^N}(u, v)\cos\theta$. Defining $\eta = \frac{r}{\text{rch}(\mathcal{M})}$, Lemma 8.2.14 yields

$$d_{\mathcal{M}}(u, v) \leq d_{\mathbb{R}^N}(u, v)(1 + 4\eta), \quad (8.8)$$

and so

$$d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) \geq \frac{d_{\mathcal{M}}(u, v) \cos\theta}{1 + 4\eta}.$$

Using Lemma 8.2.10, we find $\sin\theta_1 \leq \eta$, and Lemma 8.2.12, yields $\sin\theta_2 \leq 6\eta$. Therefore, since $\sin\theta \leq \sin\theta_1 + \sin\theta_2$, we have $\cos\theta = (1 - \sin^2\theta)^{1/2} \geq 1 - \sin\theta \geq 1 - 7\eta$ and we get

$$\begin{aligned} d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) &\geq d_{\mathcal{M}}(u, v) \left(\frac{1 - 7\eta}{1 + 4\eta} \right) \\ &\geq d_{\mathcal{M}}(u, v)(1 - 7\eta)(1 - 4\eta) \\ &\geq d_{\mathcal{M}}(u, v)(1 - 11\eta). \end{aligned}$$

Using Eq. (8.8) we find $d_{\mathcal{M}}(u, v) \leq \frac{208r}{100}$, so $d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) \geq d_{\mathcal{M}}(u, v) - 23\frac{r^2}{\text{rch}(\mathcal{M})}$, and the result follows since $d_{\mathcal{M}}(u, v) \geq d_{\mathbb{R}^N|_{\mathcal{M}}}(u, v) \geq d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v))$. $\square$

Our sampling radius is constrained by the size of a Euclidean ball that can be contained in an admissible neighbourhood. The following lemma gives a convenient expression for this:

Lemma 8.3.9 If $1 < a \leq 10^4$ and $a\epsilon \leq \frac{\text{rch}(\mathcal{M})}{100}$, and $U = B_{\mathbb{R}^m}(p, (a-1)\epsilon)$, then $\psi_p(U) = W \subseteq B_{\mathbb{R}^N|_{\mathcal{M}}}(p, a\epsilon)$.

Proof Using Lemma 8.2.10, we have that $B_{\mathbb{R}^m}(p, r) \subseteq \pi_p(B_{\mathbb{R}^N|_{\mathcal{M}}}(p, a\epsilon))$ if

$$\begin{aligned} r^2 &\leq a^2\epsilon^2 - \left(\frac{a^2\epsilon^2}{2\text{rch}(\mathcal{M})} \right)^2 \\ &= a^2\epsilon^2 \left(1 - \left(\frac{a\epsilon}{2\text{rch}(\mathcal{M})} \right)^2 \right) \\ &\leq a^2\epsilon^2 \left(1 - \left(\frac{1}{200} \right)^2 \right) \end{aligned}$$

Thus we require $r \leq \sqrt{\frac{200^2-1}{200^2}}a\epsilon$, which is satisfied by $r = (a-1)\epsilon$ if $a \leq 79999$. $\square$

Lemmas 8.3.8 and 8.3.9 lead to a sampling radius which allows us to employ Theorem 8.3.2, and so obtain an equivalence between Delaunay structures:

Proposition 8.3.10 (Equating Delaunay complexes) *Suppose $\mathcal{P} \subset \mathcal{M}$ is an ϵ -sample set with respect to $d_{\mathbb{R}^N|_{\mathcal{M}}}$, and that for every $p \in \mathcal{P}$, in the local Euclidean metric on $W = B_{\mathbb{R}^N|_{\mathcal{M}}}(p, 5\epsilon)$, every m -simplex in $\text{star}(p; \text{Del}(\mathcal{P}_W))$ is secure, where $\mathcal{P}_W = \mathcal{P} \cap W$, and $\delta = \nu_0\epsilon$. If*

$$\epsilon \leq \frac{\Upsilon_0 \mu_0 \nu_0 \text{rch}(\mathcal{M})}{20700}$$

then

$$\text{star}(p; \text{Del}(\mathcal{P}_W)) = \text{star}(p; \text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P}_W)) = \text{star}(p; \text{Del}_{\mathcal{M}}(\mathcal{P}_W)). \quad (8.9)$$

Thus

$$\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P}) = \text{Del}_{\mathcal{M}}(\mathcal{P}),$$

and they are manifold complexes.

Proof As usual, let $U = \pi_p(W)$. Then by Lemma 8.3.9 $B_{\mathbb{R}^m}(p, 4\epsilon) \subseteq U$, and thus $d_{\mathbb{R}^m}(q, \partial U) \geq 2\epsilon$ for any vertex q of a simplex in $\text{star}(p; \text{Del}(\mathcal{P}_W))$. Thus Lemma 8.3.8 allows us to apply Theorem 8.3.2 provided

$$\frac{23a^2\epsilon^2}{\text{rch}(\mathcal{M})} \leq \frac{\Upsilon_0 \mu_0 \nu_0 \epsilon}{36},$$

when $a = 5$, and we obtain the required bound on ϵ . Thus the star of every vertex in $\text{Del}_{\mathcal{M}}(\mathcal{P})$ is equal to the star of that point in the local Euclidean metric, and likewise for $\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$. The claim follows since $\sigma \in \text{Del}_{\mathcal{M}}(\mathcal{P})$ if and only if it is in the local Euclidean Delaunay triangulation of every one of its vertices, and likewise for the simplices in $\text{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$. $\square$

8.3.3 The protected tangential complex

We obtain Theorem 8.3.6 by means of Theorem 8.3.2 via the observation that power protection of the ambient Delaunay balls translates into protection in the local Euclidean metrics. We must distinguish between the geometry of a simplex defined with respect to the Euclidean metric $d_{\mathbb{R}^N}$ of the ambient space, as opposed to a local Euclidean metric $d_{\mathbb{R}^m}$. In general, we use a tilde to indicate simplices in the ambient space, and their properties.

Lemma 8.3.11 (Protection under projection) *Suppose $\mathcal{P} \subset \mathcal{M}$ and that $\tilde{\sigma} \in \text{Del}_{\mathbb{R}^N}(\mathcal{P})$ is an $\tilde{\Upsilon}_0$ -thick m -simplex, with $L_{\tilde{\sigma}} \geq \tilde{\mu}_0\epsilon$ and $B_{\mathbb{R}^N}(C, R)$ is a $\tilde{\delta}^2$ -power-protected empty ball for $\tilde{\sigma}$, with respect to the metric $d_{\mathbb{R}^N}$, where $\tilde{\delta}^2 \geq \delta_0^2 \tilde{\mu}_0^2 \epsilon^2$. Suppose also that $C \in T_p \mathcal{M}$, for some vertex $p \in \tilde{\sigma}$.*

If $R < \epsilon$, with

$$\epsilon \leq \frac{\tilde{\Upsilon}_0^2 \tilde{\mu}_0^3 \delta_0^2 \text{rch}(\mathcal{M})}{512}, \quad (8.10)$$

then $\sigma = \pi_p(\tilde{\sigma})$ has a δ -protected Delaunay ball $B_{\mathbb{R}^m}(c, r)$ with respect to the local Euclidean metric $d_{\mathbb{R}^m}$ for p on any admissible neighbourhood W that contains $B_{\mathbb{R}^N|_{\mathcal{M}}}(p, 3\epsilon)$, and $\delta = \nu_0\epsilon$, with

$$\nu_0 = \frac{\delta_0^2 \tilde{\mu}_0^2}{8}. \quad (8.11)$$

Proof We first find a bound for $d_{\mathbb{R}^m}(C, c)$ and r . Let $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_m]$, and $\sigma = [p_0, \dots, p_m]$ so that $\pi_p(\tilde{p}_i) = p_i$, and $p = p_0 = \tilde{p}_0$. We will first show that, near C , there is a circumcentre c for σ in the metric $d_{\mathbb{R}^m}$. For any $p_i \in \sigma$, $d_{\mathbb{R}^N}(p, p_i) < 2R$, and so by Lemma 8.2.10 we have

$$d_{\mathbb{R}^N}(\tilde{p}_i, p_i) \leq \frac{2R^2}{\text{rch}(\mathcal{M})} < \frac{2\epsilon^2}{\text{rch}(\mathcal{M})}.$$

In order to apply Lemma 8.2.2 we require $\frac{2\epsilon^2}{\text{rch}(\mathcal{M})} \leq \frac{\tilde{\Upsilon}_0^2 \tilde{\mu}_0 \epsilon}{28}$, or

$$\epsilon \leq \frac{\tilde{\Upsilon}_0^2 \tilde{\mu}_0 \text{rch}(\mathcal{M})}{56},$$

which is satisfied by Eq. (8.10). Since $\text{aff}(\sigma) = T_p \mathcal{M}$, the circumcentre $c \in T_p \mathcal{M}$ is the closest point in N_σ to C , Lemma 8.2.2 yields

$$|R - r| \leq d_{\mathbb{R}^m}(C, c) = d_{\mathbb{R}^N}(C, c) < \frac{16\epsilon^2}{\tilde{\Upsilon}_0 \tilde{\mu}_0 \text{rch}(\mathcal{M})}.$$

Now we seek a lower bound on the protection of $B_{\mathbb{R}^m}(c, r)$. Suppose $\tilde{q} \in \mathcal{P} \setminus \tilde{\sigma}$. We wish to establish a lower bound on $d_{\mathbb{R}^m}(c, q) - r$, where $q = \pi_p(\tilde{q})$. We may assume that $d_{\mathbb{R}^N}(C, \tilde{q}) < 3\epsilon$, since otherwise q will lie outside of our region of interest.

Let $z = \frac{(3\epsilon)^2}{2\text{rch}(\mathcal{M})}$ be the upper bound on $d_{\mathbb{R}^N}(\tilde{q}, q)$ given by Lemma 8.2.10. Then $d_{\mathbb{R}^m}(C, q)^2 \geq d_{\mathbb{R}^N}(C, \tilde{q})^2 - z^2 > R^2 + \delta^2 - z^2$. Thus

$$d_{\mathbb{R}^m}(C, q) - R > \frac{\delta^2 - z^2}{d_{\mathbb{R}^m}(C, q) + R} > \frac{\delta^2 - z^2}{4\epsilon},$$

since $R < \epsilon$. Then $d_{\mathbb{R}^m}(c, q) - r \geq (d_{\mathbb{R}^m}(C, q) - d_{\mathbb{R}^m}(C, c)) - (R + |R - r|) > \frac{\delta^2 - z^2}{4\epsilon} - 2d_{\mathbb{R}^m}(C, c)$. Putting this together, using $\delta^2 \geq \delta_0^2 \tilde{\mu}_0^2 \epsilon^2$, we get

$$d_{\mathbb{R}^m}(c, q) - r > \left(\frac{1}{4} \delta_0^2 \tilde{\mu}_0^2 - \frac{81\epsilon^2}{16\text{rch}(\mathcal{M})^2} - \frac{32\epsilon}{\tilde{\Upsilon}_0 \tilde{\mu}_0 \text{rch}(\mathcal{M})} \right) \epsilon.$$

In order to simplify away the final term, we demand

$$\frac{32\epsilon}{\tilde{\Upsilon}_0 \tilde{\mu}_0 \text{rch}(\mathcal{M})} \leq \frac{1}{16} \delta_0^2 \tilde{\mu}_0^2,$$

which is satisfied by Eq. (8.10). Under this constraint, the central term is also seen to be less than $\frac{1}{16} \delta_0^2 \tilde{\mu}_0^2$, and we obtain

$$\delta \geq \frac{1}{8} \delta_0^2 \tilde{\mu}_0^2 \epsilon.$$

□

Proposition 8.3.10 requires a thickness Υ_0 and shortest edge bound $\mu_0\epsilon$ for the simplex $\sigma \subset \mathbb{R}^m$, but Lemma 8.3.11 is expressed in terms of the corresponding quantities $\tilde{\Upsilon}_0$ and $\tilde{\mu}_0\epsilon$ for the corresponding simplex $\tilde{\sigma} \subset \mathbb{R}^N$.

Lemma 8.3.12 (Simplex distortion under projection) *Let $\tilde{\sigma} \in \text{Del}_{\mathbb{R}^N}(\mathcal{P})$ be an m -simplex as described in Lemma 8.3.11, and let $\sigma = \pi_p(\tilde{\sigma})$ be its projection in the local Euclidean metric for p on any admissible neighbourhood that contains $B_{\mathbb{R}^N|_{\mathcal{M}}}(p, 2\epsilon)$, where p is a vertex of $\tilde{\sigma}$. If ϵ satisfies Eq. (8.10), and $\delta_0^2 \tilde{\mu}_0^2 \leq \frac{1}{7}$, then $L_\sigma > \mu_0\epsilon$, where*

$$\mu_0 = \frac{20}{21} \tilde{\mu}_0,$$

and $\Upsilon_\sigma > \Upsilon_0$, where

$$\Upsilon_0 = \frac{6}{49} \tilde{\Upsilon}_0.$$

Proof Since $\pi_p(B_{\mathbb{R}^N|_{\mathcal{M}}}(p, 2\epsilon)) \subseteq B_{\mathbb{R}^m}(p, 2\epsilon)$, it is sufficient to apply the Metric distortion lemma 8.3.9 with $a = 3$.

For the shortest edge length, we find

$$L_\sigma \geq L_{\tilde{\sigma}} - \frac{3^2 \times 23\epsilon^2}{\text{rch}(\mathcal{M})} = \tilde{\mu}_0 \left(1 - \frac{207\tilde{\Upsilon}_0^2 \delta_0^2 \tilde{\mu}_0^2}{512} \right) \epsilon > \tilde{\mu}_0 \left(1 - \frac{\tilde{\Upsilon}_0^2 \delta_0^2 \tilde{\mu}_0^2}{3} \right) \epsilon > \frac{20}{21} \tilde{\mu}_0 \epsilon.$$

For the thickness bound, in order to apply Lemma 8.2.1, using $\tilde{\eta} = (1 - \eta)$, we require

$$\frac{207\tilde{\Upsilon}_0^2 \delta_0^2 \tilde{\mu}_0^3}{512} \leq \frac{\tilde{\eta} \tilde{\Upsilon}_0^2 \tilde{\mu}_0}{14},$$

which is satisfied if we choose

$$\tilde{\eta} > 6\delta_0^2 \tilde{\mu}_0^2.$$

Then Lemma 8.2.1 yields

$$\Upsilon_\sigma \geq \frac{6}{7} (1 - 6\delta_0^2 \tilde{\mu}_0^2) \Upsilon_{\tilde{\sigma}} > \frac{6}{49} \Upsilon_{\tilde{\sigma}}.$$

□

We can now express the sampling conditions in terms of the output parameters of the tangential complex algorithm, and this allows us to apply Proposition 8.3.10 and obtain our main structural result:

Proof of Theorem 8.3.6 We first translate the sampling requirements of Proposition 8.3.10 in terms of properties of simplices in the ambient metric $d_{\mathbb{R}^N}$. Using Lemma 8.3.12, together with Eq. (8.11), the upper bound on the sampling radius demanded by Proposition 8.3.10 becomes

$$\epsilon \leq \frac{20 \times 6 \tilde{\Upsilon}_0 \delta_0^2 \tilde{\mu}_0^3 \text{rch}(\mathcal{M})}{21 \times 49 \times 8 \times 20700}.$$

We obtain the stated sampling radius bound after multiplying by $\tilde{\Upsilon}_0$ in order to ensure that the demand of Eq. (8.10) is also met. Thus the stated sampling radius satisfies the requirements of both Lemma 8.3.11 and Proposition 8.3.10.

The fact that the structures are isomorphic follows from the fact that they are all locally isomorphic to the Delaunay triangulation in the local Euclidean metric. To see that $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P})) \cong \text{star}(p; \text{Del}(\mathcal{P}_W))$, observe that Lemma 8.3.11 implies that there is an injective simplicial map $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P})) \rightarrow \text{star}(p; \text{Del}(\mathcal{P}_W))$. The isomorphism is established by Lemma 7.2.6, once it is established that $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ is a triangulation at p . In fact $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ is isomorphic to the star of p in a regular triangulation of the projected points $\mathcal{P}_W$; it is a *weighted Delaunay triangulation* [BG10b], and with our choice of W , the point p is an interior point in this triangulation [BG10b, Lemma 2.7(1)]. Thus $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ is a triangulation at p , and it follows that

$$\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P})) \cong \text{star}(p; \text{Del}(\mathcal{P}_W)).$$

The equality of the Delaunay complexes now follows from Proposition 8.3.10 and Eq. (8.9).

The homeomorphism assertion follows from previous Theorem 4.0.3. $\square$

8.4 Algorithm

In this section we introduce a Delaunay refinement algorithm which, while constructing a tangential Delaunay complex, will transform the input sample set into one which meets the requirements of Theorem 8.3.6. In particular we wish to construct a tangential Delaunay complex in which every m -simplex σ is $\tilde{\Upsilon}_0$ -thick and for every $p \in \sigma$, there is a $\check{\delta}^2$ -power-protected Delaunay ball for σ centred on $T_p\mathcal{M}$. We demand $\check{\delta} \geq \delta_0 \tilde{\mu}_0 \epsilon$, where ϵ provides a strict upper bound on the radius of these Delaunay balls, and $\tilde{\mu}_0 \epsilon$ provides a lower bound on the shortest edge length of any simplex in $\text{Del}_{T\mathcal{M}}(\mathcal{P})$. The constants δ_0 and $\tilde{\mu}_0$ are both positive and smaller than one.

The algorithm is in the same vein as the one given in Chapter 5, which is in turn an adaptation of the algorithm introduced by Li [Li03b]. It is described in Section 8.4.2, after we introduce terminology and constructs which are used in the algorithm in Section 8.4.1.

8.4.1 Components of the algorithm

We now introduce the primary concepts that are used as building blocks of the algorithm.

8.4.1.1 Elementary weight functions

Elementary weight functions are a convenient device to facilitate the identification of simplices σ that are not $\check{\delta}^2$ -power-protected for $\check{\delta} = \delta_0 L_\sigma$.

In order to emphasise that we are considering a function defined only on the set of vertices of a simplex, we denote by $\circ$ the set $\{p_0, \dots, p_k\}$ of vertices of $\sigma = [p_0 \dots p_k]$. We will call $\omega_\sigma : \circ \rightarrow [0, \infty)$ an *elementary weight function* if it satisfies the following conditions:

1. There exists $p_i \in \circ$ such that $\omega_\sigma(p_i) \in [0, \delta_0 L_\sigma]$, and
2. for all $p_j \in \circ \setminus p_i$, $\omega_\sigma(p_j) = 0$.

For a given $\sigma = [p_0, \dots, p_k]$ and elementary weight function ω_σ , we define $N(\sigma, \omega_\sigma)$ as the set of solutions to the following system of k equations:

$$\|x - p_i\|^2 - \|x - p_0\|^2 = \omega_\sigma(p_i)^2 - \omega_\sigma(p_0)^2.$$

In direct analogy with the space N_σ of centres of σ , the set $N(\sigma, \omega_\sigma)$ is an affine space of dimension $m - \dim \text{aff}(\sigma)$ that is orthogonal to $\text{aff}(\sigma)$. We denote by $C(\sigma, \omega_\sigma)$ the unique point in $N(\sigma, \omega_\sigma) \cap \text{aff}(\sigma)$, and we define

$$R(\sigma, \omega_\sigma)^2 = \|p_0 - C(\sigma, \omega_\sigma)\|^2 - \omega_\sigma(p_0)^2,$$

where the notation is chosen to emphasise the close relationship with the circumcentre c_σ and circumradius R_σ . The following lemma exposes some properties of $R(\sigma, \omega_\sigma)$ in this spirit:

Lemma 8.4.1 *For a given $\sigma = [p_0, \dots, p_k]$, with $k \geq 1$, and elementary weight function ω_σ , we have:*

1. *If $\sigma_1 \leq \sigma$ then $\omega_{\sigma_1} = \omega_\sigma|_{\sigma_1}$ is an elementary weight function, and*

$$R(\sigma_1, \omega_{\sigma_1}) \leq R(\sigma, \omega_\sigma).$$

2. *$\Delta_\sigma \leq \frac{2}{1-\delta_0^2} R(\sigma, \omega_\sigma)$.*

3. *If $\Upsilon_\sigma > 0$, then*

$$1 - \eta \leq \frac{R(\sigma, \omega_\sigma)}{R_\sigma} \leq 1 + \eta,$$

$$\text{with } \eta = \frac{\delta_0^2}{\Upsilon_\sigma}.$$

Proof 1. That ω_{σ_1} is an elementary weight function follows from the observation that $L_{\sigma_1} \geq L_\sigma$. Since $N(\sigma, \omega_\sigma) \subseteq N(\sigma_1, \omega_{\sigma_1})$, the projection of $C(\sigma, \omega_\sigma)$ into $\text{aff}(\sigma_1)$ is $C(\sigma_1, \omega_{\sigma_1})$. The result then follows from the Pythagorean theorem.

2. Let $e = [p_0, p_1]$ be the longest edge of σ , and let c denote the projection of $C(\sigma, \omega_\sigma)$ onto $\text{aff}(e)$. Without loss of generality we assume that $\omega(p_0) = 0$.

We have

$$\begin{aligned} \|p_0 - c\|^2 &= \|p_1 - c\|^2 - \omega_\sigma(p_1)^2 \\ &= \|(p_1 - p_0) - (c - p_0)\|^2 - \omega_\sigma(p_1)^2 \\ &= \Delta_\sigma^2 - 2(p_1 - p_0) \cdot (c - p_0) + \|p_0 - c\|^2 - \omega_\sigma(p_1)^2. \end{aligned}$$

Since p_0, p_1 , and c are colinear, we have $2(p_1 - p_0) \cdot (c - p_0) = 2\Delta_\sigma \|p_0 - c\|$, and using the fact that $\omega_\sigma(p_1) \leq \delta_0 L_\sigma$, we get

$$\begin{aligned} \|p_0 - c\| &= \frac{\Delta_\sigma}{2} \left(1 - \frac{\omega_\sigma(p_1)^2}{\Delta_\sigma^2} \right) \\ &\geq \frac{(1 - \delta_0^2)\Delta_\sigma}{2}. \end{aligned}$$

The result follows from the fact that $R(\sigma, \omega_\sigma) \geq \|p_0 - c\|$.

3. Using the fact that $\omega_\sigma(p) = 0$ for all vertices $p \in \sigma$, except at most one, we get $p_i \in \partial B_{\mathbb{R}^k}(C(\sigma, \omega_\sigma), R(\sigma, \omega_\sigma))$ for all $p_i \in \sigma$ except at most one.

Let $\eta = \|C(\sigma, \omega_\sigma) - c_\sigma\|$, and assume, without loss of generality, that the vertex $p_0 \in \partial B(C(\sigma, \omega_\sigma), R(\sigma, \omega_\sigma))$. Therefore,

$$\begin{aligned} \|c_\sigma - p_0\| - \eta &\leq \|C(\sigma, \omega_\sigma) - p_0\| \leq \|c_\sigma - p_0\| + \eta \\ R_\sigma + \eta &\leq R(\sigma, \omega_\sigma) \leq R_\sigma + \eta. \end{aligned} \quad (8.12)$$

Since the point in N_σ that is closest to $C(\sigma, \omega_\sigma)$ is c_σ , and $\Upsilon_\sigma > 0$, we obtain the following bound using Lemma 4.1 from [?]:

$$\begin{aligned} \eta &\leq \frac{\delta_0^2 L_\sigma^2}{2\Upsilon_\sigma \Delta_\sigma} \\ &\leq \frac{\delta_0^2}{\Upsilon_\sigma} R_\sigma, \quad \text{since } L_\sigma \leq \Delta_\sigma \leq 2R_\sigma. \end{aligned} \quad (8.13)$$

The result now follows from Eq. (8.12) and (8.13). $\square$

If $\sigma = p * \sigma_p$, and ω_σ is an elementary weight function that vanishes on σ_p , then $N(\sigma, \omega_\sigma) \subseteq N_{\sigma_p}$, but no point in $N(\sigma, \omega_\sigma)$ can be the centre of a δ^2 -power-protected Delaunay ball for σ_p for any $\delta \geq \delta_0 L_\sigma$. In other words, σ and ω_σ define a quasi-cospherical configuration that is an obstruction to the power protection of σ_p at all points in $N(\sigma, \omega_\sigma)$.

8.4.1.2 Quasicospherical configurations

We now define the family of simplices that our algorithm must eliminate in order to ensure that the final point set has the desired protection properties.

Recalling the definition (8.5) of $\text{star}(p)$, we have the following result from Chapter 5 (Lemma 5.2.3):

Lemma 8.4.2 *Let $\mathcal{P} \subset \mathcal{M}$ satisfy a sampling radius of ϵ with respect to $d_{\mathbb{R}^N}$ such that $\epsilon \leq \text{rch}(\mathcal{M})/16$. Then for all $x \in \text{Vor}_{\mathbb{R}^N}(p) \cap T_p \mathcal{M}$, we have $\|p - x\| \leq 4\epsilon$. In particular, for all $p \in \mathcal{P}$, and every m -simplex $\sigma \in \text{star}(p)$, we have $R_p(\sigma) \leq 4\epsilon$.*

Since by Lemma 8.4.2, the Voronoi cell of p restricted to $T_p \mathcal{M}$ is bounded, we get:

Lemma 8.4.3 *If $\epsilon \leq \frac{\text{rch}(\mathcal{M})}{16}$, then the combinatorial dimension of the maximal simplices in $\text{star}(p)$ is at least m .*

We will always assume that $\mathcal{P}$ satisfies a sampling radius of $\epsilon \leq \frac{\text{rch}(\mathcal{M})}{16}$. If σ is a maximal simplex in $\text{star}(p)$, then $\text{Vor}_{\mathbb{R}^N}(\sigma)$ intersects $T_p \mathcal{M}$ at a single point. Indeed, since $\text{Vor}_{\mathbb{R}^N}(\sigma) \subset \text{Vor}_{\mathbb{R}^N}(p)$, by Lemma 8.4.2 the convex set $\text{Vor}_{\mathbb{R}^N}(\sigma) \cap T_p \mathcal{M}$ is bounded, and if it had a nonempty interior, then σ would not be maximal. Let σ be a maximal simplex in $\text{star}(p)$. Then, for all $\sigma^m \leq \sigma$, the unique point in $\text{Vor}_{\mathbb{R}^N}(\sigma) \cap T_p \mathcal{M}$ will be denoted by $c_p(\sigma^m)$. We denote the radius of the circumscribing ball centred at $c_p(\sigma^m)$ by $R_p(\sigma^m)$, i.e., $R_p(\sigma^m) = \|p - c_p(\sigma^m)\|$.

In our algorithm we will use the following complex, whose definition employs a particular elementary weight function:

$$\text{cosph}^{\delta_0}(p) = \left\{ \begin{array}{l} \sigma^{m+1} = p_{m+1} * \sigma^m \mid \sigma^m \in \text{star}(p), R_p(\sigma^m) < \epsilon, \\ \sigma^m \text{ is } \Gamma_0\text{-good, and } \exists \omega_{\sigma^{m+1}} \text{ with } \omega_{\sigma^{m+1}}|_{\sigma^m} = 0 \\ \text{and } c_p(\sigma^m) \in N(\sigma^{m+1}, \omega_{\sigma^{m+1}}) \end{array} \right\}. \quad (8.14)$$

The $(m+1)$ -dimensional simplices in $\text{cosph}^{\delta_0}(p)$ are analogous to inconsistent configurations defined in Chapters 3 and 5.

Unless otherwise stated, whenever $\sigma^{m+1} = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$, with $\sigma^m \in \text{star}(p)$, the mention of $\omega_{\sigma^{m+1}}$ will refer to the elementary weight function identified in Eq. (8.14). In particular,

$$\omega_{\sigma^{m+1}}(p_i) = 0 \text{ for all } p_i \in \sigma^{m+1} \setminus p_{m+1},$$

and

$$\omega_{\sigma^{m+1}}(p_{m+1}) \in [0, \delta_0 L_{\sigma^{m+1}}]$$

satisfies

$$\|c_p(\sigma^m) - p\|^2 = \|c_p(\sigma^m) - p_{m+1}\|^2 - \omega_{\sigma^{m+1}}(p_{m+1})^2.$$

We will exploit the following observations:

Lemma 8.4.4 *If $\sigma^{m+1} = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$ with $\sigma^m \in \text{star}(p)$, then*

$$R(\sigma^{m+1}, \omega_{\sigma^{m+1}}) \leq R_p(\sigma^m)$$

and

$$\Delta_{\sigma^{m+1}} \leq \frac{2}{1 - \delta_0^2} R_p(\sigma^m)$$

Proof Since $c_p(\sigma^m) \in N(\sigma^{m+1}, \omega_{\sigma^{m+1}})$, it follows that $C(\sigma^{m+1}, \omega_{\sigma^{m+1}})$ is the projection of $c_p(\sigma^m)$ into $\text{aff}(\sigma^{m+1})$, and therefore $R_p(\sigma^m) \geq R(\sigma^{m+1}, \omega_{\sigma^{m+1}})$. The bound on $\Delta_{\sigma^{m+1}}$ now follows directly from Lemma 8.4.1. $\square$

Lemma 8.4.2 implies that we can compute $\text{star}(p)$ by computing a weighted Delaunay triangulation on $T_p\mathcal{M}$ of the points obtained by projecting $\mathcal{P}$ onto $T_p\mathcal{M}$. Once $\text{star}(p)$ has been computed, we can compute $\text{cosph}^{\delta_0}(p)$ by a simple distance computation.

The importance of $\text{cosph}^{\delta_0}(p)$ lies in the observation that if an m -simplex $\sigma^m \in \text{star}(p)$ is not sufficiently power-protected, then there will be a simplex in $\text{cosph}^{\delta_0}(p)$ that is a witness to this. It is a direct consequence of the definitions, but we state it explicitly for reference:

Lemma 8.4.5 *If $\mathcal{P}$ is $\tilde{\mu}_0\epsilon$ -sparse, and $\text{cosph}^{\delta_0}(p) = \emptyset$, then every $\sigma^m \in \text{star}(p)$ is $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power protected on $T_p\mathcal{M}$.*

8.4.1.3 Unfit configurations and the picking region

The refinement algorithm, at each step, kills an *unfit configuration* by inserting a new point $x = \psi_p(x')$ where x' belongs to the so-called *picking region* of the unfit configuration, and ψ_p is the inverse projection defined in Eq. (8.7). We use the term unfit configuration to distinguish the elements under consideration from other simplices. An unfit configuration ϕ may be one of two types:

Big configuration: An m -simplex $\phi = \sigma^m$ in $\text{star}(p)$ is a *big configuration* if $R_p(\sigma^m) \geq \epsilon$.

Bad configuration: A simplex ϕ is a *bad configuration* if it is Γ_0 -bad and it is either an m -simplex $\phi = \sigma^m \in \text{star}(p)$ that is not a big configuration, or it is an $(m+1)$ -simplex $\phi = \sigma^{m+1} \in \text{cosph}^{\delta_0}(p)$.

We will show in Section 8.5.2, Lemma 8.5.13, that in fact *every* $(m+1)$ -simplex in $\text{cosph}^{\delta_0}(p)$ is a bad configuration.

The size of the picking region is governed by a positive parameter $\alpha < 1$ called the *picking ratio*.

Definition 8.4.6 (Picking region) *The picking region of a bad configuration, $\sigma^m \in \text{star}(p)$ or $p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$ with $\sigma^m \in \text{star}(p)$, denoted by $P(\sigma^m, p)$ and $P(\sigma^{m+1}, p)$ respectively, is defined to be the m -dimensional ball*

$$B_{\mathbb{R}^N}(c_p(\sigma^m), \alpha R_p(\sigma^m)) \cap T_p \mathcal{M}.$$

We choose a point in the picking region so as to minimize the introduction of new unfit configurations. We are able to avoid creating new bad configurations provided that the radius of the potential configuration is not too large. To this end, we introduce the parameter $\beta > 1$.

Definition 8.4.7 (Hitting sets and good points) *Let $\phi = \sigma^m \in \text{star}(p)$ or $\phi = q * \sigma^m \in \text{cosph}^{\delta_0}(p)$ with $\sigma^m \in \text{star}(p)$, and $x = \psi_p(y)$ where $y \in P(\phi, p)$. A set $\sigma \subset \mathcal{P}$ of size k , with $k \leq m+1$, is called a *hitting set* of x if*

a. $\tau = x * \sigma$ is a k -dimensional Γ_0 -flake

and there exists an elementary weight function ω_τ satisfying the following condition:

b. $R(\tau, \omega_\tau) < \beta R_p(\sigma^m)$

The elementary weight function ω_τ is called a *hitting map*, and we sometimes say σ hits x .

A point $x = \psi_p(y)$, where $y \in P(\phi, p)$, is said to be a *good point* if it is not hit by any set $\sigma \subset \mathcal{P}$ with $|\sigma| \leq m+1$.

A simplex σ which defines a hitting set of x , is necessarily Γ_0 -good. This follows from the requirement that $x * \sigma$ be a Γ_0 -flake.

8.4.2 The refinement algorithm

In this section, we show that we can refine an ϵ -net of $\mathcal{M}$ so that the simplices of the Delaunay tangential complex of the refined sample $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ are power-protected. An ϵ -net is a point sample $\mathcal{P} \subset \mathcal{M}$ that is an ϵ -sparse ϵ -sample set of $\mathcal{M}$ for the metric $d_{\mathbb{R}^N}$. One can obtain an ϵ -net by using a farthest point strategy to select a subset of a sufficiently dense sample set. We will assume that we know the dimension m of the submanifold $\mathcal{M}$ and the tangent space $T_p \mathcal{M}$ at any point p in $\mathcal{M}$.

The algorithm takes as input $\mathcal{P}_0$, an ϵ -net of $\mathcal{M}$, and the positive input parameters ϵ , Γ_0 , $\alpha < \frac{1}{2}$, $\beta > 1$ and $\delta_0 < \frac{1}{4}$. The algorithm refines the input point sample such that:

- (1) The output sample $\mathcal{P} \supseteq \mathcal{P}_0$ is an $\tilde{\mu}_0\epsilon$ -sparse ϵ -sample set of $\mathcal{M}$ with respect to $d_{\mathbb{R}^N}$, where $\mu_0 = \frac{1}{9}$.
- (2) For all $p \in \mathcal{P}$, every m -simplex $\sigma^m \in \text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$, σ^m is Γ_0 -good and $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power protected on $T_p\mathcal{M}$.

Algorithm 4 Refinement algorithm

Input ϵ -net $\mathcal{P}_0$ of $\mathcal{M}$, and input parameters Γ_0 , α and δ_0 ;
Initialize $\mathcal{P} \leftarrow \mathcal{P}_0$, and calculate $\text{Del}_{T\mathcal{M}}(\mathcal{P})$;
Rule (1) *Big configuration* (ϵ -big radius):
 if $\exists p \in \mathcal{P}$ such that $\exists \sigma^m \in \text{star}(p)$ with $R_p(\sigma^m) \geq \epsilon$,
 then Insert($\psi_p(c_p(\sigma^m))$);
Rule (2) *Bad configuration* (Γ_0 -bad):
 if $\exists p \in \mathcal{P}$ and $\exists \sigma^m \in \text{star}(p)$ s.t. σ^m is Γ_0 -bad,
 then Insert(Pick_valid(σ^m, p));
 if $\exists p \in \mathcal{P}$ and $\exists \sigma^{m+1} \in \text{cosph}^{\delta_0}(p)$ s.t. σ^{m+1} is Γ_0 -bad,
 then Insert(Pick_valid(σ^{m+1}, p));
Output $\text{Del}_{T\mathcal{M}}(\mathcal{P}) = \cup_{p \in \mathcal{P}} \text{star}(p)$;

The algorithm, described in Algorithm 4, applies two rules with a priority order: Rule (2) is applied only if Rule (1) cannot be applied. The algorithm ends when no rule applies any more. Each rule inserts a new point to kill an unfit configuration: either a big configuration or a bad configuration.

A crucial procedure, that selects the location of the point to be inserted, is Pick_valid, given in Algorithm 5. Pick_valid(ϕ, p) returns a good point $\psi_p(y)$ where $y \in P(\phi, p)$.

Algorithm 5 Pick_valid(σ, p)

// Assume that σ is either equal to $\sigma^m \in \text{star}(p)$
// or $\sigma^{m+1} = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$ with $\sigma^m \in \text{star}(p)$
Step 1. Pick randomly $y \in P(\sigma^m, p)$ (or $P(\sigma^{m+1}, p)$);
// Recall that ψ_p projects points from $T_p\mathcal{M}$ onto $\mathcal{M}$ along $N_p\mathcal{M}$
Step 2. $x \leftarrow \psi_p(y)$;
Step 3. *Avoid hitting sets*:
 // $|\tilde{\sigma}|$ denotes the cardinality of $\tilde{\sigma}$
 if $\exists \tilde{\sigma} \subset \mathcal{P}$, with $|\tilde{\sigma}| \leq m+1$, which is a *hitting set* of x ,
 then discard x , and go back to Step 1;
Step 4. Return x ;

The refinement algorithm will also use the procedure Insert(p), given in Algorithm 6.

Algorithm 6 Insert(p)

-
- Step 1. Add p to $\mathcal{P}$;
 Step 2. Compute $\text{star}(p)$ and $\text{cospH}^{\delta_0}(p)$;
 Step 3. For all $x \in \mathcal{P} \setminus \{p\}$, update $\text{star}(x)$ and $\text{cospH}^{\delta_0}(x)$;
-

8.5 Analysis of the algorithm

We now turn to the demonstration of the correctness of Algorithm 4. In Section 8.5.1 we show that the algorithm must terminate, and in Section 8.5.2 we show that the output of the algorithm meets the requirements of Theorem 8.3.6. In order to complete the demonstrations we impose a number of requirements on the input parameters, listed as Hypotheses $\mathcal{H}0$ to $\mathcal{H}5$ below.

Recall that our input parameters are the following positive numbers: ϵ , which is the sampling radius and sparsity bound satisfied by $\mathcal{P}_0$, the input ϵ -net sample set; δ_0 , which is used to describe the amount of power-protection enjoyed by the m -simplices in the final complex; Γ_0 , which is used to quantify the quality of the output simplices; β , which is used to describe an upper bound on the radius of the bad configurations that we will avoid; and α , which governs the relative size of the picking region.

It is often convenient to represent the sampling radius by a dimension-free parameter that has the reach of the manifold factored out. We define

$$\tilde{\epsilon} = \frac{\epsilon}{\text{rch}(\mathcal{M})}.$$

The volume of the m -dimensional Euclidean unit-ball is denoted ϕ_m . In order to state the hypotheses on the input parameters, we use some additional symbols:

$$\begin{aligned}\tilde{\epsilon}_0 &= \frac{1}{2^4(2^4 + 1)^2}, \\ B &= 4 + 2(1 + 2^7 3^2 \beta^2)^2, \\ \beta' &= \frac{\beta}{1 - 2^4 \tilde{\epsilon}_0},\end{aligned}$$

as well as ξ , E , and D . The term ξ is introduced in Lemma 8.5.5, and depends on m and $\text{rch}(\mathcal{M})$, and the term E , defined in Eq. (8.18), depends on ξ and β . The symbol D is introduced in Lemma 8.5.8, where it is said to depend on m and β .

In order to guarantee termination, we demand the following hypotheses on the input parameters:

$$\mathcal{H}0. \quad \alpha < 1/2$$

$$\mathcal{H}1. \quad \beta \geq \frac{2}{(1 - \delta_0^2)(1 - \alpha - 4.5 \tilde{\epsilon}_0)}$$

$$\mathcal{H}2. \quad \Gamma_0 < \min \left\{ \frac{\phi_m \alpha^m}{E^{m+1} \beta^m D}, \frac{1}{B+1} \right\}$$

$$\mathcal{H}3. \quad \delta_0^2 \leq \Gamma_0^{m+1}$$

$$\mathcal{H}4. \quad \tilde{\epsilon} \leq \min \left\{ \frac{\xi}{2(\beta + \beta')}, \frac{\Gamma_0^{m+1}}{8\beta} \right\}$$

To meet the quality requirements of Theorem 8.3.6 we demand an additional constraint on the sampling radius:

$$\mathcal{H}5. \quad \tilde{\epsilon} \leq \frac{\delta_0^2 \Gamma_0^{2m}}{1.1 \times 10^9}$$

The make use of the following observation:

Lemma 8.5.1 *From hypotheses $\mathcal{H}0$ to $\mathcal{H}4$ we have $\tilde{\epsilon} < \tilde{\epsilon}_0$ and $\delta_0^2 < 2^4 \tilde{\epsilon}_0$, and*

$$\frac{(1 - \delta_0^2)(1 - \alpha - 4.5\tilde{\epsilon}_0)\epsilon}{4} > \frac{\epsilon}{9} \stackrel{\text{def}}{=} \tilde{\mu}_0 \epsilon. \quad (8.15)$$

Proof From $\mathcal{H}1$ we have $\beta > 2$ and using the fact that $B > \beta^4$ and $\mathcal{H}2$ we have $\Gamma_0 < \frac{1}{2^4+1}$. And using the fact, from $\mathcal{H}4$, that

$$\tilde{\epsilon} \leq \frac{\Gamma_0^{m+1}}{8\beta} \leq \frac{\Gamma_0^2}{8\beta} < \frac{1}{2^4(2^4+1)^2} = \tilde{\epsilon}_0.$$

Similarly the bound on δ_0^2 follows from $\mathcal{H}3$.

Inequality (8.15) follows from $\mathcal{H}0$ and the definition of $\tilde{\epsilon}_0$. $\square$

From Eq. (8.15) we can see that we require $\beta \geq 4.5$. Given α satisfying $\mathcal{H}0$, and a valid choice for β , the hypotheses $\mathcal{H}2$ to $\mathcal{H}4$ sequentially yield upper bounds on the parameters Γ_0 , δ_0 , and $\tilde{\epsilon}$; we are able to choose parameters that satisfy all of the hypotheses.

The main result of this section can now be summarised:

Theorem 8.5.2 (Algorithm guarantee) *If the input parameters satisfy hypotheses $\mathcal{H}0$ to $\mathcal{H}5$, then Algorithm 4 terminates after producing an intrinsic Delaunay complex $\text{Del}_{\mathcal{M}}(\mathcal{P})$ that triangulates $\mathcal{M}$.*

8.5.1 Termination of the algorithm

This subsection is devoted to the proof of the following theorem:

Theorem 8.5.3 (Algorithm termination) *Under hypotheses $\mathcal{H}0$ to $\mathcal{H}4$, the application of Rule (1) or Rule (2) on a big or a bad configuration ϕ always leaves the interpoint distance greater than*

$$\tilde{\mu}_0 \epsilon = \frac{\epsilon}{9},$$

and if ϕ is a bad configuration then there exists $x \in P(\phi, p)$ such that $\psi_p(x)$ is a good point. Since $\mathcal{M}$ is a compact manifold this implies that the refinement algorithm terminates and returns a point sample $\mathcal{P}$ which is an $\tilde{\mu}_0 \epsilon$ -sparse ϵ -sample of the manifold $\mathcal{M}$.

We will prove that at every step the algorithm maintains the following two *invariants*:

Sparsity: Whenever a refinement rule inserts a new point $x = \psi_p(y)$, the distance between x and the existing point set $\mathcal{P}$ is greater than $\tilde{\mu}_0 \epsilon$.

Good points: For a bad configuration ϕ refined by Rule (2), there exists a set of positive volume $G \subseteq P(\phi, p)$ such that if $x \in G$, then $\psi_p(x)$ is a good point.

The Termination Theorem 8.5.3 is a direct consequence of these two algorithmic invariants. We first prove the sparsity invariant in Section 8.5.1.1, using an induction argument that relies on the fact that the algorithm only inserts good points. The existence of good points is then established in Section 8.5.1.2, using the sparsity invariant and a volumetric argument. Termination must follow since $\mathcal{M}$ is compact and therefore can only support a finite number of sample points satisfying a minimum interpoint distance.

8.5.1.1 The sparsity invariant

The proof of the sparsity invariant employs the following observation, which serves to bound the distance between a point inserted by Rule (2) and the existing point set:

Lemma 8.5.4 *Assume Hypotheses $\mathcal{H}0$ to $\mathcal{H}4$. Let $\phi = \sigma^m \in \text{star}(p)$ or $\phi = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$ be a bad configuration being refined by Rule (2). Then for all $x \in P(\phi, p)$ we have*

$$d_{\mathbb{R}^N}(c_p(\sigma^m), \psi_p(x)) < (\alpha + 4.5\tilde{\epsilon}_0)R_p(\sigma^m)$$

and

$$d_{\mathbb{R}^N}(\psi_p(x), \mathcal{P}) > (1 - \alpha - 4.5\tilde{\epsilon}_0)R_p(\sigma^m) > \frac{R_p(\sigma^m)}{3}.$$

Proof Using the facts that $\alpha < \frac{1}{2}$, and $\tilde{\epsilon} < \tilde{\epsilon}_0$, and $R_p(\sigma^m) < \epsilon$, we have that for all $x \in P(\phi, p)$

$$\|p - x\| < (1 + \alpha)R_p(\sigma^m) < \frac{3\epsilon}{2} < \frac{3\tilde{\epsilon}_0}{2} < \frac{1}{4},$$

and so we may apply Lemma 8.2.11 to get

$$\|x - \psi_p(x)\| \leq \frac{2\|p - x\|^2}{\text{rch}(\mathcal{M})} \leq \frac{2(1 + \alpha)^2 R_p(\sigma^m)^2}{\text{rch}(\mathcal{M})} \leq 4.5\tilde{\epsilon}_0 R_p(\sigma^m),$$

and

$$\|c_p(\sigma^m) - \psi_p(x)\| \leq \|c_p(\sigma^m) - x\| + \|x - \psi_p(x)\| \leq (\alpha + 4.5\tilde{\epsilon}_0) R_p(\sigma^m). \quad (8.16)$$

Let $S_p = \partial B_{\mathbb{R}^N}(c_p(\sigma^m); R_p(\sigma^m))$. From Eq. (8.16) we have for $x \in P(\phi, p)$

$$d_{\mathbb{R}^N}(\psi_p(x); \mathcal{P}) \geq d_{\mathbb{R}^N}(\psi_p(x), S_p) > (1 - \alpha - 4.5\tilde{\epsilon}_0) R_p(\sigma^m) > \frac{R_p(\sigma^m)}{3},$$

where the final inequality follows from $\mathcal{H}0$ and the definition of $\tilde{\epsilon}_0$. $\square$

We introduce some additional terminology to facilitate the demonstration of the sparsity invariant. An abstract simplex in the initial sample set $\sigma \subset \mathcal{P}_0$ is called an *original simplex*, otherwise $\sigma \subset \mathcal{P}$ is called a *created simplex*.

Let ϕ be an unfit configuration that was refined by inserting a point x . We say that x *created* σ if $x \in \sigma$ and x is the last inserted vertex of the simplex σ , i.e., $\sigma \setminus \{x\}$ already existed just before the refinement of the unfit configuration ϕ . The unfit configuration ϕ is called the *parent of* σ and will be denoted $p(\sigma)$.

Let σ denote the simplex being refined by the refinement algorithm. We will denote by $e(\sigma)$ the distance between the point newly inserted to refine σ and the current sample set.

The sparsity invariant is demonstrated by induction. We use a case analysis according to the type of unfit configuration being refined; it is necessary to consider sub-cases. The induction hypothesis is employed only in the sub-case **Case 2(b)(ii)** and the implicit similar **Case 3(b)(ii)**. The base for the induction hypothesis, i.e., the insertion of the first point, cannot involve **Case 2(b)** or **Case 3(b)**.

Case 1. Let $\phi = \sigma^m \in \text{star}(p)$ be a big configuration being refined by Rule (1).

Since $\mathcal{P}_0 (\subseteq \mathcal{P})$ is an ϵ -net, we have from the fact that $\tilde{\epsilon} \leq \tilde{\epsilon}_0 < \frac{1}{16}$ and Lemma 8.4.2, $R_p(\sigma^m) \leq 4\epsilon$. Rule (1) will refine σ by inserting $\psi_p(c_p(\sigma^m))$. Using the fact that $\tilde{\epsilon} < \tilde{\epsilon}_0 < \frac{1}{16}$, $R_p(\sigma^m) \leq 4\epsilon$ and $R_p(\sigma^m) \geq \epsilon$ (since σ^m is being refined by Rule (1)), and Lemma 8.2.11, the distance between $\psi_p(c_p(\sigma^m))$ and any vertex inserted before $\psi_p(c_p(\sigma^m))$ is not less than

$$R_p(\sigma) - \|c_p(\sigma) - \psi_p(c_p(\sigma))\| \geq R_p(\sigma) - \frac{2R_p(\sigma)^2}{\text{rch}(\mathcal{M})} > (1 - 8\tilde{\epsilon}_0)\epsilon > \frac{\epsilon}{2},$$

which establishes the sparsity invariant for this case.

Case 2. Consider now the case where $\phi = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$, with $\sigma^m \in \text{star}(p)$, is being refined by Rule (2). In this case, recalling Lemma 8.2.4, we have

- $R_p(\sigma^m) < \epsilon$, and
- there exists a face of ϕ that is a Γ_0 -flake.

Let $\sigma_1 \subseteq \phi$ denote a face of ϕ that is a Γ_0 -flake. We have to now consider two cases:

- (a) σ_1 is an original simplex
- (b) σ_1 is a created simplex

Case 2(a). If σ_1 is an original simplex then $\sigma_1 \subseteq \mathcal{P}_0$, and since $\mathcal{P}_0$ is an ϵ -net, $L_{\sigma_1} \geq \epsilon$. Since a flake must have at least three vertices, σ_1 and σ^m must share at least two vertices, and therefore $R_{\sigma^m} \geq \epsilon/2$.

Let $x = \psi_p(x')$ be point inserted to refine ϕ where $x' \in P(\phi, p)$. Using Lemma 8.5.4 and the fact that $R_{\sigma^m} \geq \epsilon/2$, we therefore have

$$\begin{aligned} d_{\mathbb{R}^N}(x, \mathcal{P}) &> (1 - \alpha - 4.5\tilde{\epsilon}_0)R_p(\sigma^m) \\ &\geq (1 - \alpha - 4.5\tilde{\epsilon}_0)R_{\sigma^m} \geq \frac{(1 - \alpha - 4.5\tilde{\epsilon}_0)\epsilon}{2} > \tilde{\mu}_0\epsilon. \end{aligned}$$

where the final inequality follows from Inequality (8.15). Hence the sparsity invariant is maintained on the refinement of ϕ if σ_1 is an original simplex.

Case 2(b) We will now consider the case when σ_1 is a created simplex. We denote by $p(\sigma_1)$ the parent simplex whose refinement gave birth to σ_1 .

We will bound the distance between $x = \psi_p(x')$, where $x' \in P(\phi, p)$, and the point set $\mathcal{P}$. Let x^* denote the point whose insertion killed $p(\sigma_1)$. By definition x^* is a vertex of σ_1 , and hence also of ϕ since $\sigma_1 \subseteq \phi$. We distinguish the following two cases:

Case 2(b)(i) Suppose $p(\sigma_1)$ was a big configuration refined by the application of Rule (1). According to **Case 1**, the lengths of the edges incident to x^* in σ_1 are greater than $\epsilon/2$. Therefore

$$\begin{aligned}
d_{\mathbb{R}^N}(x, \mathcal{P}) &\geq (1 - \alpha - 4.5 \tilde{\epsilon}_0) R_p(\sigma^m) && \text{by Lemma 8.5.4} \\
&\geq \frac{(1 - \delta_0^2)(1 - \alpha - 4.5 \tilde{\epsilon}_0) \Delta_\phi}{2} && \text{by Lemma 8.4.4} \\
&\geq \frac{(1 - \delta_0^2)(1 - \alpha - 4.5 \tilde{\epsilon}_0) \Delta_{\sigma_1}}{2} && \text{since } \sigma_1 \leq \phi \\
&> \frac{(1 - \delta_0^2)(1 - \alpha - 4.5 \tilde{\epsilon}_0) \epsilon}{4} && \text{by Case 1} \\
&> \tilde{\mu}_0 \epsilon && \text{Inequality 8.15,}
\end{aligned}$$

and the sparsity invariant is maintained.

Case 2(b)(ii) Suppose $p(\sigma_1)$ was a bad configuration refined by Rule (2). Thus $p(\sigma_1)$ was either an m -simplex $\sigma_2^m \in \text{star}(q)$ or an $(m+1)$ -simplex $q_{m+1} * \sigma_2^m \in \text{cosph}^{\delta_0}(q)$ with $\sigma_2^m \in \text{star}(q)$.

Consider the elementary weight function $\omega_{\sigma_1} = \omega_\phi|_{\sigma_1}$, where ω_ϕ is the weight function (8.14) identifying ϕ as a member of $\text{cosph}^{\delta_0}(p)$. From Lemma 8.4.1(1), and Lemma 8.4.4 we have that $R_p(\sigma^m) \geq R(\sigma_1, \omega_{\sigma_1})$. We also have that $R(\sigma_1, \omega_{\sigma_1}) \geq \beta R_q(\sigma_2^m)$. Indeed, otherwise $\sigma_1 \setminus \{x^*\}$ would be a hitting set for x^* , contradicting the hypothesis that $p(\sigma_1)$ was refined according to Rule (2) by the insertion of a good point x^* . Thus we have

$$\begin{aligned}
d_{\mathbb{R}^N}(x, \mathcal{P}) &> (1 - \alpha - 4.5 \tilde{\epsilon}_0) R_p(\sigma^m) \\
&\geq (1 - \alpha - 4.5 \tilde{\epsilon}_0) R(\sigma_1, \omega_{\sigma_1}) \\
&\geq (1 - \alpha - 4.5 \tilde{\epsilon}_0) \beta R_q(\sigma_2^m) \\
&\geq \frac{(1 - \delta_0^2)(1 - \alpha - 4.5 \tilde{\epsilon}_0) \beta \Delta_{\sigma_2^m}}{2} && \text{from Lemma 8.4.4} \\
&> \Delta(\sigma_2^m) && \text{from Hypotheses } \mathcal{H}1 \text{ on } \beta \\
&> \tilde{\mu}_0 \epsilon,
\end{aligned}$$

where the last inequality follows from the induction hypothesis. Again the sparsity invariant is maintained after refinement of ϕ .

Case 3 The proof for the case of a bad configuration $\phi = \sigma^m \in \text{star}(p)$ to be refined by Rule (2) is similar to **Case 2**, and the lower bound on the interpoint distances is the same.

This completes the demonstration of the sparsity invariant.

8.5.1.2 The good points invariant

We will now show that the good point invariant is maintained if ϕ is a bad configuration being refined by Rule (2). Without loss of generality, we will assume that ϕ is either equal to $\sigma^m \in \text{star}(p)$ or to $q * \sigma^m \in \text{cosph}^{\delta_0}(p)$, with $\sigma^m \in \text{star}(p)$.

Recall the picking region $P(\phi, p)$ introduced in Definition 8.4.6. We will show that there exists $y \in P(\phi, p)$ such that $x = \psi_p(y)$ is a good point. Let $Y \subseteq P(\phi, p)$ be the set of points that ψ_p maps to a point with a hitting set:

$$Y = \{y \in P(\phi, p) \mid \psi_p(y) \text{ is not a good point}\}.$$

We will show that the volume of $P(\phi, p)$ exceeds the volume of Y . To this end, we will first bound the number of simplices that could hit some point in $\psi_p(Y)$. Then we will bound the volume that each potential hitting set can contribute to Y .

In order to bound the number of hitting sets, we will use the sparsity invariant together with the following result from Chapter 5 (Lemma 5.4.7) to bound the number of points that can be a vertex of a hitting set:

Lemma 8.5.5 (Bound on sparse points) *For a point $p \in \mathcal{M}$ and $R > 0$, let V be a maximal set of points in $B_{\mathbb{R}^N|\mathcal{M}}(p, R)$ such that the smallest interpoint distance is not less than $2r$. There exists ξ that depends only on m , and A that depends on m , such that if $R + r \leq \xi \text{rch}(\mathcal{M})$, then*

$$|V| \leq \frac{1 + A\xi}{1 - A\xi} \left(\frac{R}{r} + 1 \right)^m.$$

We obtain the following bound on the number of hitting sets:

Lemma 8.5.6 *Let $\mathbf{S}(\phi)$ denote the set of simplices contained in $\mathbf{B}^+ \cap \mathcal{P}$ that can hit a point in $\psi_p(Y)$. Then*

$$|\mathbf{S}(\phi)| \leq \frac{E^{m+1}}{2^m}, \quad (8.17)$$

where

$$E \stackrel{\text{def}}{=} 2 \left(\frac{1 + A\xi}{1 - A\xi} \right) (18(\alpha + 2\beta' + 6.5\tilde{\epsilon}_0) + 1)^m. \quad (8.18)$$

Proof Suppose $\sigma \subseteq \mathcal{P}$ is a hitting set of a point $x = \psi_p(y)$, where $y \in P(\phi, p)$, with $|\sigma| = k$ and $k \leq m + 1$. Let $\tilde{\sigma} = x * \sigma$, and let $\omega_{\tilde{\sigma}}$ denote the corresponding hitting map (see Definition 8.4.7). Therefore, we have $R(\tilde{\sigma}, \omega_{\tilde{\sigma}}) < \beta R_p(\sigma^m)$, and it follows from Lemma 8.4.1(2) that $\Delta_{\tilde{\sigma}} \leq \frac{2\beta}{1 - \delta_0^2} R_p(\sigma^m)$. Thus from Lemma 8.5.4 and the Triangle inequality we have $\sigma \subset \mathbf{B}^- \stackrel{\text{def}}{=} B_{\mathbb{R}^N}(c_p(\sigma^m), r^-)$, where

$$r^- = 4.5\tilde{\epsilon}_0 R_p(\sigma^m) + \alpha R_p(\sigma^m) + \frac{2\beta}{1 - \delta_0^2} R_p(\sigma^m).$$

Let $c = \psi_p(c_p(\sigma^m))$. Then using Lemma 8.2.11 and the fact that $R_p(\sigma^m) < \epsilon < \tilde{\epsilon}_0 \text{rch}(\mathcal{M})$ we have

$$\|c_p(\sigma^m) - c\| \leq \frac{2R_p(\sigma^m)^2}{\text{rch}(\mathcal{M})} < 2\tilde{\epsilon} R_p(\sigma^m) \leq 2\tilde{\epsilon}_0 R_p(\sigma^m).$$

Using $\delta_0^2 < 2^4 \tilde{\epsilon}_0$ from Lemma 8.5.1, and $\beta' = \frac{\beta}{1 - 2^4 \tilde{\epsilon}_0}$, and $R_p(\sigma^m) < \epsilon$, we find

$$\begin{aligned} \|c_p(\sigma^m) - c\| + r^- &\leq \alpha R_p(\sigma^m) + 6.5\tilde{\epsilon}_0 R_p(\sigma^m) + \frac{2\beta}{1 - \delta_0^2} R_p(\sigma^m) \\ &\leq (\alpha + 2\beta' + 6.5\tilde{\epsilon}_0) \epsilon \stackrel{\text{def}}{=} R. \end{aligned}$$

Thus $\mathbf{B}^- \subseteq \mathbf{B}^+ \stackrel{\text{def}}{=} B_{\mathbb{R}^N}(c, R)$, and $y \in Y$ if and only if there exists $\sigma \in \mathbf{B}^+ \cap \mathcal{P}$ such that σ hits $\psi_p(y)$.

Using Lemma 8.5.5 we will bound the number of sample points in $\mathbf{B}^+ \cap \mathcal{P}$. Set $r = \frac{\tilde{\mu}_0 \epsilon}{2} = \frac{\epsilon}{18}$ and observe that

$$R + r = \left(\frac{1}{18} + \alpha + 2\beta' + 6.5\tilde{\epsilon}_0 \right) \epsilon \leq (2\beta' + 1)\epsilon \leq \xi \text{rch}(\mathcal{M}),$$

by Hypothesis $\mathcal{H}4$. The sparsity invariant and Lemma 8.5.5 then yields

$$|\mathbf{B}^+ \cap \mathcal{P}| \leq \frac{1 + A\xi}{1 - A\xi} \times \left(\frac{(\alpha + 2\beta' + 6.5\tilde{\epsilon}_0)}{1/18} + 1 \right)^m = \frac{E}{2}.$$

Since the number of k -simplices is less than $\left(\frac{E}{2}\right)^{k+1}$, and the maximum dimension of a hitting set is m , we have $|\mathbf{S}(\phi)| \leq \frac{E^{m+1}}{2^m}$. $\square$

We now turn to the problem of bounding the volume of Y . We will consider the contribution of each $\sigma \in \mathbf{S}(\phi)$. The following definition characterises the set of points in $\mathcal{M}$ that can be hit by σ :

Definition 8.5.7 (Forbidden region) *For a k -simplex σ with vertices in $\mathcal{M}$ with $k \leq m$ and parameter $t < \epsilon$, the forbidden region, $F(\sigma, t)$, is the set of points $x \in \mathcal{M}$ such that $\sigma_1 = x * \sigma$ satisfies the following conditions:*

- $L_{\sigma_1} > \frac{t}{9}$
- σ_1 is a Γ_0 -flake
- there exists an elementary weight function ω_{σ_1} s.t. $R(\sigma_1, \omega_{\sigma_1}) < \beta t$

We will use the following lemma, which is proved in Appendix C.1. It bounds the volume of the set of points that can be hit by a given simplex:

Lemma 8.5.8 (Volume of forbidden region) *Let σ be a k -simplex with vertices on $\mathcal{M}$ and $k \leq m$. If*

1. $\Gamma_0 \leq \frac{1}{B+1}$,
2. $\tilde{\epsilon} \leq \min\{\frac{\xi}{4\beta}, \frac{\Gamma_0^{m+1}}{8\beta}\}$ and
3. $\delta_0^2 \leq \min\{\Gamma_0^{m+1}, \frac{1}{4}\}$,

then

$$\text{vol}(F(\sigma, t)) \leq D \Gamma_0 R_{\sigma}^m,$$

where D depends on m and β .

Lemma 8.5.8, together with Lemma 8.5.6, yields a bound on the set of points Y in the picking region that do not map to a good point:

Lemma 8.5.9 *The volume of the set $Y \subset P(\phi, p)$ of points that do not map to a good point is bounded as follows:*

$$\text{vol}(Y) \leq E^{m+1} \beta^m D \Gamma_0 R_p(\sigma^m)^m.$$

Proof Let $t_0 = R_p(\sigma^m) < \epsilon$. For a given $\sigma \in \mathbf{S}(\phi)$, let $Y_\sigma \subseteq Y$ be the set of points y for which σ hits $x = \psi_p(y)$. Then from Hypotheses $\mathcal{H}0$ to $\mathcal{H}4$ and Lemma 8.5.8, we have

$$\begin{aligned} \text{vol}(Y_\sigma) &\leq \text{vol}(\pi_p(F(\sigma, t_0))) \\ &\leq \text{vol}(F(\sigma, t_0)) && \text{since } \pi_p \text{ is a projection map on } T_p \mathcal{M} \\ &\leq D \Gamma_0 R_\sigma^m. \end{aligned} \tag{8.19}$$

Let $\sigma_1 = x * \sigma$, and let ω_{σ_1} be the corresponding hitting map. From the definition of hitting sets and hitting maps, we have $R_p(\sigma^m) < \epsilon$, and $R(\sigma_1, \omega_{\sigma_1}) < \beta R_p(\sigma^m)$ and σ is Γ_0^k -thick. Define $\omega_\sigma = \omega_{\sigma_1} \upharpoonright_\sigma$. Then, using Lemma 8.4.1 (3) and the fact that $R(\sigma, \omega_\sigma) \leq R(\sigma_1, \omega_{\sigma_1}) < \beta R_p(\sigma^m)$, we have

$$\begin{aligned} R_\sigma &\leq R(\sigma, \omega_\sigma) \left(1 - \frac{\delta_0^2}{\Upsilon(\sigma)}\right)^{-1} \\ &\leq R(\sigma, \omega_\sigma) \left(1 - \frac{\delta_0^2}{\Gamma_0^m}\right)^{-1} && \text{since } \Upsilon(\sigma) \geq \Gamma_0^k \geq \Gamma_0^m \\ &\leq 2R(\sigma, \omega_\sigma) && \text{since } \frac{\delta_0^2}{\Gamma_0^m} \leq \Gamma_0 < \frac{1}{2} \text{ from Hyp. } \mathcal{H}2, \mathcal{H}3 \\ &< 2\beta R_p(\sigma^m). \end{aligned} \tag{8.20}$$

The inequalities (8.19) and (8.20) together yield

$$\text{vol}(Y_\sigma) \leq 2^m \beta^m D \Gamma_0 R_p(\sigma^m)^m,$$

and so using Lemma 8.5.6 we have

$$\text{vol}(Y) = \text{vol}\left(\bigcup_{\sigma \in \mathbf{S}(\phi)} Y_\sigma\right) \leq \sum_{\sigma \in \mathbf{S}(\phi)} \text{vol}(Y_\sigma) \leq E^{m+1} \beta^m D \Gamma_0 R_p(\sigma^m)^m.$$

□

By the definition of the picking region, we have that

$$\text{vol}(P(\phi, p)) = \phi_m \alpha^m R_p(\sigma^m)^m.$$

By Hypothesis $\mathcal{H}2$, $E^{m+1} \beta^m D \Gamma_0 R_p(\sigma^m)^m$ is less than $\text{vol}(P(\phi, p))$, the volume of the picking region of ϕ . Thus with Lemma 8.5.9, this proves the existence of points y in the picking region $P(\phi, p)$ of ϕ such that $\psi_p(y)$ is a good point.

The proof of Theorem 8.5.3 is complete.

8.5.2 Output quality

We will now show that if Hypothesis $\mathcal{H}5$ is satisfied, in addition to Hypotheses $\mathcal{H}0$ to $\mathcal{H}4$, then the output to the refinement algorithm will meet the demands imposed by Theorem 8.3.6, thus yielding Theorem 8.5.2.

The main task is to ensure that every m -simplex in $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ has, for each vertex, a $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power-protected Delaunay ball centred on the tangent space of that vertex. This is achieved in two steps. First we establish conditions to ensure that $\text{cosp}^{\delta_0}(p) = \emptyset$ for every $p \in \mathcal{P}$. As noted by Lemma 8.4.5, this ensures that every simplex in $\text{star}(p)$ has a $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power-protected Delaunay ball centred on $T_p\mathcal{M}$. Next we show conditions such that if $\sigma^m \in \text{Del}_{T\mathcal{M}}(\mathcal{P})$, then $\sigma^m \in \text{star}(p)$ for every vertex $p \in \sigma^m$. In each step the required conditions impose an additional constraint on the sampling radius, and this leads to Hypothesis $\mathcal{H}5$.

As a starting point, we observe the following direct consequence of the Termination Theorem 8.5.3:

Corollary 8.5.10 *Under Hypotheses $\mathcal{H}0$ to $\mathcal{H}4$, for all $p \in \mathcal{P}$, the output of the algorithm satisfies the following:*

1. $\sigma^m \in \text{star}(p) \implies R_p(\sigma^m) < \epsilon$ and σ^m is a Γ_0 -good simplex, and
2. all $\sigma^{m+1} \in \text{cosp}^{\delta_0}(p)$ are Γ_0 -good.

We will show that for an appropriate sampling radius, there cannot be a Γ_0 -good simplex in $\text{cosp}^{\delta_0}(p)$. We exploit the following bound on the thickness of a small $(m+1)$ -simplex:

Lemma 8.5.11 (Small $(m+1)$ -simplices are not thick) *Let σ^{m+1} be an $(m+1)$ -simplex with vertices in $\mathcal{M}$ and $\Delta_{\sigma^{m+1}} < \text{rch}(\mathcal{M})$. For distinct vertices $p, q \in \sigma^{m+1}$ define $\theta = \angle(\text{aff}(\sigma_q), T_p\mathcal{M})$. Then*

$$\Upsilon_{\sigma^{m+1}} \leq \left(\frac{\Delta_{\sigma^{m+1}}}{2 \text{rch}(\mathcal{M})} + \sin \theta \right).$$

Proof We will bound the altitude $D_{\sigma^{m+1}}(q)$. Let ℓ be the line through p and q . Using Lemma 8.2.10 and the fact that $\angle(\text{aff}(\sigma_q), T_p\mathcal{M}) = \theta$, we get

$$\begin{aligned} D_{\sigma^{m+1}}(q) &= d_{\mathbb{R}^N}(q, \text{aff}(\sigma_q)) \\ &= \sin \angle(\ell, \text{aff}(\sigma_q)) \times d_{\mathbb{R}^N}(p, q) \\ &\leq (\sin \angle(\ell, T_p\mathcal{M}) + \sin \angle(\text{aff}(\sigma_q), T_p\mathcal{M})) \times d_{\mathbb{R}^N}(p, q) \\ &\leq \left(\frac{d_{\mathbb{R}^N}(p, q)}{2 \text{rch}(\mathcal{M})} + \sin \theta \right) \times d_{\mathbb{R}^N}(p, q) \\ &\leq \left(\frac{\Delta_{\sigma^{m+1}}}{2 \text{rch}(\mathcal{M})} + \sin \theta \right) \times \Delta_{\sigma^{m+1}}. \end{aligned}$$

Therefore we have

$$\Upsilon_{\sigma^{m+1}} \leq \left(\frac{\Delta_{\sigma^{m+1}}}{2 \text{rch}(\mathcal{M})} + \sin \theta \right).$$

□

Also, Whitney's Lemma 7.2.1 implies that a Γ_0 -good simplex in $\text{star}(p)$ makes a small angle with the tangent space at p :

Lemma 8.5.12 *If $\sigma^m \in \text{star}(p)$ is Γ_0 -good with $R_p(\sigma^m) < \epsilon$, then*

$$\sin \theta < \frac{2\epsilon}{\Gamma_0^m \text{rch}(\mathcal{M})},$$

where $\theta = \angle(\text{aff}(\sigma^m), T_p \mathcal{M})$.

Proof Let $\zeta = \max_{x \in \sigma^m} d_{\mathbb{R}^N}(x, T_p \mathcal{M})$ where x is a vertex of σ^m . From Lemma 8.2.10, we have

$$\zeta = \max_{x \in \sigma^m} d_{\mathbb{R}^N}(x, T_p \mathcal{M}) \leq \max_{x \in \sigma^m} \frac{d_{\mathbb{R}^N}(p, x)^2}{2 \text{rch}(\mathcal{M})} \leq \frac{\Delta_{\sigma^m}^2}{2 \text{rch}(\mathcal{M})}.$$

Using Lemma 7.2.1 and the facts that $R_{\sigma^m} \leq R_p(\sigma^m) < \epsilon$ and $\Upsilon_{\sigma^m} \geq \Gamma_0^m$ (since σ^m is a Γ_0 -good simplex), we have

$$\sin \theta \leq \frac{2\zeta}{\Upsilon_{\sigma^m} \Delta_{\sigma^m}} \leq \frac{\Delta_{\sigma^m}}{\Upsilon_{\sigma^m} \text{rch}(\mathcal{M})} < \frac{2\epsilon}{\Gamma_0^m \text{rch}(\mathcal{M})}.$$

The last inequality follows from the fact that $\Delta_{\sigma^m} \leq 2R_{\sigma^m} < 2\epsilon$. $\square$

Using Lemmas 8.5.11 and 8.5.12 we get that no $(m+1)$ -dimensional simplices in $\text{cosph}^{\delta_0}(p)$ can be Γ_0 -good when ϵ is sufficiently small:

Lemma 8.5.13 ($\text{cosph}^{\delta_0}(p)$ simplices are Γ_0 -bad) *Let $\sigma^{m+1} = p_{m+1} * \sigma^m \in \text{cosph}^{\delta_0}(p)$ with $\sigma^m \in \text{star}(p)$. If*

$$\tilde{\epsilon} \leq \frac{\Gamma_0^{2m+1}}{4},$$

and $\delta_0^2 \leq \frac{1}{2}$, then $\Upsilon_{\sigma^{m+1}} < \Gamma_0^{m+1}$.

Proof By Lemma 8.4.4

$$\Delta_{\sigma^{m+1}} \leq \frac{2}{1 - \delta_0^2} R_p(\sigma^m) < 4\epsilon,$$

since $R_p(\sigma^m) < \epsilon$ and $\delta_0^2 \leq \frac{1}{2}$.

Then from the fact that $\Gamma_0 < 1$, and Lemmas 8.5.11 and 8.5.12, we get

$$\Upsilon_{\sigma^{m+1}} \leq \left(\frac{\Delta_{\sigma^{m+1}}}{2 \text{rch}(\mathcal{M})} + \sin \theta \right) \leq \frac{2\epsilon}{\text{rch}(\mathcal{M})} \left(1 + \frac{1}{\Gamma_0^m} \right) < \frac{4\epsilon}{\Gamma_0^m \text{rch}(\mathcal{M})}.$$

From the hypothesis on $\tilde{\epsilon}$, we get $\Upsilon_{\sigma^{m+1}} < \Gamma_0^{m+1}$. $\square$

We emphasise the consequence of Lemma 8.5.13:

Corollary 8.5.14 *If $\delta_0^2 < \frac{1}{2}$ and*

$$\tilde{\epsilon} \leq \frac{\Gamma_0^{2m+1}}{4},$$

and all the simplices in $\text{cosph}^{\delta_0}(p)$ are Γ_0 -good, then

$$\text{cosph}^{\delta_0}(p) = \emptyset.$$

Now we proceed to the second step of the analysis. Assuming that $\text{cosph}^{\delta_0}(p) = \emptyset$ for all p in $\mathcal{P}$, the following lemma says that if $\sigma \in \text{star}(p)$, then also $\sigma \in \text{star}(q)$ for every vertex $q \in \sigma$, provided the appropriate constraints are met.

Lemma 8.5.15 *Let $\mathcal{P}$ be a $\tilde{\mu}_0\epsilon$ -sparse ϵ -sample of $\mathcal{M}$ with $\tilde{\mu}_0 \leq 1$ independent of ϵ . We further assume $\delta_0 \leq 1$ and*

- (1) *for all $p \in \mathcal{P}$, every $\sigma^m \in \text{star}(p)$ is a Γ_0 -good simplex with $R_p(\sigma^m) < \epsilon$, and*
- (2) *for all $p \in \mathcal{P}$, $\text{cosph}^{\delta_0}(p) = \emptyset$.*

If

$$\tilde{\epsilon} \leq \frac{\delta_0^2 \tilde{\mu}_0^2 \Gamma_0^m}{36},$$

then $\text{star}(p) = \text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ for all p in $\mathcal{P}$.

Proof For $p \in \mathcal{P}$, let $\sigma^m \in \text{star}(p)$ and $q (\neq p)$ be a vertex of σ^m . We will show that σ^m is also in $\text{star}(q)$.

Let $\theta = \max \angle(\text{aff}(\sigma^m), T_x \mathcal{M})$ where the max is taken over the vertices x of σ^m . Since

$$\tilde{\epsilon} \leq \frac{\delta_0^2 \tilde{\mu}_0^2 \Gamma_0^m}{36} < \frac{\Gamma_0^m}{4},$$

Lemma 8.5.12 yields

$$\sin \theta \leq \frac{2\tilde{\epsilon}}{\Gamma_0^m} \stackrel{\text{def}}{=} c_1 \tilde{\epsilon} < \frac{1}{2}.$$

It follows that $\cos \theta > \sqrt{3}/2$ and so

$$\tan \theta \leq 2c_1 \tilde{\epsilon}.$$

Recall that N_{σ^m} denotes the affine space orthogonal to $\text{aff}(\sigma^m)$ and passing through c_{σ^m} . Let c be the unique point in $N_{\sigma^m} \cap T_q \mathcal{M}$, and let $R = d_{\mathbb{R}^N}(c, p)$.

Using the fact that $\angle(\text{aff}(\sigma^m), T_q \mathcal{M}) \leq \theta$, we have

$$d_{\mathbb{R}^N}(c_{\sigma^m}, c) \leq R_{\sigma^m} \tan \theta \leq 2c_1 \tilde{\epsilon} R_{\sigma^m},$$

and likewise

$$d_{\mathbb{R}^N}(c_{\sigma^m}, c_p(\sigma^m)) \leq 2c_1 \tilde{\epsilon} R_{\sigma^m}.$$

It follows that $R \leq (1 + 2c_1 \tilde{\epsilon}) R_{\sigma^m}$, and $d_{\mathbb{R}^N}(c_p(\sigma^m), c) \leq 4c_1 \tilde{\epsilon} R_{\sigma^m}$. From the above observations, and using the fact that $R_{\sigma^m} \leq R_p(\sigma^m) < \epsilon$, we get

$$\begin{aligned} B_{\mathbb{R}^N}(c, R) &\subseteq B_{\mathbb{R}^N}(c_p(\sigma^m), (1 + 6c_1 \tilde{\epsilon}) R_{\sigma^m}) \\ &\subseteq B_{\mathbb{R}^N}(c_p(\sigma^m), R_p(\sigma^m) + 6c_1 \tilde{\epsilon} \epsilon). \end{aligned}$$

Since $\text{cosph}^{\delta_0}(p) = \emptyset$, and $\mathcal{P}$ is $\tilde{\mu}_0\epsilon$ -sparse, we have that σ^m is $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power protected on $T_p \mathcal{M}$ (Lemma 8.4.5). This means that

$$B_{\mathbb{R}^N}(c_p(\sigma^m), R_p(\sigma^m) + \Delta) \cap (\mathcal{P} \setminus \sigma^m) = \emptyset,$$

where

$$\begin{aligned}
\Delta &= \sqrt{R_p(\sigma^m)^2 + \delta_0^2 \tilde{\mu}_0^2 \epsilon^2} - R_p(\sigma^m) \\
&= \frac{\delta_0^2 \tilde{\mu}_0^2 \epsilon^2}{\sqrt{R_p(\sigma^m)^2 + \delta_0^2 \tilde{\mu}_0^2 \epsilon^2} + R_p(\sigma^m)} \\
&> \frac{\delta_0^2 \tilde{\mu}_0^2 \epsilon}{\sqrt{1 + \delta_0^2 \tilde{\mu}_0^2} + 1} \\
&> \frac{\delta_0^2 \tilde{\mu}_0^2 \epsilon}{3} \stackrel{\text{def}}{=} c_2 \epsilon.
\end{aligned}$$

Since $6c_1 \tilde{\epsilon} \epsilon \leq c_2 \epsilon$, by our hypothesis on $\tilde{\epsilon}$, we have

$$B_{\mathbb{R}^N}(c, R) \subset B_{\mathbb{R}^N}(c_p(\sigma^m), R_p(\sigma^m) + \Delta),$$

and thus the m -simplex σ^m belongs to $\text{star}(q)$. $\square$

The consequence of Lemma 8.5.15, together with Lemma 8.4.5 is that every m -simplex in $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ has, for each vertex, a $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power-protected Delaunay ball centred on the tangent space of that vertex:

Corollary 8.5.16 *Let $\mathcal{P}$ be a $\tilde{\mu}_0 \epsilon$ -sparse ϵ -sample of $\mathcal{M}$ with $\tilde{\mu}_0$ being independent of ϵ . Under the hypotheses in Lemma 8.5.15, for all $p \in \mathcal{P}$, all the m -simplices σ^m in $\text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ are $\delta_0^2 \tilde{\mu}_0^2 \epsilon^2$ -power protected on $T_p \mathcal{M}$. I.e, for all $\sigma^m \in \text{star}(p; \text{Del}_{T\mathcal{M}}(\mathcal{P}))$ there exists a $c_p(\sigma^m) \in N_{\sigma^m} \cap T_p \mathcal{M}$ such that for all $q \in \mathcal{P} \setminus \sigma^m$*

$$d_{\mathbb{R}^N}(q, c_p(\sigma^m))^2 > d_{\mathbb{R}^N}(p, c_p(\sigma^m))^2 + \delta_0^2 \tilde{\mu}_0^2 \epsilon^2.$$

We are now in a position to show that Hypothesis $\mathcal{H}5$, when added to Hypotheses $\mathcal{H}0$ to $\mathcal{H}4$, results in the output of the algorithm meeting the demands of Theorem 8.3.6.

Recalling that $\tilde{\mu}_0 = \frac{1}{9}$, Hypotheses $\mathcal{H}3$ yields the following consequence of $\mathcal{H}5$:

$$\tilde{\epsilon} \leq \frac{\delta_0^2 \Gamma_0^{2m}}{1.1 \times 10^9} \leq \min \left\{ \frac{\Gamma_0^{2m+1}}{4}, \frac{\delta_0^2 \tilde{\mu}_0^2 \Gamma_0^m}{36}, \frac{\delta_0^2 \tilde{\mu}_0^3 \Gamma_0^{2m}}{1.5 \times 10^6} \right\}.$$

In other words, the sampling radius bounds demanded by Corollary 8.5.14, Lemma 8.5.15, and Theorem 8.3.6 are all simultaneously satisfied. Corollary 8.5.10 together with Corollary 8.5.14 ensure that the hypotheses of Lemma 8.5.15 are satisfied, and so it follows that the m -simplices of $\text{Del}_{T\mathcal{M}}(\mathcal{P})$ are power-protected as described by Corollary 8.5.16. Thus all the requirements of Theorem 8.3.6 are satisfied, and we obtain Theorem 8.5.2.

8.6 Summary

We have described an algorithm which meshes a manifold according to extrinsic sampling conditions which guarantee that the intrinsic Delaunay complex coincides with the restricted Delaunay complex, and that it is homeomorphic to the manifold. The algorithm constructs the tangential Delaunay complex, which is also shown to be equal to the

intrinsic Delaunay complex, and in this way we are able to exploit existing structural results from Chapter 4 to obtain the homeomorphism guarantee.

This approach relies on an embedding of $\mathcal{M}$ in $\mathbb{R}^N$. In future work we aim to develop algorithms and structural results which enable the construction of an intrinsic Delaunay triangulation in the absence of an embedding in Euclidean space.

Chapter 9

Conclusion

In this thesis, we mainly worked on the algorithmic questions that arise in the field of piecewise linear approximation of smooth submanifolds of Euclidean space. More specifically, we developed algorithms for reconstruction of manifolds from a dense point sample, and sampling and meshing of manifolds. Along the way, we also studied the stability of *Delaunay-type* structures.

We developed algorithms, unlike the current algorithms in manifold reconstruction and meshing, whose complexity depends exponentially on the intrinsic dimension of the manifold rather than the ambient dimension. The central idea behind the algorithms we developed to solve this problem can best be summarized as building *locally and fitting globally*. Rather than building the whole structure that we want at once, we build parts of the main structure locally and hopefully efficiently, and fit these local parts in a correct way globally. Note that this concept has been successfully applied before, see [BWY08, She05]. We applied this principle, in the context of manifold reconstruction and meshing, by defining Delaunay triangulations locally and glueing this local triangulations together by removing inconsistencies among the local triangulation. We give the first algorithm with theoretical guarantee, i.e. the output is homeomorphic to the manifold, for manifold reconstruction whose complexity depends exponentially on the intrinsic dimension, linearly on the ambient dimension and quadratically on the size of the input sample. We also developed the first algorithm for sampling and meshing a submanifold of Euclidean space, according to a prescribed sampling parameter ε , whose asymptotic complexity is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$. The algorithm generates a dense ε -sample of the manifold and a simplicial approximation which is homeomorphic to the manifold.

Liebon and Letscher [LL00] claimed that sampling density of point sample on a manifold alone can guarantee existence of intrinsic Delaunay triangulation. We give an explicit counterexample which shows that density of the sample points alone cannot guarantee that the intrinsic Delaunay complex is homeomorphic to the manifold. Cheng et al. [CDR05a] and Boissonnat et al. [BGO09, Lemma 3.1] presented an example which showed that the restricted Delaunay complex is not homeomorphic to the original manifold even if the point sample is dense and well separated. We build on this example but from the perspective of the intrinsic metric of the manifold, to develop our counterexample.

Our study on the stability of Delaunay-type structure was motivated by the counterexample to Liebon and Letscher [LL00] claimed results. Liebon and Letscher [LL00] missed the problem that arises when we have close to degenerate configuration in the

point sample. The key observation in our work was to realize that Delaunay triangulation is unstable around cospherical configurations which makes it difficult to extend Delaunay triangulations on domains where the metric is unstable. To address this problem we introduce parametrized notion of genericity for Delaunay triangulation and using this framework studied the stability of Delaunay triangulations under perturbations of the metric and of the vertex positions. We then show that, for any sufficiently regular submanifold of Euclidean space, and appropriate ϵ and δ , any sample set which meets a localized δ -generic ϵ -dense sampling criteria yields an intrinsic Delaunay triangulation. Finally, we give an algorithm for generating δ -generic point sets.

We will now discuss some of the open questions and extensions of this work. In Chapter 3, we assume that the underlying manifold that we are reconstructing from the point sample has no boundaries and the point sample either lies on the manifold or very close to the manifold (small Hausdorff noise). Other works on manifold reconstruction also suffer from the same flaws [BG10a, BGO09, CL08, CDR05a]. The main hurdle in extending the current methods comes from the fact that it is difficult to identify the points that are outliers or the points that lie on the boundary. Recent progress has been made on both fronts. In [DLRW09], Dey, Li, Ramos and Wenger gave the first provably correct algorithm for the reconstruction of two-dimensional surfaces with boundaries embedded in $\mathbb{R}^3$. Even though their reconstruction algorithm is quite elegant but it relies very heavily on the fact the codimension of two-dimensional surface is one. And this hinders its extension to the general case. The notion of distance to a probability distribution in $\mathbb{R}^d$, introduced in [CCSM11], have been successfully used to recover geometric and topological features from the point sample [CCSM11]. We think that the notion of distance to a probability distribution in $\mathbb{R}^d$ can be used to have a provably correct reconstruction algorithm that will take care of both the presence of outliers and boundaries.

In Chapter 5, we gave an algorithm for meshing submanifold of Euclidean space with positive reach. The requirement that the reach of the submanifold to be greater than zero can be relaxed to include Lipschitz manifolds. This has already been for Lipschitz surfaces in $\mathbb{R}^3$ [BO06]. The asymptotic time complexity of our meshing algorithm is $O(\epsilon^{-k^2-k})$ but the algorithm begins by first computing a crude sample of the manifold by using a grid (other methods can also be used to compute this sample), which will add a constant to the asymptotic complexity which depends exponentially on the ambient dimension, see Section 5.3.2 of Chapter 5. We expect that this additive constant to the asymptotic time complexity can be removed from the exact time complexity.

In the same vein as Gruber [Gru93], Clarkson [Cla06] claimed for a general manifold, under some general conditions, a tight bound on the minimum Hausdorff distance for a mesh with n simplices to the manifold when $n \rightarrow \infty$. The construction in the paper [Cla06] was for a limited setting and relied on an invalid implication of Leifson and Letscher's work [LL00]. In Chapter 6, we gave a counterexample which showed that the claimed results in [LL00] are false. Therefore the problem of Hausdorff distance between $\mathcal{M}$ and a mesh approximating $\mathcal{M}$, in the limiting case, addressed in [Cla06] still remains open.

In Chapter 8, we provide sampling conditions for smooth submanifold of Euclidean space, that depends on the reach (which is an extrinsic property), under which the intrinsic Delaunay triangulation is well defined and homeomorphic to the manifold. The proofs use the fact that the manifold is a smooth submanifold of Euclidean space. Since

we are interested in the intrinsic Delaunay triangulations, which depends on the intrinsic metric, these sampling conditions that depend on extrinsic property of the manifold are not satisfactory. We want to get sampling conditions that depend on the intrinsic properties of the manifold, like injectivity radius or sectional curvature.

Appendix A

Appendix for Chapter 2

A.1 Whitney's proof of Lemma 2.3.2

In this appendix we will give the details of the Whitney's proof of Lemma 2.3.2. Rather than using *alternating forms*, as in Whitney's proof, we will be working with determinants.

We will be using the following result to prove Lemma A.1.1.

Lemma A.1.1 *Let $\tau = [p_0, p_1, \dots, p_j]$ be a j -dimensional simplex, and let $u_i = \frac{p_i - p_0}{\|p_i - p_0\|}$ for all $i \in S = \{1, \dots, j\}$. Then $\left\| \sum_{i=1}^j \lambda_i u_i \right\| \geq j! \Theta_\tau \times \max_{i \in S} (|\lambda_i|)$.*

Proof Before we prove the result we need the following claim.

Claim A.1.2 *Let $u_1, \dots, u_j$ are unit vectors, $\lambda_1, \dots, \lambda_j$ are real numbers, and k -simplex $\tau = [0, q_1, \dots, q_j]$ where $q_i = 0 + u_i$ and $i \in \{1, \dots, j\}$. Then $\left\| \sum_{i=1}^j \lambda_i u_i \right\| \geq |\det(u_1 \dots u_j)| \times \max_i (|\lambda_i|) = j! \text{vol}(\tau) \times \max_i (|\lambda_i|)$.*

Proof We will prove this result by contradiction. Lets assume that there exist unit vectors $u_1, \dots, u_j$ and real numbers $\lambda_1, \dots, \lambda_j$ s.t $\left\| \sum_{i=1}^j \lambda_i u_i \right\| < |\det(u_1 \dots u_j)| \times \max_i (|\lambda_i|)$.

Without loss of generality assume that $j = \arg \max_i (|\lambda_i|)$ and set $w = \sum_{i=1}^{j-1} \mu_i u_i + \lambda_j u_j$ where $\mu_1, \dots, \mu_{j-1}$ are chosen so as to minimize $\|w\|$. Then

$$\|w\| = \left\| \sum_{i=1}^j \lambda_i u_i \right\| \stackrel{\text{hyp}}{<} \lambda |\det(u_1 \dots u_j)| = |\det(u_1 \dots w)| \leq \prod_{i=1}^{j-1} \|u_i\| \times \|w\| = \|w\|,$$

a contradiction. Using the fact that $\text{vol}(\tau) = \frac{|\det(u_1 \dots u_j)|}{j!}$ we get the full lemma. $\square$

Without loss of generality we can assume that $\tau = [p_0, \dots, p_j]$ is embedded in $\mathbb{R}^j$. Let $\lambda = \max_i (|\lambda_i|)$. Using Claim A.1.2 we get

$$\left\| \sum_{i=1}^j \lambda_i u_i \right\| \geq \lambda |\det(u_1 \dots u_j)| = \frac{\lambda |\det(v_1 \dots v_j)|}{\|v_1\| \dots \|v_j\|} \geq \frac{\lambda j! \text{vol}(\tau)}{\Delta^j(\tau)} = \lambda j! \Theta_\tau.$$

$\square$




Figure A.1

Using Lemma A.1.1, we can now get a proof of Lemma 2.3.2.

Proof of Lemma 2.3.2 Let π_H denote the orthogonal projection onto H . Write $v_i = p_i - p_0$, $u_i = \frac{v_i}{\|v_i\|}$, and let $\bar{u}_i$ denote the orthogonal projection of u_i onto H , for $i \in \{0, \dots, p_j\}$. We have

$$\begin{aligned} \|u_i - \bar{u}_i\| &\leq \frac{\|p_i - p_0 - (\pi_H(p_i) - \pi_H(p_0))\|}{\|v_i\|} \\ &\leq \frac{\|p_i - \pi_H(p_i)\| + \|p_0 - \pi_H(p_0)\|}{\|v_i\|} \leq \frac{2\eta}{L_\tau}. \end{aligned}$$

The last inequality follows from the fact that $\|\pi_H(p_i) - p_i\| \leq \eta$ for all $i \in \{0, \dots, j\}$ and $\|v_i\| = \|p_i - p_0\| \geq L_\tau$.

Let $u = \sum_{i=1}^j \lambda_i u_i$. We deduce from Lemma A.1.1 that $|\lambda_i| \leq \frac{\|u\|}{j! \Theta_\tau} = \frac{1}{j! \Theta_\tau}$ since $\|u\| = 1$. Hence

$$\left\| u - \sum_{i=1}^j \lambda_i \bar{u}_i \right\| = \left\| \sum_{i=1}^j \lambda_i (u_i - \bar{u}_i) \right\| \leq \frac{j}{j! \Theta_\tau} \times \frac{2\eta}{L_\tau} = \frac{2\eta}{(j-1)! \Theta_\tau L_\tau}. \quad (\text{A.1})$$

Let u_H be the unit vector along $\sum_{i=1}^j \lambda_i \bar{u}_i$. Then from inequality A.1 we have $\sin \angle(u, u_H) = \frac{2\eta}{(j-1)! \Theta_\tau L_\tau}$. $\square$

Appendix B

Appendix for Chapter 5

B.1 Proof of Lemmas 5.2.3 and 5.2.5

Proof of Lemma 5.2.3 Proof is similar to Lemma 3.4.2 in Chapter 3. Assume for a contradiction that there exists a point $x \in \text{Vor}(p) \cap T_p \mathcal{M}$ s.t. $\|p - x\| > 4\varepsilon \text{rch}(\mathcal{M})$. Let q be a point on the line segment $[px]$ s.t. $\|p - q\| = 2\varepsilon \text{rch}(\mathcal{M})$. Let q' be the nearest to q on $\mathcal{M}$. From Lemma 2.2.2 (2), we have $\|q - q'\| \leq 8\varepsilon^2 \text{rch}(\mathcal{M})$. Since P is an ε -sample, there exists a point $t \in P$, s.t. $\|q' - t\| \leq \varepsilon \text{rch}(\mathcal{M})$. We thus have

$$\|q - t\| \leq \|q - q'\| + \|q' - t\| \leq 8\varepsilon^2 \text{rch}(\mathcal{M}) + \varepsilon \text{rch}(\mathcal{M}) < 2\varepsilon \text{rch}(\mathcal{M}),$$

the last inequality follows from the fact that $\varepsilon \leq 1/8$.

From Eq. (B.1) we get $p \neq t$, as $\|p - q\| = 2\varepsilon \text{rch}(\mathcal{M})$ and $\|t - q\| < 2\varepsilon \text{rch}(\mathcal{M})$. We can see that $\angle ptx > \pi/2$. This implies that

$$\|x - p\|^2 - \|x - t\|^2 > \|p - t\|^2 > 0,$$

the last inequality follows from the fact that $p \notin t$. This implies $x \notin \text{Vor}(p)$, which contradicts our initial assumption. We conclude that $\text{Vor}(p) \cap T_p \mathcal{M} \subseteq B(p, 4\varepsilon \text{rch}(\mathcal{M}))$ (i). (ii) and (iii) are easy consequences of (i). $\square$

Proof of Lemma 5.2.5 1. We have $c_p(\tau) = \text{Vor}(\tau) \cap T_p \mathcal{M}$, $c_q(\tau) = \text{aff}(\text{Vor}(\tau)) \cap T_q \mathcal{M}$, and $R_p(\tau) = \|c_p(\tau) - p\|$ and $R_q(\tau) = \|c_q(\tau) - q\|$. Since $\theta = \max_x \theta_x$ where $\theta_x = \angle(\text{aff}(\tau), T_x \mathcal{M})$ and x is a vertex of τ , we have $R_{p'}(\tau) \leq R_\tau / \cos \theta$ and $\|c_{p'}(\tau) - c_\tau\| \leq R_\tau \tan \theta$, for $p' \in \{p, q\}$. As $i_\phi \in [c_p(\tau), c_q(\tau)]$, we have $\|i_\phi - c_\tau\| \leq R_\tau \tan \theta$. Then, by Pythagoras theorem, we have

$$\widetilde{R}_\phi = \sqrt{R_\tau^2 + \|i_\phi - c_\tau\|^2} \leq R_\tau \sqrt{1 + \tan^2 \theta} = R_\tau / \cos \theta.$$

2. We will now bound $\Theta_\phi = \frac{\text{vol}(\phi)}{\Delta_\phi^{k+1}}$. We use $\text{vol}(\phi) = \frac{D_\phi(r) \text{vol}(\tau)}{k+1}$ and bound $D_\phi(r)$ and $\text{vol}(\tau)$.

Using the fact that $\Delta_\phi \leq 2R_\phi \leq \frac{2R_\tau}{\cos\theta}$ from 1, we have

$$\begin{aligned}
 D_\phi(r) &= \text{dist}(r, \text{aff}(\tau)) \\
 &= \sin \angle(pr, \text{aff}(\tau)) \times \|p - r\| \\
 &\leq (\sin \angle(pr, T_p\mathcal{M}) + \sin \angle(\text{aff}(\tau), T_p\mathcal{M})) \times \Delta_\phi \\
 &\leq \left(\frac{\|p - r\|}{2\text{rch}(\mathcal{M})} + \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau \text{rch}(\mathcal{M})} \right) \Delta_\phi \\
 &\leq \frac{\Delta_\phi^2}{2\text{rch}(\mathcal{M})} \left(1 + \frac{4\rho_\tau}{\Theta_\tau} \right)
 \end{aligned} \tag{B.1}$$

From the definition of fatness of a simplex and Lemma 2.3.1 (1), we get

$$\text{vol}(\tau) = \Theta_\tau \Delta_\tau^k \leq \frac{\Delta_\tau^k}{k!}. \tag{B.2}$$

Using inequalities (B.1) and (B.2), and $\Delta_\tau \leq \Delta_\phi \leq 2R_\phi \leq \frac{2R_\tau}{\cos\theta}$, we get

$$\begin{aligned}
 \Theta_\phi &= \frac{\text{vol}(\phi)}{\Delta_\phi^{k+1}} \\
 &= \frac{D_\phi(r) \text{vol}(\tau)}{k+1} \times \frac{1}{\Delta_\phi^{k+1}} \\
 &\leq \frac{\Delta_\phi^2}{2\text{rch}(\mathcal{M})} \left(1 + \frac{4\rho_\tau}{\Theta_\tau} \right) \times \frac{\Delta_\tau^k}{(k+1)! \Delta_\phi^{k+1}} \\
 &\leq \frac{R_\tau}{\text{rch}(\mathcal{M}) \cos\theta} \left(1 + \frac{4\rho_\tau}{\Theta_\tau} \right)
 \end{aligned}$$

□

B.2 Proof of Lemma 5.4.1

B.2.1 Geodesic curves and balls

Recall that the *geodesic (intrinsic) distance* $d_{\mathcal{M}}(p, q)$ between points $p, q \in \mathcal{M}$ is $\inf |\gamma_{pq}|$ where the infimum is taken over all the geodesic curves γ_{pq} connecting p and q , and recall that the *intrinsic ball* of radius r at a point $p \in \mathcal{M}$ is defined as

$$B_{\mathcal{M}}(p, r) = \{q \in \mathcal{M} : d_{\mathcal{M}}(p, q) < r\}.$$

We get from Proposition 6.3 in [NSW08b] the following lemma.

Lemma B.2.1 *Let $p, q \in \mathcal{M}$, $\|p - q\| \leq t \text{rch}(\mathcal{M})$, and $t \leq \frac{1}{2}$. Then,*

$$d_{\mathcal{M}}(p, q) \leq \frac{\|p - q\|}{1 - t},$$

and this implies $B(p, r) \cap \mathcal{M} \subseteq B_{\mathcal{M}}(p, r/(1 - t))$.

Proof From Proposition 6.3 of [NSW08b], we have for $\|p - q\| \leq \text{rch}(\mathcal{M})/2$

$$d_{\mathcal{M}}(p, q) \leq \text{rch}(\mathcal{M}) \times \left(1 - \sqrt{1 - \frac{2\|p - q\|}{\text{rch}(\mathcal{M})}} \right). \quad (\text{B.3})$$

Using the fact that $\|p - q\| \leq t \text{rch}(\mathcal{M})$ and inequality (B.3), we get

$$d_{\mathcal{M}}(p, q) \leq \frac{2\|p - q\|}{1 + \sqrt{1 - 2t}} \leq \frac{\|p - q\|}{1 - t}$$

The second statement of the lemma is a direct consequence of the first one. $\square$

B.2.2 Injectivity radius and reach

Let γ be a geodesic curve starting at a point $p \in \mathcal{M}$. A *cut point* on γ is the first point of γ where γ stops minimizing the distance to p . The *cut locus* $CL(p)$ of a point p is the set of cut points of all geodesic curves of $\mathcal{M}$ starting at p . The injectivity radius $\text{inj}(p)$ at point p is defined as

$$\text{inj}(p) = \inf_{q \in CL(p)} d_{\mathcal{M}}(p, q). \quad (\text{B.4})$$

The *injectivity radius* $\text{inj}(\mathcal{M})$ of $\mathcal{M}$ is defined as

$$\text{inj}(\mathcal{M}) = \inf_{p \in \mathcal{M}} \text{inj}(p). \quad (\text{B.5})$$

In this section, we will bound the injectivity radius $\text{inj}(\mathcal{M})$ in terms of the reach $\text{rch}(\mathcal{M})$ of the manifold. We need first to recall the definition of the sectional curvature of a manifold. Given a point $p \in \mathcal{M}$ and two linearly independent vectors $u, v \in T_p\mathcal{M}$, the *sectional curvature* is defined as

$$\mathcal{K}(p, u, v) = \frac{\langle R(u, v)v, u \rangle}{\langle u, u \rangle \langle v, v \rangle - \langle u, v \rangle^2} \quad (\text{B.6})$$

where $\langle \cdot, \cdot \rangle$ is the metric tensor and $R(\cdot)$ is the Riemann curvature tensor.

The following theorem is due to Cheeger et al. [CGT82, Theorem 4.7]. See also [AM97].

Theorem B.2.2 *Assume that $\mathcal{M}$ is a connected, complete Riemannian k -manifold such that $\lambda \leq \mathcal{K}(p, u, v) \leq \Lambda$ for all $p \in \mathcal{M}$ and independent vectors u and v in $T_p\mathcal{M}$. If $\Lambda > 0$ and $0 < r < \pi/(4\sqrt{\Lambda})$, then*

$$\text{inj}(p) \geq r \frac{\text{vol}(B_{\mathcal{M}}(p, r))}{\text{vol}(B_{\mathcal{M}}(p, r)) + V_{\lambda}^k(2r)}, \quad (\text{B.7})$$

where $V_{\lambda}^k(\rho)$ denotes the volume of a ball of radius ρ in the k -dimensional space M_{λ}^k with constant sectional curvature λ .

In order to apply this theorem, we need to bound $\mathcal{K}(p, u, v)$, $V_{\lambda}^k(2r)$ and $\text{vol}(B_{\mathcal{M}}(p, r))$. This will be done in Lemmas B.2.3, B.2.4 and B.2.5 respectively.

Lemma B.2.3 ([DL09]) *If $\mathcal{M}$ is a submanifold of $\mathbb{R}^d$ with reach $\text{rch}(\mathcal{M})$, then*

$$\sup_{p,u,v} |\mathcal{K}(p,u,v)| \leq \frac{2}{\text{rch}^2(\mathcal{M})} \stackrel{\text{def}}{=} \lambda_0.$$

Lemma B.2.4 *For $\lambda \geq 0$ and $r \leq \frac{1}{\sqrt{\lambda}}$, we have $V_{-\lambda}^k(r) \leq \phi_k(1 + a\lambda r^2)^{k-1} r^k$, where a is an absolute constant.*

Proof 1. It is known (see [Ber90]) that

$$V_{-\lambda}^k(r) = k \phi_k \int_0^r s(\lambda, x)^{k-1} dx \quad (\text{B.8})$$

where

$$s(\lambda, x) = \begin{cases} \frac{\sin(x\sqrt{\lambda})}{\sqrt{\lambda}} & \text{if } \lambda > 0; \\ t & \text{if } \lambda = 0; \\ \frac{\sinh(x\sqrt{|\lambda|})}{\sqrt{|\lambda|}} & \text{if } \lambda < 0. \end{cases}$$

and ϕ_k is the volume of the k -dimensional unit Euclidean ball.

2. For $0 \leq x \leq 1$, we have

$$\begin{aligned} \frac{\sinh(x)}{x} &= \sum_{i=0}^{\infty} \frac{x^{2i}}{(2i+1)!} \leq 1 + x^2 \sum_{i=1}^{\infty} \frac{1}{(2i+1)!} \\ &= 1 + x^2(\sinh(1) - 1) \stackrel{\text{def}}{=} 1 + ax^2 \end{aligned} \quad (\text{B.9})$$

3. Observing that $r\sqrt{\lambda} \leq 1$ by assumption, we deduce from equation (B.8) and inequality (B.9)

$$\begin{aligned} V_{-\lambda}^k(r) &= k \phi_k \int_0^r \left(\frac{\sinh(\sqrt{\lambda} x)}{\sqrt{\lambda}} \right)^{k-1} dx \\ &\leq k \phi_k \int_0^r (1 + a\lambda x^2)^{k-1} x^{k-1} dx \\ &\leq \phi_k (1 + a\lambda r^2)^{k-1} r^k \end{aligned}$$

□

Lemma B.2.5 ([NSW08b]) *Let $\mathcal{M}$ be a k -dimensional submanifold of $\mathbb{R}^d$ with reach $\text{rch}(\mathcal{M})$ and let p be a point on $\mathcal{M}$. Then, for $r < \frac{\text{rch}(\mathcal{M})}{2}$, we have $\text{vol}(B(p, r) \cap \mathcal{M}) \geq \phi_k r^k \cos^k \theta$ where $\theta = \arcsin\left(\frac{r}{2\text{rch}(\mathcal{M})}\right)$.*

Using the three above lemmas and Theorem B.2.2, we get a lower bound of $\text{inj}(\mathcal{M})$ in terms of $\text{rch}(\mathcal{M})$ as stated in the following lemma.

Lemma B.2.6 *Let $\mathcal{M}$ be a k -dimensional submanifold of $\mathbb{R}^d$ with reach $\text{rch}(\mathcal{M})$. Then $\text{inj}(\mathcal{M}) \geq \xi_1 \text{rch}(\mathcal{M})$, where ξ_1 only depends on k .*

Proof From Lemma B.2.3, we have, for all $p \in \mathcal{M}$ and independent vectors u and v in $T_p\mathcal{M}$, $-\lambda_0 \leq \mathcal{K}(p, u, v) \leq \lambda_0$, where $\lambda_0 = \frac{\sqrt{2}}{\text{rch}(\mathcal{M})}$. Let $r = t \text{rch}(\mathcal{M})$ with $t \leq 1/2$ and observe that $2r\sqrt{\lambda_0} \leq \sqrt{2}$.

1. We can apply Lemma B.2.4 to get

$$\begin{aligned} V_{-\lambda_0}^k(2r) &\leq \phi_k(1 + 4a\lambda_0 r^2)^{k-1}(2r)^k \\ &\leq 2^k(1 + 2a)^{k-1}\phi_k r^k \stackrel{\text{def}}{=} \zeta' r^k \end{aligned} \quad (\text{B.10})$$

2. By Lemma B.2.1, we have for any point $p \in \mathcal{M}$, $B(p, (1-t)r) \cap \mathcal{M} \subseteq \mathcal{B}_{d_{\mathcal{M}}}(p, r)$. It follows that

$$\begin{aligned} \text{vol}(B_{\mathcal{M}}(p, r)) &\geq \text{vol}(B(p, (1-t)r) \cap \mathcal{M}) \\ &\geq \text{vol}(B(p, \frac{3r}{4}) \cap \mathcal{M}) \\ &\geq \frac{\phi_k 3^k r^k \cos^k \theta'}{4^k} \stackrel{\text{def}}{=} \zeta r^k \end{aligned}$$

where $\theta' = \arcsin\left(\frac{3r/4}{2\text{rch}(\mathcal{M})}\right) < \arcsin\left(\frac{3}{16}\right)$.

3. Since $r \leq \frac{\text{rch}(\mathcal{M})}{2} \leq \frac{\sqrt{2}}{2\sqrt{\lambda_0}} < \frac{\pi}{4\sqrt{\lambda_0}}$, we have by Theorem B.2.2, Lemma B.2.5 and inequality (B.10)

$$\text{inj}(p) \geq \frac{\text{rch}(\mathcal{M})}{4} \left(1 + \frac{\zeta'}{\zeta}\right)^{-1} \stackrel{\text{def}}{=} \xi_1 \text{rch}(\mathcal{M}).$$

The same lower bound plainly holds for $\text{inj}(\mathcal{M}) = \inf_{p \in \mathcal{M}} \text{inj}(p)$. $\square$

B.2.3 Proof of Lemma 5.4.1

Once the injectivity radius of $\mathcal{M}$ is bounded, we can apply the following theorem from Differential Geometry that bounds the volume of intrinsic balls. Refer to [Gra90].

Theorem B.2.7 (The Bishop-Günther inequalities) *Let $\mathcal{M}$ be a complete k -dimensional Riemannian manifold and assume that $r \leq \text{inj}(\mathcal{M})$. Assume that there exists two constants λ and Λ such that $\lambda \leq \mathcal{K}(p, u, v) \leq \Lambda$ for all $p \in \mathcal{M}$ and independent vectors u and v in $T_p\mathcal{M}$. Then*

$$V_{\Lambda}^k(r) \leq \text{vol}(B_{\mathcal{M}}(p, r)) \leq V_{\lambda}^k(r).$$

We can now prove Lemma 5.4.1 using Lemmas B.2.1, B.2.3 and Theorem B.2.7.

Proof of Lemma 5.4.1 Let $r = t \text{rch}(\mathcal{M})$ and $t \leq \min(\frac{1}{1+\sqrt{2}}, \frac{\xi_1}{2}) \leq \frac{1}{2}$, where ξ_1 is the constant defined in Lemma B.2.6.

1. From Lemma B.2.5, we have

$$\begin{aligned} \text{vol}(B(p, r) \cap \mathcal{M}) &\geq \phi_k r^k \left(1 - \frac{r^2}{4\text{rch}(\mathcal{M})^2}\right)^{\frac{k}{2}} \\ &= \phi_k r^k \left(1 - \frac{t^2}{4}\right)^{\frac{k}{2}} \geq \phi_k r^k \left(1 - \frac{k}{8} t^2\right) \end{aligned} \quad (\text{B.11})$$

2. Since $r \leq \frac{r}{1-t} < 2r = 2t \text{rch}(\mathcal{M}) \leq \xi \text{rch}(\mathcal{M}) \leq \text{inj}(\mathcal{M})$ (from Lemma B.2.6), we can apply Theorem B.2.7. We can also apply Lemma B.2.4 since $\frac{r\sqrt{\lambda_0}}{1-t} = \frac{t\sqrt{2}}{1-t} \leq 1$.

$$\begin{aligned} \text{vol}(B(p, r) \cap \mathcal{M}) &\leq \text{vol}(B_{\mathcal{M}}(p, r/(1-t))) && \text{Lemma B.2.1} \\ &\leq V_{-\lambda_0}^k(r/(1-t)) && \text{Theorem B.2.7} \\ &\leq \phi_k \left(1 + a\lambda_0 \frac{r^2}{(1-t)^2}\right)^{k-1} \frac{r^k}{(1-t)^k} && \text{Lemma B.2.4} \\ &= \phi_k \left(1 + \frac{2at^2}{(1-t)^2}\right)^{k-1} \frac{r^k}{(1-t)^k}. \end{aligned} \quad (\text{B.12})$$

Observe that from inequalities (B.11) and (B.12), we deduce that there exists ξ and A that depends only on k such that for $t \leq \xi$, we have

$$0 < 1 - At \leq \frac{\text{vol}(B(p, r) \cap \mathcal{M})}{\phi_k r^k} \leq 1 + At.$$

□

Appendix C

Appendix for Chapter 8

C.1 Forbidden volume calculation

In this appendix we demonstrate:

Lemma 8.5.8 (Volume of forbidden region) *Let σ be a k -simplex with vertices on $\mathcal{M}$ and $k \leq m$. If*

1. $\Gamma_0 \leq \frac{1}{B+1}$,
2. $\tilde{\epsilon} \leq \min\{\frac{\xi}{4\beta}, \frac{\Gamma_0^{m+1}}{8\beta}\}$ and
3. $\delta_0^2 \leq \min\{\Gamma_0^{m+1}, \frac{1}{4}\}$,

then

$$\text{vol}(F(\sigma, t)) \leq D \Gamma_0 R_\sigma^m,$$

where D depends on m and β .

We will use the following lemmas in the proof of Lemma 8.5.8:

Lemma C.1.1 (Triangle altitude bound) *For any non-degenerate triangle $\sigma = [p, q, r]$, we have*

$$D_\sigma(p) = \frac{\|p - q\| \|p - r\|}{2R_\sigma}.$$

Proof Let $\alpha = \angle prq$ and observe that

$$\sin \alpha = \frac{\|p - q\|}{2R_\sigma}.$$

Since $D_\sigma(p) = \|p - r\| \sin \alpha$, the result follows. $\square$

Lemma C.1.2 *Let $\sigma = [p_0 \dots p_k] \subset \mathbb{R}^N$ be a k -simplex with $1 \leq k \leq m < N$. Suppose $p_{k+1} \in \mathbb{R}^N$ is such that $\sigma_1 = p_{k+1} * \sigma$ admits an elementary weight function $\omega_{\sigma_1} : \mathring{\sigma}_1 \rightarrow [0, \infty)$, and the following conditions are satisfied:*

1. $L_{\sigma_1} > \frac{t}{9}$,

2. $R(\sigma_1, \omega_{\sigma_1}) < \beta t$,
3. σ_1 is a Γ_0 -flake, and
4. $\delta_0^2 \leq \min\{\Gamma_0^{m+1}, \frac{1}{4}\}$.

Then

$$d_{\mathbb{R}^N}(p_{k+1}; \partial S') \leq B \Gamma_0 R_\sigma$$

where $S' = B_{\mathbb{R}^N}(c_\sigma, R_\sigma) \cap \text{aff}(\sigma)$ and

$$B \stackrel{\text{def}}{=} 4 + 96\beta(1 + 2^7 3^2 \beta^2).$$

Proof Let $\omega_\sigma = \omega_{\sigma_1}|_{\sigma}$. Note that $\omega_\sigma : \sigma \rightarrow [0, \infty)$ is an elementary weight function, and $C(\sigma, \omega_\sigma)$ is the orthogonal projection of $C(\sigma_1, \omega_{\sigma_1})$ onto $\text{aff}(\sigma)$.

From Lemma 8.4.1 (2) and the fact that $\delta_0^2 \leq \frac{1}{4}$, we have

$$\Delta_{\sigma_1} \leq \frac{2}{1 - \delta_0^2} R(\sigma_1, \omega_{\sigma_1}) < \frac{8}{3} R(\sigma_1, \omega_{\sigma_1}). \quad (\text{C.1})$$

and

$$\begin{aligned} \frac{R(\sigma, \omega_\sigma)}{R(\sigma_1, \omega_{\sigma_1})} &\geq \frac{(1 - \delta_0^2) \Delta_\sigma}{2R(\sigma_1, \omega_{\sigma_1})} && \text{from Lemma 8.4.1 (2)} \\ &\geq \frac{3L_\sigma}{8R(\sigma_1, \omega_{\sigma_1})} && \text{as } \delta_0^2 \leq \frac{1}{4} \text{ and } L_\sigma \leq \Delta_\sigma \\ &\geq \frac{1}{24\beta} && (\text{C.2}) \end{aligned}$$

Therefore, from Lemma 8.2.6, we have

$$\begin{aligned} \frac{D_{\sigma_1}(p_{k+1})}{\Delta_\sigma} &< \left(1 + \frac{1}{k}\right) \Gamma_0 \times \frac{\Delta(\sigma_1)^2}{L_{\sigma_1} \Delta_\sigma} \\ &\leq 2\Gamma_0 \times \frac{\Delta_{\sigma_1}^2}{L_{\sigma_1}^2} && \text{from } k \geq 1 \text{ and } L_{\sigma_1} \leq \Delta_\sigma \\ &< \frac{128\Gamma_0}{9} \times \frac{R(\sigma_1, \omega_{\sigma_1})^2}{L_{\sigma_1}^2} && \text{from Eq. (C.1)} \\ &< 2^7 3^2 \beta^2 \times \Gamma_0 && \text{from hyp. 1 \& 2} \quad (\text{C.3}) \end{aligned}$$

Let p be the point closest to p_{k+1} in $\partial B_{\mathbb{R}^N}(C, R)$ where $C = C(\sigma_1, \omega_{\sigma_1})$ and $R = R(\sigma_1, \omega_{\sigma_1})$. We have

$$\|p - p_{k+1}\| = \sqrt{R^2 + \omega_{\sigma_1}(p_{k+1})^2} - R \leq \omega_{\sigma_1}(p_{k+1}) \leq \delta_0 L(\sigma_1) \quad (\text{C.4})$$

Let q be the point closest to p on $\partial B_{\mathbb{R}^N}(C; R) \cap \text{aff}(\sigma)$, p' be the projection of p onto $\text{aff}(\sigma)$, and let r denotes the intersection of the line $\text{aff}([q C(\sigma, \omega_\sigma)])$ with $\partial B_{\mathbb{R}^N}(C; R)$. Note that $C(\sigma_1, \omega_{\sigma_1})$, $C(\sigma, \omega_\sigma)$, p_{k+1} , p , p' , q and r lie on the same 2-dimensional affine space.

Using the fact that $\|p - p_{k+1}\| \leq \delta_0 L_{\sigma_1}$, we get

$$\|p - p'\| \leq D(p_{k+1}, \sigma_1) + \delta_0 L_{\sigma_1} \quad (\text{C.5})$$

We will now consider the triangle $\sigma_2 = [pqr]$. Note that $C(\sigma_1, \omega_{\sigma_1})$, $R(\sigma_1, \omega_{\sigma_1})$ are the circumcenter and radius of σ_2 respectively. Also, $C(\sigma, \omega_\sigma)$ is the midpoint of the line segment $[qr]$ with $2R(\sigma, \omega_\sigma) = \|q - r\|$ and $D_{\sigma_2}(p) = \|p - p'\|$. From the definition of q , we have $\|p - r\| \geq \|p - q\|$. Using the fact $\|q - r\| = 2R(\sigma, \omega_\sigma)$, we have

$$\|p - r\| \geq \frac{\|q - r\|}{2} = R(\sigma, \omega_\sigma).$$

This implies from Lemma C.1.1

$$\begin{aligned} \|p - q\| &= \frac{2R_{\sigma_2} D_{\sigma_2}(p)}{\|p - r\|} \\ &\leq \frac{2R(\sigma_1, \omega_{\sigma_1}) D_{\sigma_2}(p)}{R(\sigma, \omega_\sigma)} \quad \text{as } R_{\sigma_2} \leq R(\sigma_1, \omega_{\sigma_1}) \text{ and } \|p - r\| \geq R(\sigma, \omega_\sigma) \\ &\leq 48\beta D_{\sigma_2}(p) = 48\beta \|p - p'\| \quad \text{as Eq. (C.2)} \end{aligned} \quad (\text{C.6})$$

From Eq. (C.4), (C.5) and (C.6)

$$\begin{aligned} \|p_{k+1} - q\| &\leq \|p_{k+1} - p\| + \|p - q\| \\ &\leq \delta_0 L_{\sigma_1} + 48\beta (D_{\sigma_1}(p_{k+1}) + \delta_0 L_{\sigma_1}) \\ &\stackrel{\text{def}}{=} \eta_1 \end{aligned} \quad (\text{C.7})$$

Using the fact that σ is Γ_0^k -thick (since σ_1 is a Γ_0 -flake), and the bound $\delta_0^2 L_{\sigma_1}^2$ on the differences of the squared distances between $C(\sigma, \omega_\sigma)$ and the vertices of σ , we obtain a bound [?, Lemma 4.1] on the distance from $C(\sigma, \omega_\sigma)$ to c_σ :

$$\begin{aligned} \|c_\sigma - C(\sigma, \omega_\sigma)\| &\leq \frac{\delta_0^2 L_{\sigma_1}^2}{2\Upsilon(\sigma)\Delta_\sigma} \\ &\leq \frac{\delta_0^2 R_\sigma}{\Upsilon_\sigma} \quad \text{as } L_{\sigma_1} \leq \Delta_\sigma \leq 2R_\sigma \\ &\leq \frac{\delta_0^2 R_\sigma}{\Gamma_0^k} \quad \text{since } \sigma \text{ is } \Gamma_0^k\text{-thick, } \Upsilon_\sigma \geq \Gamma_0^k \\ &\leq \frac{\delta_0^2 R_\sigma}{\Gamma_0^m} \quad \text{as } \Gamma_0 \leq 1 \\ &\stackrel{\text{def}}{=} \eta_2 \end{aligned} \quad (\text{C.8})$$

Since $k \geq 1$, there exists $p_i \in \circ$ such that

$$p_i \in B_{\mathbb{R}^N}(c_\sigma, R_\sigma) \cap B_{\mathbb{R}^N}(C(\sigma, \omega_\sigma), R(\sigma, \omega_\sigma)) \cap \text{aff}(\sigma).$$

Also, $\|c_\sigma - p_i\| = R_\sigma$ and $\|C(\sigma, \omega_\sigma) - p_i\| = R(\sigma, \omega_\sigma)$.

Using the facts that $R_\sigma = \|c_\sigma - p_i\|$ and $R(\sigma, \omega_\sigma) = \|C(\sigma, \omega_\sigma) - p_i\|$, and the Triangle inequality, we get

$$\begin{aligned} R_\sigma - \|c_\sigma - C(\sigma, \omega_\sigma)\| &\leq \|C(\sigma, \omega_\sigma) - p_i\| \leq R_\sigma + \|c_\sigma - C(\sigma, \omega_\sigma)\| \\ R_\sigma - \eta_2 &\leq R(\sigma, \omega_\sigma) \leq R_\sigma + \eta_2 \end{aligned} \quad (\text{C.9})$$

The last equation follows from Eq. (C.8).

Let S' and S denote $B_{\mathbb{R}^N}(c_\sigma, R_\sigma) \cap \text{aff}(\sigma)$ and $B_{\mathbb{R}^N}(C(\sigma, \omega_\sigma), R(\sigma, \omega_\sigma))$ respectively. From Eq. (C.8) and (C.9), we have $d_{\mathbb{R}^N}(\partial S', \partial S) \leq \eta_1 + 2\eta_2$. This implies that there exists $q' \in \partial S'$ such that

$$\|q' - q\| \leq 2\eta_2. \quad (\text{C.10})$$

Therefore from Eq. (C.7) and (C.10), we get

$$\|p_{k+1} - q'\| \leq \|p_{k+1} - q\| + \|q' - q\| \leq \eta_1 + 2\eta_2$$

Using the facts that $\delta_0^2 \leq \Gamma_0^{m+1} \leq \Gamma_0^2$ (from hyp. 4 of the lemma and $\Gamma_0 \leq 1$), $L_{\sigma_1} \leq L_\sigma \leq \Delta_\sigma \leq 2R_\sigma$ and $\frac{D_{\sigma_1}(p_{k+1})}{\Delta_\sigma} \leq 2^7 3^2 \beta^2 \Gamma_0$ (from Eq. (C.3)), and Eq. (C.7) and (C.8), we get

$$\begin{aligned} d_{\mathbb{R}^N}(p_{k+1}; \partial S') &\leq \|p_{k+1} - q'\| \\ &\leq \eta_1 + 2\eta_2 \\ &\leq \delta_0 L_{\sigma_1} + 48\beta \left(D_{\sigma_1}(p_{k+1}) + \delta_0 L_{\sigma_1} \right) + \frac{2\delta_0^2 R_\sigma}{\Gamma_0^m} \\ &\leq B\Gamma_0 R_\sigma \end{aligned}$$

where

$$B = 4 + 96\beta(1 + 2^7 3^2 \beta^2).$$

□

We will now restate Lemma 5.4.1 from Chapter 5 in terms of m .

Lemma C.1.3 *Let p be a point on $\mathcal{M}$. There exists ξ that depends only on m , and A that depends only on m such that, for all $r = t \leq \xi \text{rch}(\mathcal{M})$, we have*

$$0 < 1 - \frac{At}{\text{rch}(\mathcal{M})} \leq \frac{\text{vol}(B_{\mathbb{R}^N}(p, r) \cap \mathcal{M})}{\phi_m r^k} \leq 1 + \frac{At}{\text{rch}(\mathcal{M})}$$

where ϕ_m is the volume of the m -dimensional unit Euclidean ball.

This result will be used to bound the volume of $F(\sigma, t)$.

Proof of Lemma 8.5.8 For the rest of the proof we define

$$\tilde{t} = \frac{t}{\text{rch}(\mathcal{M})}$$




Figure C.1: Proof of Lemma 8.5.8.

Consider the following elementary weight function: $\omega_\sigma = \omega_{\sigma_1} \mid_{\sigma^\circ}$. Using the facts that $R(\sigma, \omega_\sigma) \leq R(\sigma_1, \omega_{\sigma_1})$, $R(\sigma, \omega_\sigma) < \beta t \text{rch}(\mathcal{M})$, and Lemma 8.4.1 (3)

$$\begin{aligned}
R_\sigma &\leq R(\sigma, \omega_\sigma) \left(1 - \frac{\delta_0^2}{\Upsilon_\sigma}\right)^{-1} \\
&\leq R(\sigma, \omega_\sigma) \left(1 - \frac{\delta_0^2}{\Gamma_0^m}\right)^{-1} && \text{since } \Upsilon_\sigma \geq \Gamma_0^k \geq \Gamma_0^m \\
&\leq 2\beta \tilde{t} \text{rch}(\mathcal{M})
\end{aligned} \tag{C.11}$$

Let p be a vertex of σ . Let c be the point closest to c_σ on $T_p\mathcal{M}$ and c^* be the point closest to c on $\mathcal{M}$ (see Fig. C.1).

From Lemma 8.2.10, we have for all $q \in \mathring{\sigma}$

$$d_{\mathbb{R}^m}(q, T_p \mathcal{M}) \leq \frac{\|p - q\|^2}{2\text{rch}(\mathcal{M})} \leq \frac{\Delta_\sigma^2}{2\text{rch}(\mathcal{M})} \stackrel{\text{def}}{=} \eta$$

From Lemma 7.2.1, and the facts that $\Upsilon_\sigma \geq \Gamma_0^m$ and $\Delta_\sigma \leq 2R_\sigma \leq 4\beta t$, we have

$$\sin \angle(T_p \mathcal{M}, \text{aff}(\sigma)) \leq \frac{2\eta}{\Upsilon_\sigma \Delta_\sigma} \leq \frac{\Delta_\sigma}{\Upsilon_\sigma \text{rch}(\mathcal{M})} \leq \frac{4\beta \tilde{t}}{\Gamma_0^m}$$

Therefore

$$\|c - c_\sigma\| \leq \sin \angle(T_p \mathcal{M}, \text{aff}(\sigma)) \times R_\sigma \leq \left(\frac{4\beta \tilde{t}}{\Gamma_0^m} \right) R_\sigma, \quad (\text{C.12})$$

and from Lemma 8.4.2

$$\|c - c^*\| \leq \frac{2\|c - p\|^2}{\text{rch}(\mathcal{M})} \leq 4\beta\tilde{t}R_\sigma. \quad (\text{C.13})$$

Let $x \in F(\sigma, t)$ and x^* be the point closest to x on $\partial B_{\mathbb{R}^N}(c_\sigma, R_\sigma) \cap \text{aff}(\sigma)$. Then from Lemma C.1.2, we have

$$\|x - x^*\| < B\Gamma_0 R_\sigma \quad (\text{C.14})$$

Using the fact that $\|c_\sigma - x^*\| = R_\sigma$, we get

$$\begin{aligned}
\|c^* - x\| &\leq \|c^* - c\| + \|c - C(\sigma)\| + \|C(\sigma) - x^*\| + \|x^* - x\| \\
&< R_\sigma \left(1 + B\Gamma_0 + 4\beta t \left(\frac{1}{\Gamma_0^m} + 1 \right) \right) && \text{from Eq. (C.12), (C.13), (C.14)} \\
&\leq R_\sigma \left(1 + B\Gamma_0 + \frac{8\beta \tilde{t}}{\Gamma_0^m} \right) && \text{since } \Gamma_0 \leq 1 \\
&\leq R_\sigma (1 + (B+1)\Gamma_0) && \text{from hyp. 2 of the lemma.}
\end{aligned}$$

Similarly we can show that

$$\|c^* - x\| < R_\sigma (1 - (B+1)\Gamma_0)$$

Therefore

$$F(\sigma, t) \subseteq (B_{\mathbb{R}^N}(c^*, (1+\zeta)R_\sigma) \setminus B_{\mathbb{R}^N}(c^*, (1-\zeta)R_\sigma)) \cap \mathcal{M}$$

where $\zeta = (B+1)\Gamma_0$.

Observe that Lemma C.1.3 can be applied since

$$\begin{aligned}
R_\sigma(1+\zeta) &\leq 2R_\sigma && \text{since } \zeta \leq 1 \text{ from hyp. 1} \\
&\leq 4\beta t && \text{from Eq. (C.11)} \\
&\leq \xi \text{rch}(\mathcal{M}) && \text{because } t < \epsilon.
\end{aligned}$$

Therefore

$$\begin{aligned}
\frac{\text{vol}(F(\sigma, t))}{\phi_m} &\leq \frac{\text{vol}(B_{\mathbb{R}^N}(c^*, R_\sigma(1+\zeta)) \cap \mathcal{M} \setminus B_{\mathbb{R}^N}(c^*, R_\sigma(1-\zeta)) \cap \mathcal{M})}{\phi_m} \\
&\leq (1 + A(1+\zeta)\tilde{t})R_\sigma^m(1+\zeta)^m - (1 - A(1-\zeta)\tilde{t})R_\sigma^m(1-\zeta)^m \\
&\leq R_\sigma^m((1+\zeta)^m - (1-\zeta)^m) + A\tilde{t}R_\sigma^m((1+\zeta)^m + (1-\zeta)^m) \\
&\leq 2^m\zeta R_\sigma^m + A(2^{m+1} + 1)\tilde{t}R_\sigma^m && \text{(C.15)}
\end{aligned}$$

The last inequality follows from the fact that $(1+x)^m - (1-x)^m \leq 2^m x$ for all $x \in [0, 1]$.

From hyp. 2 and the fact that $\Gamma_0 < 1$, we have

$$\tilde{t} \leq \tilde{\epsilon} \leq \frac{\Gamma_0^{m+1}}{8\beta} < \Gamma_0. \quad \text{(C.16)}$$

The lemma now follows from Eq. (C.15) and (C.16). $\square$

Bibliography

- [AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1999. (Cited on page 10.)
- [ACDL02a] N. Amenta, S. Choi, T. K. Dey, and Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. *Internat. Journal of Comput. Geom. and Applications*, 12(1-2):125–141, 2002. (Cited on page 71.)
- [ACDL02b] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Intl. Journal of Computational Geometry and Application*, 12:125–141, 2002. (Cited on pages 45 and 51.)
- [AE84] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984. (Cited on page 47.)
- [AGG⁺10] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin, and M. Sharir. Kinetic stable Delaunay graphs. In *Symp. Comp. Geom.*, pages 127–136, 2010. (Cited on page 106.)
- [AM97] U. Abresch and W. T. Meyer. Injectivity Radius Estimates and Sphere Theorems. In K. Grove and P. Petersen, editors, *Comparison Geometry*. Mathematical Sciences Research Institute Publications, 1997. (Cited on page 177.)
- [Aur87] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1):78–96, 1987. (Cited on pages 16, 47, 50 and 98.)
- [Ber90] M. Berger. *Geometry 2*. Universitext, Springer Verlag, 1990. (Cited on page 178.)
- [BF04] J.-D. Boissonnat and J. Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004. (Cited on pages i, 3, 5, 25, 26, 71, 72 and 195.)
- [BFC01] J.-D. Boissonnat and F. F. Cazals. Natural neighbour coordinates of points on a surface. *Computational Geometry Theory and Applications*, 19(2):155–173, 2001. (Cited on page 141.)
- [BG10a] J.-D. Boissonnat and A. Ghosh. Manifold Reconstruction using Tangential Delaunay Complexes. In *Symp. Comp. Geom.*, pages 324–333, 2010. (Cited on page 170.)

- [BG10b] J.-D. Boissonnat and A. Ghosh. Triangulating smooth submanifolds with light scaffolding. *Mathematics in Computer Science*, 4(4):431–461, 2010. (Cited on pages 140, 143 and 150.)
- [BG11] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. Technical Report N° 7142 v.3, INRIA, 2011. (Cited on pages 9, 70, 71, 72, 134 and 140.)
- [BGO09] J.-D. Boissonnat, L. J. Guibas, and S. Y. Oudot. Manifold Reconstruction in Arbitrary Dimensions using Witness Complexes. *Discrete and Computational Geometry*, 42(1):37–70, 2009. (Cited on pages 4, 5, 16, 25, 26, 45, 98, 102, 169 and 170.)
- [BN02] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, volume 14, pages 585–591, 2002. (Cited on pages 4 and 26.)
- [BNN10] J.-D. Boissonnat, F. Nielsen, and R. Nock. Bregman Voronoi Diagrams. *Discrete and Computational Geometry*, 44(2):281–307, 2010. (Cited on pages 45 and 98.)
- [BO05] J.-D. Boissonnat and S. Oudot. Provably Good Sampling and Meshing of Surfaces. *Graphical Models*, 67:405–451, 2005. (Cited on pages 71 and 94.)
- [BO06] J.-D. Boissonnat and S. Y. Oudot. Provably Good Sampling and Meshing of Lipschitz Surfaces. In *Proc. ACM Symp. on Computational Geometry*, pages 337–346, 2006. (Cited on pages 94 and 170.)
- [Bre94] G. E. Bredon. *Topology and Geometry*. Graduate Text in Mathematics, Springer, 1994. (Cited on page 20.)
- [BS04] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: nearest neighbor relations for imprecise points. In *SODA*, pages 410–419, 2004. (Cited on page 106.)
- [BSW08] M. Belkin, J. Sun, and Y. Wang. Discrete Laplace operator on meshed surfaces. In *Proc. ACM Symp. on Computational Geometry*, pages 278–287, 2008. (Cited on page 45.)
- [BWC02] P. Baniramka, R. Wenger, and R. Crawfis. Isosurface construction in any dimension using convex hulls. *IEEE Trans. Vis. Comput. Graph*, 10(2):130–141, 2002. (Cited on pages 6 and 70.)
- [BWY08] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. ACM Symp. on Computational Geometry*, pages 270–277, 2008. (Cited on pages 3, 5, 26, 70, 94 and 169.)
- [BWY11] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. Research Report RR-7712, INRIA, 2011. (Cited on pages 108 and 134.)

- [BY98] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic geometry*. Cambridge University Press, 1998. (Cited on page 44.)
- [Cai61] S. S. Cairns. A simple triangulation method for smooth manifolds. *Bulletin of the American Mathematical Society*, 67(4):389–390, 1961. (Cited on pages 5 and 69.)
- [CCSM11] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011. (Cited on page 170.)
- [CDE⁺00a] S-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S-H. Teng. Sliver Exudation. *Journal of ACM*, 47:883–904, 2000. (Cited on pages 3, 5, 16, 17, 19, 25, 26 and 33.)
- [CDE⁺00b] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S. H. Teng. Sliver exudation. *Journal of the ACM*, 47(5):883–904, 2000. (Cited on pages 108 and 134.)
- [CDR05a] 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. (Cited on pages 4, 5, 15, 16, 25, 26, 27, 34, 71, 169 and 170.)
- [CDR05b] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *Symp. on Discrete Algorithms*, pages 1018–1027, 2005. (Cited on pages 98, 106, 133 and 134.)
- [CG06] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J. D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curve and Surfaces*. Springer, 2006. (Cited on pages 4, 6, 25 and 70.)
- [CG12] G. D. Cañas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are triangulations. In *Symp. Comp. Geom.*, 2012. (Cited on page 106.)
- [CGT82] J. Cheeger, M. Gromov, and M. Taylor. Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete riemannian manifolds. *Journal of Differential Geometry*, 17:15–53, 1982. (Cited on page 177.)
- [Cha93] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993. (Cited on pages 28 and 44.)
- [Cha06] I. Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006. (Cited on page 100.)
- [Che97] L. P. Chew. Guaranteed-Quality Delaunay Meshing in 3D. In *Proc. ACM Symp. on Computational Geometry*, pages 391–393, 1997. (Cited on page 71.)
- [CIdSZ08] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, 2008. (Cited on page 45.)

- [CL08] F. Chazal and A. Lieutier. Smooth Manifold Reconstruction from Noisy and Non-uniform Approximation with Guarantees. *Comp. Geom: Theory and Applications*, 40:156–170, 2008. (Cited on pages 4, 25 and 170.)
- [Cla06] K. Clarkson. Building triangulations using ε -nets. In *Proc. ACM Symp. on Theory of Computing (STOC)*, pages 326–335, 2006. (Cited on pages 5, 6, 70, 94, 104 and 170.)
- [CO08] F. Chazal and S. Y. Oudot. Towards Persistence-Based Reconstruction in Euclidean Spaces. In *Proc. ACM Symp. on Computational Geometry*, pages 232–241, 2008. (Cited on page 26.)
- [CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989. (Cited on page 28.)
- [CSD04] D. Cohen-Steiner and T. K. F. Da. A greedy Delaunay Based Surface Reconstruction Algorithm. *The Visual Computer*, 20:4–16, 2004. (Cited on page 27.)
- [CWW08] S-W. Cheng, Y. Wang, and Z. Wu. Provable Dimension Detection using Principle Component Analysis. *Intl. Journal of Computational Geometry and Application*, 18(5):415–440, 2008. (Cited on pages 31 and 45.)
- [Dan00] E. N. Dancer. Degree theory on convex sets and applications to bifurcation. In *Calculus of Variations and Partial Differential Equations*, pages 185–225. Springer-Verlag, 2000. (Cited on page 127.)
- [dC92] M. P. do Carmo. *Riemannian Geometry*. Birkhäuser, 1992. (Cited on page 104.)
- [Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. (Cited on pages 98, 100, 101, 105, 116 and 129.)
- [Dey06] T. K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2006. (Cited on pages 4 and 25.)
- [DG03] D. L. Donohu and C. Grimes. Hessian eigenmaps: new locally linear embedding techniques for high dimensional data. *Proceedings of the Natural Academy of Sciences*, 100:5591–5596, 2003. (Cited on pages 4 and 26.)
- [DL09] T. K. Dey and K. Li. Topology from Data via Geodesic complexes. *Tech Report OSU-CISRC-3/09-TR05*, 2009. (Cited on page 178.)
- [DLRW09] T. K. Dey, K. Li, E. A. Ramos, and R. Wenger. Isotopic reconstruction of surfaces with boundaries. *Comput. Graph. Forum*, 28(5):1371–1382, 2009. (Cited on page 170.)
- [DMT92] O. Devillers, S. Meiser, and M. Teillaud. The space of spheres, a geometric tool to unify duality results on Voronoi diagrams. Technical Report RR-1620, INRIA, 1992. (Cited on page 98.)

- [dS08] V. de Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 135:39–64, 2008. (Cited on page 130.)
- [Dud74] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Journal of Approximation Theory*, 10:227–236, 1974. (Cited on pages 5 and 70.)
- [Dug66] J. Dugundji. *Topology*. Allyn and Bacon, Inc., Boston, 1966. (Cited on pages 20 and 113.)
- [DZM08] R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symp. Geometry Processing)*, 27(5):1393–1402, 2008. (Cited on pages 104 and 106.)
- [Ede01] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, 2001. (Cited on page 71.)
- [Ehr74] P. Ehrlich. Continuity properties of the injectivity radius function. *Composito Mathematica*, 29:151–178, 1974. (Cited on page 104.)
- [ELS⁺00] H. Edelsbrunner, X-Y. Li, A. Stathopoulos, D. Talmor, S-H. Teng, A. ÜNgÖR, and N. Walkington. Smoothing and cleaning up slivers. In *Proc. ACM Symp. on Theory of Computing*, pages 273–277, 2000. (Cited on pages 9 and 12.)
- [ES97] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. Comput. Geometry Appl.*, 7(4):365–378, 1997. (Cited on pages 104, 106 and 133.)
- [Fed59] H. Federer. Curvature Measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. (Cited on pages 5, 69 and 140.)
- [Fed69] H. Federer. *Geometric measure theory*. Springer New York, 1969. (Cited on pages 5, 48 and 69.)
- [FKMS05] S. Funke, C. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for Delaunay triangulations. In *Symp. on Discrete Algorithms*, pages 1047–1056, 2005. (Cited on page 106.)
- [Flö03] J. Flötotto. *A coordinate system associated to a point cloud issued from a manifold: definition, properties and applications*. PhD thesis, Université of Nice Sophia-Antipolis, 2003. (Cited on pages i, 3, 5, 25, 26 and 195.)
- [FR02] S. Funke and E. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pages 781–780, 2002. (Cited on page 45.)
- [Fre02] D. Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002. (Cited on pages i, 3, 5, 25, 26 and 195.)
- [Fu93] J. H. G. Fu. Convergence of curvature in secant approximations. *Journal of Differential Geometry*, 37:117–190, 1993. (Cited on page 15.)

- [GKS00] M. Gopi, S. Khrisnan, and C. T. Silva. Surface Reconstruction based on Lower Dimensional Localized Delaunay Triangulation. In *Proc. Eurographics*, pages 363–371, 2000. (Cited on page 27.)
- [Gra90] A. Gray. *Tubes*. Addison-Wesley, Reading, MA, 1990. (Cited on page 179.)
- [Gru93] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Mathematicum*, 5:281–297, 1993. (Cited on pages 5, 70 and 170.)
- [Gru04] P. M. Gruber. Optimum quantization and its applications. *Advances in Mathematics*, 186:456–497, 2004. (Cited on pages 5 and 70.)
- [GW04a] J. Giesen and U. Wagner. Shape Dimension and Intrinsic Metric from Samples of Manifolds. *Discrete and Computational Geometry*, 32(2):245–267, 2004. (Cited on pages 11, 31, 45, 65 and 67.)
- [GW04b] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete Comput. Geom.*, 32:245–267, 2004. (Cited on page 140.)
- [Hat02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. (Cited on page 21.)
- [Hen02] M. E. Henderson. Multiple parameter continuation: computing implicitly defined k -manifolds. *Int. Journal of Bifurcation and Chaos*, 12(3):451–476, 2002. (Cited on pages 6, 70 and 74.)
- [Kam08] G. K. Kamenev. The initial convergence rate of adaptive methods for polyhedral approximation of convex bodies. *Computational Mathematics and Mathematical Physics*, 48(5):724–738, 2008. (Cited on pages 6 and 70.)
- [Lei99] G. Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999. arXiv:math/0011016v1. (Cited on page 104.)
- [Li00] X-Y. Li. *Sliver-Free three dimensional Delaunay mesh generation*. PhD thesis, University of Illinois at Urbana-Champaign, 2000. (Cited on page 9.)
- [Li03a] X-Y. Li. Generating Well-Shaped d -dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003. (Cited on pages 3, 9, 15, 71 and 74.)
- [Li03b] X-Y. Li. Generating well-shaped d -dimensional Delaunay meshes. *Theoretical Computer Science*, 296(1):145–165, 2003. (Cited on pages 108 and 150.)
- [LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Symp. Comp. Geom.*, pages 341–349, 2000. (Cited on pages i, 4, 6, 7, 98, 100, 104, 106, 132, 133, 145, 169, 170 and 195.)
- [LL06] S. Lafon and A. B. Lee. Diffusion Maps and Coarse-Graining: A Unified Framework for Dimensionality Reduction, Graph Partitioning, and Data Set Parameterization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1393–1403, 2006. (Cited on pages 4 and 26.)

- [LS03] F. Labelle and J. R. Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *Symp. Comp. Geom.*, pages 191–200, 2003. (Cited on pages 7, 106 and 133.)
- [LT01] X-Y. Li and S-H. Teng. Generating well-shaped delaunay meshed in 3d. In *Proc. ACM-SIAM Symp. on Discrete Algorithms*, pages 28–37, 2001. (Cited on page 9.)
- [Mas67] W. S. Massey. *Algebraic Topology : An Introduction*. Springer-Verlang, GTM 56, 1967. (Cited on page 21.)
- [Min03] C. Min. Simplicial Isosurfacing in Arbitrary Dimension and Codimension. *Journal of Computational Physics*, 190:295–310, 2003. (Cited on pages 6 and 70.)
- [Mun66] J. R. Munkres. *Elementary Differential Topology*. Annals of Mathematics Studies, Princeton University Press, 1966. (Cited on pages 5 and 69.)
- [Mun68] J. R. Munkres. *Elementary differential topology*. Princeton University press, second edition, 1968. (Cited on pages 102 and 108.)
- [Mun84] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984. (Cited on page 111.)
- [Mun00] J. R. Munkres. *Topology*. Prentice-Hall, 2nd edition, 2000. (Cited on page 20.)
- [NLCK05] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis. Diffusion Maps, Spectral Clustering and Eigenfunctions of Fokker-Planck Operators. *Neural Information Processing Systems*, 18, 2005. (Cited on pages 4 and 26.)
- [NSW08a] 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. (Cited on pages 4, 25 and 51.)
- [NSW08b] 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. (Cited on pages 80, 87, 88, 176, 177 and 178.)
- [NSW08c] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008. (Cited on pages 140, 141 and 145.)
- [PC05] G. Peyré and L. Cohen. Geodesic computations for fast and accurate surface remeshing and parameterization. In *Progress in Nonlinear Differential Equations and Their Applications*, volume 63, pages 157–171. Birkhäuser Verlag Basel/Switzerland, 2005. (Cited on pages 6 and 70.)
- [RS00] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000. (Cited on pages 4 and 26.)

- [Rup95] J. Ruppert. A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation. *Journal of Algorithms*, 18(3):548–585, 1995. (Cited on page 71.)
- [Sak83] T. Sakai. On continuity of injectivity radius function. *Mathematical Journal of Okayama University*, 25(1):91–97, 1983. (Cited on page 104.)
- [She05] J. Shewchuk. Star Splaying: An Algorithm for Repairing Delaunay Triangulations and Convex Hulls. In *Proc. ACM Symp. on Computational Geometry*, pages 237–246, 2005. (Cited on pages 3, 5, 26, 28 and 169.)
- [SL00] H. S. Seung and D. D. Lee. The manifold ways of perception. *Science*, 290:2268–2269, 2000. (Cited on page 26.)
- [TB97] L.N. Trefethen and D. Bau. *Numerical linear algebra*. Society for Industrial Mathematics, 1997. (Cited on page 109.)
- [TdSL00] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000. (Cited on pages 4 and 26.)
- [Whi40] J. H. C. Whitehead. On C^1 -complexes. *Annals of Mathematics*, 41:809–824, 1940. (Cited on pages 5 and 69.)
- [Whi57a] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957. (Cited on pages 5, 9, 14, 21, 51, 54, 57, 58, 59, 60 and 69.)
- [Whi57b] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957. (Cited on page 108.)
- [Zee66] E. C. Zeeman. *Seminar on Combinatorial Topology*. Institut des Hautes Études Scientifiques (Paris) and University of Warwick (Coventry), Notes, 1963–1966. (Cited on pages 49 and 57.)
- [ZZ04] Z. Zhang and H. Zha. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. *SIAM Journal of Scientific Computing*, 26(1):313–338, 2004. (Cited on pages 4 and 26.)

Piecewise linear reconstruction and meshing of submanifolds of Euclidean space

Abstract. In this thesis we address some of the problems in the field of piecewise linear approximation of k -dimensional smooth submanifolds of Euclidean space $\mathbb{R}^d$. The main goal of this thesis was to develop algorithms that solve these problems with *theoretical guarantees*, i.e. the output being homeomorphic to the submanifold, and also have *intrinsic dimension sensitive complexity*, i.e. time and space complexity depend exponentially on the intrinsic dimension k of the submanifold and linearly on the ambient Euclidean dimension d .

The two standard questions in this field are the following:

- **Manifold reconstruction.** From a dense point sample $P \subset \mathbb{R}^d$, from an unknown smooth k -dimensional submanifold $\mathcal{M}$ of $\mathbb{R}^d$, we want to build a simplicial approximation $\hat{\mathcal{M}} \subset \mathbb{R}^d$ of $\mathcal{M}$ with theoretical guarantees.
- **Sampling and meshing manifolds.** For a given parameter ε and a k -dimensional smooth submanifold, known through some standard oracles, we want to generate a dense sample $P \subset \mathcal{M}$, according to the prescribed parameter ε , and build a simplicial approximation $\hat{\mathcal{M}}$ of $\mathcal{M}$ on top of the sample P with theoretical guarantees.

In this thesis we try to chip away at both these problems with the following results:

- For a dense point sample P of a smooth submanifold $\mathcal{M}$ of $\mathbb{R}^d$ we give sufficient conditions under which the *tangential Delaunay complex*, defined in [BF04, Flö03, Fre02], build using the point sample P is homeomorphic and a close geometric approximation of $\mathcal{M}$.
- We give an algorithm, whose complexity is intrinsic dimension sensitive, to reconstruct smooth k -dimensional submanifolds of $\mathbb{R}^d$ from a dense point sample P using tangential Delaunay complexes. We show, using the above result, that the output is homeomorphic and a close geometric approximation of $\mathcal{M}$. To the best of our knowledge, this is the first certified algorithm for manifold reconstruction whose complexity is intrinsic dimension sensitive.
- We give an algorithm to sample and mesh a k -dimensional smooth submanifold $\mathcal{M}$ of $\mathbb{R}^d$. According to the prescribed parameter ε , the algorithm generates a dense sample of $\mathcal{M}$ and a mesh with theoretical guarantees. The algorithm uses only simple numerical operations. We show that the size of the sample is $O(\varepsilon^{-k})$ and the asymptotic complexity of the algorithm is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (for fixed $\mathcal{M}$, d and k).
- We provide a counterexample to the result announced by Liebon and Letscher [LL00]. We show that density of the sample points on a manifold $\mathcal{M}$ alone cannot guarantee that the nerve of the intrinsic Voronoi diagram, i.e. the intrinsic Delaunay triangulation, is homeomorphic to the manifold $\mathcal{M}$.

- We introduce a parameterized notion of δ -generic point set for Delaunay triangulations. We show that Delaunay triangulation of a δ -generic point sample is (1) combinatorially stable under small perturbation of the underlying metric and vertex positions, and (2) simplices of Delaunay triangulation are *well shaped*.
- Using the stability results of Delaunay triangulations of δ -generic point set, we show that, for any sufficiently regular submanifold of Euclidean space, and appropriate ε and δ , any sample set which meets a localized δ -generic ε -dense sampling criteria, intrinsic Delaunay triangulation is equal to restricted Delaunay triangulation and tangential Delaunay triangulation, and intrinsic Delaunay triangulation is homeomorphic to the submanifold. We also give a refinement algorithm for generating intrinsic Delaunay triangulations of submanifolds.

Keywords. Delaunay complex, intrinsic Delaunay complex, manifold reconstruction, meshing, slivers, stability of Delaunay triangulation, Voronoi diagram, and weighted points.
