



Project number IST-25582

CGL

Computational Geometric Learning

**Manifold Reconstruction Using Tangential Delaunay
Complexes**

STREP

Information Society Technologies

Period covered: November 1, 2010–October 31, 2011
Date of preparation: October 25, 2011
Start date of project: November 1, 2010
Duration: 3 years
Project coordinator name: Joachim Giesen (FSU)
Project coordinator organisation: Friedrich-Schiller-Universität Jena
Jena, Germany

Manifold Reconstruction Using Tangential Delaunay Complexes *

Jean-Daniel Boissonnat

Arijit Ghosh

October 25, 2011

Abstract

We give a provably correct algorithm to reconstruct a k -dimensional manifold embedded in d -dimensional Euclidean space. The input to our algorithm is a point sample coming from an unknown manifold. Our approach is based on two main ideas : the notion of tangential Delaunay complex defined in [6, 24, 25], and the technique of sliver removal by weighting the sample points [17]. 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 ambient isotopic to the manifold and a close geometric approximation of the manifold.

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

*This research has been partially supported by the Agence Nationale de la Recherche (project GAIA 07-BLAN-0328-04) and by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL Computational Geometry Learning).

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 [22], and the survey by Cazals and Giesen in [13] 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 [7, 14, 18, 37].

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 [39]. 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 [41], LLE [38], Laplacian eigenmaps [4], Hessian eigenmaps [32], diffusion maps [33, 36], principal manifolds [43]. 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 investigations along this line can be found in the work of Cheng, Dey and Ramos [18], and Boissonnat, Guibas and Oudot [7]. 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 [7] and [18] 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 [15] 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 [6, 24, 25], and the technique of sliver removal by weighting the sample points [17]. 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 [25], no solution was known except for the case of curves ($k = 1$) [24]. The idea of removing inconsistencies among local triangulations that have been computed independently has already been used for maintaining dynamic meshes [40] and generating anisotropic meshes [9]. Our approach is close in spirit to the one in [9]. 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. [18]. By local, we mean that we do not compute any d -dimensional data structure like a grid or a triangulation of the ambient space. Still, the tangential complex is a subcomplex of the weighted d -dimensional Delaunay triangulation of the (weighted) data points and therefore implicitly relies on a global partition of the ambient space. This is a key to our analysis and distinguishes our method from other local algorithms that have been proposed in the surface reconstruction literature [21, 29].

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 bound its 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 have been added in the appendix as a reference for the readers.

2 Definitions and preliminaries

A j -simplex is the convex hull of $j + 1$ affinely independent points. 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 $\text{aff}(\tau)$.

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 no boundary and a positive reach (see Section 2.2). We denote by T_p the k -dimensional tangent space at point $p \in \mathcal{M}$.

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 two spaces of the same dimension. The angle between U and V is defined as

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

where u and v are vectors in U and V respectively.

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

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

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

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

$$\tilde{\omega} = \max_{p \in \mathcal{P}, q \in \mathcal{P} \setminus \{p\}} \frac{\omega(p)}{\|p - q\|}. \quad (1)$$

In the paper, we make the following hypothesis.

Hypothesis 2.2 $\tilde{\omega} \leq \omega_0$, for some constant $\omega_0 \in [0, 1/2)$

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

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 $\text{aff}(\tau)$.

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

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

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} [3].

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 will assume that the points of \mathcal{P} are in general position. Then, $\text{Vor}^\omega(\tau) = \emptyset$ when $|\tau| > d + 1$, and the collection of all simplices $\text{conv}(\tau)$ such that $\text{Vor}^\omega(\tau) \neq \emptyset$ constitutes a triangulation 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.

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 real positive weights, except in Lemma 2.3.

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. [3]. 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 2.3). The other main reason is that some flat simplices named slivers can be removed from a Delaunay triangulation by weighting the vertices (see [7, 17, 18] and Section 3).




Figure 1: Refer to Lemma 2.3. The grey line denotes the k -dimensional plane H and the black line denotes $\text{Vor}^\omega(p_i p_j)$.

Lemma 2.3 *Let H be a k -dimensional affine space of \mathbb{R}^d . The restriction of the weighted Voronoi diagram of \mathcal{P} to H is identical to the k -dimensional weighted Voronoi diagram of \mathcal{P}' in H , where \mathcal{P}' is the orthogonal projection of \mathcal{P} onto H and the squared weight of p'_i is $\omega(p_i)^2 - \|p_i - p'_i\|^2$.*

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 of the weighted points $(p'_i, \xi_i) \in H$ where $\xi_i^2 = -\|p_i - p'_i\|^2 + \omega(p_i)^2$. \square

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} [1]. As is well known and can be easily proved, lfs is *Lipschitz continuous* i.e. $\text{lfs}(x) \leq \text{lfs}(y) + \|x - y\|$. The infimum of lfs over \mathcal{M} is called the *reach* of \mathcal{M} . In this paper, we assume that the reach of \mathcal{M} is (strictly) positive.

2.2.2 (ε, δ) -sample

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

Lemma 2.4 *Given an (ε, δ) -sample \mathcal{P} of \mathcal{M} , we have*

1. $\delta \text{lfs}(p) \leq \mathfrak{m}(p) \leq \frac{2\varepsilon}{1-\varepsilon} \text{lfs}(p)$.
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_τ , Δ_τ , L_τ , V_τ and $\rho_\tau = R_\tau/L_\tau$ the circumradius, the longest edge length, the shortest edge length, the j -dimensional volume, and the radius-edge ratio of τ respectively.

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

$$\Theta_\tau = \begin{cases} 1 & \text{if } j = 0 \\ V_\tau^{1/j} / \Delta_\tau & \text{if } j > 0 \end{cases} \quad (3)$$

The following important lemma is due to Whitney [42].

Lemma 2.5 *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 [26] and Lemma 16 in [18].

Corollary 2.6 (Tangent 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*

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

Proof. It suffices to take $H = T_p$ and to use $\eta = \Delta_\tau^2/2 \text{lfs}(p)$ (from Lemma 2.4 (2)) and $R_\tau/\rho_\tau = L_\tau \leq \Delta_\tau \leq 2R_\tau$. Hence

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

□

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 3.

The following definition is a variant of a definition given in [34].

Definition 2.7 (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.

2.4 Tangential Delaunay complex and inconsistent configurations




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, $\text{star}(u) = \{uv, uw\}$.

Let 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 $\text{Del}_u^\omega(\mathcal{P})$ is in general a k -dimensional triangulation. Since this situation can always be ensured by applying some infinitesimal perturbation on \mathcal{P} , we will assume, in the rest of the paper, that all $\text{Del}_u^\omega(\mathcal{P})$ are k -dimensional triangulations. Finally, write $\text{star}(u)$ for the *star* of u in $\text{Del}_u^\omega(\mathcal{P})$, i.e. the set of simplices that are incident to u in $\text{Del}_u^\omega(\mathcal{P})$ (see Figure 2).

We denote by *tangential Delaunay complex* or *tangential complex* for short, the simplicial complex $\{\tau : \tau \in \text{star}(u), u \in \mathcal{P}\}$. 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 $\text{Del}^\omega(\mathcal{P})$.

By duality, computing $\text{star}(u)$ is equivalent to computing the restriction to T_u of the (weighted) Voronoi cell of u , which, by Lemma 2.3, 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 [16, 20]. It follows that the tangential complex can




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 .

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 $\text{star}(u)$ but not in $\text{star}(v)$ (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. 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 2.8 (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 $\text{star}(u)$ but not in $\text{star}(v)$.*
- *$\text{Vor}^\omega(w)$ is one of the first weighted Voronoi cells of $\text{Vor}^\omega(\mathcal{P})$, other than the weighted Voronoi cells of the vertices of 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.*




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).

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 $\text{Vor}^\omega(\phi)$.

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 simplex hence 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 $\text{Del}(\mathcal{P})$ 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* 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 4.10 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 just above, all inconsistencies in the tangential complex will then also be removed. All simplices being consistent, the resulting weighted tangential Delaunay complex $\hat{\mathcal{M}}$ output by the algorithm will be a simplicial k -manifold that approximates \mathcal{M} well, as will be shown in Section 5.

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

3.1 Algorithm

Let \mathcal{M} be a differentiable submanifold of positive reach, and let \mathcal{P} be an (ε, δ) -sample of \mathcal{M} . \mathcal{M} , ε , δ are unknown and the input to the algorithm consists only of the sample \mathcal{P} and an upper bound η_0 on the sampling ratio ε/δ of \mathcal{P} . As shown in [19, 28], 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 k and T_p , for any point $p \in \mathcal{P}$, are known.

The algorithm fixes ω_0 , the bound on the relative amplitude of the weight assignment, in the interval $[0, 1/2)$ (Hypothesis 2.2). The algorithm also fixes Θ_0 to a constant defined in Theorem 4.16, that depends on k , ω_0 and η_0 .

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

$$LN(p) = \{q \in \mathcal{P} : |B(p, \|p - q\|) \cap \mathcal{P}| \leq N\}. \quad (4)$$

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 4.13, 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.

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
    ω(p_i) ← 0
end for
// Build the full unweighted completed complex C^ω(P)
C^ω(P) ← update_completed_complex(P, ω)
// Weight assignment to remove inconsistencies
for i = 1 to n do
    ω(p_i) ← weight(p_i, ω)
    update_completed_complex(LN(p_i), ω)
end for
// Output
output : M̂ ← Del_{T_M}^ω(P)

```

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, and `build_inconsistent_configurations`(p, τ) and `build_star`(p).

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 `star`(u). More precisely, for each vertex $v \neq u$ of τ such that $\tau \notin \text{star}(v)$, 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.

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  $\text{star}(q)$  do
    if  $\tau$  is  $\Theta_0$ -fat and  $\exists v \in \tau, \tau \not\subseteq \text{star}(v)$  then
      //  $\tau$  is inconsistent
      build_inconsistent_configurations( $q, \tau$ )
    end if
  end for
end for

```

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 [17] 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 $\text{star}(p)$.

Let τ be a candidate simplex of p that is a Θ_0 -sliver. We associate to τ a forbidden interval $W(\tau)$ 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  $p$ 
 $J(p) \leftarrow [0, \omega_0^2 \mathfrak{m}(p)^2] \setminus \bigcup_{\tau \in S(p)} W(\tau)$ 
 $\omega(p)^2 \leftarrow$  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 $W(\tau)$. 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 $W(\tau)$ on the fly.

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} .

For this section, the following hypothesis is assumed to be satisfied as well as Hypothesis 2.2.

Hypothesis 4.1 \mathcal{P} is an (ε, δ) -sample of \mathcal{M} of sampling ratio $\varepsilon/\delta \leq \eta_0$ for some positive constant η_0 .

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

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 [17].

Lemma 4.2 *Assume that Hypothesis 2.2 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. $R_\tau \leq \frac{\Phi_\tau}{\sqrt{1 - 4\omega_0^2}}$.
3. $\forall z \in \text{aff}(\text{Vor}^\omega(\tau)), \sqrt{\|z - p\|^2 - \omega^2(p)} \geq \Phi_\tau$.




Figure 5: For the proofs of Lemmas 4.2 and 4.4.

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. Recall that c_τ and R_τ denote the circumcenter and circumradius of τ . For any vertex p of τ , $\mathfrak{m}(p) = \min_{x \in \mathcal{P}, x \neq p} \|x - p\| \leq 2R_\tau$ and $\omega(p) \leq 2\omega_0 R_\tau$. From the definition, we have

$$\Phi_\tau \geq \min_{p \in \tau} \sqrt{\|c_\tau - p\|^2 - \omega(p)^2} \geq \sqrt{1 - 4\omega_0^2} R_\tau.$$

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

$$\|z - p\|^2 - \omega(p)^2 = \|p - o_\tau\|^2 + \|o_\tau - z\|^2 - \omega(p)^2 = \Phi_\tau^2 + \|o_\tau - z\|^2 \geq \Phi_\tau^2. \quad \square$$

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

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 bound from Lemma 4.3 (1) and the definition of fatness, we get

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

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

$$\frac{D_\tau(p)}{\Delta_\tau} = \frac{j V_\tau}{\Delta_\tau V_{\tau_p}} = j \frac{\Theta_\tau^j \Delta_\tau^{j-1}}{\Theta_{\tau_p}^{j-1} \Delta_{\tau_p}^{j-1}} \leq j \frac{\Theta_\tau^j \Delta_\tau^{j-1}}{\Theta_{\tau_p}^{j-1} L_{\tau_p}^{j-1}} \leq j 2^{j-1} \rho_\tau^{j-1} \times \frac{\Theta_\tau^j}{\Theta_{\tau_p}^{j-1}}. \quad \square$$

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 bounds the excentricity of a simplex.

Lemma 4.4 *Assume that Hypothesis 2.2 is satisfied. Let τ be a simplex of $\text{Del}^\omega(\mathcal{P})$, p a vertex of τ , and α_1, α_2 and α_3 three scalars such that*

1. *There exists $z \in \text{aff}(\text{Vor}^\omega(\tau))$ s.t. $\|z - p\| \leq \alpha_1$,*
2. *$\|p - q\| \leq \alpha_2$ for all vertices q of τ ,*
3. *$\Phi_\tau \leq \alpha_3 L_\tau$.*

Then $|H_\tau(q, \omega(q))| = \text{dist}(o_\tau, \text{aff}(\tau_q)) \leq \alpha_1 + (1 + \alpha_3 + \omega_0)\alpha_2$ for all vertices q of τ .

Proof. $|H_\tau(q, \omega(q))|$ is equal to $\|o_\tau - o_{\tau_q}\|$. Let s be a vertex of τ_q . Using the facts that $\|o_\tau - p\| \leq \|z - p\|$, $\Phi_{\tau_q} \leq \Phi_\tau \leq \alpha_2\alpha_3$ and $\omega(s) \leq \omega_0\|p - s\| \leq \omega_0\alpha_2$, we have

$$\begin{aligned}
\|o_\tau - o_{\tau_q}\| &\leq \|o_\tau - p\| + \|p - s\| + \|s - o_{\tau_q}\| \\
&\leq \|z - p\| + \|p - s\| + \|s - o_{\tau_q}\| \\
&\leq \alpha_1 + \alpha_2 + \sqrt{\Phi_{\tau_q}^2 + \omega(s)^2} \\
&\leq \alpha_1 + \alpha_2 + \sqrt{\Phi_\tau^2 + \omega(s)^2} \\
&\leq \alpha_1 + \alpha_2 + \sqrt{\alpha_2^2\alpha_3^2 + \omega_0^2\alpha_2^2} \\
&\leq \alpha_1 + (1 + \alpha_3 + \omega_0)\alpha_2.
\end{aligned}$$

□

The following lemma is a generalization of Claim 13 from [17]. It bounds the excentricity of a simplex τ as a function of $\omega(p)$ where p is a vertex of τ .

Lemma 4.5 *Let τ be a simplex of $\text{Del}^\omega(\mathcal{P})$ and let p be any vertex of τ . We have $H_\tau(p, \omega(p)) = H_\tau(p, 0) - \frac{\omega(p)^2}{2D_\tau(p)}$.*




Figure 6: For the proof of Lemma 4.5.

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. 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 [18]. The first lemma states that the restriction of a (weighted) Voronoi cell to a tangent space is small.

Lemma 4.6 *Assume that Hypotheses 2.2 and 4.1 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

$$C_1(1 - C_1\varepsilon) > 2 + C_1\varepsilon(1 + C_1\varepsilon) \quad (5)$$

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 2.4, we have $\|q - q'\| \leq C_1^2 \varepsilon^2 \text{lfs}(p)/2$.

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. (5),

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

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 Fig. 7), that $\angle ptx > \pi/2$, which implies that $\|x - p\|^2 - \|x - t\|^2 - \|p - t\|^2 > 0$. Hence,

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




Figure 7: Refer to Lemma 4.6. 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 $\angle ptx \geq \angle ptx' > \pi/2$.

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. (5) is satisfied, which is true for $C_1 \stackrel{\text{def}}{=} 3 + \sqrt{2} \approx 4.41$ and $\varepsilon < 0.09$. \square

The following lemma states that, under Hypotheses 2.2 and 4.1, the simplices of $\text{Del}_p^\omega(\mathcal{P})$ are small, have a good radius-edge ratio and a small excentricity.

Lemma 4.7 *Assume that Hypotheses 2.2 and 4.1 are satisfied. There exists positive constants C_2, C_3 and C_4 that depend on ω_0 and η_0 such that, if $\varepsilon < \frac{1}{2C_2}$, the following holds:*

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

Proof. 1a. Consider first the case where pq is an edge of $\text{Del}_p^\omega(\mathcal{P})$. Then $T_p \cap \text{Vor}^\omega(pq) \neq \emptyset$. Let $x \in T_p \cap \text{Vor}^\omega(pq)$. From Lemma 4.6, we have $\|p - x\| \leq C_1 \varepsilon \text{lfs}(p)$. By Lemma 4.2, $\|q - x\| \leq \frac{\|p - x\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C_1 \varepsilon \text{lfs}(p)}{\sqrt{1 - 4\omega_0^2}}$. Hence, $\|p - q\| \leq C'_1 \varepsilon \text{lfs}(p)$ where $C'_1 \stackrel{\text{def}}{=} C_1(1 + 1/\sqrt{1 - 4\omega_0^2})$.

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

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

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

2. Assume that $\tau \in \text{star}(p)$ for some vertex p of τ . Let $z \in \text{Vor}^\omega(\tau) \cap T_p$, 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 . From Lemma 4.2, we have $r_z \geq \Phi_\tau$. Hence it suffices to prove that there exists a constant C_3 such that $r_z \leq C_3 L_\tau$. Since $z \in \text{Vor}^\omega(\tau) \cap T_p$, we deduce from Lemma 4.6 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 2.4 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 Lemma 4.2 (2), we have $R_\tau \leq \frac{\Phi_\tau}{\sqrt{1-4\omega_0^2}}$. Therefore

$$\rho_\tau = R_\tau / L_\tau = \frac{\Phi_\tau}{L_\tau \sqrt{1-4\omega_0^2}} \leq \frac{r_z}{L_\tau \sqrt{1-4\omega_0^2}} \leq \frac{2C_1 \eta_0}{\sqrt{1-4\omega_0^2}} \stackrel{\text{def}}{=} C_3.$$

3. Assume that $\tau \in \text{Del}_q^\omega(\mathcal{P})$ for some vertex q of τ . Let $z \in \text{Vor}^\omega(\tau) \cap T_q$. From Lemma 4.6, we have $\|q - z\| \leq C_1 \varepsilon \text{lfs}(q)$. We have for all vertices p of τ , $\|p - q\| \leq C_2 \varepsilon \text{lfs}(p)$ (by part 1) and, using the fact that lfs is 1-Lipschitz and $2C_2 \varepsilon < 1$, we get

$$\|q - z\| \leq C_1 \varepsilon (\text{lfs}(p) + \|p - q\|) \leq C_1 \varepsilon (1 + C_2 \varepsilon) \text{lfs}(p) < \frac{3C_1}{2} \varepsilon \text{lfs}(p).$$

We can now apply Lemma 4.4 with $\alpha_1 = \frac{3C_1}{2} \varepsilon \text{lfs}(p)$, $\alpha_2 = C_2 \varepsilon \text{lfs}(p)$, and $\alpha_3 = C_3$ (by part 2). We get

$$|H_\tau(p, \omega(p))| = \text{dist}(o_\tau, \text{aff}(\tau_p)) \leq \frac{3C_1}{2} \varepsilon \text{lfs}(p) + (1 + C_3 + \omega_0) C_2 \varepsilon \text{lfs}(p)$$

Setting $C_4 \stackrel{\text{def}}{=} \frac{3C_1}{2} + (1 + C_3 + \omega_0) C_2$, we get the result. \square

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 4.7 applied to inconsistent configurations. Differently from Lemma 4.7, we need to use Corollary 2.6 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 4.8 $2A\varepsilon < 1$ where $A \stackrel{\text{def}}{=} 4C_2 C_3 / \Theta_0^k$, and C_2, C_3 are the constants defined in Lemma 4.7.

Lemma 4.9 Assume that Hypotheses 2.2, 4.1 and 4.8 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 $\rho_\sigma = R_\sigma / L_\sigma \leq C'_3$.
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 4.6, $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$. Then, by Lemma 4.2, we have for all vertices $p \in \tau$

$$\|p - o_\tau\| \leq \frac{\|u - o_\tau\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{\|u - m_u\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C_1 \varepsilon \text{lfs}(u)}{\sqrt{1 - 4\omega_0^2}}. \quad (6)$$

Let $C'_2 \stackrel{\text{def}}{=} C_2 / \sqrt{1 - 4\omega_0^2} > C_2$. We now use the facts that the k -dimensional simplex τ is a Θ_0 -fat simplex (by definition of an inconsistent configuration), $\Delta_\tau \leq 2C_2 \varepsilon \text{lfs}(p)$ (Lemma 4.7 (1) and the Triangle Inequality), $\Delta_\tau \leq 2R_\tau \leq 2\rho_\tau L_\tau$ and $\rho_\tau \leq C_3$ (Lemma 4.7 (2)). Then Corollary 2.6 yields

$$\sin \angle(\text{aff}(\tau), T_p) \leq \frac{\Delta_\tau^2}{\Theta_0^k L_\tau \text{lfs}(p)} \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_0^k \text{lfs}(p)} \leq \frac{4C_2 C_3 \varepsilon}{\Theta_0^k} = A \varepsilon \quad (7)$$

for all vertices p of τ . Which implies, together with $2A \varepsilon < 1$ (Hypothesis 4.8),

$$\tan^2 \angle(\text{aff}(\tau), T_u) \leq \frac{A^2 \varepsilon^2}{1 - A^2 \varepsilon^2} < 4A^2 \varepsilon^2, \quad (8)$$

Observing again that $\|u - m_u\| \leq C_1 \varepsilon \text{lfs}(u)$ (from Lemma 4.6) and Eq. 7, we deduce

$$\|m_u - o_\tau\| \leq \|m_u - u\| \sin \angle(\text{aff}(\tau), T_u) \leq A C_1 \varepsilon^2 \text{lfs}(u). \quad (9)$$

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 (6) and (8),

$$\|m_v - o_\tau\| \leq \|v - o_\tau\| \tan \angle(\text{aff}(\tau), T_v) < \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \text{lfs}(u) \quad (10)$$

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 (9) and (10) that

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

1. Using Lemma 4.6, and the facts that $2A \varepsilon < 1$ and $\|u - o_\tau\| \leq \|u - m_u\|$, we get

$$\begin{aligned} \|u - i_\phi\| &\leq \|u - o_\tau\| + \|o_\tau - i_\phi\| \\ &\leq \|u - m_u\| + \|o_\tau - i_\phi\| \\ &\leq \left(C_1 \varepsilon + \frac{2A C_1 \varepsilon^2}{\sqrt{1 - 4\omega_0^2}} \right) \text{lfs}(u) \\ &\leq \frac{C_2}{4} \varepsilon \text{lfs}(u) \end{aligned} \quad (11)$$

where C_2 is the constant introduced in Lemma 4.7. Eq. (11), together with Lemma 4.2 and $C'_2 = C_2 / \sqrt{1 - 4\omega_0^2}$, yields $\|p - i_\phi\| \leq \frac{\|u - i_\phi\|}{\sqrt{1 - 4\omega_0^2}} \leq \frac{C'_2}{4} \varepsilon \text{lfs}(u)$ 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 $C'_2 \varepsilon < 1$:

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

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 part 1 of this lemma, we have

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

3. If σ belongs to $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, the result has been proved in Lemma 4.7(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 4.2, we have $r_\phi \geq \Phi_\sigma$. Hence it suffices to show that there exists a constant C'_3 such that $r_\phi \leq C'_3 L_\sigma$. Using (11), 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 4.1,

$$\Phi_\sigma \leq r_\phi \leq \frac{C'_2}{2} \varepsilon \text{lfs}(q) = \frac{C'_2}{2} \left(\frac{\varepsilon}{\delta} \times \delta \right) \times \text{lfs}(q) \leq \frac{C'_2}{2} \eta_0 L_\sigma.$$

From Lemma 4.2, we have $R_\sigma \leq \Phi_\sigma / \sqrt{1 - 4\omega_0^2}$. Therefore, $\frac{R_\sigma}{L_\sigma} \leq \frac{C'_2 \eta_0}{2\sqrt{1 - 4\omega_0^2}} \stackrel{\text{def}}{=} C'_3$. Note that

$$\frac{C'_3}{C_3} = \frac{C'_2}{4C_1} = 1 + \frac{2}{\sqrt{1 - \omega_0^2}} > 1.$$

4. If σ belongs to $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$, the result has been proved in Lemma 4.7(4) with $C'_4 = C_4$. From part 1, we know that $\|i_\phi - p\| \leq \frac{C'_2}{2} \varepsilon \text{lfs}(p)$. Hence, $\|q - p\| \leq C'_2 \varepsilon \text{lfs}(p)$ for all $p, q \in \phi$ and, by part 3, $\frac{\Phi_\sigma}{L_\sigma} \leq C'_3$. Using the above facts, Lemma 4.4 (with $2\alpha_1 = \alpha_2 = C'_2 \varepsilon \text{lfs}(p)$, and $\alpha_3 = C'_3$) and Lemma 4.7 (1), we get

$$\begin{aligned} |H_\sigma(p, \omega(p))| &= \text{dist}(o_\sigma, \text{aff}(\sigma_p)) \leq \text{dist}(i_\phi, \text{aff}(\sigma_p)) \\ &\leq \frac{C'_2}{2} \varepsilon \text{lfs}(p) + (1 + C'_3 + \omega_0) C'_2 \varepsilon \text{lfs}(p) \\ &\leq \left(\frac{3}{2} + C'_3 + \omega_0 \right) C'_2 \varepsilon \text{lfs}(p) \stackrel{\text{def}}{=} C'_4 \varepsilon \text{lfs}(p). \end{aligned}$$

□

The next crucial lemma bounds the fatness of inconsistent configurations.

Lemma 4.10 *Assume Hypotheses 2.2, 4.1 and 4.8. The fatness Θ_ϕ of an inconsistent configuration ϕ is at most*

$$\left(C'_2 \varepsilon \left(1 + \frac{4C_3}{\Theta_0^k} \right) \right)^{\frac{1}{k+1}}$$

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 $\text{star}(u)$ and τ is a Θ_0 -fat simplex. As in the proof of Lemma 4.9, we have $\sin \angle(T_u, \text{aff}(\tau)) \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(u)}$ (refer to Eq. (7)) by Corollary 2.6. Also, from Lemma 2.4, $\sin \angle(uw, T_u) \leq \frac{\|u-w\|}{2\text{lfs}(u)} \leq \frac{\Delta_\phi}{2\text{lfs}(u)}$. Using the fact that τ is a Θ_0 -fat simplex and $\|w - q\| \leq C'_2 \varepsilon \text{lfs}(q)$ for any vertex q of τ (Lemma 4.9 (2)), we can bound the altitude $D_\phi(w)$ of w in ϕ

$$\begin{aligned} D_\phi(w) &= \text{dist}(w, \text{aff}(\tau)) \\ &= \sin \angle(uw, \text{aff}(\tau)) \times \|u - w\| \\ &\leq (\sin \angle(uw, T_u) + \sin \angle(\text{aff}(\tau), T_u)) \times \Delta_\phi \\ &\leq \left(\frac{\Delta_\phi}{2\text{lfs}(u)} + \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau^k \text{lfs}(u)} \right) \Delta_\phi \\ &\leq \frac{\Delta_\phi^2}{2\text{lfs}(u)} \left(1 + \frac{4\rho_\tau}{\Theta_\tau^k} \right). \end{aligned} \tag{12}$$

From the definition of fatness of a simplex and Lemma 4.3(1), we get

$$V_\tau = \Theta_\tau^k \Delta_\tau^k \leq \frac{\Delta_\tau^k}{k!}. \quad (13)$$

We deduce

$$\begin{aligned} \Theta_\phi^{k+1} &= \frac{V_\phi}{\Delta_\phi^{k+1}} \\ &= \frac{D_\phi(w) V_\tau}{(k+1)} \times \frac{1}{\Delta_\phi^{k+1}} \\ &\leq \frac{\Delta_\phi^2}{2 \text{lfs}(u)} \left(1 + \frac{4\rho_\tau}{\Theta_\tau^k}\right) \times \frac{\Delta_\tau^k}{(k+1)! \Delta_\phi^{k+1}} \quad (\text{using Eq. (12) and (13)}) \\ &\leq C'_2 \varepsilon \left(1 + \frac{4C_3}{\Theta_0^k}\right). \end{aligned}$$

The last inequality comes from the facts that $\Delta_\phi \leq C'_2 \varepsilon \text{lfs}(u)$ (from Lemma 4.9(2)) and $\rho_\tau \leq C_3$ (from Lemma 4.7(3)). □

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

Hypothesis 4.11 $C'_2 \varepsilon \left(1 + \frac{4C_3}{\Theta_0^k}\right) < \Theta_0^{k+1}$

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 Hypothesis 4.8, we assumed that $2C_2\varepsilon < 1$ in Lemma 4.7 and $C'_2\varepsilon < 1$ in Lemma 4.9. We will make another hypothesis that subsumes these two previous conditions.

Hypothesis 4.12 $C'_2(1 + C'_2 \eta_0) \varepsilon < 1/2$.

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

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.

Lemma 4.13 *Assume the Hypotheses 2.2, 4.1, 4.8 and 4.12 are satisfied and let $N \stackrel{\text{def}}{=} (4C'_2\eta_0 + 6)^k$, where the constant C'_2 is defined in Lemma 4.9. The set*

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

includes all the points of \mathcal{P} that can form an edge with p in $C^\omega(\mathcal{P})$.

Proof. Lemmas 4.7 and 4.9 show that, in order to construct $\text{star}(p)$ 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 $\text{lfs}(\cdot)$ is a 1-Lipschitz function, we have

$$\text{lfs}(p) (1 - C'_2 \varepsilon) \leq \text{lfs}(x), \text{lfs}(y) \leq \text{lfs}(p) (1 + C'_2 \varepsilon). \quad (14)$$

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 $\overline{B}_x = B_x \cap T_p$, $\overline{B}_y = B_y \cap T_p$ and $\overline{B}_p^+ = B_p^+ \cap T_p$. From Lemma 2.4(2), the distance from x to T_p is

$$\text{dist}(x, T_p) = \|p - x\| \times \sin(\angle px, T_p) \leq C'_2{}^2 \varepsilon^2 \text{lfs}(p)/2. \quad (15)$$

Using Eq.s (14), (15) and the fact that $\varepsilon/\delta \leq \eta_0$, we see that \overline{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

$$\begin{aligned} \left(\frac{r^+}{r^-} \right)^k &= \left(\frac{(C'_2 \varepsilon + (1 + C'_2 \varepsilon) \delta)^2}{\delta^2/4 \times ((1 - C'_2 \varepsilon)^2 - C'_2{}^4 \eta_0^2 \varepsilon^2)} \right)^{k/2} \\ &\leq \left(\frac{4(C'_2 \eta_0 + (1 + C'_2 \varepsilon))^2}{(1 - C'_2 \varepsilon)^2 - C'_2{}^4 \eta_0^2 \varepsilon^2} \right)^{k/2} \\ &= \left(\frac{4(C'_2 \eta_0 + (1 + C'_2 \varepsilon))^2}{(1 - C'_2 \varepsilon - C'_2{}^2 \eta_0 \varepsilon)(1 - C'_2 \varepsilon + C'_2{}^2 \eta_0 \varepsilon)} \right)^{k/2} \\ &\leq (4C'_2 \eta_0 + 6)^k \stackrel{\text{def}}{=} N \quad (\text{using Hypothesis 4.12}) \end{aligned}$$

And the result follows. \square

4.5 Correctness of the algorithm

Definition 4.14 (Sliverity range) Let ω be a weight assignment satisfying Hypothesis 2.2. 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 4.15 Under Hypotheses 2.2, 4.1, 4.8, 4.11 and 4.12, 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 4.9(4), we have

$$|H_\tau(p, \omega(p))| \leq C'_4 \varepsilon \text{lfs}(p) \stackrel{\text{def}}{=} D \quad (16)$$

Using Lemma 4.3 (2), we have

$$D_p(\tau) \leq j 2^{j-1} \rho_\tau^{j-1} \Delta_\tau \times \frac{\Theta_\tau^j}{\Theta_{\tau_p}^{j-1}} \leq 2^k (k+1) C_3' \Theta_0 \Delta_\tau \stackrel{\text{def}}{=} E \quad (17)$$

The last inequality follows from the facts that $j \leq k+1$, $\rho_\tau \leq C_3'$ (from Lemmas 4.7 (2) and 4.9 (3)) and τ is a Θ_0 -sliver. Moreover, from Lemma 4.5,

$$H_\tau(p, \omega(p)) = H_\tau(p, 0) - \frac{\omega(p)^2}{2D_p(\tau)}. \quad (18)$$

It then follows from Eqs. (16), (17) and (18) 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.s (17) and (18), the measure of the set of weights for which τ belongs to $C^\omega(\mathcal{P})$ is at most

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

Let q_1 and q_2 be two vertices of τ such that $\Delta_\tau = \|q_1 - q_2\|$. Using Lemma 4.9 (2), we get

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

Using this inequality, $\text{lfs}(p) \leq \mathfrak{m}(p)/\delta$ (Lemma 2.4) and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 4.1), the sliverity range of τ is at most $2^{k+3}(k+1)C_3'^k C_2' C_4' \Theta_0 \eta_0^2 \mathfrak{m}(p)^2 = C_5 \Theta_0 \mathfrak{m}(p)^2$ with $C_5 \stackrel{\text{def}}{=} 2^{k+3}(k+1)C_3'^k C_2' C_4' \eta_0^2$. By Lemma 4.13, the number of j -simplices that are incident to p is at most N^j . Hence, the sliverity range of p is at most

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

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

Theorem 4.16 *Let \mathcal{P} be an (ε, δ) -sample of \mathcal{M} , $\varepsilon/\delta \leq \eta_0$ and $\Theta_0 = \frac{\omega_0^2}{2N^{k+1}C_5}$. If Hypotheses 2.2, 4.1, 4.8, 4.11 and 4.12 are satisfied, the simplicial complex $\hat{\mathcal{M}} = \text{Del}_{T, \mathcal{M}}^\omega(\mathcal{P})$ 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 4.15. 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 4.10 and Hypothesis 4.11, all inconsistent configurations in $C^\omega(\mathcal{P})$ 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 4.17 *Assume that Hypotheses 2.2, 4.1, 4.8, 4.11 and 4.12 are satisfied. Then the space complexity of the algorithm is*

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

and the time complexity is

$$O(d)|\mathcal{P}|^2 + d2^{O(k^2)}|\mathcal{P}|.$$

Proof. **1. 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 $LN(p)$.

By Lemma 2.3, each $\text{star}(p)$, $p \in \mathcal{P}$, has the same combinatorial complexity as a Voronoi cell in the k -dimensional flat T_p . The sites needed to compute this Voronoi cell all belong to $LN(p)$. Using the fact that $|LN(p)| = N$ for all $p \in \mathcal{P}$ and from the Upper Bound Theorem of convex geometry, see e.g. [10], the number of k -simplices in 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 $\text{star}(p)$, we can have from Lemmas 4.9 (2) and 4.13, at most $k|LN(p)| = kN$ different inconsistent configurations. Hence, the number of inconsistent configurations to be stored in the completed complex $C^\omega(\mathcal{P})$ is at most $O(k^2N^{\lfloor k/2 \rfloor + 1})|\mathcal{P}|$.

With $N = O(2^k)$ (refer to Lemma 4.13), we conclude that the total space complexity of the algorithm is

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

2. 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 $\text{star}(p)$ is dominated by the time to compute the cell of p in the weighted k -dimensional Voronoi diagram of the projected points of $LN(p)$ onto T_p . Since, by definition, $|LN(p)| = N$, 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 , which is

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

see e.g. [16, 10]. The factor $O(kd)$ appears in the first term because to calculate the projection of a point in \mathbb{R}^d on a k -flat we have to do k inner products. The $O(k^3)$ factor comes from the fact that the basic operation we need to perform when computing a weighted Voronoi cell in \mathbb{R}^k 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 $\text{star}(u)$. For each vertex $v (\neq u)$ of τ with $\tau \notin \text{star}(v)$, we need to compute the inconsistent configurations of the form $\phi = [p_0, \dots, p_k, w]$ witnessed by (u, v, w) where $w \in LN(p) \setminus \tau$. The number of such inconsistent configurations is therefore less than $|LN(p)| = N$. 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

$$O(dkN) \tag{19}$$

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

Hence, the time complexity for building the inconsistent configurations of $C^\omega(\mathcal{P})$ 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 $O(dk N^2 + (k^3 + dk)N^{k+1}) |\mathcal{P}|$.

Combining the time complexities for all the steps of the algorithm and using $N = O(2^k)$ (refer to Lemma 4.13), we get the total time complexity of the algorithm

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

□

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 will prove that the simplicial complex $\hat{\mathcal{M}}$ output by the algorithm in Section 3 is ambient isotopic to and a close geometric approximation of the manifold \mathcal{M} . We assume for the rest of this section that the Hypotheses 2.2, 4.1, 4.8, 4.11 and 4.12 are satisfied. Therefore, $\hat{\mathcal{M}}$ satisfies the following properties from Theorem 4.16 :

1. For all simplices τ in $\hat{\mathcal{M}}$, $\Theta_\tau \geq \Theta_0$.
2. $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ have no inconsistent simplex.

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 [23].

Lemma 5.1 *Let \mathcal{M} be a smooth submanifold of \mathbb{R}^d without boundary.*

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 $d - k$.*

We will now state the main result of this section.

Theorem 5.2 *Under the conditions of Theorem 4.16 and for ε sufficiently small, the simplicial complex $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ output by the algorithm does not contain slivers nor inconsistent configurations, and the map π restricted to $\hat{\mathcal{M}}$ has the following properties :*

P1 PL k -manifold without boundary : $\hat{\mathcal{M}}$ is a piecewise linear (PL) k -manifold without boundary;

P2 Tangent space approximation : Let τ be a k -simplex in $\hat{\mathcal{M}}$. For all vertices p of τ , we have $\sin \angle(T_p, \text{aff}(\tau)) = \sin \angle(\text{aff}(\tau), T_p) = O(\varepsilon)$, where the constant in the big- O depends on k, ω_0, η_0 and Θ_0 ;

P3 Homeomorphism : The restriction of π to $\hat{\mathcal{M}}$ provides a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} ;

P4 Ambient isotopy : There exists an ambient 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 $F(\hat{\mathcal{M}}, 1) = \mathcal{M}$;

P5 Pointwise approximation : $\forall x \in \mathcal{M}, \text{dist}(x, \pi^{-1}(x)) = O(\varepsilon^2 \text{lfs}(x))$ where the constant in the big- O depends on k, ω_0, η_0 and Θ_0 .

In the rest of this section, we prove the Theorem 5.2.

We first define three maps. For a point $p \in \mathcal{P}$, let the map

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

denote the orthogonal projection of \mathbb{R}^d onto T_p .

We also define the map

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

which maps each point $x \in \mathbb{R}^d \setminus \mathcal{O}$ to the point of intersection of T_p and $N_{\pi(x)}$.

Finally, we define the map

$$\psi_p : T_p \setminus \mathcal{O} \longrightarrow \mathcal{M}$$

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

Proof of Property P1

Lemma 5.3 (PL k -manifold without boundary) *Assume that Hypotheses 2.2 and 4.1 are satisfied. If $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ has no inconsistent simplex, then $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ is a piecewise linear k -manifold without boundary.*

Proof. The $\text{star}(p)$ of a point $p \in \mathcal{P}$ is a set of simplices incident to p in the k -dimensional triangulation $\text{Del}_p^\omega(\mathcal{P})$, and from Lemma 4.6, we know that $\text{Vor}^\omega(p) \cap T_p$ is bounded. Hence $\text{star}(p)$ is a PL topologically closed k -ball with point p in its interior, i.e. $p \in \text{star}(p) \setminus \partial \text{star}(p)$.

Let x be a point of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$. Let $\sigma \in \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ denote the minimal simplex containing x , i.e. if τ is a simplex in $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ containing x then $\tau \supseteq \sigma$. Let q be a vertex of σ . Since there is no inconsistency, all the simplices of $\text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ that are incident to q are in $\text{star}(q)$ and, in particular, $\sigma \in \text{star}(q)$. The point $x \notin \partial \text{star}(q)$, since σ is the minimal simplex containing x . This implies that the point x belongs to the interior of the topologically closed PL k -ball $\text{star}(q)$. \square

Property **P1** then follows from the fact that $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ has no inconsistent simplex (Theorem 4.16).

Proof of Property P2

Lemma 5.4 (Tangent approximation) *Let τ be a k -simplex in $\hat{\mathcal{M}}$. For all vertices p of τ , we have $\sin \angle(\text{aff}(\tau), T_p) = \sin \angle(T_p, \text{aff}(\tau)) \leq A\varepsilon$. The constant A is defined in Hypothesis 4.8 and depends on k, ω_0, η_0 and Θ_0 .*

Proof. Since $\dim(\text{aff}(\tau)) = \dim(T_p) (= k)$, we have from Lemma 2.1, $\angle(\text{aff}(\tau), T_p) = \angle(T_p, \text{aff}(\tau))$.

All the simplices in $\hat{\mathcal{M}} = \text{Del}_{T\mathcal{M}}^\omega(\mathcal{P})$ are Θ_0 -fat (from Theorem 4.16), hence $\Theta_\tau \geq \Theta_0$. Using Corollary 2.6, and the facts that $\Delta_\tau \leq 2R_\tau \leq 2\rho_\tau L_\tau$, $\Delta_\tau \leq 2C_2\varepsilon \text{lfs}(p)$ (from Lemma 4.7 (1) and the Triangle Inequality), and $\rho_\tau \leq C_3$ (from Lemma 4.7 (2)), we get

$$\sin \angle(\text{aff}(\tau), T_p) \leq \frac{\Delta_\tau^2}{\Theta_\tau^k L_\tau \text{lfs}(p)} \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_0^k \text{lfs}(p)} \leq \frac{4C_2 C_3 \varepsilon}{\Theta_0^k} = A\varepsilon$$

□

Proof of Property P3

Outline of the proof. The main ingredients of the proof are due to Whitney [42]. For the reader convenience, we have recalled in Appendix B the basic definitions and results that will be used. Here is a brief outline of the proof. We will first show in Lemma 5.12 that π_p^* restricted to the *open star of p*, $\text{star}(p) \setminus \partial \text{star}(p)$, is injective. Then, using this result, we will show in Lemma 5.13 that the restriction of π to the open star of p is also injective. It will follow that the restriction of π to $\hat{\mathcal{M}}$, denoted by $\pi|_{\hat{\mathcal{M}}}$, is a covering space of $\pi(\hat{\mathcal{M}})$. In order to show that the covering consists of a single sheet, we show in Lemma 5.14 that $\pi_{\hat{\mathcal{M}}}^{-1}(p)$ is reduced to p for all p in \mathcal{P} . This proves that $\pi|_{\hat{\mathcal{M}}}$ is injective. It will then be easy to prove that $\pi|_{\hat{\mathcal{M}}}$ provides a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} from Theorem A.2. This will finish the proof of the Property **P3**.

To prove the main lemmas just mentioned, we will need some technical lemmas. The following one is a generalization of Proposition 6.2 in [37] which bounds the variation of the angle between tangent spaces between two points on the manifold \mathcal{M} .

Lemma 5.5 (Tangent variation) *Let $p, q \in \mathcal{M}$ and $\|p - q\| = t \text{lfs}(p)$. Then, $\sin \angle(T_p, T_q) = O(t)$, where the constant in the big-O is an absolute constant.*

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

$$(1 - t) \text{lfs}(p) \leq \text{lfs}(q) \leq (1 + t) \text{lfs}(p) \quad (20)$$

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

For a unit vector u in T_p , let p_u be a point in T_p such that

$$p_u = p + t \text{lfs}(p) \cdot u$$

Let v denote the unit vector in T_q which makes the smallest angle with the unit vector u .

Let p'_u denote the point closest to p_u on \mathcal{M} . Then, from Lemma 2.4 (3) and Eq. (20), we have

$$\begin{aligned} \|q - p'_u\| &\leq \|q - p\| + \|p - p_u\| + \|p_u - p'_u\| \\ &\leq 2t(1 + t) \text{lfs}(p) \leq 2t \left(\frac{1+t}{1-t} \right) \text{lfs}(q). \end{aligned} \quad (21)$$

Using Lemma 2.4 (2) and Eq. (20) and 21, we have

$$\text{dist}(p, T_q) \leq \|p - q\| \sin \angle(pq, T_q) \leq \frac{t^2}{2(1-t)^2} \text{lfs}(q) \quad (22)$$

Using Lemma 2.4 (3), and Eq. (20), we have

$$\begin{aligned}
\text{dist}(p_u, T_q) &\leq \text{dist}(p'_u, T_q) + \|p_u - p'_u\| \\
&\leq 2t^2 \left(\frac{1+t}{1-t} \right)^2 \text{lfs}(q) + 2t^2 \text{lfs}(p) \\
&\leq 2t^2 \left(\frac{1+t}{1-t} \right)^2 \text{lfs}(q) + \frac{2t^2}{1-t} \text{lfs}(q) \\
&= \frac{2t^2}{1-t} \left(\frac{(1+t)^2}{1-t} + 1 \right) \text{lfs}(q)
\end{aligned} \tag{23}$$

Let $\eta = \max\{\text{dist}(p, T_q), \text{dist}(p_u, T_q)\}$. From Eq. (22) and (23), we have

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

Therefore,

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

□

We will show that the map ψ_p is a diffeomorphism using Lemma 5.5.

Lemma 5.6 (ψ_p is a C^1 -diffeomorphism) *Let $p \in \mathcal{M}$. There exists an absolute constant t_0 (we can assume $t_0 \leq 1/4$) such that for $t \leq t_0$, we have the following:*

1. *The map ψ_p restricted to $B(p, t \text{lfs}(p)) \cap T_p$ is a C^1 -diffeomorphism.*
2. *$B(p, (1-2t)t \text{lfs}(p)) \cap \mathcal{M} \subseteq \psi_p(B(p, t \text{lfs}(p)) \cap T_p)$.*

Proof. 1. By Lemma 5.1, ψ_p is a C^1 -function.

For all $x \in B(p, t \text{lfs}(p)) \cap T_p$, we have from Lemma 2.4 (3)

$$\|p - \psi_p(x)\| \leq \|p - x\| + \|x - \psi_p(x)\| \leq (1+2t)t \text{lfs}(p). \tag{24}$$

From Eq. (24), and Lemmas 5.1 (2) and 5.5, we have that $d\psi_p(x)$, the derivative of ψ_p , for all $x \in B(p, t \text{lfs}(p)) \cap T_p$ is non-singular. Indeed for $x \in B(p, t \text{lfs}(p)) \cap T_p$, we have from Lemmas 2.1 and 5.5 and Eq. (24),

$$\sin \angle(T_p, T_{\psi_p(x)}) = O(t)$$

Since the constant in the big- O is an absolute constant, there exists an absolute constant t_0 such that for $t \leq t_0$, we have

$$\sin \angle(T_p, T_{\psi_p(x)}) < 1. \tag{25}$$

Plainly we can take $t_0 \leq 1/4$.

The map ψ_p is injective. Indeed, otherwise there exists x, y ($x \neq y$) $\in B(p, t_0 \text{lfs}(p)) \cap T_p$ such that $\psi_p(x) = \psi_p(y)$. This implies that the line segment $[x, y] \in T_p$ is orthogonal to $T_{\psi_p(x)}$. We have reached a contradiction since $\sin \angle(T_p, T_{\psi_p(x)}) < 1$ for all $x \in B(p, t_0 \text{lfs}(p)) \cap T_p$ from Eq. (25).

Since ψ_p is injective and the derivative of ψ_p is non-singular, ψ_p is a diffeomorphism from the Inverse Function Theorem.

2. The fact that $B(p, (1 - 2t)t \text{lfs}(p)) \cap \mathcal{M} \subseteq \psi_p(B(p, t \text{lfs}(p)) \cap T_p)$ follows from Lemma 2.4 (3).
 \square

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

Lemma 5.7 *Let ε be sufficiently small and $x \in \text{star}(p)\hat{\mathcal{M}}$. There exists a constant C depending on k , ω_0 and η_0 and Θ_0 such that*

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

Proof. Let τ be a k -simplex in $\text{star}(p)$. We will show that for all $x \in \tau$,

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

1. From Lemma 4.7 (1), we know that, for all vertices q of τ , $\|p - q\| \leq C_2 \varepsilon \text{lfs}(p)$. Therefore, for all $x \in \tau$, $\|p - x\| \leq C_2 \varepsilon \text{lfs}(p)$.

From Lemma 5.4, we have

$$\|\pi_p(x) - x\| \leq \|p - x\| \sin \angle(T_p, \text{aff}(\tau)) \leq AC_2 \varepsilon^2 \text{lfs}(p). \quad (26)$$

This proves the first part of the lemma.

2. Using the definition of the map π , the fact that

$$\|p - \pi_p(x)\| \leq \|p - x\| \leq C_2 \varepsilon \text{lfs}(p),$$

Lemma 2.4 (3) and Eq. (26), we have

$$\begin{aligned} \|\pi(x) - x\| &\leq \|\pi(\pi_p(x)) - x\| \leq \|\pi(\pi_p(x)) - \pi_p(x)\| + \|\pi_p(x) - x\| \\ &\leq 2 \left(\frac{\|\pi_p(x) - p\|}{\text{lfs}(p)} \right)^2 \text{lfs}(p) + AC_2 \varepsilon^2 \text{lfs}(p) \\ &\leq (2C_2 + A) C_2 \varepsilon^2 \text{lfs}(p) \end{aligned} \quad (27)$$

This proves the second part of the lemma.

3. We then have for all x in τ :

$$\|p - \pi(x)\| \leq \|p - x\| + \|x - \pi(x)\| \leq 2\|p - x\| \leq 2C_2 \varepsilon \text{lfs}(p) \quad (28)$$

We assume that ε is small enough so that, for all $x \in \tau$, $\|p - \pi(x)\| < (1 - 2t_0)t_0 \text{lfs}(p)$ where $t_0 \leq 1/4$ is the constant introduced in Lemma 5.6. From Lemma 5.6 (1), we have for all $x \in \tau$, $\psi_p^{-1}(\pi(x)) \in N_{\pi(x)} \cap T_p$. From Lemma 5.5 and ε sufficiently small, we have for all $x \in \tau$,

$$\sin \angle(T_p, T_{\pi(x)}) = O\left(\frac{\|\pi(x) - p\|}{\text{lfs}(p)}\right) = O(\varepsilon) < 1$$

where the constant in the last bi- gO depends on C_2 and therefore on ω_0 and η_0 . This implies that $N_{\pi(x)} \cap T_p$ consists of a single point. Therefore, writing $y = \pi_p^*(x)$, we have $\pi(x) = \psi_p(y)$. Since $\|p - \pi(x)\| < (1 - 2t_0)t_0 \text{lfs}(p)$, we get from Lemma 5.6 (2), that $y = \psi_p(\pi(x)) \in B(p, t_0 \text{lfs}(p)) \cap T_p$, and since ψ_p is the restriction of π to $T_p \setminus \mathcal{O}$, we have from Lemma 2.4 (3)

$$\|y - \pi(x)\| = \|y - \psi_p(y)\| \leq 2t_0^2 \text{lfs}(p). \quad (29)$$

Using the fact that $t_0 \leq 1/4$ and the above inequalities, we have

$$\begin{aligned} \frac{1}{2} \|p - y\| &\leq \frac{t_0 \text{lfs}(p)}{2} < (t_0 - 2t_0^2) \text{lfs}(p) \leq \|p - y\| - \|y - \pi(x)\| && \text{using (29)} \\ &\leq \|p - \pi(x)\| \leq 2C_2 \varepsilon \text{lfs}(p) && \text{using (28)} \end{aligned}$$

This implies $t_0 \leq 4C_2\varepsilon$. We then have

$$\begin{aligned}
\|x - \pi_p^*(x)\| &\leq \|x - \pi(x)\| + \|\pi(x) - y\| && \text{(since } y = \pi_p^*(x)\text{)} \\
&\leq (2C_2 + A)C_2\varepsilon^2 \text{ lfs}(p) + 2t_0^2 \text{ lfs}(p) && \text{(from Eq. (27) and (29))} \\
&\leq (2C_2 + A)C_2\varepsilon^2 \text{ lfs}(p) + 32C_2^2\varepsilon^2 \text{ lfs}(p) && \text{(since } t < 4C_2\varepsilon\text{)} \\
&= (A + 34C_2)C_2\varepsilon^2 \text{ lfs}(p)
\end{aligned}$$

This proves the third part of the lemma. \square

The next lemma uses the notion of C^r -embedding of a simplex in \mathcal{M} recalled in Definition A.5 in Appendix A.

Lemma 5.8 (π C^1 -embeds $\tau \in \text{star}(p)$) *Let $\tau = [p, p_1, \dots, p_k]$ be a k -simplex in $\text{star}(p)$ and let $\sigma = [q_0, \dots, q_k]$ be a k -simplex with $q_0 = p$ and $\|q_i - p_i\| \leq C\varepsilon^2 \text{ lfs}(p)$ for all $i \in \{1, \dots, k\}$. Then, the map π C^1 -embeds the simplex σ into \mathcal{M} .*

Proof. The proof of this lemma is the same as the proof of Lemma 21a in Chapter 4 of [42] and Lemma 5.6. For completeness, we recall the main arguments. Since σ is obtained by a minute perturbation of τ and since $\Theta_\tau \geq \Theta_0$, we have $\Theta_\sigma \geq \Theta_0/2$ for a sufficiently small ε . Therefore $\text{aff}(\sigma)$ approximates the tangent space T_p at p . Specifically, from Corollary 2.6, we have $\sin \angle(\text{aff}(\sigma), T_p) = O(\varepsilon)$ where the constant in the big- O depends on k, ω_0, η_0 and Θ_0 . It follows, using Lemma 5.5, that the restriction of π to σ is injective. Similarly, using Lemmas 5.1 (2) and 5.5, and the fact that $\sin \angle(\text{aff}(\sigma), T_p) = O(\varepsilon)$, we can show that the derivative $d\pi(x)$ is non-singular when π is restricted to σ . This will complete the proof of the lemma. \square

Orientation of $T_p, \text{star}(p)$ and $\pi_p(\text{star}(p))$. For all points p in \mathcal{P} , fix an orientation of T_p , and orient all the k -simplices of $\text{star}(p)$, and $\pi_p(\text{star}(p))$ accordingly. Note that from Lemma 2.3, π_p restricted to $\text{star}(p)$ gives an isomorphic simplicial mapping between the simplicial complex $\text{star}(p)$ and $\pi_p(\text{star}(p))$. Therefore $\text{star}(p)$ (and $\pi_p(\text{star}(p))$) is a PL oriented k -manifold, i.e., $\text{star}(p)$ is a PL k -manifold and each k -dimensional simplex in $\text{star}(p)$ is oriented.

The following lemma is similar to Lemma 23a in Chapter 4 of [42]. It will be used for proving that π_p^* restricted to $\text{star}(p) \setminus \partial \text{star}(p)$ is injective in Lemma 5.12. We give a proof for completeness. The lemma uses the notion of a simplexwise positive mapping recalled in Definition A.6 of Appendix A.

Lemma 5.9 *For ε sufficiently small, π_p^* is a simplexwise positive mapping of $\text{star}(p)$ into T_p .*

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

It then follows from Lemma 5.7 that

$$\forall j \in \{1, \dots, k\}, \|p_j - p_{jt}\| \leq C\varepsilon^2 \text{ lfs}(p) \quad (30)$$

where C is the constant defined in Lemma 5.7. We deduce

Claim 5.10 *For ε sufficiently small, π_p^* C^1 -embeds the simplex σ_t in T_p .*

Proof. For ε sufficiently small, $\pi(\sigma_t) \subset B_{\mathcal{M}}(p, r)$ where $r = (1 - 2t_0)t_0 \text{ lfs}(p)$. It follows that $\pi_p^* = \psi_p^{-1} \circ \pi$ as in part 3 of the proof of Lemma 5.7. Refer to Fig. 8. The fact that the map π_p^* C^1 -embeds the simplex σ_t into T_p follows from the four following observations : 1. $\pi(\sigma_t) \subset B_{\mathcal{M}}(p, r)$;




Figure 8: Refer to the proof of the Claim 5.10.

2. ψ_p^{-1} is a C^1 -diffeomorphism when restricted to $B_{\mathcal{M}}(p, r) \subseteq \psi_p(B(p, t_0 \text{lfs}(p))) \cap T_p$ (from Lemma 5.6); 3. π C^1 -embeds the simplex σ_t into \mathcal{M} (from Lemma 5.8); and 4. $\pi_p^* = \psi_p^{-1} \circ \pi$. \square

Since the simplex σ_1 is in T_p , π_p^* is the identity on σ_1 . Note σ_1 belongs to the PL oriented k -manifold $\pi_p(\text{star}(p))$; therefore (by convention) σ_1 and T_p have the same orientation. This implies $\det(J(\pi_p^*)) > 0$ in σ_1 . Since, from Claim 5.10, $\det(J(\pi_p^*)) \neq 0$ in σ_t for all $t \in [0, 1]$, we also have $\det(J(\pi_p^*)) > 0$ in σ_0 . This concludes the proof of the lemma. \square

We will need the following standard lemma from convex geometry which bounds the distance between an interior point and a point on the boundary of a simplex. See, e.g., Lemma 14b from Chapter 4 in [42] for a proof.

Lemma 5.11 *For any i -simplex $\sigma = [p_0, \dots, p_i]$ and point $p = \sum_{j=0}^i \mu_j p_j$ in σ , $\text{dist}(p, \partial\sigma) \geq i! \Theta_\sigma^i L_\sigma \inf\{\mu_0, \dots, \mu_i\}$.*

Using Lemmas 5.9 and 5.11, we can now prove that π_p^* restricted to the open star of p , $\text{star}(p) \setminus \partial \text{star}(p)$, is injective.

Lemma 5.12 (Injectivity of π_p^*) *Let ε be sufficiently small. For each point $p \in \mathcal{P}$, the map π_p^* is injective on $\text{star}(p) \setminus \partial \text{star}(p)$.*

Proof. For convenience, rewrite $f = \pi_p^*$ and $S = \text{star}(p) \setminus \partial \text{star}(p)$. By Property **P1**, $\text{star}(p)$ is a piecewise linear k -manifold and, by Lemma 5.9, $f = \pi_p^*$ is a simplexwise positive mapping of $\text{star}(p)$ into 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 R be any connected open subset of T_p . By a standard theorem in piecewise linear topology (see, e.g., Appendix II of [42, Lemma 15a]), any two points of $R \setminus S_{k-1}$ are covered the same number of times. If this number is 1, then f , restricted to the open subset $f^{-1}(R)$ of $\text{star}(p)$, is injective onto R .

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

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

Using Lemma 5.11 and the facts that $\Theta_\sigma \geq \Theta_0$ (since the simplices of $\hat{\mathcal{M}}$ are Θ_0 -fat), $R_\sigma/L_\sigma \leq C_3$

(Lemma 4.7), $R_\sigma \geq \delta \text{ lfs}(p)/2$ and $\varepsilon/\delta \leq \eta_0$ (Hypothesis 4.1), we have

$$\begin{aligned} \text{dist}(z, \partial\sigma) &\geq \frac{k! \Theta_\sigma^k L_\sigma}{k+1} \geq \frac{k! \Theta_0^k}{k+1} \times \frac{R_\sigma}{C_3} \\ &\geq \frac{k! \Theta_0^k}{k+1} \times \frac{\delta \text{ lfs}(p)}{2C_3} \geq \frac{k! \Theta_0^k}{k+1} \times \frac{\varepsilon \text{ lfs}(p)}{2C_3 \eta_0} = C' \varepsilon \text{ lfs}(p) \end{aligned} \quad (31)$$

where $C' = \frac{k! \Theta_0^k}{2(k+1) C_3 \eta_0}$.

Also, from Lemma 14b of Chapter 4 of [42], Eq. (31) and using the fact that $\|\pi_p(x) - x\| \leq C \varepsilon^2 \text{ lfs}(p)$ for all $x \in \sigma$ (Lemma 5.7), we have

$$\text{dist}(\pi_p(z), \partial\pi_p(\sigma)) \geq C' \varepsilon \text{ lfs}(p) - 2C \varepsilon^2 \text{ lfs}(p), \quad (32)$$

Since π_p embeds $\text{star}(p)$ into T_p (Lemma 2.3), Eq. (32) implies that, for all $x \in \text{star}(p) \setminus \sigma$

$$\|\pi_p(z) - \pi_p(x)\| \geq C' \varepsilon \text{ lfs}(p) - 2C \varepsilon^2 \text{ lfs}(p). \quad (33)$$

From Eq. (33), we have for all $x \in \text{star}(p) \setminus \sigma$

$$\begin{aligned} \|\pi_p^*(x) - \pi_p^*(z)\| &\geq \|\pi_p(x) - \pi_p(z)\| - (\|\pi_p(x) - x\| + \|\pi_p^*(x) - x\|) \\ &\quad - (\|\pi_p(z) - z\| + \|\pi_p^*(z) - z\|) \\ &\geq C' \varepsilon \text{ lfs}(p) - 6C \varepsilon^2 \text{ lfs}(p) > 0 \end{aligned} \quad (34)$$

The last inequality holds for a sufficiently small ε .

Using the fact that π_p^* is injective (from Claim 5.10) together with Eq. (34) show that $\pi_p^*(x) \neq \pi_p^*(z)$ for all $x \in \text{star}(p) \setminus \{z\}$. Hence, z is not covered by any other point of S , and the lemma follows. \square

We will now prove that π is also injective when restricted to $\text{star}(p) \setminus \partial \text{star}(p)$.

Lemma 5.13 (Injectivity of π) *Let ε be sufficiently small. For all p in \mathcal{P} , the map π restricted to $\text{star}(p) \setminus \partial \text{star}(p)$ is injective.*

Proof. To reach a contradiction, assume that there exist x_1, x_2 ($x_1 \neq x_2$) in $\text{star}(p) \setminus \partial \text{star}(p)$ such that $\pi(x_1) = \pi(x_2)$. Then $\pi_p^*(x_1) = \pi_p^*(x_2) = N_{\pi(x_1)} \cap T_p$. Which contradicts the fact that π_p^* is injective when restricted to $\text{star}(p) \setminus \partial \text{star}(p)$ from Lemma 5.12. \square

We will now show that for all $p \in \mathcal{P}$, we have $\pi^{-1}(p) = \{p\}$ when π is restricted to $\hat{\mathcal{M}}$. The following lemma will be used to show that π restricted to $\hat{\mathcal{M}}$ is a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} in Lemma 5.15.

Lemma 5.14 *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 5.4, we have $\sin \angle(\text{aff}(\tau), T_p) \leq A\varepsilon$, which is strictly less than 1 for ε sufficiently.

Case 2. p is not a vertex of τ . Since $\hat{\mathcal{M}}$ has no inconsistencies, $\tau \in \text{star}(q_i)$ for all $i \in \{0, \dots, k\}$.

We will divide the k -simplex τ into a union of $k + 2$ sets,

$$\tau = S \cup \left(\bigcup_{i=0}^k S_i \right) \quad (35)$$

where $S_i = B(q_i, \lambda_i) \cup \tau$ and $S = \tau \setminus \bigcup_{i=0}^k S_i$. The exact value of λ_i will be defined later in the proof but for the time being we assume that $\lambda_i < \mathfrak{m}(q_i)/2$. We prove that x cannot belong to S_i nor to S .

For ε sufficiently small, $x \notin S_i$. Indeed, using Lemmas 2.4 (1), 5.7 and $\varepsilon < \frac{1}{2C\eta_0}$,

$$\begin{aligned} \|q_i - p\| &\leq \|q_i - x\| + \|x - \pi(x)\| \\ &\leq \lambda_i + C\varepsilon^2 \text{fs}(q_i) < \frac{\mathfrak{m}(q_i)}{2} + \frac{\delta \text{fs}(q_i)}{2} < \mathfrak{m}(q_i) \end{aligned}$$

We have reached a contradiction.

Let us prove now that $x \notin S$ under the assumption (proved below) that the distance of any point in the set S to p is $\Omega(\varepsilon \text{fs}(q_i))$, for all q_i , where the constant in the big- Ω depends on ω_0 and η_0 . If there exists $x \in S$ such that $\pi(x) = p$, then using the fact that $\tau \in \text{star}(q_i)$ and Lemma 5.7, we have

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

for all vertices q_i of τ . This contradicts the assumption made above for ε sufficiently small.

We will now show that the assumption is satisfied and prove that for all $x \in S$ and for all q_i , $\|x - p\| = \Omega(\varepsilon \text{fs}(q_i))$.

Since $\hat{\mathcal{M}}$ has no inconsistent configuration, τ is in $\text{star}(q_0)$. Let $m = \text{Vor}^\omega(\tau) \cap T_{q_0}$. From Lemma 4.6, we have

$$R = \sqrt{\|m - q_0\|^2 - \omega(q_0)^2} \leq \|m - q_0\| \leq C_1 \varepsilon \text{fs}(q_0). \quad (36)$$

Using the facts that for all vertices q_i and q_j of τ , $\|q_i - q_j\| \leq C_2 \varepsilon \text{fs}(q_i)$ (from Lemma 4.7 (1)), fs is 1-Lipschitz and ε sufficiently small, we have

$$\text{fs}(q_j) \geq \text{fs}(q_i) - \|q_i - q_j\| \geq (1 - C_2 \varepsilon) \text{fs}(q_i) \geq \frac{\text{fs}(q_i)}{2} \quad (37)$$

Therefore using Eq. (36) and (37), we have for all vertices q_i of τ

$$R \leq C_1 \varepsilon \text{fs}(q_0) \leq 2C_1 \varepsilon \text{fs}(q_i). \quad (38)$$

Since $m \in \text{Vor}(\tau)$, we have

$$\|p - m\| \geq R^2 + \omega(p)^2 \geq R^2. \quad (39)$$

Consider the edge $q_i q_j$ of τ and let c_{ij} be the projection of m onto the line segment $[q_i q_j]$. Observe that the ball of radius $r_{ij} = \sqrt{\|c_{ij} - q_i\|^2 - \omega(q_i)^2} \leq R$ centered at c_{ij} is orthogonal to the balls $B(q_i, \omega(q_i))$ and $B(q_j, \omega(q_j))$. Using the fact that \mathcal{P} is an (ε, δ) -sample of \mathcal{M} and Lemma 4.2 (2), we have

$$r_{ij} \geq \frac{\sqrt{1 - 4\omega_0^2} \|q_i - q_j\|}{2} \geq B\delta \max\{\text{fs}(q_i), \text{fs}(q_j)\} \quad (40)$$

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




Figure 9: Refer to the proof of Lemma 5.14, Case 2.

Let $\lambda_i = \max\{\omega(q_i), B\delta \text{lfs}(q_i)/4\}$ for $q_i \in \{q_0, \dots, q_k\}$. Since $\omega(q_i) < \mathfrak{m}(q_i)/2$ and $B\delta \text{lfs}(q_i)/4 \leq \mathfrak{m}(q_i)/8$ (Lemma 2.4 (1)), we have $\lambda_i < \mathfrak{m}(q_i)/2$. Using Lemma 2.4 (1) with $\varepsilon \leq 1/2$,

$$\lambda_i < \frac{\mathfrak{m}(q_i)}{2} \leq \frac{\varepsilon}{1-\varepsilon} \text{lfs}(q_i) \leq 2\varepsilon \text{lfs}(q_i). \quad (41)$$

Let $x_{ij} = [c_{ij} q_i] \cap \partial B(c_{ij}, r_{ij})$ and $y_{ij} = [c_{ij} q_i] \cap \partial B(q_i, \lambda_i)$. Note that $x_{ij} = \partial B(m, R) \cap [c_{ij} q_i]$, see Figure 9. Therefore

$$\begin{aligned} \|x_{ij} - y_{ij}\| &= \|c_{ij} - x_{ij}\| + \|q_i - y_{ij}\| - \|q_i - c_{ij}\| \\ &= r_{ij} + \lambda_i - \sqrt{r_{ij}^2 + \omega(q_i)^2} \\ &= \frac{2r_{ij}\lambda_i + \lambda_i^2 - \omega(q_i)^2}{r_{ij} + \lambda_i + \sqrt{r_{ij}^2 + \omega(q_i)^2}} \\ &\geq \frac{2r_{ij}\lambda_i}{R + \lambda_i + \sqrt{R^2 + \lambda_i^2}} \quad (\text{since } \lambda_i \geq \omega(q_i) \text{ and } r_{ij} \leq R) \\ &\geq \frac{2r_{ij}\lambda_i}{(2C_1\varepsilon + 2\varepsilon + \sqrt{4C_1^2\varepsilon^2 + 4\varepsilon^2}) \text{lfs}(q_i)} \quad (\text{from Eq. (38) and (41)}) \\ &\geq \frac{B^2\delta^2 \text{lfs}(q_i)}{4\varepsilon(C_1 + 1 + \sqrt{C_1^2 + 1})} \quad (\text{using } \lambda_i \geq B\delta \text{lfs}(q_i)/4 \text{ and Eq. (40)}) \\ &\geq \frac{B^2\delta \text{lfs}(q_i)}{4\eta_0(C_1 + 1 + \sqrt{C_1^2 + 1})} \quad (\text{since } \varepsilon/\delta \leq \eta_0 \text{ from Hypothesis 4.1}) \end{aligned}$$


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

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

where the constant in the big- Ω that depends on ω_0 and η_0 .

Let z_{ij} denote the point closest to y_{ij} on $\partial B(m, R)$, see Figure 9.

r0.55



Intersecting Chords Theorem.

From the Intersecting Chords Theorem (see Figure 5) for circles, we have :

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

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

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

where the constant in the big- Ω depends on ω_0 and η_0 . The last equality follows from the facts that $A_{ij} = \Omega(\delta \text{lfs}(q_i))$, $R \leq 2C_1 \varepsilon \text{lfs}(q_i)$ (from Eq. (38)) and $\varepsilon/\delta \leq \eta_0$.

Using the fact that ε is sufficiently small and Eq. (37), we have for all vertices q of τ

$$\|z_{ij} - y_{ij}\| = \Omega(\varepsilon \text{lfs}(q_i)) = \Omega(\varepsilon \text{lfs}(q)) \quad (42)$$

where the constant in the big- Ω depends on ω_0 and η_0 .

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

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

From the definition of S , we have $S \subset \text{conv}(\tilde{S}) \subset B(m, R)$. This implies

$$\text{dist}(\text{conv}(\tilde{S}), \partial B(m, R)) \leq \text{dist}(S, \partial B(m, R)). \quad (43)$$

Using convexity and Eq. (42), we have

$$\text{dist}(\text{conv}(\tilde{S}), \partial B(m, R)) = \min_{i, j (\neq i) \in \{0, \dots, k\}} \|z_{ij} - y_{ij}\| = \Omega(\varepsilon \text{lfs}(q)) \quad (44)$$

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

Using Eq. (43) and (44), and the facts that $S \subset B(m, R)$ and $p \notin B(m, R)$ (from Eq. (39)), we have

$$\text{dist}(p, S) \geq \text{dist}(S, \partial B(m, R)) = \Omega(\varepsilon \text{lfs}(q))$$

for all vertices q of τ . This proves that our assumption was valid and achieves the proof of the lemma. \square

We will now show that π restricted to $\hat{\mathcal{M}}$ gives a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} .

Lemma 5.15 (Homeomorphism) *For ε sufficiently small, the restriction $\pi|_{\hat{\mathcal{M}}}$ of the map π to $\hat{\mathcal{M}}$ is a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} .*

Proof. From Lemma 4.7, we know that $\hat{\mathcal{M}} \subset \mathbb{R}^d \setminus \mathcal{O}$. Therefore, according to Lemma 5.1, $\pi|_{\hat{\mathcal{M}}}$ is continuous. As $\hat{\mathcal{M}}$ is compact and \mathcal{M} is a Hausdorff space, $\pi|_{\hat{\mathcal{M}}}$ is a homeomorphism if it is injective and surjective (see Theorem A.2 in Appendix A). We assume in the rest of the proof that \mathcal{M} is connected. Arguments for one connected component can be extended to the case where \mathcal{M} has more than one connected component if the sample \mathcal{P} contains at least one point from each connected component of \mathcal{M} .

The proof of injectivity is similar to the proof of Lemma 18 in [2]. Using the fact that π is injective in $\text{star}(p) \setminus \partial \text{star}(p)$ for all p in \mathcal{P} (Lemma 5.13), that \mathcal{M} together with $\pi|_{\hat{\mathcal{M}}}$ forms a covering space of the image $\pi(\hat{\mathcal{M}})$ (see Definition A.11 in Appendix A).

Let $\pi(\hat{\mathcal{M}}) = \bigcup_i C_i$ where C_i are the maximal connected components of $\pi(\hat{\mathcal{M}})$.

Claim 5.16 $C_i \cap \pi(\mathcal{P}) \neq \emptyset$.

Proof. Let x be a point in $\hat{\mathcal{M}}$ such that $\pi(x) \in C_i$, and let $p \in \mathcal{P}$ such that $x \in \text{star}(p)$. Since π is a continuous map and $\text{star}(p)$ is a connected space, $\pi(\text{star}(p)) \subset \mathcal{M}$ is also connected by Proposition A.3. Since C_i is a maximal connected component, $\pi(\text{star}(p)) \subset C_i$. This implies $\pi(p) \in C_i$. \square

Let p_i be a point of \mathcal{P} such that $\pi(p_i) \in C_i$. Since $|\pi^{-1}(p_i)| = 1$ by Lemma 5.14, we have from Lemma A.12 in Appendix A, $|\pi^{-1}(x)| = |\pi^{-1}(p_i)| = 1$ for all $x \in C_i$. This implies $|\pi^{-1}(x)| = 1$ for all $x \in \pi(\hat{\mathcal{M}}) = \bigcup_i C_i$, i.e. π restricted to $\hat{\mathcal{M}}$ is injective.

Surjectivity of $\pi|_{\hat{\mathcal{M}}}$ follows from the fact that π embeds $\hat{\mathcal{M}}$ in \mathcal{M} . As $\hat{\mathcal{M}}$ is a compact topological k -manifold without boundary, $\pi(\hat{\mathcal{M}})$ is a k -dimensional manifold without boundary. Since \mathcal{M} is a connected topological k -manifold without boundary, we conclude that $\pi(\hat{\mathcal{M}}) = \mathcal{M}$. \square

Proof of Property P4

Lemma 5.17 (Ambient isotopy) *For ε sufficiently small there exists an ambient isotopy*

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

such that the map $F(\cdot, 0)$ restricted to $\hat{\mathcal{M}}$ is an identity map on $\hat{\mathcal{M}}$ and $F(\hat{\mathcal{M}}, 1) = \mathcal{M}$.

Proof. Let

$$f : \hat{\mathcal{M}} \times [0, 1] \longrightarrow \mathbb{R}^d, \quad (x, t) \mapsto x + t(\pi(x) - x)$$

Note that $f(\cdot, 0)$ is an identity map on $\hat{\mathcal{M}}$ and $f(\cdot, 1)$ is the map π restricted to $\hat{\mathcal{M}}$. The map f is an isotopy because the maps

$$f_t : \hat{\mathcal{M}} \longrightarrow \mathbb{R}^d, \quad x \mapsto f(x, t)$$

are homeomorphisms between $\hat{\mathcal{M}}$ and $f_t(\hat{\mathcal{M}})$.

By Theorem A.10, there exists an ambient isotopy $F : \mathbb{R}^d \times [0, 1] \longrightarrow \mathbb{R}^d$ such that $F(\cdot, 0)|_{\hat{\mathcal{M}}} = f(\cdot, 0)$ and $F(\cdot, 1)|_{\hat{\mathcal{M}}} = f(\cdot, 1)$. \square

Proof of Property P5

Lemma 5.18 (Pointwise approximation) *For ε sufficiently small and the map π restricted to $\hat{\mathcal{M}}$, we have $\text{dist}(x, \pi^{-1}(x)) = O(\varepsilon^2 \text{lfs}(x))$, where the constant in big- O depends on k , ω_0 , η_0 and Θ_0 .*

Proof. From Lemma 5.15, π restricted to $\hat{\mathcal{M}}$ is a homeomorphism between $\hat{\mathcal{M}}$ and \mathcal{M} . In this proof we will only consider π restricted to $\hat{\mathcal{M}}$.

For $x \in \mathcal{M}$, let $x' = \pi^{-1}(x) \in \hat{\mathcal{M}}$ and let p be a point in \mathcal{P} such that $x' \in \text{star}(p)$. Using the facts that $\|p - x'\| \leq C_2\varepsilon \text{lfs}(p)$ (from Lemma 4.7 (1)) and $\|x - x'\| \leq C\varepsilon^2 \text{lfs}(p)$ (from Lemma 5.7), we have

$$\|p - x\| \leq \|p - x'\| + \|x - x'\| \leq (C_2\varepsilon + C\varepsilon^2)\text{lfs}(p) \quad (45)$$

We will now bound $\text{lfs}(p)$ as a function of $\text{lfs}(x)$ using the 1-Lipschitz property of lfs . Hence, for ε sufficiently small, we have

$$\text{lfs}(x) \geq \text{lfs}(p) - \|p - x\| \geq (1 - C_2\varepsilon - C\varepsilon^2)\text{lfs}(p) \geq \text{lfs}(p)/2 \quad (46)$$

Using Eq. (45) and (46), we get $\|x - x'\| \leq 2C\varepsilon^2 \text{lfs}(x)$. \square

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 [19, 28] 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 [7, 27] 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 [5]. It easily follows from [8] 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 [12].

Acknowledgments

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

References

- [1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1999.
- [2] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *International Journal of Computational Geometry and Application*, 12:125–141, 2002.
- [3] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1):78–96, 1987.
- [4] 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.
- [5] M. Belkin, J. Sun, and Y. Wang. Discrete laplace operator on meshed surfaces. In *Proc. ACM Symp. on Computational Geometry*, pages 278–287, 2008.
- [6] J.-D. Boissonnat and J. Flötotto. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004.
- [7] 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.
- [8] J.-D. Boissonnat, F. Nielsen, and R. Nock. Bregman Voronoi Diagrams. *Discrete and Computational Geometry*, 44(2):281–307, 2010.
- [9] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. ACM Symp. on Computational Geometry*, pages 270–277, 2008.
- [10] J.-D. Boissonnat and M. Yvinec. *Algorithmic geometry*. Cambridge University Press, 1998.
- [11] G. E. Bredon. *Topology and Geometry*. Springer-Verlag, GTM, 1993.
- [12] 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.
- [13] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curve and Surfaces*. Springer, 2006.
- [14] 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.
- [15] F. Chazal and S. Y. Oudot. Towards Persistence-Based Reconstruction in Euclidean Spaces. In *Proc. ACM Symp. on Computational Geometry*, pages 232–241, 2008.
- [16] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.
- [17] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver Exudation. *Journal of ACM*, 47:883–904, 2000.
- [18] 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.
- [19] S.-W. Cheng, Y. Wang, and Z. Wu. Provable Dimension Detection using Principle Component Analysis. *International Journal of Computational Geometry and Application*, 18(5):415–440, 2008.

- [20] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [21] D. Cohen-Steiner and T. K. F. Da. A greedy Delaunay Based Surface Reconstruction Algorithm. *The Visual Computer*, 20:4–16, 2004.
- [22] T. K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2006.
- [23] H. Federer. *Geometric measure theory*. Springer New York, 1969.
- [24] 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.
- [25] D. Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002.
- [26] J. H. G. Fu. Convergence of curvature in secant approximations. *Journal of Differential Geometry*, 37:117–190, 1993.
- [27] S. Funke and E. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pages 781–780, 2002.
- [28] J. Giesen and U. Wagner. Shape Dimension and Intrinsic Metric from Samples of Manifolds. *Discrete and Computational Geometry*, 32(2):245–267, 2004.
- [29] M. Gopi, S. Khrisnan, and C. T. Silva. Surface Reconstruction based on Lower Dimensional Localized Delaunay Triangulation. In *Proc. Eurographics*, pages 363–371, 2000.
- [30] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [31] M. W. Hirsh. *Differential Topology*. Springer-Verlag, New York, NY, 1976.
- [32] D. D. L. 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.
- [33] 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.
- [34] X.-Y. Li. Generating Well-Shaped d -dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.
- [35] W. S. Massey. *Algebraic Topology : An Introduction*. Springer-Verlang, GTM 56, 1967.
- [36] 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.
- [37] 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.
- [38] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [39] H. S. Seung and D. D. Lee. The manifold ways of perception. *Science*, 290:2268–2269, 2000.
- [40] J. Shewchuk. Star splaying: An algorithm for repairing delaunay triangulations and convex hulls. In *Proc. ACM Symp. on Computational Geometry*, pages 237–246, 2005.

- [41] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.
- [42] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957.
- [43] 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.

A Definitions and results from topology

In this section, we will give the definitions and results from topology used in the paper.

Definition A.1 (Homeomorphism) *Two topological spaces X and Y are homeomorphic if there exists a continuous bijective map $f : X \rightarrow Y$ such that the inverse f^{-1} is also continuous. The map f is said to be a homeomorphism between X and Y .*

The following standard result on homeomorphisms is proved, e.g., in Chapter 1 of [11, Theorem 7.8].

Theorem A.2 *A continuous bijective map f from a compact space to a Hausdorff space is a homeomorphism.*

For a proof of the following standard result from topology see, e.g., Chapter 1 of [11, Proposition 4.6].

Proposition A.3 *Let $f : X \rightarrow Y$ be a continuous map and let X be a connected space. Then, $f(X)$ is a connected space.*

Definition A.4 (C^r -diffeomorphism) *Let $U \subseteq \mathbb{R}^k$ and $V \subseteq \mathbb{R}^d$. A bijective C^r -function $f : U \rightarrow V$ is called C^r -diffeomorphism if f^{-1} is a C^r -function.*

Definition A.5 (C^r -embedding of simplices in \mathcal{M}) *Let σ be an i -simplex, and let $f : \sigma \rightarrow \mathcal{M}$ be a C^r -function. The simplex σ is C^r -embedded by f in \mathcal{M} if f is an injective mapping and for all $x \in \sigma$, the rank of the linear map $df(x) : \mathbb{R}^i \rightarrow T_{f(x)}$ is i , where $T_{f(x)}$ is the tangent space to \mathcal{M} at $f(x)$.*

Definition A.6 (Simplexwise positive map) *Let σ be an i -simplex, 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 lemma is a special case of a standard result from piecewise linear topology. See, e.g., Appendix II of [42, Lemma 15a].

Lemma A.7 *Let K be a d -dimensional piecewise linear manifold with all the d -dimensional simplices of K be oriented, and let the continuous map $f : K \rightarrow \mathbb{R}^d$ be a simplexwise positive map for all the d -simplices in K . Then for any connected open subset R of $\mathbb{R}^d \setminus f(\partial K)$, any two points of R not in $f(K^{d-1})$, where K^{d-1} is the $d-1$ skeleton of K , are covered the same number of times. If this number is 1, then f , restricted to the open subset $R' = f^{-1}(R)$ of K , is injective.*

Definition A.8 (Isotopy) *Let X, Y be topological spaces. The map $F : X \times [0, 1] \rightarrow Y$ is called an isotopy of X if for any $t \in [0, 1]$*

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

is a homeomorphism between X and $F_t(X)$.

Definition A.9 (Ambient isotopy) *An isotopy $F : X \times [0, 1] \rightarrow Y$ is called an ambient isotopy if $X = Y$.*

The following theorem is a special case of a more general result in topology. See, e.g., Theorem 1.3 from Chapter 8 of [31].

Theorem A.10 *Let M be a manifold, $V \subset M$ be a compact submanifold and let $F : V \times [0, 1] \rightarrow M$ be an isotopy of V . If $F(V \times [0, 1]) \subset M \setminus \partial M$, then F extends to an ambient isotopy of M .*

We will now recall the definition of a covering space. See, e.g., [11, 30, 35].

Definition A.11 (Covering space) *Let X be a topological space. A covering space of X is a space \tilde{X} together with a continuous surjective map $f : \tilde{X} \rightarrow X$ satisfying the following condition: There exists an open cover $\{U_\alpha\}$ of X such that, for each α , $f^{-1}(U_\alpha)$ is a disjoint union of open sets in \tilde{X} , each of which is mapped homeomorphically onto U_α by f .*

The following lemma follows directly from the above definition. See, e.g., [11, 30, 35].

Lemma A.12 *Let $f : \tilde{X} \rightarrow X$ be a covering map, and let $X = \bigcup_i X_i$ where X_i are the connected components of X . Then, the cardinality of $f^{-1}(x)$ is constant for all $x \in X_i$.*

B Main notations

General notations

$B(c, r) = \{x \in \mathbb{R}^d \mid \|c - x\| < r\}$

$\bar{B}(c, r) = \{x \in \mathbb{R}^d \mid \|c - x\| \leq r\}$

$B_{\mathcal{M}}(c, r) = B(c, r) \cap \mathcal{M}$

$\bar{B}_{\mathcal{M}}(c, r) = \bar{B}(c, r) \cap \mathcal{M}$

$df(x)$ derivative of the function f at x

$\dim(U)$ dimension of the affine space U

$\text{conv}(S)$ smallest convex set containing the set S

Submanifold

d dimension of the ambient space

k dimension of \mathcal{M}

\mathcal{M} manifold

N_x normal space at x

\mathcal{P} sample

$|\mathcal{P}|$ cardinality of the sample \mathcal{P}

T_x tangent space at x

Weights (Section 2.1.1)

$p^\omega = (p, \omega(p))$ weighted point

ω weight assignment

$\tilde{\omega}$ relative amplitude of ω

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

$LN(p)$ local neighborhood of p (Section 3.1)

$\mathfrak{m}(p)$ distance of p to its nearest neighbor

Shape measure of simplex τ (Section 2.3)

$\text{aff}(\tau)$ affine hull of τ_p

c_τ circumcenter

$D_\tau(p)$ altitude (Section 4.1)

Δ_τ length of longest edge (diameter)

$H_\tau(p, \omega(p))$ excentricity (Section 4.1)

L_τ length of shortest edge

N_τ affine space orthogonal to $\text{aff}(\tau)$

o_τ orthocenter
 Φ_τ orthoradius
 R_τ circumradius
 $\rho_\tau = R_\tau/L_\tau$ radius-edge ratio
 $\tau_p = \tau \setminus \{p\}$ face of τ opposite to p (Section 4.1)
 Θ_τ fatness
 V_τ volume