

Manifold reconstruction using tangential Delaunay complexes

Jean-Daniel Boissonnat, Arijit Ghosh

► **To cite this version:**

Jean-Daniel Boissonnat, Arijit Ghosh. Manifold reconstruction using tangential Delaunay complexes. Discrete and Computational Geometry, Springer Verlag, 2014, 51 (1), pp.221-267. <10.1007/s00454-013-9557-2>. <hal-00932209>

HAL Id: hal-00932209

<https://hal.inria.fr/hal-00932209>

Submitted on 16 Jan 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.

Manifold reconstruction using tangential Delaunay complexes

Jean-Daniel Boissonnat *

INRIA, Geometrica

Sophia-Antipolis, France

Jean-Daniel.Boissonnat@inria.fr

Arijit Ghosh *†

Indian Statistical Institute, ACM Unit

Kolkata, India

arijitiiiiitkgpster@gmail.com

Abstract

We give a provably correct algorithm to reconstruct a k -dimensional smooth 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⁺00]. 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. We also prove that for a dense enough sample the output of our algorithm is isotopic to the manifold and a close geometric approximation of the manifold. A preliminary version of this paper appeared in the proceedings of the Symposium on Computational Geometry, SoCG 2010.

Keywords: Tangential Delaunay complex, weighted Delaunay triangulation, manifold reconstruction, manifold learning, sampling conditions, sliver exudation.

Contents

1	Introduction	2
2	Definitions and preliminaries	4
2.1	Weighted Delaunay triangulation	6
2.1.1	Weighted points	6
2.1.2	Weighted Voronoi diagram and Delaunay triangulation	6
2.2	Sampling conditions	8
2.2.1	Local feature size	8
2.2.2	Sampling parameters	8
2.3	Fat simplices	9
2.4	Tangential Delaunay complex and inconsistent configurations	10

*This research has been partially supported by the Agence Nationale de la Recherche (project GAIA 07-BLAN-0328-04).

†This work was done when the author was a PhD student at Geometrica, INRIA Sophia Antipolis.

3	Manifold reconstruction	12
4	Analysis of the algorithm	15
4.1	First lemmas	16
4.2	Properties of the tangential Delaunay complex	18
4.3	Properties of inconsistent configurations	21
4.4	Number of local neighbors	25
4.5	Correctness of the algorithm	26
4.6	Time and space complexity	27
5	Topological and geometric guarantees	29
5.1	Piecewise-linear k -manifold	31
5.2	Homeomorphism	31
5.2.1	Basic lemmas	32
5.2.2	Properties of π_p^*	33
5.2.3	Properties of π	36
5.3	Pointwise approximation	37
5.4	Isotopy	37
6	Conclusion	38
A	Appendix to Section 5.2	39
	Bibliography	42

1 Introduction

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 \mathcal{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, CDR05, NSW08].

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 breakthrough, along this line, was the work of Cheng, Dey and Ramos [CDR05] and this was followed by the work of 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 paper, 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 [CDR05] 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⁺00]. 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. [CDR05]. 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 the other local algorithms that have been proposed in the surface reconstruction literature [CSD04, GKS00].

Organisation of the paper

In Section 2, we introduce the basic concepts used in this paper. We recall the notion of weighted Voronoi (or power) diagrams and Delaunay triangulations in Section 2.1 and define sampling conditions in Section 2.2. We introduce various quantities to measure the shape of simplices in Section 2.3 and, in particular, the central notion of fatness. In Section 2.4, we define the two main notions of this paper: the tangential complex and inconsistent configurations.

The algorithmic part of the paper is given in Section 3.

The main structural results are given in Section 4. Under some sampling condition, we bound the shape measure of the simplices of the tangential complex in Section 4.2 and of inconsistent configurations in Section 4.3. 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 4.4. In Sections 4.5 and 4.6, we prove the correctness of the algorithm, and space and time complexity respectively. In Section 5, we prove that the simplicial complex output by the algorithm is indeed a good approximation of the sampled manifold. Finally, in Section 6, we conclude with some possible extensions.

The list of main notations has been added at the end of the paper as a reference for the readers.

2 Definitions and preliminaries

The standard Euclidean distance between two points $x, y \in \mathbb{R}^d$ will be denoted by $\|x - y\|$. *Distance* between a point $x \in \mathbb{R}^d$ and a set $X \subset \mathbb{R}^d$, is defined as

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

We will denote the topological boundary and interior of a set $X \subseteq \mathbb{R}^d$ by ∂X and $\text{int } X$ respectively.

For a map $f : X \rightarrow Y$ and $X_1 \subseteq X$, $f|_{X_1}$ denotes the restriction of map f to the subset X_1 .

A j -simplex is the convex hull of $j + 1$ affinely independent points. If τ is a j -simplex with vertices $\{x_0, \dots, x_j\}$, we will also denote τ as $[x_0, \dots, x_j]$. For convenience, we often identify a simplex and the set of its vertices. Hence, if τ is a simplex, $p \in \tau$ means that p is a vertex of τ . If τ is a j -simplex, $\text{aff}(\tau)$ denotes the j -dimensional affine hull of τ and N_τ denotes the $(d - j)$ -dimensional normal space of $\text{aff}(\tau)$.

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} , $\overset{\circ}{\text{st}}(\sigma, \mathcal{K}) = \text{st}(\sigma, \mathcal{K}) \setminus \text{lk}(\sigma, \mathcal{K})$.

For all x, y in \mathbb{R}^d , $[xy]$ denotes the line segment joining the points x and y .

In this paper, \mathcal{M} denotes a differentiable manifold of dimension k embedded in \mathbb{R}^d and $\mathcal{P} = \{p_1, \dots, p_n\}$ a finite sample of points from \mathcal{M} . We will further assume that \mathcal{M} has positive reach (see Section 2.2). We denote by T_p and N_p the k -dimensional tangent space and $(d - k)$ -dimensional normal space at point $p \in \mathcal{M}$ respectively.

For a given $p \in \mathbb{R}^d$ and $r \geq 0$, $B(p, r)$ ($\bar{B}(p, r)$) denotes the d -dimensional Euclidean open (close) ball centered at p of radius r , and $B_{\mathcal{M}}(p, r)$ ($\bar{B}_{\mathcal{M}}(p, r)$) denotes $B(p, r) \cap \mathcal{M}$ ($\bar{B}(p, r) \cap \mathcal{M}$). For a given $p \in \mathcal{P}$, $\mathfrak{m}(p)$ denotes the distance of p to its nearest neighbor in $\mathcal{P} \setminus \{p\}$, i.e.,

$$\mathfrak{m}(p) = \min_{x \in \mathcal{P}, x \neq p} \|x - p\|.$$

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

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

where u and v are vectors in U and V respectively. This is the largest *principal angle* between U and V . In case of affine spaces, the angle between two of them is defined as the angle between the corresponding parallel vector subspaces.

The following lemma is a direct consequence of the above definition. We have included the proof for completeness.

Lemma 2.1 *Let U and V be affine spaces of \mathbb{R}^d with $\dim(U) \leq \dim(V)$, and let U^\perp and V^\perp are 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.,

$$\angle(v_*, u_*) = \frac{\pi}{2} - \angle(v_*, u). \quad (1)$$

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 Eq. (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

2.1 Weighted Delaunay triangulation

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

We define the *relative amplitude* of ω as

$$\tilde{\omega} = \max_{p \in \mathcal{P}, q \in \mathcal{P} \setminus \{p\}} \frac{\omega(p)}{\|p - q\|}. \quad (2)$$

In the paper, we will assume that $\tilde{\omega} \leq \omega_0$, for some constant $\omega_0 \in [0, 1/2)$ (see Hypothesis 3.2). Under this hypothesis, all the balls bounded by weighted spheres are disjoint.

Given a subset τ of $d + 1$ weighted points whose centers are affinely independent, there exists a unique sphere orthogonal to the weighted points of τ . The sphere is called the *orthosphere* of τ and its center o_τ and radius Φ_τ are called the *orthocenter* and the *orthoradius* of τ . If the weights of the vertices of τ are 0 (or all equal), then the orthosphere is simply the *circumscribing sphere* of τ whose center and radius are respectively called *circumcenter* and *circumradius*. If τ is a j -simplex, $j < d$, the orthosphere of τ is the smallest sphere that is orthogonal to the (weighted) vertices of τ . Its center o_τ lies in $\text{aff}(\tau)$. Note that a simplex τ may have an orthoradius which is imaginary, i.e. $\Phi_\tau^2 < 0$. This situation however cannot happen if the relative amplitude of the weight function is $< 1/2$.

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

2.1.2 Weighted Voronoi diagram and Delaunay triangulation

Let ω be a weight function defined on \mathcal{P} . We define the weighted Voronoi cell of $p \in \mathcal{P}$ as

$$\text{Vor}^\omega(p) = \{x \in \mathbb{R}^d : \|p - x\|^2 - \omega(p)^2 \leq \|q - x\|^2 - \omega(q)^2, \forall q \in \mathcal{P}\}. \quad (3)$$

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 \mathcal{P} [Aur87].

Let τ be a subset of points of \mathcal{P} and write

$$\text{Vor}^\omega(\tau) = \bigcap_{x \in \tau} \text{Vor}^\omega(x).$$

We say \mathcal{P}^ω is in *general position*¹ if $\text{Vor}^\omega(\tau) = \emptyset$ when $|\tau| > d + 1$. The collection of all simplices $\text{conv}(\tau)$ such that $\text{Vor}^\omega(\tau) \neq \emptyset$ constitutes a triangulation of $\text{conv}(\mathcal{P})$ if $\dim \text{aff}(\mathcal{P}) = d$, and this triangulation is called the weighted Delaunay triangulation $\text{Del}^\omega(\mathcal{P})$. The mapping that associates to the face $\text{Vor}^\omega(\tau)$ of $\text{Vor}^\omega(\mathcal{P})$ the face $\text{conv}(\tau)$ of $\text{Del}^\omega(\mathcal{P})$ is a *duality*, i.e. a bijection that reverses the inclusion relation. See [Del34, Aur87, BDG12].

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

Observe that the definition of weighted Voronoi diagrams makes sense if, for some $p \in \mathcal{P}$, $\omega(p)^2 < 0$, 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 positive weight functions with relative amplitude $< 1/2$.

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]. In this paper, 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 2.2). 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⁺00, CDR05] and Section 3).

Lemma 2.2 *Let H be a k -dimensional affine space of \mathbb{R}^d . The restriction of the weighted Voronoi diagram of $\mathcal{P} = \{p_0, \dots, p_m\}$ to H is identical to the k -dimensional weighted Voronoi diagram of $\mathcal{P}' = \{p'_0, \dots, p'_m\}$ in H , where p'_i denotes the orthogonal projection of p_i onto H and the squared weight of p'_i is*

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

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

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

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

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 $p_i \in \mathcal{P}$ onto H . See Figure 1. Hence the restriction of $\text{Vor}^\omega(\mathcal{P})$ to H is the weighted Voronoi diagram $\text{Vor}^\xi(\mathcal{P}')$ of the points \mathcal{P}' with the weight function: $\xi : \mathcal{P}' \rightarrow [0, \infty)$, with $\xi(p'_i)^2 = -\|p_i - p'_i\|^2 + \omega(p_i)^2 + \lambda^2$ where $\lambda = \max_{p_j \in \mathcal{P}} \|p_j - p'_j\|$. \square

¹This is an extension of the general position notion introduced by Delaunay [Del34] for Delaunay triangulations. See also [BDG12].




Figure 1: Refer to Lemma 2.2. The grey line denotes the k -dimensional plane H and the black line denotes $\text{Vor}^\omega(p_i p_j)$.

2.2 Sampling conditions

2.2.1 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 1-*Lipschitz*, 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 paper, we assume that \mathcal{M} is a compact submanifold of \mathbb{R}^d of (strictly) positive reach. Note that the class of submanifolds with positive reach includes C^2 -smooth submanifolds [Fed59].

2.2.2 Sampling parameters

The point sample \mathcal{P} is said to be an (ε, δ) -sample (where $0 < \delta < \varepsilon < 1$) if (1) for any point $x \in \mathcal{M}$, there exists a point $p \in \mathcal{P}$ such that $\|x - p\| \leq \varepsilon \text{lfs}(x)$, and (2) for any two distinct points $p, q \in \mathcal{P}$, $\|p - q\| \geq \delta \text{lfs}(p)$. The parameter ε is called the *sampling rate*, δ the *sparsity*, and ε/δ the *sampling ratio* of the sample \mathcal{P} .

The following lemma, proved in [GW04], states basic properties of manifold samples. As before, we write $\mathbf{m}(p)$ for the distance between $p \in \mathcal{P}$ and its nearest neighbor in $\mathcal{P} \setminus \{p\}$.

Lemma 2.3 *Given an (ε, δ) -sample \mathcal{P} of \mathcal{M} , we have*

1. For all $p \in \mathcal{P}$, we have

$$\delta \text{lfs}(p) \leq \mathbf{m}(p) \leq \frac{2\varepsilon}{1 - \varepsilon} \text{lfs}(p).$$

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

$$\|x - x'\| \leq 2t^2 \text{lfs}(p).$$

2.3 Fat simplices

Consider a j -simplex τ , where $1 \leq j \leq k+1$. We denote by $R_\tau, \Delta_\tau, L_\tau, V_\tau$ and $\Gamma_\tau = \Delta_\tau/L_\tau$ the circumradius, the longest edge length, the shortest edge length, the j -dimensional volume, and *spread* i.e. the longest edge to shortest edge ratio of τ respectively.

We define the *fatness* of a j -dimensional simplex τ as

$$\Theta_\tau = \begin{cases} 1 & \text{if } j = 0 \\ V_\tau^{1/j}/\Delta_\tau & \text{if } j > 0 \end{cases} \quad (4)$$

The following important lemma is due to Whitney [Whi57, Chap. II].

Lemma 2.4 *Let $\tau = [p_0, \dots, p_j]$ be a j -dimensional simplex and let H be an affine flat such that τ is contained in the offset of H by η (i.e. any point of τ is at distance at most η from H). If u is a unit vector in $\text{aff}(\tau)$, then there exists a unit vector u_H in H such that*

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

We deduce from the above lemma the following corollary. See also Lemma 1 in [Fu93] and Lemma 16 in [CDR05].

Corollary 2.5 (Tangent space approximation) *Let τ be a j -simplex, $j \leq k$, with vertices on \mathcal{M} , and let p be vertex of τ . Assuming that $\Delta_\tau < \text{lfs}(p)$, we have*

$$\sin \angle(\text{aff}(\tau), T_p) \leq \frac{\Delta_\tau^2}{\Theta_\tau^k L_\tau \text{lfs}(p)} = \frac{\Gamma_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(p)}.$$

Proof. It suffices to take $H = T_p$ and to use $\eta = \Delta_\tau^2/2 \text{lfs}(p)$ (from Lemma 2.3 (2)) and $\Gamma_\tau = \Delta_\tau/L_\tau$. Hence

$$\sin \angle(\text{aff}(\tau), T_p) \leq \frac{2\eta}{(j-1)! \Theta_\tau^j L_\tau} \leq \frac{2\eta}{\Theta_\tau^j L_\tau} \leq \frac{\Delta_\tau^2}{\Theta_\tau^j L_\tau \text{lfs}(p)} = \frac{\Gamma_\tau \Delta_\tau}{\Theta_\tau^j \text{lfs}(p)}$$

□

A sliver is a special type of flat simplex. The property of being a sliver is defined in terms of a parameter Θ_0 , to be fixed later in Section 4.

The following definition is a variant of a definition given in [Li03].

Definition 2.6 (Θ_0 -fat simplices and Θ_0 -slivers) *Given a positive parameter Θ_0 , a simplex τ is said to be Θ_0 -fat if the fatness of τ and of all its subsimplices is at least Θ_0 .*

A simplex of dimension at least 2 which is not Θ_0 -fat but whose subsimplices are all Θ_0 -fat is called a Θ_0 -sliver.




Figure 2: \mathcal{M} is the black curve. The sample \mathcal{P} is the set of small circles. The tangent space at u is denoted by T_u . The Voronoi diagram of the sample is in grey. The edges of the Delaunay triangulation $\text{Del}(\mathcal{P})$ are the line segments between small circles. In bold, $\text{st}(u) = \{uv, uw\}$.

2.4 Tangential Delaunay complex and inconsistent configurations

Let \mathcal{P} be a point sample of \mathcal{M} and u be a point of \mathcal{P} . We denote by $\text{Del}_u^\omega(\mathcal{P})$ the weighted Delaunay triangulation of \mathcal{P} restricted to the tangent space T_u . Equivalently, the simplices of $\text{Del}_u^\omega(\mathcal{P})$ are the simplices of $\text{Del}^\omega(\mathcal{P})$ whose Voronoi dual faces intersect T_u , i.e., $\tau \in \text{Del}_u^\omega(\mathcal{P})$ iff $\text{Vor}^\omega(\tau) \cap T_u \neq \emptyset$. Observe that, from Lemma 2.2, $\text{Del}_u^\omega(\mathcal{P})$ is in general isomorphic to a weighted Delaunay triangulation of a k -dimensional convex hull. Since this situation can always be ensured by applying some infinitesimal perturbation on \mathcal{P} or of the weights, in addition to the general position assumption of \mathcal{P}^ω we will also assume, in the rest of the paper, that all $\text{Del}_u^\omega(\mathcal{P})$ are isomorphic to a weighted Delaunay triangulation of a k -dimensional convex hull. We will discuss this in more details at the beginning of Section 5. Finally, write $\text{st}(u)$ for the *star* of u in $\text{Del}_u^\omega(\mathcal{P})$, i.e., $\text{st}(u) = \text{st}(u, \text{Del}_u^\omega(\mathcal{P}))$ (see Figure 2).

We denote by *tangential Delaunay complex* or *tangential complex* for short, the simplicial complex $\{\tau : \tau \in \text{st}(u), u \in \mathcal{P}\}$. We denote it by $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$. By our assumption above, $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ is a k -dimensional subcomplex of $\text{Del}^\omega(\mathcal{P})$. Note that $\text{st}(u) = \text{st}(u, \text{Del}_u^\omega(\mathcal{P}))$ and $\text{st}(u, \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P}))$, the star of u in the complex $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, are in general different.

By duality, computing $\text{st}(u)$ is equivalent to computing the restriction to T_u of the (weighted) Voronoi cell of u , which, by Lemma 2.2, reduces to computing a cell in a k -dimensional weighted Voronoi diagram embedded in T_u . To compute such a cell, we need to compute the intersection of $|\mathcal{P}| - 1$ halfspaces of T_u where $|\mathcal{P}|$ is the cardinality of \mathcal{P} . Each halfspace is bounded by the bisector consisting of the points of T_u that are at equal weighted distance from u^ω and some other point in \mathcal{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} .

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 τ of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ with two vertices u and v such that τ is in $\text{st}(u)$ but not in $\text{st}(v)$ (refer to Figure 3). We write $B_u(\tau)$ (and $B_v(\tau)$) for the open ball centered on T_u (and T_v) that is orthogonal to the (weighted) vertices of τ , and denote by $m_u(\tau)$ (and $m_v(\tau)$), or m_u (m_v) for short, its center.

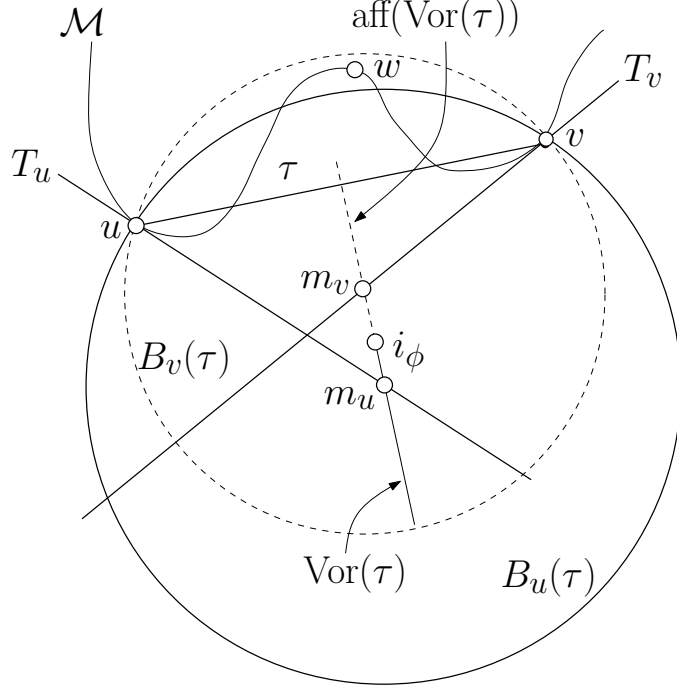


Figure 3: An inconsistent configuration in the unweighted case. Edge $[uv]$ is in $\text{Del}_u(\mathcal{P})$ but not in $\text{Del}_v(\mathcal{P})$ since $\text{Vor}(uv)$ intersects T_u but not T_v . This happens because $[m_u m_v]$ penetrates (at i_ϕ) the Voronoi cell of a point $w \neq u, v$, therefore creating an inconsistent configuration $\phi = [u, v, w]$. Also note that i_ϕ is the center of an empty sphere circumscribing simplex ϕ .

According to our definition, τ is inconsistent iff $B_u(\tau)$ is further than orthogonal from all weighted points in $\mathcal{P}^\omega \setminus \tau$ while there exists a weighted point in $\mathcal{P}^\omega \setminus \tau$ that is closer than orthogonal from $B_v(\tau)$. 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 $\mathcal{P}^\omega \setminus \tau$.

We now formally define an inconsistent configuration as follows.

Definition 2.7 (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 -simplex $\tau = \phi \setminus \{w\}$ is in $\text{st}(u)$ but not in $\text{st}(v)$.*
- *$\text{Vor}^\omega(w)$ is one of the first weighted Voronoi cells of $\text{Vor}^\omega(\mathcal{P})$, other than the weighted Voronoi cells of the vertices of τ , that is intersected by the line segment $[m_u m_v]$ oriented from m_u to m_v . Here $m_u = T_u \cap \text{Vor}^\omega(\tau)$ and $m_v = T_v \cap \text{aff}(\text{Vor}^\omega(\tau))$. Let i_ϕ denote the point where the oriented segment $[m_u m_v]$ first intersects $\text{Vor}^\omega(w)$.*
- *τ is a Θ_0 -fat simplex.*

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




Figure 4: In Figure (a), \mathcal{M} is the black curve, the sample \mathcal{P} is the set of small circles, the tangent space at a point $x \in \mathcal{P}$ is denoted by T_x and the Voronoi diagram of the sample is in grey and $\text{Del}_{T,\mathcal{M}}(\mathcal{P})$ is the line segments between the sample points. In dashed lines, are the inconsistent simplices in $\text{Del}_{T,\mathcal{M}}(\mathcal{P})$. In Figure (b), the grey triangles denote the inconsistent configurations corresponding to the inconsistent simplices in Figure (a).

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

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 $\text{Inc}^\omega(\mathcal{P})$ for the subcomplex of $\text{Del}^\omega(\mathcal{P})$ consisting of all the Θ_0 -inconsistent configurations of $\text{Del}_{T,\mathcal{M}}^\omega(\mathcal{P})$ and their subfaces. We also define the *completed complex* $C^\omega(\mathcal{P})$ as

$$C^\omega(\mathcal{P}) = \text{Del}_{T,\mathcal{M}}^\omega(\mathcal{P}) \cup \text{Inc}^\omega(\mathcal{P}).$$

Refer to Figure 4.

An important observation, stated as Lemma 4.9 in Section 4.3, 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 Manifold reconstruction

The algorithm removes all Θ_0 -slivers from $C^\omega(\mathcal{P})$ by weighting the points of \mathcal{P} . By the observation mentioned 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 Section 5.

In this section, we describe the algorithm. Its analysis is deferred to Section 4.

We will make the two following hypotheses.

Hypothesis 3.1 \mathcal{M} is a compact smooth submanifold of \mathbb{R}^d without boundary, and \mathcal{P} is an (ε, δ) -sample of \mathcal{M} of sampling ratio $\varepsilon/\delta \leq \eta_0$ for some positive constant η_0 .

As shown in [CWW08, GW04], we can estimate the tangent space T_p at each sample point p and also the dimension k of the manifold from \mathcal{P} and η_0 . We assume now that T_p , for any point $p \in \mathcal{P}$, k and η_0 are known.

The algorithm fixes a bound ω_0 on the relative amplitude of the weight assignment :

Hypothesis 3.2 $\tilde{\omega} \leq \omega_0$, for some constant $\omega_0 \in [0, 1/2)$

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

The algorithm also fixes Θ_0 to a constant defined in Theorem 4.15, that depends on k , ω_0 and η_0 .

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

$$LN(p) = \{q \in \mathcal{P} : |B(p, \|p - q\|) \cap \mathcal{P}| \leq N\}. \quad (5)$$

where N is a constant that depends on k , ω_0 and η_0 to be defined in Section 4.4. We will show in Lemma 4.12, that $LN(p)$ includes all the points of \mathcal{P} that can form an edge with p in $C^\omega(\mathcal{P})$. In fact, the algorithm can use instead of $LN(p)$ any subset of \mathcal{P} that contains $LN(p)$. This will only affect the complexity of the algorithm, not the output.

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

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

The function **build_inconsistent_configurations**(u, τ) adds to $C^\omega(\mathcal{P})$ all the inconsistent configurations of the form $\phi = \tau \cup \{w\}$ where τ is an inconsistent simplex of $\text{st}(u)$. More precisely, for each vertex $v \neq u$ of τ such that $\tau \not\subseteq \text{st}(v)$, we calculate the points $w \in LN(p)$, such that (u, v, w) witnesses the inconsistent configuration $\phi = \tau \cup \{w\}$. Specifically, we compute the restriction of the Voronoi diagram of the points in $LN(u)$ to the line segment $[m_u m_v]$, where $m_u = T_u \cap \text{aff}(\text{Vor}^\omega(\tau))$ and $m_v = T_v \cap \text{aff}(\text{Vor}^\omega(\tau))$. 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 $\phi = \tau \cup \{w\}$ 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⁺00] for

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

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

Algorithm 2 Function `update_completed_complex`(Q, ω)

```
for each point  $q \in Q$  do
  build_star( $q$ )
end for
for each  $q \in Q$  do
  for each  $k$ -simplex  $\tau$  in  $\text{st}(q)$  do
    if  $\tau$  is  $\Theta_0$ -fat and  $\exists v \in \tau, \tau \notin \text{st}(v)$  then
      //  $\tau$  is inconsistent
      build_inconsistent_configurations( $q, \tau$ )
    end if
  end for
end 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 $C^\omega(\mathcal{P})$ that becomes incident to p when the weight of p is varied from 0 to $\omega_0 \mathfrak{m}(p)$, keeping the weights of all the points in $\mathcal{P} \setminus \{p\}$ fixed. Note that a candidate simplex of p is incident to p for some weight $\omega(p)$ but does not necessarily belong to $\text{st}(p)$.

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

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  $C^\omega(\mathcal{P})$  contains
// no  $\Theta_0$ -sliver incident to  $p$ 
 $J(p) \leftarrow [0, \omega_0^2 \mathfrak{m}(p)^2] \setminus \bigcup_{\tau \in S(p)} W(p, \tau)$ 
 $\omega(p)^2 \leftarrow$  a squared weight from  $J(p)$ 
return  $\omega(p)$ 

```

The function `candidate_slivers`(p, ω) varies the weight of p and computes all the candidate slivers of p and their corresponding weight intervals $W(p, \tau)$. 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 $\omega_0 \mathfrak{m}(p)$, 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 $0 = P_0 < P_1 < \dots < P_{n-1} < P_n = \omega_0 \mathfrak{m}(p)$.

2. We determine the next weight assignment of p in the following way. For each new simplex τ currently incident to p , we keep it in a priority queue ordered by the weight of p at which τ 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 τ of p which is detected, we compute $W(p, \tau)$ on the fly. The candidate slivers are stored in a list with their corresponding forbidden intervals.

4 Analysis of the algorithm

The analysis of the algorithm relies on structural results that will be proved in Sections 4.2, 4.3 and 4.4. We will then prove that the algorithm is correct and analyze its complexity in Sections 4.5 and 4.6. In Section 5, we will show that the output $\hat{\mathcal{M}}$ of the reconstruction algorithm is a good approximation of \mathcal{M} .

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.

4.1 First lemmas

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⁺00].

Lemma 4.1 *Assume that Hypothesis 3.2 is satisfied. If τ is a simplex of $\text{Del}^\omega(\mathcal{P})$ 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 5: For the proofs of Lemma 4.1.

Proof. Refer to Figure 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

Altitude and fatness

If $p \in \tau$, we define $\tau_p = \tau \setminus \{p\}$ to be the $(j-1)$ -face of τ opposite to p . We also write $D_\tau(p)$ for the distance from p to the affine hull of τ_p . $D_\tau(p)$ will be called the *altitude* of p in τ .

From the definition of fatness, we easily derive the following lemma.

Lemma 4.2 *Let $\tau = [p_0, \dots, p_j]$ be a j -dimensional simplex and p be a vertex of τ .*

1. $\Theta_\tau^j \leq \frac{1}{j!}$
2. $j! \Theta_\tau^j \leq \frac{D_\tau(p)}{\Delta_\tau} \leq j \Gamma_\tau^{j-1} \times \frac{\Theta_\tau^j}{\Theta_{\tau_p}^{j-1}}$

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

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

2. Using the above inequality and the definition of fatness, we get

$$D_\tau(p) = \frac{j V_\tau}{V_{\tau_p}} \geq \frac{j \Theta_\tau^j \Delta_\tau^j}{\frac{\Delta_\tau^{j-1}}{(j-1)!}} \geq j! \Theta_\tau^j \Delta_\tau.$$

Moreover, using $\Delta_{\tau_p} \geq L_{\tau_p} \geq L_\tau$ and $\Gamma_\tau = \frac{\Delta_\tau}{L_\tau}$,

$$\frac{D_\tau(p)}{\Delta_\tau} = \frac{j V_\tau}{\Delta_\tau V_{\tau_p}} = j \frac{\Theta_\tau^j \Delta_\tau^{j-1}}{\Theta_{\tau_p}^{j-1} \Delta_{\tau_p}^{j-1}} \leq j \frac{\Theta_\tau^j \Delta_\tau^{j-1}}{\Theta_{\tau_p}^{j-1} L_\tau^{j-1}} = j \Gamma_\tau^{j-1} \times \frac{\Theta_\tau^j}{\Theta_{\tau_p}^{j-1}}.$$

□

Excentricity

Let τ be a simplex and p be a vertex of τ . We define the *excentricity* $H_\tau(p, \omega(p))$ of τ with respect to p as the signed distance from o_τ to $\text{aff}(\tau_p)$. Hence, $H_\tau(p, \omega(p))$ is positive if o_τ and p lie on the same side of $\text{aff}(\tau_p)$ and negative if they lie on different sides of $\text{aff}(\tau_p)$.

The following lemma is a generalization of Claim 13 from [CDE⁺00]. It bounds the excentricity of a simplex τ as a function of $\omega(p)$ where p is a vertex of τ . The proof is included for completeness even though it is exactly the same as the one given in [CDE⁺00].

Lemma 4.3 *Let τ be a simplex of $\text{Del}^\omega(\mathcal{P})$ and let p be any vertex of τ . We have*

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

Proof. Refer to Figure 6. For convenience, we write $R = \Phi_\tau$ for the orthoradius of τ and $r = \Phi_{\tau_p}$ for the orthoradius of τ_p . The orthocenter o_{τ_p} of τ_p is the projection of o_τ onto τ_p . When the weight $\omega(p)$ varies while the weights of other points remain fixed, o_τ moves on a (fixed) line L that passes through o_{τ_p} . Now, let p' and p'' be the projections of p onto L and $\text{aff}(\tau_p)$ respectively.




Figure 6: For the proof of Lemma 4.3.

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 $\text{aff}(\tau)$, $\|p' - o_{\tau_p}\| = \|p - p''\| = D_\tau(p)$.

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

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

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

4.2 Properties of the tangential Delaunay complex

The following two lemmas are slight variants of results of [CDR05]. The first lemma states that the restriction of a (weighted) Voronoi cell to a tangent space is small.

Lemma 4.4 *Assume that Hypotheses 3.1 and 3.2 are satisfied. There exists a positive constant C_1 such that for all $T_p \cap \text{Vor}^\omega(p)$, $\|x - p\| \leq C_1 \varepsilon \text{lfs}(p)$.*

Proof. Assume for a contradiction that there exists a point $x \in \text{Vor}^\omega(p) \cap T_p$ s.t. $\|x - p\| > C_1 \varepsilon \text{lfs}(p)$ with

$$C_1(1 - C_1\varepsilon) > 2 + C_1\varepsilon(1 + C_1\varepsilon) \quad (6)$$

Let q be a point on the line segment $[px]$ such that $\|p - q\| = C_1 \varepsilon \text{lfs}(p)/2$. Let q' be the point closest to q on \mathcal{M} . From Lemma 2.3, we have $\|q - q'\| \leq C_1^2 \varepsilon^2 \text{lfs}(p)/2$.




Figure 7: Refer to Lemma 4.4. x' is a point on the line segment such that $\|p - x'\| = C_1 \varepsilon \text{lfs}(p)$, $L = C_1 \varepsilon \text{lfs}(p)$, $\angle pax' = \pi/2$ and $\angle ptx \geq \angle ptx' > \pi/2$.

Hence,

$$\|p - q'\| \leq \|p - q\| + \|q - q'\| < \frac{C_1}{2} \varepsilon (1 + C_1 \varepsilon) \text{lfs}(p)$$

Since \mathcal{P} is an ε -sample, there exists a point $t \in \mathcal{P}$ s.t. $\|q' - t\| \leq \varepsilon \text{lfs}(q')$. Using the fact that lfs is 1-Lipschitz and Eq. (6),

$$\text{lfs}(q') \leq \text{lfs}(p) + \|p - q'\| < (1 + \frac{C_1}{2} \varepsilon (1 + C_1 \varepsilon)) \text{lfs}(p) < \frac{C_1}{2} (1 - C_1 \varepsilon) \text{lfs}(p),$$

which yields $\|q' - t\| < \frac{C_1}{2} \varepsilon (1 - C_1 \varepsilon) \text{lfs}(p)$. We thus have

$$\|q - t\| \leq \|q - q'\| + \|q' - t\| < \frac{C_1}{2} \varepsilon \text{lfs}(p).$$

It follows (see Figure 7), that $\angle ptx > \pi/2$, which implies that

$$\|x - p\|^2 - \|x - t\|^2 - \|p - t\|^2 > 0.$$

Hence,

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

This implies $x \notin \text{Vor}^\omega(p)$, which contradicts our initial assumption. We conclude that $\text{Vor}^\omega(p) \cap T_p \subseteq B(p, C_1 \varepsilon \text{lfs}(p))$ if Eq. (6) 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 3.2 and 3.1, the simplices of $\text{Del}_p^\omega(\mathcal{P})$ are small, have a good radius-edge ratio and a small excentricity.

Lemma 4.5 *Assume that Hypotheses 3.1 and 3.2 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 τ is a simplex of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, then $\Phi_\tau \leq C_3L_\tau$ and $\Gamma_\tau = \Delta_\tau/L_\tau \leq C_3$.*
3. *If τ is a simplex of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ and p a vertex of τ , the eccentricity $|H_\tau(p, \omega(p))|$ is at most $C_4\varepsilon\text{lfs}(p)$.*

Proof. 1. We will consider the following two cases:

- (a) Consider first the case where pq is an edge of $\text{Del}_p^\omega(\mathcal{P})$. Then $T_p \cap \text{Vor}^\omega(pq) \neq \emptyset$. Let $x \in T_p \cap \text{Vor}^\omega(pq)$. From Lemma 4.4, we have $\|p - x\| \leq C_1\varepsilon\text{lfs}(p)$. By Lemma 4.1,

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

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

- (b) From the definition of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, there exists a vertex r of τ such that $[pq] \in \text{st}(r)$. From (a), $\|r - p\|$ 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 1 (a)), 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 $\tau_1 \in \text{st}(p)$ be a simplex incident to p and $\tau \subseteq \tau_1$. Let $z \in \text{Vor}^\omega(\tau_1) \cap T_p$, 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 τ_1 . From Lemma 4.1 (2) and (3), we have $r_z \geq \Phi_{\tau_1}$ and $\Phi_\tau \leq \Phi_{\tau_1}$. Hence it suffices to prove that there exists a constant C_3 such that $r_z \leq C_3L_\tau$. Since $z \in \text{Vor}^\omega(\tau) \cap T_p$, we deduce from Lemma 4.4 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 τ , 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 τ , we have, using Lemma 2.3 and Hypothesis 1,

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

From part (1) of the lemma we have $\Delta_{\tau_1} \leq 2C_2\varepsilon\text{lfs}(p)$. Therefore

$$\Gamma_\tau = \frac{\Delta_\tau}{L_\tau} \leq \frac{\Delta_{\tau_1}}{L_\tau} \leq \frac{2C_2\varepsilon\text{lfs}(p)}{\delta\text{lfs}(q)} \leq 4C_2\eta_0 \stackrel{\text{def}}{=} C_3.$$

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 $H_\tau(p, \omega(p))$, we have for all vertices $q \in \tau_p$

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

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

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

□

4.3 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 paper. The first lemma is the analog of Lemma 4.5 applied to inconsistent configurations. Differently from Lemma 4.5, we need to use Corollary 2.5 to control the orientation of the facets of $\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ and require the following additional hypothesis relating the sampling rate ε and the fatness bound Θ_0 .

Hypothesis 4.6 $2A\varepsilon < 1$ where $A \stackrel{\text{def}}{=} 2C_2C_3/\Theta_0^k$, and C_2 and C_3 are the constants defined in Lemma 4.5.

Lemma 4.7 (Tangent approximation) *Let τ be a k -simplex of $\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ that is Θ_0 -fat. For all vertices p of τ , we have $\sin \angle(\text{aff}(\tau), T_p) = \sin \angle(T_p, \text{aff}(\tau)) \leq A\varepsilon$. The constant A , defined in Hypothesis 4.6, depends on k , ω_0 , η_0 and Θ_0 .*

Proof. Since $\dim(\tau) = \dim(T_p) (= k)$, we have from Lemma 2.1, $\angle(\text{aff}(\tau), T_p) = \angle(T_p, \text{aff}(\tau))$. From Lemma 4.5 (1) and the Triangle inequality, we get $\Delta_\tau \leq 2C_2\varepsilon \text{lfs}(p)$, and, from Lemma 4.5 (2), we get $\Gamma_\tau \leq C_3$. Using Corollary 2.5, we obtain

$$\sin \angle(\text{aff}(\tau), T_p) \leq \frac{\Gamma_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(p)} \leq \frac{2C_2C_3\varepsilon}{\Theta_0^k} = A\varepsilon$$

□

Lemma 4.8 *Assume that Hypotheses 3.1, 3.2 and 4.6 are satisfied. Let $\phi \in \text{Inc}^\omega(\mathcal{P})$ 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 $m_u \in (\text{Vor}^\omega(\tau) \cap T_u)$ denote, as in Section 2.4, the point of T_u 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 4.4, $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$. Then, by Lemma 4.1, we have for all vertices $p \in \tau$

$$\|p - o_\tau\| \leq \frac{\|u - o_\tau\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{\|u - m_u\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C_1 \varepsilon \text{lfs}(u)}{\sqrt{1 - 4\omega_0^2}}. \quad (7)$$

Since τ is a Θ_0 -fat simplex (by definition of an inconsistent configuration), we get from Lemma 4.7 that

$$\sin \angle(\text{aff}(\tau), T_p) \leq A \varepsilon \quad (8)$$

for all vertices p of τ . This implies, together with $2A \varepsilon < 1$ (Hypothesis 4.6),

$$\tan^2 \angle(\text{aff}(\tau), T_u) \leq \frac{A^2 \varepsilon^2}{1 - A^2 \varepsilon^2} < 4A^2 \varepsilon^2, \quad (9)$$

Observing again that $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$ (from Lemma 4.4) and Eq. 8, we deduce

$$\|m_u - o_\tau\| \leq \|m_u - u\| \sin \angle(\text{aff}(\tau), T_u) \leq A C_1 \varepsilon^2 \text{lfs}(u). \quad (10)$$

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.s (7) and (9),

$$\|m_v - o_\tau\| \leq \|v - o_\tau\| \tan \angle(\text{aff}(\tau), T_v) < \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \text{lfs}(u) \quad (11)$$

Let i_ϕ denote, as in Section 2.4, the first point of the line segment $[m_u m_v]$ that is in $\text{Vor}^\omega(\phi)$. We get from Eq.s (10) and (11) that

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

1. Using Lemma 4.4, and the facts that $2A \varepsilon < 1$ and $\|u - o_\tau\| \leq \|u - m_u\|$, we get

$$\begin{aligned} \|u - i_\phi\| &\leq \|u - o_\tau\| + \|o_\tau - i_\phi\| \\ &\leq \|u - m_u\| + \|o_\tau - i_\phi\| \\ &\leq \left(C_1 \varepsilon + \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \right) \text{lfs}(u) \\ &\leq \left(C_1 \varepsilon + \frac{C_1 \varepsilon}{\sqrt{1 - 4\omega_0^2}} \right) \text{lfs}(u) \\ &\leq \frac{C_2}{4} \varepsilon \text{lfs}(u) \end{aligned} \quad (12)$$

where C_2 is the constant introduced in Lemma 4.5. Eq. (12), together with Lemma 4.1 and $C'_2 \stackrel{\text{def}}{=} C_2/\sqrt{1-4\omega_0^2} > C_2$, yields

$$\|p - i_\phi\| \leq \frac{\|u - i_\phi\|}{\sqrt{1-4\omega_0^2}} \leq \frac{C'_2}{4} \varepsilon \text{lfs}(u) \quad (13)$$

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 $C'_2 \varepsilon < 1$:

$$|\text{lfs}(p) - \text{lfs}(u)| \leq \|p - u\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(u) \leq \frac{1}{2} \text{lfs}(u). \quad (14)$$

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. (13) and (14), 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 $\text{Del}_{\mathcal{T}\mathcal{M}}^\omega(\mathcal{P})$, the result has been proved in Lemma 4.5 (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 Lemma 4.1 (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. (12), 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 3.1,

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

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

$$\Gamma_\sigma = \frac{\Delta_\sigma}{L_\sigma} \leq \frac{\Delta_\phi}{L_\sigma} \leq 2C'_2 \eta_0$$

We get the result by taking $C'_3 \stackrel{\text{def}}{=} \max\{4C_2\eta_0, 2C'_2\eta_0\}$.

4. Using the same arguments as in the proof of Lemma 4.5 (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 $\|p - q\| \leq 2C'_2 \varepsilon \text{lfs}(p)$ (from part 2), we have

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

□

The next crucial lemma bounds the fatness of inconsistent configurations.

Lemma 4.9 *Assume Hypotheses 3.1, 3.2 and 4.6, and $\varepsilon < 1/2C'_2$. The fatness Θ_ϕ of an inconsistent configuration ϕ is at most*

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

Proof. Let ϕ be witnessed by (u, v, w) . From the definition of inconsistent configurations, the k -dimensional simplex $\tau = \phi \setminus \{w\}$ belongs to $\text{st}(u)$ and τ is a Θ_0 -fat simplex. Using Lemma 4.8 (2), we have $\Delta_\phi \leq 2C'_2 \varepsilon \text{lfs}(u)$. As in the proof of Lemma 4.8, we have $\sin \angle(T_u, \text{aff}(\tau)) \leq \frac{\Gamma_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(u)}$ (refer to Eq. (8)) by Corollary 2.5 and the fact that $\Delta_\phi < \text{lfs}(u)$. Also, using the fact that $\|u - w\| < C'_2 \varepsilon \text{lfs}(u) < \text{lfs}(u)$ (from Lemma 4.5 (2) and $\varepsilon < 1/2C'_2$) and Lemma 2.3 (2), we have

$$\sin \angle(uw, T_u) \leq \frac{\|u - w\|}{2 \text{lfs}(u)} \leq \frac{\Delta_\phi}{2 \text{lfs}(u)}.$$

We can bound the altitude $D_\phi(w)$ of w in ϕ

$$\begin{aligned} D_\phi(w) &= \text{dist}(w, \text{aff}(\tau)) \\ &= \sin \angle(uw, \text{aff}(\tau)) \times \|u - w\| \\ &\leq (\sin \angle(uw, T_u) + \sin \angle(\text{aff}(\tau), T_u)) \times \Delta_\phi \\ &\leq \left(\frac{\Delta_\phi}{2 \text{lfs}(u)} + \frac{\Gamma_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(u)} \right) \Delta_\phi \\ &\leq \frac{\Delta_\phi^2}{2 \text{lfs}(u)} \left(1 + \frac{2\Gamma_\tau}{\Theta_\tau^k} \right). \end{aligned} \tag{15}$$

From the definition of fatness of a simplex and Lemma 4.2 (1), we get

$$V_\tau = \Theta_\tau^k \Delta_\tau^k \leq \frac{\Delta_\tau^k}{k!}. \tag{16}$$

We deduce

$$\begin{aligned} \Theta_\phi^{k+1} &= \frac{V_\phi}{\Delta_\phi^{k+1}} \\ &= \frac{D_\phi(w) V_\tau}{(k+1)} \times \frac{1}{\Delta_\phi^{k+1}} \\ &\leq \frac{\Delta_\phi^2}{2 \text{lfs}(u)} \left(1 + \frac{2\Gamma_\tau}{\Theta_\tau^k} \right) \times \frac{\Delta_\tau^k}{(k+1)! \Delta_\phi^{k+1}} \quad \text{using Eq. (15) and (16)} \\ &\leq \frac{C'_2 \varepsilon}{(k+1)!} \left(1 + \frac{2C_3}{\Theta_0^k} \right). \end{aligned}$$

The last inequality comes from the facts that σ is Θ_0 -fat, $\Delta_\phi \leq 2C'_2 \varepsilon \text{lfs}(u)$ (from Lemma 4.8 (2)) and $\Gamma_\tau \leq C_3$ (from Lemma 4.5 (3)). □

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

Hypothesis 4.10 $\frac{C'_2 \varepsilon}{(k+1)!} \left(1 + \frac{2C_3}{\Theta_0^k}\right) < \Theta_0^{k+1}$

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 3.1, 3.2 and 4.6, we assumed that $2C_2\varepsilon < 1$ in Lemma 4.5, $C'_2\varepsilon < 1$ in Lemma 4.8 and $2C'_2\varepsilon < 1$ in Lemma 4.9. We will make another hypothesis that subsumes these two previous conditions.

Hypothesis 4.11 $C'_2(1 + C'_2\eta_0)\varepsilon < 1/2$.

Observe that this hypothesis implies $C'_2(1 + C'_2)\varepsilon < 1/2$ since $\eta_0 > 1$. Observe also that Hypotheses 4.6, 4.10 and 4.11 are satisfied for fixed ω_0 , η_0 , Θ_0 and a sufficiently small ε .

4.4 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 4.8.

Lemma 4.12 *Assume Hypotheses 3.1, 3.2, 4.6 and 4.11. The set*

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

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

Proof. Lemmas 4.5 and 4.8 show that, in order to construct $\text{st}(p)$ and search for inconsistencies involving p , it is enough to consider the points of \mathcal{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 $B_p \cap \mathcal{P}$.

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

$$\text{lfs}(p)(1 - C'_2\varepsilon) \leq \text{lfs}(x), \text{lfs}(y) \leq \text{lfs}(p)(1 + C'_2\varepsilon). \quad (17)$$

By definition of an (ε, δ) -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 $\bar{B}_x = B_x \cap T_p$, $\bar{B}_y = B_y \cap T_p$ and $\bar{B}_p^+ = B_p^+ \cap T_p$. From Lemma 2.3(2), the distance from x to T_p is

$$\text{dist}(x, T_p) = \|p - x\| \times \sin(\angle px, T_p) \leq C_2'^2 \varepsilon^2 \text{lfs}(p)/2. \quad (18)$$

Using Eq.s (17), (18) and the fact that $\varepsilon/\delta \leq \eta_0$, we see that \bar{B}_x is a k -dimensional ball of squared radius

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

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

$$\begin{aligned} \left(\frac{r^+}{r^-}\right)^k &= \left(\frac{(C'_2\varepsilon + (1 + C'_2\varepsilon)\delta)^2}{\delta^2/4 \times ((1 - C'_2\varepsilon)^2 - C_2'^4 \eta_0^2 \varepsilon^2)}\right)^{k/2} \\ &\leq \left(\frac{4(C'_2\eta_0 + (1 + C'_2\varepsilon))^2}{(1 - C'_2\varepsilon)^2 - C_2'^4 \eta_0^2 \varepsilon^2}\right)^{k/2} \\ &= \left(\frac{4(C'_2\eta_0 + (1 + C'_2\varepsilon))^2}{(1 - C'_2\varepsilon - C_2'^2 \eta_0 \varepsilon)(1 - C'_2\varepsilon + C_2'^2 \eta_0 \varepsilon)}\right)^{k/2} \\ &\leq (4C'_2\eta_0 + 6)^k \stackrel{\text{def}}{=} N \quad (\text{using Hypothesis 4.11}) \end{aligned}$$

And the result follows. \square

4.5 Correctness of the algorithm

Definition 4.13 (Sliverity range) *Let ω be a weight assignment satisfying Hypothesis 3.2. The weight of all the points in $\mathcal{P} \setminus \{p\}$ are fixed and the weight $\omega(p)$ of p is varying. The sliverity range $\Sigma(p)$ of a point $p \in \mathcal{P}$ 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 $C^\omega(\mathcal{P})$.*

Lemma 4.14 *Under Hypotheses 3.1, 3.2, 4.6, 4.10 and 4.11, the sliverity range satisfies*

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

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

Proof. Let τ be a j -dimensional simplex of $C^\omega(\mathcal{P})$ incident on p (with $2 \leq j \leq k+1$). assume that τ is a Θ_0 -sliver. If $\omega(p)$ is the weight of p , we write $H_\tau(p, \omega(p))$ for the excentricity of τ with respect to p and parameterized by $\omega(p)$. From Lemma 4.8(4), we have

$$|H_\tau(p, \omega(p))| \leq C'_4 \varepsilon \text{ifs}(p) \stackrel{\text{def}}{=} D \quad (19)$$

Using Lemma 4.2 (2), we have

$$D_p(\tau) \leq j \Gamma_\tau^{j-1} \Delta_\tau \times \frac{\Theta_\tau^j}{\Theta_{\tau_p}^{j-1}} \leq (k+1) C_3^k \Theta_0 \Delta_\tau \stackrel{\text{def}}{=} E \quad (20)$$

The last inequality follows from the facts that $j \leq k+1$, $\Gamma_\tau \leq C'_3$ (from Lemmas 4.5 (2) and 4.8 (3)) and τ is a Θ_0 -sliver. Moreover, from Lemma 4.3,

$$H_\tau(p, \omega(p)) = H_\tau(p, 0) - \frac{\omega(p)^2}{2D_p(\tau)}. \quad (21)$$

It then follows from Eq. (19), (20) and (21) that the set of squared weights of p for which τ belongs to $C^\omega(\mathcal{P})$ is a subset of the following interval

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

where $\beta = 2DE$. Therefore, from Eq. (20) and (21), the measure of the set of weights for which τ belongs to $C^\omega(\mathcal{P})$ is at most

$$2\beta = 4DE = 4(k+1)C_3'^k \Theta_0 \Delta_\tau C_4' \varepsilon \text{lfs}(p).$$

Let q_1 and q_2 be two vertices of τ such that $\Delta_\tau = \|q_1 - q_2\|$. Using Lemmas 4.5 (1) and 4.8 (2), we get

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

Using this inequality, $\text{lfs}(p) \leq \mathfrak{m}(p)/\delta$ (Lemma 2.3) and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 3.1), the sliverity range of τ is at most $8(k+1)C_3'^k C_2' C_4' \Theta_0 \eta_0^2 \mathfrak{m}(p)^2 = C_5 \Theta_0 \mathfrak{m}(p)^2$ with $C_5 \stackrel{\text{def}}{=} 8(k+1)C_3'^k C_2' C_4' \eta_0^2$. By Lemma 4.12, the number of j -simplices that are incident to p is at most N^j . Hence, the sliverity range of p is at most

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

The last inequality follows from the fact that $\sum_{j=3}^{k+1} N^j < 2N^{k+1}$ (as $N = (4C_2' \eta_0 + 6)^k > 2$). \square

Theorem 4.15 *Take $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$ and assume that Hypotheses 3.1, 3.2, 4.6, 4.10 and 4.11 are satisfied. Then, 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 $2N^{k+1}C_5\Theta_0\mathfrak{m}(p)^2$ from Lemma 4.14. Since $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$, $\Sigma(p)$ is less than the total range of possible squared weights $\omega_0^2 \mathfrak{m}(p)^2$. Hence, Function **weight** (p, ω) will always find a weight for any point $p \in \mathcal{P}$ and any weight assignment of relative amplitude at most ω_0 for the points of $\mathcal{P} \setminus \{p\}$.

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

By Lemma 4.9 and Hypothesis 4.10, all the inconsistent configurations in $C^\omega(\mathcal{P})$ are either Θ_0 -slivers or contain a subfaces which are Θ_0 -slivers. It follows that $\hat{\mathcal{M}}$ has no inconsistency since, when the algorithm terminates, all simplices of $C^\omega(\mathcal{P})$ are Θ_0 -fat. \square

4.6 Time and space complexity

Theorem 4.16 *Assume that Hypotheses 3.1, 3.2, 4.6, 4.10 and 4.11 are satisfied. Then the space complexity of the algorithm is*

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

and the time complexity is

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

where $|\mathcal{P}|$ denotes the cardinality of \mathcal{P} .

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

Rather than maintaining $C^\omega(\mathcal{P})$ as one data structure, for all $p \in \mathcal{P}$, we maintain $\text{st}(p)$ and $\text{cosph}(p)$, where

$$\text{cosph}(p) = \left\{ \begin{array}{l} \text{inconsistent configuration } \phi \mid \\ \phi \text{ is witnessed by } (p, q, r), \text{ for some } q, r \in \mathcal{P} \setminus p \end{array} \right\}$$

Note that

$$C^\omega(\mathcal{P}) = \bigcup_{p \in \mathcal{P}} (\text{st}(p) \cup \text{cosph}(p)).$$

By Lemma 2.2, each $\text{st}(p)$, $p \in \mathcal{P}$, has the same combinatorial complexity as a Voronoi cell in the k -dimensional flat T_p . Since the sites needed to compute this Voronoi cell all belong to $LN(p)$, there number is at most N by Lemma 4.12. 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 $O(kN^{\lfloor k/2 \rfloor})|\mathcal{P}|$.

For a given inconsistent Θ_0 -fat k -simplex in $\text{st}(p)$, we can have from Lemmas 4.8 (2) and 4.12, at most $k|LN(p)| \leq kN$ different inconsistent configurations. Hence, the number of inconsistent configurations to be stored in the completed complex $C^\omega(\mathcal{P})$ is at most $O(k^2N^{\lfloor k/2 \rfloor + 1})|\mathcal{P}|$.

The main loop of the algorithm processes each point $p \in \mathcal{P}$ in turn, and finds a weight assignment $\omega(p)$ using the function **weight**(,). In the function **weight**, we maintain a list of candidate simplices of p with their corresponding forbidden intervals and a priority queue of candidate simplices of p . By Lemmas 4.5 (1), 4.8 (2) and 4.12, the space complexities of both these data structures are bounded by $O(kN^{k+1})$.

With $N = 2^{O(k)}$ (refer to Lemma 4.12), we conclude that the total space complexity of the algorithm is

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

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

Then the algorithm builds $C^\omega(\mathcal{P})$ for the zero weight assignment. The time to compute $\text{st}(p)$ 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 . 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 calculating the projection of a point in \mathbb{R}^d on a k -flat requires to compute 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 \mathcal{P} is

$$O(kdN + k^3(N \log N + N^{\lfloor k/2 \rfloor}))|\mathcal{P}|.$$

Let $\tau = [p_0, \dots, p_k]$ be a Θ_0 -fat k -simplex in $\text{st}(u)$. For each vertex $v (\neq u)$ of τ with $\tau \notin \text{st}(v)$, we need to compute the inconsistent configurations of the form $\phi = [p_0, \dots, p_k, w]$ witnessed by

(u, v, w) where $w \in LN(p) \setminus \tau$. 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 τ and w a point in $LN(u) \setminus \tau$ is

$$O(dkN) \tag{22}$$

The time complexity to build all the inconsistent configurations corresponding to $st(u)$ 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 $C^\omega(\mathcal{P})$ is $O(dN + kN^{\lfloor k/2 \rfloor + 1})|\mathcal{P}|$. Therefore the total time complexity of the initialization phase is

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

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

Combining the time complexities for all the steps of the algorithm and using $N = O2^{O(k)}$ (refer to Lemma 4.12), we get the total time complexity of the algorithm

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

□

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

5 Topological and geometric guarantees

In this section we give conditions under which $\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ is isotopic to and a close geometric approximation of the manifold \mathcal{M} . The results in this section will be used to prove the theoretical guarantees on the quality of the simplicial complex $\hat{\mathcal{M}}$ output by the Algorithm 1 in Section 3.

General position assumption

We assume that the weighted points in \mathcal{P}^ω are in general position in \mathbb{R}^d and this will imply $\text{Del}^\omega(\mathcal{P})$ is a simplicial complex, see Section 2.1.2. Therefore $\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ is a simplicial complex.

We know from Lemma 2.2 that $\text{Vor}^\omega(\mathcal{P}) \cap T_p$ is identical to the weighted Voronoi diagram $\text{Vor}^{\xi_p}(\mathcal{P}_p)$ of \mathcal{P}_p in T_p , where \mathcal{P}_p is obtained by projecting \mathcal{P} orthogonally onto T_p and ξ_p is the weight assignment defined in Lemma 2.2.

Under Hypotheses 3.1 and 3.2, $\text{Vor}^{\xi_p}(p)$ is bounded (Lemma 4.4) and therefore $\dim \text{aff}(\mathcal{P}_p) = k$. In addition to the general position assumption of \mathcal{P}^ω , we will also assume that the weighted points $\mathcal{P}_p^{\xi_p}$ are in *general position* on T_p , meaning that there exists no orthosphere centered on T_p that is orthogonal to $k + 2$ weighted points of $\mathcal{P}_p^{\xi_p}$. See Sections 2.1.1 and 2.1.2. As we have already mentioned in Section 2.4, the general position assumption can be fulfilled by applying an infinitesimal perturbation to the point sample or to the weight assignment ω .

Under the general position assumption above, $\text{Del}^{\xi_p}(\mathcal{P}_p)$ is a triangulation of $\text{conv}(\mathcal{P}_p)$ [Aur87, AE84, BDG12]. Moreover, since $\text{Vor}^{\xi_p}(p)$ is bounded (from Lemma 4.4), p is an interior vertex of $\text{Del}^{\xi_p}(\mathcal{P}_p)$. This implies the following result.

Lemma 5.1 *For all $p \in \mathcal{P}$, $\text{st}(p)$ is isomorphic to the star of an interior vertex of a compact triangulated k -dimensional polytope.*

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 5.2 *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 $d\pi(x) : \mathbb{R}^d \rightarrow T_{\pi(x)}$, where $d\pi(x)$ denotes the derivative of π at x , is parallel to $N_{\pi(x)}$ and has dimension $d - k$.*

Main result

The main result of this section is the following theorem.

Theorem 5.3 *Assume that Hypotheses 3.1 and 3.2 are satisfied as well as 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 $\sin \angle(T_p, \text{aff}(\tau)) = \sin \angle(\text{aff}(\tau), T_p) = O(\varepsilon)$. Note that this property directly follows from Lemma 4.7.*

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: $\forall x \in \mathcal{M}$, $\text{dist}(x, \pi|_{\hat{\mathcal{M}}}^{-1}(x)) = O(\varepsilon^2 \text{ifs}(x))$;

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 .

Since Conditions **C1** and **C2** are satisfied when the conditions of Theorem 4.15 are satisfied, it follows that the simplicial complex $\hat{\mathcal{M}}$ output by Algorithm 1 satisfies the above properties **P1-P5** provided that the additional condition on ε in the above theorem is satisfied.

The rest of the section will be devoted to the proof of Theorem 5.3.

5.1 Piecewise-linear k -manifold

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

Lemma 5.4 *Under condition **C1** of Theorem 5.3 and for a sufficiently small ε , $\text{st}(p) = \text{st}(p, \hat{\mathcal{M}})$ for all $p \in \mathcal{P}$.*

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

Lemma 5.5 (PL k -manifold without boundary) *Assume condition **C1** of Theorem 5.3. If ε is sufficiently small, then $\hat{\mathcal{M}}$ is a PL k -manifold without boundary.*

Proof. Since $\text{Del}^{\xi_p}(\mathcal{P}_p)$ is a triangulation of a k -dimensional polytope, $\text{Del}^{\xi_p}(\mathcal{P}_p)$ is a PL k -manifold [Zee66, Lem. 8 of Chap. 3] and, since p is an interior vertex of $\text{Del}^{\xi_p}(\mathcal{P}_p)$, $\text{lk}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$ is a PL $(k-1)$ -sphere [Zee66, Proof of Lem. 9 of Chap. 3].

From Lemma 5.4, we have $\text{st}(p) = \text{st}(p, \hat{\mathcal{M}})$. By Lemma 5.1, $\text{st}(p, \hat{\mathcal{M}})$ 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. This directly implies that $\hat{\mathcal{M}}$ is a PL k -manifold with no boundaries. \square

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

5.2 Homeomorphism

Following lemma establishes Property **P3** of Theorem 5.3, and the rest of the section is devoted to its proof.

Lemma 5.6 (Homeomorphism) *Assume conditions **C1** and **C2** of Theorem 5.3. For ε sufficiently small, the map π restricted to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} .*

In the broad outline, the proof of Lemma 5.6 is similar to the proof of homeomorphism given in [ACDL02] but the technical details are quite different in most places. The difficulties that arose in getting results analogous to the ones in [ACDL02] were handled using ideas from [Whi57, Chap. II].

In the rest of this paper, we assume that $\omega_0, \eta_0, \Theta_0$ are fixed constants, and that ε is small enough. We use the asymptotic notations $O()$ and $\Omega()$ hiding the constants.

Remark 5.7 *For the details of the missing proofs of Lemmas 5.8, 5.9, 5.10, and 5.11 refer to [Gho12, Chap. 4].*

5.2.1 Basic lemmas

For a point $p \in \mathcal{P}$, let

$$\pi_p : \mathbb{R}^d \rightarrow T_p$$

denotes the orthogonal projection of \mathbb{R}^d onto T_p , and

$$\tilde{\pi}_p : T_p \setminus \mathcal{O} \rightarrow \mathcal{M}$$

denotes the restriction of π to $T_p \setminus \mathcal{O}$.

We define in addition

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

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

The following lemma is a generalization of Proposition 6.2 in [NSW08] which bounds the variation of the angle between tangent spaces between two points on the manifold \mathcal{M} .

Lemma 5.8 (Tangent variation) *Let $p, q \in \mathcal{M}$ and $\|p - q\| = t \text{lfs}(p)$ with $t \leq 1$. Then,*

$$\sin \angle(T_q, T_p) = \sin \angle(T_p, T_q) < O(t).$$

Using Lemma 5.8, we can show that the derivative of the restriction $\tilde{\pi}_p$ of π_p to a neighborhood of p is nonsingular and that $\tilde{\pi}_p$ is injective. Using the Inverse function theorem this implies that $\tilde{\pi}_p$ is a diffeomorphism.

Lemma 5.9 ($\tilde{\pi}_p$ is a C^1 -diffeomorphism) *There exists absolute constants r_1 and r_2 such that $\tilde{\pi}_p$ restricted to $\tilde{T}_p \stackrel{\text{def}}{=} B(p, r_1 \text{lfs}(p)) \cap T_p$ is a diffeomorphism onto $\tilde{\pi}_p(\tilde{T}_p)$ and $\tilde{\pi}_p(\tilde{T}_p) \supseteq \tilde{\mathcal{M}}_p \stackrel{\text{def}}{=} B_{\mathcal{M}}(p, r_2 \text{lfs}(p))$.*

Note that $\tilde{\mathcal{M}}_p \subseteq \tilde{\pi}_p(\tilde{T}_p)$ follows from Lemma 2.3 (3).

The next lemma follows directly from Lemmas 2.3 4.7, 5.8 and 5.9.

Lemma 5.10 *Under Conditions **C1** and **C2** of Theorem 5.3, for any $x \in \text{st}(p, \hat{\mathcal{M}})$.*

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

We will need the following technical lemma bounding the fatness of a perturbed simplex. A similar inequality can be found in the proof of Lemma 14c in [Whi57, Chapter 4].

Lemma 5.11 *Let $\sigma = [p_0, \dots, p_j]$ be a j -dimensional simplex of \mathbb{R}^d , and let $\sigma' = [p'_0, \dots, p'_j]$ where $\|p_i - p'_i\| \leq \rho \Delta_\sigma$ with $\rho < 1/2$, for $0 \leq i \leq j$. Then,*

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

This directly implies that, if $\rho \leq \frac{\Theta_\sigma^{2j}}{j^{2j+1}} (1 - \frac{1}{2j})$, then $\Theta_{\sigma'} \geq \Theta_\sigma/2$.

We will also 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. The proof can be found in [Whi57, Lem. 14b of Chap. 4].

Lemma 5.12 *Let $\sigma = [p_0, \dots, p_j]$ be a j -simplex and let p be a point of sigma. Writing $p = \sum_{i=0}^j \mu_i p_i$ with $\sum_{i=0}^j \mu_i = 1$ and $\mu_i \geq 0$, we have*

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

We will use the standard notion of *oriented simplices*, *positively oriented simplices* and *oriented PL manifolds* (see [Zee66, RS72] or [Whi57, Appendix II]).

Recall that the restriction of π_p to $\text{st}(p)$ provides a simplicial isomorphism between $\text{st}(p)$ and $\text{st}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$ (Lemma 2.2), and that $\text{st}(p) = \text{st}(p, \hat{\mathcal{M}})$ (by Lemma 5.4). We orient T_p and orient positively the k -dimensional simplices of $\text{st}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$. We then orient the k -dimensional simplices of $\text{st}(p, \hat{\mathcal{M}})$ isomorphically, using the orientation of the corresponding k -dimensional simplices of $\text{st}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$.

Since $\text{lk}(p, \hat{\mathcal{M}})$ is a PL $(j-1)$ -sphere (see the proof of Lemma 5.5), we immediately get that $\text{st}(p, \hat{\mathcal{M}})$ is an oriented PL k -ball.

Definition 5.13 (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 following result is a special case of a standard result from PL-topology (see, e.g., Appendix II of [Whi57, Lemma 15a]).

Lemma 5.14 *Let \mathcal{K} be a j -dimensional PL-manifold whose j -simplices are oriented, and let the continuous map $f : \mathcal{K} \rightarrow \mathbb{R}^j$ be a simplexwise positive map. Then, for any connected open subset W of $\mathbb{R}^j \setminus f(\partial\mathcal{K})$, any two points of W not in $f(\mathcal{K}^{-1})$ are covered the same number of times. If this number is 1, then f , restricted to the open subset $W' = f^{-1}(W)$ of \mathcal{K} , is injective.*

5.2.2 Properties of π_p^*

Let $\sigma = [p, p_1, \dots, p_k]$ be a k -simplex of $\text{st}(p, \hat{\mathcal{M}})$. Now, for all points $q \in \sigma$, let $\pi_{p,t}(q) = (1-t)q + t\pi_p(q)$, and $\sigma_t = \pi_{p,t}(\sigma)$. 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 5.10 that

$$\forall j \in \{1, \dots, k\}, \|p_j - p_{jt}\| = O(\varepsilon^2) \text{ lfs}(p). \quad (23)$$

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

Lemma 5.15 (π C^1 -embeds σ_t) *Assume condition **C1** and **C2** of Theorem 5.3. If ε is sufficiently small, then the map π C^1 -embeds σ_t in \mathcal{M} , for all $t \in [0, 1]$.*

Proof. By Lemma 4.5 (1) and the definition of σ_t , we have $\sigma_t \subset B(p, O(\varepsilon) \text{lfs}(p))$. Therefore, for a sufficiently small ε , σ_t cannot intersect the medial axis of \mathcal{M} , which implies, from Lemma 5.2 (1), that the restriction of the map π to σ_t is a C^1 -function. The rest of the proof is devoted to showing that the restriction of π to σ_t is injective and that $d\pi$ is nonsingular.

Let $\rho \stackrel{\text{def}}{=} \max_{i \in \{1, \dots, k\}} (\|q_i - p_i\|) / \Delta_\sigma$. Since $\Delta_\sigma \geq \delta \text{lfs}(p)$ and $\varepsilon / \delta \leq \eta_0$, $\rho = O(\varepsilon^2) \text{lfs}(p) / \Delta_\sigma = O(\varepsilon)$. It then follows from Lemma 5.11 and the fact that $\Theta_\sigma \geq \Theta_0$, that $\Theta_{\sigma_t} \geq \Theta_0 / 2$ for ε sufficiently small. Using $\Gamma_\sigma \leq C_3$ (Lemma 4.5 (2)), we get for sufficiently small ε

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

We then deduce from Corollary 2.5 and the bounds on Θ_{σ_t} and L_{σ_t}

$$\sin \angle(\text{aff}(\sigma), \text{aff}(\sigma_t)) = \frac{O(\varepsilon^2) \text{lfs}(p)}{\Theta_{\sigma_t}^k L_{\sigma_t}} = O(\varepsilon)$$

and from Lemma 4.7

$$\sin \angle(T_p, \text{aff}(\sigma_t)) \leq \sin \angle(T_p, \text{aff}(\sigma)) + \sin \angle(\text{aff}(\sigma), \text{aff}(\sigma_t)) = O(\varepsilon). \quad (24)$$

Let x be a point of σ . We have, from Lemmas 4.5 (1) and 5.10, $\|p - \pi(x)\| = O(\varepsilon) \text{lfs}(p)$. Therefore, from Lemma 5.8 and Eq. (24),

$$\sin \angle(T_{\pi(x)}, \text{aff}(\sigma_t)) \leq \sin \angle(T_{\pi(x)}, T_p) + \sin \angle(T_p, \text{aff}(\sigma_t)) = O(\varepsilon). \quad (25)$$

Eq. (25) implies that π restricted to σ_t is injective for ε sufficiently small. Indeed otherwise, there would exist $x_1, x_2 \in \sigma_t$ such that $\pi(x_1) = \pi(x_2)$, and the line segment in $\text{aff}(\sigma_t)$ joining the points x_1 and x_2 would be parallel to $N_{\pi(x_1)}$, contradicting Eq. (25).

We deduce also from Eq. (25) and Lemma 5.2 (2) that the derivative $d\pi$ is non-singular when π is restricted to σ_t . This completes the proof of the lemma. \square

Lemma 5.16 (π^* C^1 -embeds σ_t) *For ε sufficiently small, π^* C^1 -embeds σ_t in T_p for all $t \in [0, 1]$.*

Proof. From the definition of σ_t , for all $x \in \sigma_t$, we have $\text{dist}(x, \sigma) = O(\varepsilon^2) \text{lfs}(p)$. From Lemma 5.10, we have $\|x - \pi(x)\| = O(\varepsilon^2) \text{lfs}(p)$ for all $x \in \sigma$. Combining these facts we get $\|x - \pi(x)\| = O(\varepsilon^2) \text{lfs}(p)$ for all $x \in \sigma_t$. This implies, from Lemma 4.5 (1),

$$\pi(\sigma_t) \subset B_{\mathcal{M}}(p, O(\varepsilon) \text{lfs}(p)) \subset \widetilde{\mathcal{M}}_p$$

From Lemma 5.8 and the fact that $\pi(\sigma_t) \subset B_{\mathcal{M}}(p, O(\varepsilon) \text{lfs}(p))$, we have

$$\sin \angle(T_x, T_p) < 1, \quad \forall x \in \pi(\sigma_t) \quad (26)$$

Since for all $x \in \pi(\sigma_t)$, $\sin \angle(T_x, T_p) < 1$, we get from Lemma 5.9,

$$N_x \cap T_p = \widetilde{\pi}_p^{-1}(x), \quad \forall x \in \pi(\sigma_t)$$




Figure 8: Refer to the proof of the Lemma 5.16.

Therefore, we deduce $\pi_p^*|_{\sigma_t} = \tilde{\pi}_p^{-1} \circ \pi$. Refer to Fig. 8.

Then, by Lemma 5.9, $\tilde{\pi}_p^{-1}$ is a diffeomorphism from $\tilde{\pi}_p(\tilde{T}_p) \supseteq \tilde{\mathcal{M}}_p$ to \tilde{T}_p , and, by Lemma 5.15, π C^1 -embeds σ_t into \mathcal{M} . We conclude that $\pi_p^*|_{\sigma_t} = \tilde{\pi}_p^{-1} \circ \pi$ C^1 -embeds σ_t into T_p . \square

The proof of the following lemma is similar to Lemma 23a in Chapter 4 from [Whi57] and it will be used for proving that π_p^* restricted to open star $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is injective in Lemma 5.18. We give a proof for completeness.

The proof of the following lemma is similar to Lemma 23a in Chapter 4 from [Whi57]. We give a proof for completeness.

Lemma 5.17 *Assume conditions C1 and C2 of Theorem 5.3. For ε sufficiently small, π_p^* is a simplexwise positive mapping of $\text{st}(p, \hat{\mathcal{M}})$ into T_p .*

Proof. We use the same notation σ_t as in the beginning of this section. We first observe that, since the simplex σ_1 is in T_p , π_p^* is the identity on σ_1 . Therefore, $\det(J(\pi_p^*)) > 0$ on σ_1 . Moreover, since from Lemma 5.16, $\det(J(\pi_p^*)) \neq 0$ in σ_t for all $t \in [0, 1]$, we also have $\det(J(\pi_p^*)) > 0$ in σ_0 . Which proves the lemma. \square

We are now ready to prove that π_p^* restricted to the open star $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is injective.

Lemma 5.18 (Injectivity of π_p^* restricted to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$) *Assume conditions C1 and C2 of Theorem 5.3 and let ε be sufficiently small. For each point $p \in \mathcal{P}$, the map π_p^* is injective on the open star $\mathring{\text{st}}(p, \hat{\mathcal{M}})$.*

Proof. Note that, since $\hat{\mathcal{M}}$ is a PL-manifold without boundary (Lemma 5.5), $\text{lk}(p, \hat{\mathcal{M}})$ is a PL $(k-1)$ -sphere, $\text{st}(p, \hat{\mathcal{M}})$ is a PL k -ball, and $\partial \text{st}(p, \hat{\mathcal{M}}) = \text{lk}(p, \hat{\mathcal{M}})$.

For convenience, rewrite $f = \pi_p^*$ and $S = \text{st}(p, \hat{\mathcal{M}})$ and orient S as indicated in Section 5.2.1. S is thus an oriented PL k -ball, and, by Lemma 5.17, $f = \pi_p^*$ is a simplexwise positive mapping from S to T_p . 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 W be any connected open subset of $T_p \setminus f(\partial S)$. From Lemma 5.14, any two points of $W \setminus f(S^{k-1})$ are covered the same number of times. If this number is 1, then f , restricted to the open subset $f^{-1}(W)$ of $\text{st}(p)$, is injective onto W .

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 5.12 and the facts that $\Theta_\sigma \geq \Theta_0$ (since the simplices of $\hat{\mathcal{M}}$ are Θ_0 -fat), $\Delta_\sigma \geq \delta \text{lfs}(p)$, and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 3.1), we have

$$\text{dist}(g, \partial\sigma) \geq \frac{k! \Theta_\sigma^k \Delta_\sigma}{k+1} \geq \frac{k! \Theta_0^k}{k+1} \times \delta \text{lfs}(p) \geq \frac{k! \Theta_0^k}{k+1} \times \frac{\varepsilon \text{lfs}(p)}{\eta_0} = \Omega(\varepsilon) \text{lfs}(p).$$

From the above equation and using the fact that $\|\pi_p(x) - x\| = O(\varepsilon^2) \text{lfs}(p)$ for all $x \in \sigma$ (Lemma 5.10), we have for ε sufficiently small

$$\text{dist}(\pi_p(g), \partial\pi_p(\sigma)) = \Omega(\varepsilon) \text{lfs}(p) - O(\varepsilon^2) \text{lfs}(p) = \Omega(\varepsilon) \text{lfs}(p). \quad (27)$$

Since π_p embeds S into T_p (from Lemma 2.2) with $\pi_p(S) = \text{st}(p, \text{Del}^{\xi_p}(\mathcal{P}_p))$ and since $\text{Del}^{\xi_p}(\mathcal{P}_p)$ is a triangulation of a k -dimensional convex hull (see the discussion on the general position assumption at the beginning of Section 5), Eq. (27) implies that, for all $x \in S \setminus \text{int } \sigma$

$$\|\pi_p(g) - \pi_p(x)\| = \Omega(\varepsilon) \text{lfs}(p) - O(\varepsilon^2) \text{lfs}(p). \quad (28)$$

From Eq. (28), we have for all $x \in S \setminus \text{int } \sigma$

$$\begin{aligned} \|\pi_p^*(x) - \pi_p^*(g)\| &\geq \|\pi_p(x) - \pi_p(g)\| - (\|\pi_p(x) - x\| + \|\pi_p^*(x) - x\|) \\ &\quad - (\|\pi_p(g) - g\| + \|\pi_p^*(g) - g\|) \\ &= \Omega(\varepsilon) \text{lfs}(p) - O(\varepsilon^2) \text{lfs}(p). \end{aligned} \quad (29)$$

which is positive for a sufficiently small ε .

Lemma 5.16 together with Eq. (29) shows 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

5.2.3 Properties of π

Lemma 5.19 (Injectivity of the restriction of π to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$) *Assume conditions C1 and C2 of Theorem 5.3. Let ε be sufficiently small. For all p in \mathcal{P} , the map π restricted to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is injective.*

Proof. To reach a contradiction, assume that there exist x_1, x_2 ($x_1 \neq x_2$) in $\text{st}(p) \setminus \partial \text{st}(p)$ such that $\pi(x_1) = \pi(x_2)$. Then $\pi_p^*(x_1) = \pi_p^*(x_2) = N_{\pi(x_1)} \cap T_p$. Which contradicts the fact that π_p^* is injective when restricted to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ (Lemma 5.18). \square

Lemma 5.20 (Surjectivity of $\pi|_{\hat{\mathcal{M}}}$) *Assume conditions C1 and C2 of Theorem 5.3. If ε is sufficiently small, then $\pi|_{\hat{\mathcal{M}}}$ is surjective on \mathcal{M} .*

Proof. Since $\hat{\mathcal{M}}$ and \mathcal{M} are topological k -manifolds without boundaries, and $\mathring{\text{st}}(p, \hat{\mathcal{M}})$, which is an open set in $\hat{\mathcal{M}}$ with $\hat{\mathcal{M}} = \bigcup_{p \in \mathcal{P}} \mathring{\text{st}}(p, \hat{\mathcal{M}})$, is mapped injectively by π into \mathcal{M} , we get from the Generalized Invariance of Domain Theorem, π is an open map and $\pi(\hat{\mathcal{M}})$ is an open set in \mathcal{M} . Also, note that since $\hat{\mathcal{M}}$ is compact and π is a continuous function, $\pi(\hat{\mathcal{M}})$ is a compact subset of \mathcal{M} [Bre94, Thm. 7.6 of Chap. 1]. As \mathcal{M} is a Hausdorff space, this implies that $\pi(\hat{\mathcal{M}})$ is a closed set of \mathcal{M} [Bre94, Thm. 7.5 of Chap. 1].

Now consider a connected component \mathcal{M}_α of \mathcal{M} with $\mathcal{V}_\alpha = \pi(\hat{\mathcal{M}}) \cap \mathcal{M}_\alpha \neq \emptyset$. From the discussion in the previous paragraph we have \mathcal{V}_α is both open and closed in \mathcal{M}_α . Since \mathcal{M}_α is connected, this implies $\mathcal{V}_\alpha = \mathcal{M}_\alpha$ [Bre94, Def. 4.1 of Chap. 1]. To complete the proof of surjectivity, we only need to show that \mathcal{P} has non empty intersection with each connected components of \mathcal{M} , since $\pi(p) = p$ for all $p \in \mathcal{P}$. By using the proof technique of Giesen and Wagner [GW04, Lem. 3], we can show that if \mathcal{P} is an (ε, δ) -sample of \mathcal{M} with $\varepsilon < 1$, then \mathcal{P} intersects all the connected components of \mathcal{M} . See, [Gho12, Chap. 4] for details. This completes the proof. \square

As discussed earlier, π is an open map $\pi|_{\mathring{\text{st}}(p, \hat{\mathcal{M}})}$ is injective (Lemma 5.19) and $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is an open set in $\hat{\mathcal{M}}$. The restriction of π to $\mathring{\text{st}}(p, \hat{\mathcal{M}})$ is thus a homeomorphism onto its image. In addition, we know from Lemma 5.20 that $\pi(\hat{\mathcal{M}}) = \mathcal{M}$. Therefore, for all $x \in \mathcal{M}$, there exists open neighborhoods U_x of x in \mathcal{M} such that $\pi^{-1}(U_x)$ is a disjoint union of open sets, each of which is contained in an open star of a vertex in $\hat{\mathcal{M}}$ and mapped homeomorphically onto U_x by π . In other words, $(\hat{\mathcal{M}}, \pi)$ is a covering space of \mathcal{M} .

A standard result in topology [Mas67, Lem. 3.4 of Chap. 5] asserts that π covers all the points of a connected component of its image \mathcal{M} the same number of times. Let us show that this number is exactly one. It cannot be zero by definition. Neither can it be more than one since, otherwise, some sample point p would be covered more than once, which would contradict Lemma A.1 from Appendix A. We have therefore proved that the restriction of π to $\hat{\mathcal{M}}$ is injective.

Lemma 5.21 (Injectivity of $\pi|_{\hat{\mathcal{M}}}$) *Under the same conditions as in Lemma 5.21, π restricted to $\hat{\mathcal{M}}$ is injective.*

Since $\hat{\mathcal{M}}$ is compact and \mathcal{M} is a Hausdorff space, Lemmas 5.20 and 5.21, and a standard result from topology [Bre94, Thm. 7.8 of Chap. 1] implies that $\pi : \hat{\mathcal{M}} \rightarrow \mathcal{M}$ is a homeomorphism.

This completes the proof of Lemma 5.6.

5.3 Pointwise approximation

Following lemma is a direct consequence is Lemmas 5.10 and 5.6, and the fact that lfs is 1-Lipschitz function. For the details of the proof, refer to [Gho12, Chap. 4].

Lemma 5.22 (Pointwise approximation) *Under conditions C1 and C2 of Theorem 5.3, and ε sufficiently, we have $\text{dist}(x, \pi|_{\hat{\mathcal{M}}}^{-1}(x)) = O(\varepsilon^2)\text{lfs}(x)$.*

5.4 Isotopy

We need the following lemma to show the isotopy property **P5** of Theorem 5.3. The proof of the lemma follows directly from [GW04, Lem. 5] and one can check the details in [Gho12, Chap. 4].

Lemma 5.23 *Assume conditions C1 and C2 of Theorem 5.3. If ε is sufficiently small, for all points x, y ($\neq p$) in $\hat{\mathcal{M}}$,*

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

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 [Bre94, Thm. 7.8 of Chap. 1], the following result:

Lemma 5.24 *Assume conditions C1 and C2 of Theorem 5.3. 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 5.25 (Isotopy) *Assume conditions C1 and C2 of Theorem 5.3. If ε is sufficiently small then*

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

is an isotopy.

6 Conclusion

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.

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 dimension of \mathcal{M} is known. If not, we can use algorithms given in [CWW08, GW04] 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 \mathcal{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].

Acknowledgements

The authors thank Mariette Yvinec and Ramsay Dyer for their careful reading of earlier drafts and insightful comments. The authors also gratefully acknowledge the work of the reviewers and thank them for their detailed comments that helped to improve the manuscript.

A Appendix to Section 5.2

This section is devoted to the proof of the following lemma.

Lemma A.1 *Let ε be sufficiently small. For all p in \mathcal{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 . But from Lemma 4.7 and ε sufficiently small, we have $\sin \angle(\text{aff}(\tau), T_p) \leq A\varepsilon < 1$.

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 5.10.

Since $\hat{\mathcal{M}}$ has no inconsistent configuration, τ is in $\text{st}(q_0)$. Let $c = \text{Vor}^\omega(\tau) \cap T_{q_0}$. From Lemma 4.4, we have

$$R = \sqrt{\|c - q_0\|^2 - \omega(q_0)^2} \leq \|c - q_0\| \leq C_1 \varepsilon \text{lfs}(q_0). \quad (30)$$

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 4.5 (1)), lfs is 1-Lipschitz and ε sufficiently small, we have

$$\text{lfs}(q_j) \geq \text{lfs}(q_i) - \|q_i - q_j\| \geq (1 - C_2 \varepsilon) \text{lfs}(q_i) \geq \frac{\text{lfs}(q_i)}{2} \quad (31)$$

Therefore using Eq. (30) and (31), we have for all vertices q_i of τ

$$R \leq C_1 \varepsilon \text{lfs}(q_0) \leq 2C_1 \varepsilon \text{lfs}(q_i). \quad (32)$$

Consider the edge $q_i q_j$ of τ and let c_{ij} be the projection of c 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) -sample of \mathcal{M} and Lemma 4.1 (2), we have

$$r_{ij} \geq \frac{\sqrt{1 - 4\omega_0^2} \|q_i - q_j\|}{2} \geq b\delta \max\{\text{lfs}(q_i), \text{lfs}(q_j)\} \quad (33)$$

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




Figure 9: Refer to the proof of Lemma A.1, Case 2.

Let $\lambda(q_i) = \max\{\omega(q_i), b\delta \text{lfs}(q_i)/4\}$ for $q_i \in \{q_0, \dots, q_k\}$. Observe that $\lambda(q_i) < \mathfrak{m}(q_i)/2$. For $\varepsilon \leq 1/2$ and Lemma 2.3 (1),

$$\lambda(q_i) < \frac{\mathfrak{m}(q_i)}{2} \leq \frac{\varepsilon}{1-\varepsilon} \text{lfs}(q_i) \leq 2\varepsilon \text{lfs}(q_i). \quad (34)$$

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 9. Therefore

$$\begin{aligned} \|x_{ij} - y_{ij}\| &= r_{ij} + \lambda(q_i) - \sqrt{r_{ij}^2 + \omega(q_i)^2} \\ &= \frac{2r_{ij}\lambda(q_i) + \lambda(q_i)^2 - \omega(q_i)^2}{r_{ij} + \lambda(q_i) + \sqrt{r_{ij}^2 + \omega(q_i)^2}} \\ &\geq \frac{2r_{ij}\lambda(q_i)}{R + \lambda(q_i) + \sqrt{R^2 + \lambda(q_i)^2}} && \text{since } \lambda(q_i) \geq \omega(q_i) \text{ and } r_{ij} \leq R \\ &\geq \frac{2r_{ij}\lambda(q_i)}{\left(2C_1\varepsilon + 2\varepsilon + \sqrt{4C_1^2\varepsilon^2 + 4\varepsilon^2}\right) \text{lfs}(q_i)} && \text{from Eq. (32) and (34)} \\ &\geq \frac{b^2\delta^2 \text{lfs}(q_i)}{4\varepsilon(C_1 + 1 + \sqrt{C_1^2 + 1})} && \text{using } \lambda(q_i) \geq \frac{b\delta \text{lfs}(q_i)}{4} \text{ and Eq. (33)} \\ &\geq \frac{b^2\delta \text{lfs}(q_i)}{4\eta_0(C_1 + 1 + \sqrt{C_1^2 + 1})} && \text{since } \varepsilon/\delta \leq \eta_0 \text{ from Hypothesis 3.1} \end{aligned}$$

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

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

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

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

From the Intersecting Chords Theorem (see Figure 10) 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. (32)) and $\varepsilon/\delta \leq \eta_0$.

Using the fact that ε is sufficiently small and Eq. (31), 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

$$\text{dist}(\text{conv}(S), \partial B(c, R)) = \min_{i, j (\neq i) \in \{0, \dots, k\}} \|z_{ij} - y_{ij}\| \geq C_7\varepsilon \text{lfs}(q) \quad (35)$$

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. (35) implies that $x \notin \text{conv}(S)$. Indeed, if $x \in \text{conv}(S)$, then from Eq. (35) 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 5.10, we have

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

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 :

$$\begin{aligned} \text{conv}(S_i) &\subset B(q_i, \lambda(q_i)) \\ \tau &= \text{conv}(S) \cup \left(\bigcup_{i \in \{0, \dots, k\}} \text{conv}(S_i) \right) \end{aligned} \quad (36)$$




Figure 10: Intersecting Chords Theorem.

If $x \in \text{conv}(S_i)$, then from Lemma 5.10, we have for ε sufficiently small

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

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

References

- [AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1999. 8
- [ACDL02] 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. 31
- [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. 30
- [Aur87] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1):78–96, 1987. 6, 7, 30
- [BDG12] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations, 2012. Submitted to *International Journal of Computational Geometry and Applications* (special issue for SoCG 2012). 7, 30
- [BF04] J.-D. Boissonnat and J. Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004. 1, 3
- [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. 2, 3, 7, 38
- [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. 3
- [BNN10] J.-D. Boissonnat, F. Nielsen, and R. Nock. Bregman Voronoi Diagrams. *Discrete and Computational Geometry*, 44(2):281–307, 2010. 38
- [Bre94] G. E. Bredon. *Topology and Geometry*. Graduate Text in Mathematics, Springer, 1994. 37, 38

- [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. 38
- [BWy08] J-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. ACM Symp. on Computational Geometry*, pages 270–277, 2008. 3
- [BY98] J-D. Boissonnat and M. Yvinec. *Algorithmic geometry*. Cambridge University Press, 1998. 28
- [CDE⁺00] S-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S-H. Teng. Sliver Exudation. *Journal of ACM*, 47:883–904, 2000. 1, 3, 7, 13, 16, 17
- [CDR05] 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. 2, 3, 7, 9, 18
- [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. 2
- [Cha93] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993. 10, 28
- [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. 38
- [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. 2
- [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. 3
- [CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989. 10
- [CSD04] D. Cohen-Steiner and T. K. F. Da. A greedy Delaunay Based Surface Reconstruction Algorithm. *The Visual Computer*, 20:4–16, 2004. 3
- [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. 13, 38
- [Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. 7
- [Dey06] T. K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2006. 2
- [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. 3

- [Fed59] H. Federer. Curvature Measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. 8
- [Fed69] H. Federer. *Geometric measure theory*. Springer New York, 1969. 30
- [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. 1, 3
- [FR02] S. Funke and E. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pages 781–780, 2002. 38
- [Fre02] D Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002. 1, 3
- [Fu93] J. H. G. Fu. Convergence of curvature in secant approximations. *Journal of Differential Geometry*, 37:117–190, 1993. 9
- [Gho12] A. Ghosh. *Piecewise linear reconstruction and meshing of submanifolds of Euclidean space*. PhD thesis, INRIA Sophia Antipolis & Université of Nice Sophia-Antipolis, 2012. 31, 37
- [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. 3
- [GW04] J. Giesen and U. Wagner. Shape Dimension and Intrinsic Metric from Samples of Manifolds. *Discrete and Computational Geometry*, 32(2):245–267, 2004. 8, 13, 37, 38
- [Li03] X-Y. Li. Generating Well-Shaped d -dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003. 9
- [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. 3
- [Mas67] W. S. Massey. *Algebraic Topology : An Introduction*. Springer-Verlang, Graduate texts in Mathematics 56, 1967. 37
- [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. 3
- [NSW08] 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. 2, 32
- [RS72] C. P. Rourke and B. J. Sanderson. *Introduction to Piecewise-Linear Topology*. Springer-Verlang, Berlin, 1972. 33

- [RS00] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000. 3
- [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. 3
- [SL00] H. S. Seung and D. D. Lee. The manifold ways of perception. *Science*, 290:2268–2269, 2000. 2
- [TdSL00] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000. 3
- [Whi57] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957. 9, 31, 32, 33, 35
- [Zee66] E. C. Zeeman. *Seminar on Combinatorial Topology*. Institut des Hautes Études Scientifiques (Paris) and University of Warwick (Coventry), Notes, 1963-1966. 31, 33
- [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. 3

Main notations

General notations (Section 2)

$B(c, r)$	is $\{x \in \mathbb{R}^d \mid \ c - x\ < r\}$
$\bar{B}(c, r)$	is $\{x \in \mathbb{R}^d \mid \ c - x\ \leq r\}$
$\dim(U)$	dimension of the affine space U
$\text{dist}(p, X)$	is $\min_{x \in X} \ p - x\ $
∂X	topological boundary of X
$\text{int } X$	topological interior of X
$df(x)$	derivative of the function f at x
$f _{X_1}$	map f restricted to the subset X_1
$[xy]$	line segment connecting points x and y
$\ x - y\ $	standard Euclidean distance between the points x and y
$\angle(u, v)$	angle between the vectors u and v
$\angle(U, V)$	angle between the affine spaces U and V

Submanifold (Section 2)

$B_{\mathcal{M}}(p, r)$	is $B(p, r) \cap \mathcal{M}$
$\bar{B}_{\mathcal{M}}(p, r)$	is $\bar{B}(p, r) \cap \mathcal{M}$
d	dimension of the ambient space
k	intrinsic dimension of \mathcal{M}
\mathcal{M}	manifold
\mathcal{P}	sample
$ \mathcal{P} $	cardinality of the sample \mathcal{P} (Theorem 4.16, page 25)
T_x	tangent space at x
N_x	normal space at x

Weights (Section 2.1.1)

p^ω	weighted point $(p, \omega(p))$
ω	weight assignment
$\tilde{\omega}$	relative amplitude
ω_0	bound on the relative amplitude

Sampling (Section 2.2)

ε	sampling rate
ε/δ	sampling ratio
η_0	bound on the sampling ratio
δ	sparsity
$\text{lfs}(p)$	local feature size at p
$\text{m}(p)$	distance of p to its nearest neighbor
N	is $(4C'_2\eta_0 + 6)^k$ (Section 4.4, page 23)
$LN(p)$	local neighborhood of p (Section 3)

Shape measure of simplex τ (Section 2.3)

c_τ	circumcenter
$D_\tau(p)$	altitude (Section 4.1, page 15)
$\text{aff}(\tau)$	affine hull of τ
$m_u(\tau)$ or m_u	see Section 2.4, page 10
i_ϕ	see Section 2.4, page 11
N_τ	affine space orthogonal to $\text{aff}(\tau)$
Δ_τ	length of longest edge (diameter)
$H_\tau(p, \omega(p))$	excentricity (Section 4.1, page 15)
L_τ	length of shortest edge
o_τ	orthocenter
Φ_τ	orthoradius
R_τ	circumradius
$\Gamma_\tau = \frac{\Delta_\tau}{L_\tau}$	longest edge to shortest-edge ratio
$\tau_p = \tau \setminus \{p\}$	face of τ opposite to p (Section 4.1)
Θ_τ	fatness
Θ_0	fatness bound
V_τ	volume

Complexes (Section 2.4)

\mathcal{K}	simplicial complex
$\text{st}(p, \mathcal{K})$	$\{\tau : \text{for some } \tau_1 \in \mathcal{K}, \sigma, \tau \subseteq \tau_1\}$
$\text{lk}(p, \mathcal{K})$	$\{\tau \in \text{st}(p, \mathcal{K}) : \tau \cap \sigma = \emptyset\}$
$\overset{\circ}{\text{st}}(p, \mathcal{K})$	$\text{st}(p, \mathcal{K}) \setminus \text{lk}(p, \mathcal{K})$
$\text{st}(p)$	$\text{st}(p, \text{Del}_p^\omega(\mathcal{P}))$ (Section 2.4)
$C^\omega(\mathcal{P})$	Complete complex (page 10)
$\text{Del}^\omega(\mathcal{P})$	Weighted Delaunay triangulation (Section 2.1.2, page 5)
$\text{Del}_p^\omega(\mathcal{P})$	restriction of $\text{Del}^\omega(\mathcal{P})$ to T_p
$\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$	Tangential Delaunay complex (page 8)
$\text{Inc}^\omega(\mathcal{P})$	contains Θ_0 -inconsistent configurations of $\text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ (page 10)

Mappings

π	projection onto \mathcal{M}
$\pi _{\hat{\mathcal{M}}}$	restriction of π to $\hat{\mathcal{M}}$
π_p	restriction of π to T_p
π_p^*	mapping onto T_p