Manifold reconstruction using Tangential Delaunay Complexes

Jean-Daniel Boissonnat, Arijit Ghosh

Manifold Reconstruction using Tangential Delaunay Complexes *

Jean-Daniel Boissonnat
Jean-Daniel.Boissonnat@sophia.inria.fr

Arijit Ghosh
Arijit.Ghosh@sophia.inria.fr

INRIA, Team Geometrica
2004 route des Lucioles
06902 Sophia-Antipolis, France

ABSTRACT

We give a provably correct algorithm to reconstruct a k-dimensional manifold embedded in d-dimensional Euclidean space. 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, 19, 20], and the technique of sliver removal by weighting the sample points [13]. Differently from previous methods, we do not construct any subdivision of the embedding d-dimensional space. As a result, the running time of our algorithm depends only linearly on the extrinsic dimension d while it depends quadratically on the size of the input sample, and exponentially on the intrinsic dimension k. To the best of our knowledge, this is the first certified algorithm for manifold reconstruction whose complexity depends linearly on the ambient dimension. We also prove that for a dense enough sample the output of our algorithm is isotopic to the manifold and a close geometric approximation of the manifold.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms and problems**]: Geometrical problems and computations; I.3.5 [**Computational Geometry and Object Modeling**]: Curve, surface, solid, and object representations.

General Terms

Algorithms, Theory.

Keywords

Tangential Delaunay complex, Weighted Delaunay triangulation, manifold reconstruction, manifold learning, sampling conditions, sliver exudation.

*This work was partially supported by the ANR project GAIA. For the full version of the paper refer to [7].

1. INTRODUCTION

Manifold reconstruction consists in computing a PL approximation of an unknown manifold $\mathbb{M} \subset \mathbb{R}^d$ from a finite sample of unorganized points $\mathcal{P}$ lying on $\mathbb{M}$ or close to $\mathbb{M}$. 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 [17], and the survey by Cazals and Giesen in [10] for recent results. The output of those methods is a triangulated surface that approximates $\mathbb{M}$. 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 has led to algorithms whose complexities depend exponentially on d [26, 11, 14].

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 [28]. 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 [30], LLE [27], Laplacian eigenmaps [4], Hessian eigenmaps [18], diffusion maps [24, 25], principal manifolds [31]. Most of those methods reduces 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 $\mathbb{M}$. Those methods come with no or very limited guarantees. For example, Isomap provides a correct embedding only if $\mathbb{M}$ is isometric to a convex open set of $\mathbb{R}^k$. To be able to better approximate the sampled manifold, another route is to extend the work on surface reconstruction and to construct a PL approximation of $\mathbb{M}$ 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 [14], and Boissonnat, Guibas and Oudot [8]. 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 developped in small dimensions and propose a way to avoid computing data structures in the ambient space. We assume that $\mathbb{M}$ is a smooth manifold of known dimension k and that we can compute the tangent space to $\mathbb{M}$ at any sample point. Under those conditions, we propose a provably correct algorithm that allows to construct a simplicial complex of dimension k that approximates $\mathbb{M}$. 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 [14] and [8] but dramatically reduces the dependance 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 [12] 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* defined in [20, 6, 19], and the technique of sliver removal by weighting the sample points [13]. 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 [20], no solution was known except for the case of curves ($k = 1$) [19]. We show that we can remove inconsistencies by weighting the sample points under appropriate sample conditions. We can then prove that the approximation returned by our algorithm is isotopic to $\mathbb{M}$, and a close geometric approximation of $\mathbb{M}$.

Our algorithm can be seen as a *local* version of the cocone algorithm of Cheng et al. [14]. 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 d-dimensional Delaunay triangulation of the data points and therefore implicitly relies on a global partition of the ambient space. This is key to our analysis and makes our method depart from other local algorithms that have been proposed in the surface reconstruction literature [16, 23].

Notations In the rest of the paper, we assume that $\mathbb{M}$ is a smooth manifold of dimension k embedded in $\mathbb{R}^d$. We call $\mathcal{P} = \{p_1, \ldots, p_n\}$ a finite sample of points from $\mathbb{M}$. We denote by T_p the k-dimensional tangent space at point $p \in \mathbb{M}$. We write $B(c, r)$ for the d-dimensional ball centered at c of radius r. We define the angle between two vector spaces U and V as

$$UV = \max_{u \in U} \min_{v \in V} \angle uv. \tag{1}$$

If τ is a j-simplex, the $(d - j)$-dimensional normal space of $\mathrm{aff}(\tau)$ is denoted by N_τ.

2. DEFINITIONS AND PRELIMINARIES

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 visualize the weighted point $(p, \omega(p))$ as 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 $\mathcal{P} = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^d$, a *weight function* on $\mathcal{P}$ is a non-negative real-valued function $\omega : \mathcal{P} \to [0, \infty)$. Write $p_i^\omega = (p_i, \omega(p_i))$ and $\mathcal{P}^\omega = \{p_1^\omega, \ldots, p_n^\omega\}$.

We define the *relative amplitude* of ω, denoted as $\tilde{\omega}$, as $\max_{p \in \mathcal{P}, q \in \mathcal{P} \setminus \{p\}} \frac{\omega(p)}{||p-q||}$. In the paper, we assume that $\tilde{\omega} \leq \omega_0 < 1/2$, for some constant ω_0 to be fixed later.

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 and radius are called the *orthocenter* and the *orthoradius* of τ. 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. Plainly, its center o_τ lies in $\mathrm{aff}(\tau)$. The radius of the orthosphere of τ is denoted by R'_τ.

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

2.1.2 Weighted Voronoi diagram and Delaunay triangulation

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

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

The weighted Voronoi cells and their k-dimensional faces, $0 \leq k \leq d$, form a cell complex, called the weighted Voronoi diagram of $\mathcal{P}$, that decomposes $\mathbb{R}^d$ into convex polyhedral cells. See [2].

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

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

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, 10]. 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

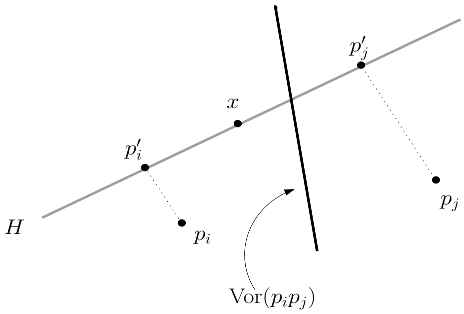

Figure 1: Refer to Lemma 1. The red line denotes the k-dimensional plane H and the black line denotes $\mathrm{Vor}^\omega(p_ip_j)$.

(see Lemma 1). The other main reason is that some flat simplices named slivers can be removed from a Delaunay triangulation by weighting the vertices (see [14, 8, 13] and Section 4).

LEMMA 1. *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 the k-dimensional weighted Voronoi diagram of $\mathcal{P}'$ where $\mathcal{P}'$ is the orthogonal projection of $\mathcal{P}$ onto H and the squared weight of p_i' is $\omega^2(p_i) - \|p_i - p_i'\|^2$.*

PROOF. By Pythagoras theorem, we have $\forall x \in H \cap \mathrm{Vor}^\omega(p_i)$, $\|x-p_i\|^2 - \omega^2(p_i) \le \|x-p_j\|^2 - \omega^2(p_j) \Leftrightarrow \|x-p_i'\|^2 + \|p_i - p_i'\|^2 - \omega^2(p_i) \le \|x-p_j'\|^2 + \|p_j - p_j'\|^2 - \omega^2(p_j)$, where p_i' denotes the orthogonal projection of $p_i \in \mathcal{P}$ onto H. Hence the restriction of $\mathrm{Vor}^\omega(\mathcal{P})$ to H is the weighted Voronoi diagram of the weighted points $(p_i', \omega_i) \in H$ where $\omega_i^2 = -\|p_i - p_i'\|^2 + \omega^2(p_i)$. $\square$

LEMMA 2 ([13]). *If τ is a simplex of $\mathrm{Del}^\omega(\mathcal{P})$, then*

1. *$\forall z \in \mathrm{aff}(\mathrm{Vor}^\omega(\tau))$ and $\forall p, q \in \tau$ we have $\|q - z\| \le \frac{\|p-z\|}{\sqrt{1-4\omega_0^2}}$.*
2. *$R_\tau \le \frac{R_\tau'}{\sqrt{1-4\omega_0^2}}$.*
3. *$\forall z \in \mathrm{aff}(\mathrm{Vor}^\omega(\tau))$ and $\forall p \in \tau$, $r_z = \sqrt{\|p-z\|^2 - \omega^2(p)} \ge R_\tau'$.*

2.2 Sampling conditions

2.2.1 Local feature size

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

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 $x \in \mathbb{M}$ there exists a point $p \in \mathcal{P}$ such that $||x-p|| \le \varepsilon\ \mathrm{lfs}(x)$, and (2) for any two distinct points $p, q \in \mathcal{P}$, $||p-q|| \ge \delta\ \mathrm{lfs}(p)$.[1] The ratio ε/δ is called the *relative density* of $\mathcal{P}$.

[1]Observe that the sparsity condition (2) is mandatory if one wants to infer the dimension of $\mathbb{M}$ from a sample [22].

We will use the following results from [22]. We write l_p for the distance between $p \in \mathcal{P}$ and its nearest neighbor in $\mathcal{P} \setminus \{p\}$.

LEMMA 3. *Given an (ε, δ)-sample $\mathcal{P}$ of $\mathbb{M}$, we have*

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

2.3 Slivers and good simplices

Consider a j-simplex τ, where $1 \le j \le k+1$. We denote by R_τ, L_τ, V_τ and $\rho(\tau) = R_\tau/L_\tau$ the circumradius, the shortest edge length, the volume, and the radius-edge ratio of τ respectively. We further define $\sigma(\tau) = \left(V_\tau/L_\tau^j\right)^{1/j}$, as the *sliverity measure* of τ. The orthocenter of τ is denoted by o_τ and its orthoradius by R_τ'.

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, and $H_\tau(p)$ for the *signed distance*[2] from o_τ to $\mathrm{aff}(\tau_p)$. We state without proof the following easy lemma.

LEMMA 4. *Let τ be a simplex of $\mathrm{Del}^\omega(\mathcal{P})$ and $p \in \tau$ s.t.*

1. *There exists $z \in \mathrm{aff}(\mathrm{Vor}^\omega(\tau))$ s.t. $\|z-p\| \le L_1$.*
2. *$\|p-q\| \le L_2$ for all vertices q of τ.*
3. *$R_\tau' \le \gamma_0 L_\tau$.*

Then $|H_\tau(x)| = \mathrm{dist}(o_\tau, \mathrm{aff}(\tau_x)) \le L_1 + (1+\gamma_0+\omega_0)L_2$ for all vertices x of τ.

LEMMA 5 ([13]). *Let τ be a simplex of $\mathrm{Del}^\omega(\mathcal{P})$. Let p be any vertex of τ and write $H_\tau(\omega(p))$ (instead of $H_\tau(p)$) for the signed distance of o_τ to $\mathrm{aff}(\tau_p)$ parametrized by the weight of p. We have $H_\tau(\omega(p)) = H_\tau(0) - \frac{\omega^2(p)}{2D_\tau(p)}$.*

Slivers are a special type of flat simplices. The property of being a sliver is measured in terms of a parameter σ_0, called the *sliverity bound*, to be fixed later in Section 4.

DEFINITION 1 (SLIVER). *Given a positive parameter σ_0, slivers are defined by induction on the dimension : (1) a simplex of dimension less than 3 is not a σ_0-sliver, and (2) for $j \ge 3$, a j-simplex τ is a σ_0-sliver if $\sigma(\tau) < \sigma_0$ and, $\forall \tau' \subset \tau$, $\sigma(\tau') \ge \sigma_0$.*

LEMMA 6. *If a j-simplex τ is a σ_0-sliver, then $D_\tau(p) < j\sigma_0 L_\tau$ for all vertices $p \in \tau$.*

PROOF. The volume of τ is

$$V_\tau = \frac{V_{\tau_p} \cdot D_\tau(p)}{j} = \frac{\sigma^{j-1}(\tau_p) L_{\tau_p}^{j-1} \cdot D_\tau(p)}{j}$$

and it is also equal to $\sigma^j(\tau)L_\tau^j$. Since τ is a σ_0-sliver, we have $\sigma(\tau) < \sigma_0$ and $\sigma(\tau_p) \ge \sigma_0$. Therefore we get

$$D_\tau(p) = j\,\frac{\sigma^j(\tau)}{\sigma^{j-1}(\tau_p)} \times \frac{L_\tau^j}{L_{\tau_p}^{j-1}} < j\sigma_0 L_\tau.$$

$\square$

[2]$H(p)$ is positive if $o_\tau \in \mathrm{aff}(\tau)$ and p lie on the same side of $\mathrm{aff}(\tau_p)$, and it is negative if they lie on different sides of $\mathrm{aff}(\tau_p)$.

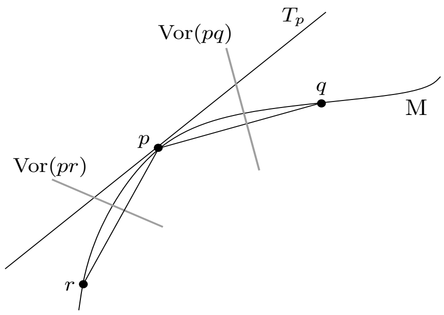

Figure 2: $[pq]$ and $[pr]$ are edges of the star of p in $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ since their dual Voronoi edges intersect the tangent space T_p at p.

Definition 2 (Good simplex). *Given two positive constants ρ_0 and σ_0, a simplex τ is called a (ρ_0, σ_0)-good simplex if $\rho(\tau) \le \rho_0$ and τ nor its subsimplices are σ_0-slivers.*

The following important lemma is known (see e.g. [14]).

Lemma 7 (Normal approximation). *Let τ be a (ρ_0, σ_0)-good j-simplex for $j \le k$ with vertices on a k-dimensional smooth manifold $\mathbb{M}$, and $p \in \tau$ s.t. the lengths of the edges of τ that are incident to p are less than $c_3\varepsilon\,\mathrm{lfs}(p)$ for $c_3\varepsilon < 1$. Then, for any normal vector n_p of $\mathbb{M}$ at p, τ has a normal n_τ such that $\angle n_p n_\tau \le \alpha_j(\sigma_0)\varepsilon$ where $\alpha_j(\sigma_0)$ depends on j, σ_0, ρ_0 and c_3, and $1/\alpha_j(\sigma_0)$ vanishes with σ_0.*

In the sequel, we will simply write $\alpha(\sigma_0)$ for $\alpha_k(\sigma_0)$.

2.4 Tangential Delaunay complex and inconsistent configurations

Let $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$ be the weighted Delaunay triangulation of $\mathcal{P}$ restricted to the tangent space T_{p_i}. Equivalently, the simplices of $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$ are the simplices of $\mathrm{Del}^{\omega}(\mathcal{P})$ whose Voronoi dual faces intersect T_{p_i}, i.e. $\tau \in \mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$ iff $\mathrm{Vor}^{\omega}(\tau) \cap T_{p_i} \neq \emptyset$. Observe that $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$ 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 $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$ are k-dimensional triangulations. Finally, write $\mathrm{star}(p_i)$ for the star of p_i in $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$, i.e. the set of simplices that are incident to p_i in $\mathrm{Del}^{\omega}_{p_i}(\mathcal{P})$.

We call *tangential Delaunay complex* or tangential complex for short, the simplicial complex $\{\tau, \tau \in \mathrm{star}(p), p \in \mathcal{P}\}$. We denote it by $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$. By our assumption above, $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ is a k-dimensional complex contained in $\mathrm{Del}^{\omega}(\mathcal{P})$.

By duality, computing $\mathrm{star}(p_i)$ is equivalent to computing the restriction of the (weighted) Voronoi cell of p_i to T_{p_i}, which, by Lemma 1, reduces to computing a cell in a k-dimensional weighted Voronoi diagram embedded in T_{p_i}. It follows that the tangential complex can be computed without constructing any data structure of dimension higher than k, the intrinsic dimension of $\mathbb{M}$.

The tangential Delaunay complex is *not* in general a triangulated manifold and therefore not a good approximation of $\mathbb{M}$. This is due to the presence of so-called inconsistencies. Consider a k-simplex τ of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ with two vertices p_i and p_j such that τ is in $\mathrm{star}(p_i)$ but not in $\mathrm{star}(p_j)$ (refer to Fig. 3). We write $B_i(\tau)$ for the open ball centered on T_{p_i} that is orthogonal to the (weighted) vertices of τ^{ω}, and

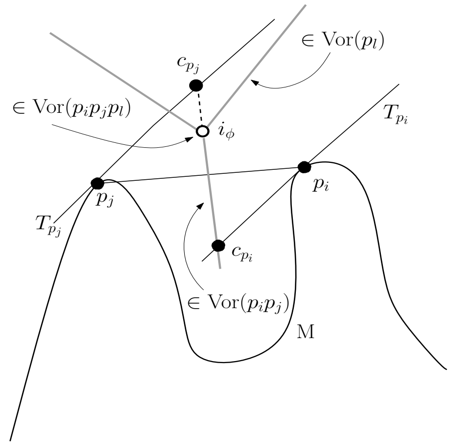

Figure 3: An inconsistent configuration in the unweighted case. Edge $[p_ip_j]$ is in $\mathrm{Del}_{p_i}(\mathcal{P})$ but not in $\mathrm{Del}_{p_j}(\mathcal{P})$ since $\mathrm{Vor}(p_ip_j)$ intersects T_{p_i} but not T_{p_j}. This happens because $[c_{p_i}c_{p_j}]$ penetrates (at i_ϕ) the Voronoi cell of a point $p_l \neq p_i, p_j$, therefore creating an inconsistent configuration $\phi = [p_i, p_j, p_l]$.

denote by c_{p_i} and r_{p_i} its center and its radius. According to our definition, τ is inconsistent iff $B_i(\tau)$ is further than orthogonal from all weighted points in $\mathcal{P}^{\omega} \setminus \tau^{\omega}$ while there exists a weighted point of $\mathcal{P}^{\omega} \setminus \tau^{\omega}$, say p_l^{ω}, that is closer than orthogonal from $B_j(\tau)$. We deduce from the above discussion that the line segment $[c_{p_i}c_{p_j}]$ has to penetrate the interior of $\mathrm{Vor}^{\omega}(p_l)$.

We formally define an inconsistent configuration as follows.

Definition 3 (Inconsistent configuration). *$\phi = [p_1, \ldots, p_{k+2}]$ is called an inconsistent configuration of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ witnessed by the triplet (p_i, p_j, p_l) if*

- *The k-simplex $\tau = \phi \setminus \{p_l\}$ is in $\mathrm{star}(p_i)$ but not in $\mathrm{star}(p_j)$.*
- *τ is a (ρ_0, σ_0)-good simplex.*
- *$\mathrm{Vor}^{\omega}(p_l)$ is the first cell of $\mathrm{Vor}^{\omega}(\mathcal{P})$ whose interior is intersected (at a point denoted by i_ϕ) by the line segment $[c_{p_i}c_{p_j}]$ oriented from c_{p_i} to c_{p_j}. Here $c_{p_i} = T_{p_i} \cap \mathrm{Vor}^{\omega}(\tau)$ and $c_{p_j} = T_{p_j} \cap \mathrm{aff}(\mathrm{Vor}^{\omega}(\tau))$.*

Note that i_ϕ is the center of a sphere that is orthogonal to the weighted vertices of τ and also to p_l^{ω}, and further than orthogonal from all the other weighted points of $\mathcal{P}^{\omega}$. Equivalenty, i_ϕ is the point on $[c_{p_i}c_{p_j}]$ that belongs to $\mathrm{Vor}^{\omega}(\phi)$.

An inconsistent configuration is thus a $(k+1)$-simplex of $\mathrm{Del}^{\omega}(\mathcal{P})$. Such a configuration or its subfaces do not necessarily belong to $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$[3] We write $\mathrm{In}^{\omega}(\mathcal{P})$ for the subcomplex of $\mathrm{Del}(\mathcal{P})$ consisting of all the inconsistent configurations of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ and their subfaces.

3. STRUCTURAL RESULTS

In the rest of the paper, we will need several constants : ω_0, ρ_0 and σ_0 (to define slivers, good simplices and inconsistent configurations), and c_3 (to be able to use lemma 7).

[3] In fact, as already noted, no $(k+1)$-simplex belongs to $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ when the points are in general position.

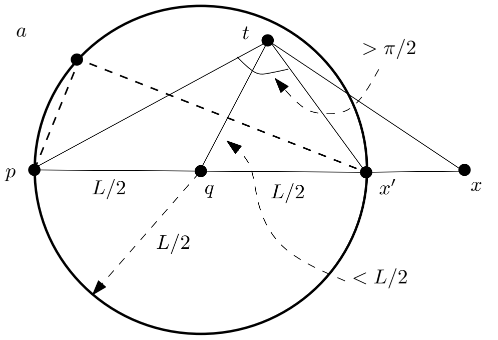

Figure 4: Refer to Lemma 8. x' is a point on the line segment such that $||p-x'|| = c_1\varepsilon\,\mathrm{lfs}(p)$, $L = c_1\varepsilon\,\mathrm{lfs}(p)$, $\angle pax' = \pi/2$ and $\angle ptx \geq \angle ptx' > \pi/2$.

These constants will be fixed in Section 3 and in Section 4. For simplicity, we will write sliver for σ_0-sliver and inconsistent configuration for (ρ_0, σ_0)-inconsistent configuration.

We give now an hypothesis which is assumed to be satisfied in the following results.

HYPOTHESIS 1. *$\mathcal{P}$ is an (ε, δ)-sample of $\mathbb{M}$ of bounded relative density, i.e. $\varepsilon/\delta \leq \eta_0$ for some positive constant η_0. We assume further that ε is small enough and that $2\alpha(\sigma_0)\varepsilon < 1$. Finally, we assume that $\tilde{\omega} \leq \omega_0$ where $\tilde{\omega}$ is the relative amplitude of the weight assignement ω and ω_0 is a positive constant less than $1/2$.*

3.1 Properties of the tangential Delaunay complex

We now give the following two lemmas which are slight variants of results of [14].

LEMMA 8. *For all $x \in T_p \cap \mathrm{Vor}^\omega(p)$, there exists a positive constant c_1 such that $||p-x|| \leq c_1\varepsilon\mathrm{lfs}(p)$.*

PROOF. Assume for a contradiction that there exists a point $x \in \mathrm{Vor}^\omega(p) \cap T_p$ s.t. $||p-x|| > c_1\varepsilon\mathrm{lfs}(p)$ with $c_1(1-c_1\varepsilon) > 2 + c_1\varepsilon(1+c_1\varepsilon)$ (*). Let q be a point on the line segment $[px]$ such that $||p-q|| = c_1\varepsilon\mathrm{lfs}(p)/2$. Let q' be the point nearest to q on $\mathbb{M}$. From Lemma 3, we have $||q-q'|| \leq c_1^2\varepsilon^2\mathrm{lfs}(p)/2$.

Hence, $||p-q'|| \leq ||p-q|| + ||q-q'|| < \frac{c_1}{2}\,\varepsilon(1+c_1\varepsilon)\mathrm{lfs}(p)$. From the 1-Lipschitz property, $\mathrm{lfs}(q') \leq \mathrm{lfs}(p) + ||p-q'|| < (1 + \frac{c_1}{2}\,\varepsilon(1+c_1\varepsilon))\,\mathrm{lfs}(p) < \frac{c_1}{2}(1-c_1\varepsilon)\,\mathrm{lfs}(p)$, which yields

$$||p-q'|| \geq ||p-q|| - ||q-q'|| > \frac{c_1}{2}\varepsilon(1-c_1\varepsilon)\,\mathrm{lfs}(p)$$

Since $\mathcal{P}$ is an ε-sample, there exists a point $t \in \mathcal{P}$, s.t. $||q'-t|| \leq \varepsilon\mathrm{lfs}(q') < \frac{c_1}{2}\varepsilon(1-c_1\varepsilon)\,\mathrm{lfs}(p)$. We thus have

$$||q-t|| \leq ||q-q'|| + ||q'-t|| < \frac{c_1}{2}\,\varepsilon\mathrm{lfs}(p).$$

From Fig. 4, we can see that $\angle ptx > \pi/2$. This implies that $||x-p||^2 - ||x-t||^2 - ||p-t||^2 > 0$. Hence,

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

This implies $x \notin \mathrm{Vor}^\omega(p)$, which contradicts our initial assumption. We conclude that $\mathrm{Vor}^\omega(p) \cap T_p \subseteq B(p, c_1\varepsilon\mathrm{lfs}(p))$ if Inequality (*) is satisfied, which is true for all $c_1 = 3+\sqrt{2} \approx 4.41$ and $\varepsilon < 0.09$. $\square$

LEMMA 9. *There exists positive constants c_2 and ρ_0' s.t.*

1. *If pq is an edge of $\mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$, then $||p-q|| \leq c_2\varepsilon\mathrm{lfs}(p)$.*
2. *If τ is a simplex of $\mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$, then $R'_\tau \leq \rho_0' L_\tau$ and $\rho(\tau) = R_\tau/L_\tau \leq \rho_0'$.*

PROOF. 1a. Consider first the case where pq is an edge of $\mathrm{Del}^\omega_p(\mathcal{P})$. Then $T_p \cap \mathrm{Vor}^\omega(pq) \neq \emptyset$. Let $x \in T_p \cap \mathrm{Vor}^\omega(pq)$. From Lemma 8, we have $||p-x|| \leq c_1\varepsilon\mathrm{lfs}(p)$. By Lemma 2, $||q-x|| \leq \frac{||p-x||}{\sqrt{1-4\omega_0^2}} \leq \frac{c_1\varepsilon\mathrm{lfs}(p)}{\sqrt{1-4\omega_0^2}}$. Hence, $||p-q|| \leq c_2'\varepsilon\mathrm{lfs}(p)$ where $c_2' = c_1(1+1/\sqrt{1-4\omega_0^2}) > 2c_1$.

1b. From the definition of $\mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$, there exists a vertex r of τ such that $[pq] \in \mathrm{star}(r)$. From 1a, $\|r-p\|$ and $\|r-q\|$ are at most $c_2'\varepsilon\mathrm{lfs}(r)$. From the 1-Lipschitz property of lfs and assuming that $2c_2'\varepsilon < 1$, we have $\mathrm{lfs}(r) \leq \mathrm{lfs}(p) + \|p-r\| \leq \frac{\mathrm{lfs}(p)}{1-c_2'\varepsilon} \leq 2\,\mathrm{lfs}(p)$. It follows that $\|p-q\| \leq \|p-r\| + \|r-q\| \leq 4c_2'\varepsilon\mathrm{lfs}(p)$. The first part of the lemma is proved by taking $c_2 = 4c_2' > 8c_1$.

2. Assume that $\tau \in \mathrm{star}(p)$. Let $z \in \mathrm{Vor}^\omega(\tau) \cap T_p$, and $r_z = \sqrt{||z-p||^2 - \omega^2(p)}$. By definition, the ball centered at z with radius r_z is orthogonal to the weighted vertices of τ. From Lemma 2, we have $r_z \geq R'_\tau$. Hence it suffices to prove $r_z \leq \rho_0' L_\tau$. Since $z \in \mathrm{Vor}^\omega(\tau) \cap T_p$, we deduce from Lemma 8 that $||z-p|| \leq c_1\varepsilon\mathrm{lfs}(p)$. Therefore

$$r_z = \sqrt{||z-p||^2 - \omega^2(p)} \leq ||z-p|| \leq c_1\varepsilon\mathrm{lfs}(p).$$

For any vertex q of τ, we have $\|p-q\| \leq c_2\varepsilon\mathrm{lfs}(p)$ (By 1.). Assuming $2c_2\varepsilon \leq 1$ and using the fact that lfs is 1-Lipschitz, $\mathrm{lfs}(p) \leq 2\mathrm{lfs}(q)$. Therefore, taking for q a vertex of the shortest edge of τ, we have, using Lemma 3,

$$\begin{aligned} r_z &\leq c_1\left(\frac{\varepsilon}{\delta}\right)\delta\,\mathrm{lfs}(p) \leq c_1\left(\frac{\varepsilon}{\delta}\right)\delta \ \times\ 2\mathrm{lfs}(q) \\ &\leq 2c_1\eta_0\,L_\tau = \rho_0'' L_\tau. \end{aligned}$$

From Lemma 2, we have $R_\tau \leq \frac{R'_\tau}{\sqrt{1-4\omega_0^2}}$. Therefore $\rho(\tau) \leq \rho_0' = \frac{\rho_0''}{\sqrt{1-4\omega_0^2}}$. $\square$

3.2 Properties of inconsistent configurations

We now give the lemmas on inconsistent configurations which are central to the proof of correctness of the reconstruction algorithm given later in the paper.

LEMMA 10. *Let $\phi \in \mathrm{In}^\omega(\mathcal{P})$ be an inconsistent configuration witnessed by (p_i, p_j, p_l). Then, there exists positive constants $c_3 > c_2$ and $\rho_0 > \rho_0'$ s.t.*

1. *$||p - \mathrm{i}_\phi|| \leq \frac{c_3}{2}\,\varepsilon\mathrm{lfs}(p)$ for all vertices p of ϕ.*
2. *If pq is an edge of ϕ then $\|p-q\| \leq c_3\varepsilon\,\mathrm{lfs}(p)$.*
3. *If τ is a j-dimensional face of ϕ ($j \leq k+1$) and R'_τ is the orthosphere of τ, then $R'_\tau \leq \rho_0 L_\tau$ and $R_\tau/L_\tau \leq \rho_0$.*

PROOF. From the definition, $\tau = \phi \setminus \{p_l\}$ belongs to $\mathrm{Del}^\omega_{p_i}(\mathcal{P})$. We first bound $\mathrm{dist}(\mathrm{i}_\phi, \mathrm{aff}(\tau)) = \|\mathrm{o}_\tau - \mathrm{i}_\phi\|$ where

o_τ is the orthocenter of τ. By Lemma 8, $\|p_i - c_{p_i}\| \leq c_1\varepsilon\mathrm{lfs}(p_i)$ and, by Lemma 2, we have

$$\|p_j - c_{p_i}\| \leq \frac{\|p_i - c_{p_i}\|}{\sqrt{1-4\omega_0^2}} \leq \frac{c_1\varepsilon\,\mathrm{lfs}(p_i)}{\sqrt{1-4\omega_0^2}}$$

Let $c_3 = c_2/\sqrt{1-4\omega_0^2} > c_2$. By Lemma 7, $\angle(\mathrm{aff}(\tau), T_{p_i}) \leq \alpha(\sigma_0)\,\varepsilon$, which implies that $\sin\angle(\mathrm{aff}(\tau), T_{p_i}) \leq \alpha(\sigma_0)\varepsilon$ and

$$\tan^2\angle(\mathrm{aff}(\tau), T_{p_i}) \leq \frac{\alpha^2(\sigma_0)\,\varepsilon^2}{1-\alpha^2(\sigma_0)\,\varepsilon^2} < 4\,\alpha^2(\sigma_0)\,\varepsilon^2,$$

since $2\alpha(\sigma_0)\varepsilon < 1$ (Hypothesis 1). Observing that $||p_i - o_\tau|| \leq ||p_i - c_{p_i}||$, since o_τ belongs to $\mathrm{aff}(\mathrm{Vor}^\omega(\tau))$ and therefore is the closest point to p_i in $\mathrm{aff}(\mathrm{Vor}^\omega(\tau))$, we deduce

$$\|c_{p_i} - o_\tau\| \leq \|c_{p_i} - p_i\| \sin\angle(\mathrm{aff}(\tau), T_{p_i}) \leq \alpha(\sigma_0)c_1\varepsilon^2\mathrm{lfs}(p_i).$$

Moreover, $||p_j - o_\tau|| \leq ||p_j - c_{p_i}||$ as o_τ is the closest point to p_j in $\mathrm{aff}(\mathrm{Vor}^\omega(\tau))$. Hence we have,

$$\begin{aligned}\|c_{p_j} - o_\tau\| &\leq \|p_j - o_\tau\| \tan\angle(\mathrm{aff}(\tau), T_{p_j}) \\ &< \frac{2\,\alpha(\sigma_0)\,c_1\,\varepsilon^2}{\sqrt{1-4\omega_0^2}}\,\mathrm{lfs}(p_i)\end{aligned}$$

As $i_\phi \in [c_{p_i} c_{p_j}]$, we conclude that

$$\|o_\tau - i_\phi\| \leq \frac{2\,\alpha(\sigma_0)\,c_1\,\varepsilon^2}{\sqrt{1-4\omega_0^2}}\,\mathrm{lfs}(p_i).$$

1. Assuming that $2\alpha(\sigma_0)\,\varepsilon < 1$ and using the facts that $\|p_i - o_\tau\| \leq \|p_i - i_\phi\|$, we get

$$\begin{aligned}\|p_i - i_\phi\| &\leq \|p_i - o_\tau\| + \|o_\tau - i_\phi\| \\ &\leq \|p_i - c_{p_i}\| + \|o_\tau - i_\phi\| \\ &\leq \left(c_1\varepsilon + \frac{2\alpha(\sigma_0)c_1\varepsilon^2}{\sqrt{1-4\omega_0^2}}\right)\mathrm{lfs}(p_i) \\ &\leq \frac{c_2}{4}\,\varepsilon\mathrm{lfs}(p_i)\end{aligned}$$

From Lemma 2, we have $\|p - i_\phi\| \leq \frac{\|p_i - i_\phi\|}{\sqrt{1-4\omega_0^2}} = \frac{c_3}{4}\,\varepsilon\mathrm{lfs}(p_i)$ for all vertices p of ϕ.

2. Using part 1 of this lemma, we have

$$\|p - q\| \leq \|p - i_\phi\| + \|q - i_\phi\| \leq \frac{c_3}{2}\,\varepsilon\mathrm{lfs}(p_i)$$

Moreover, $\mathrm{lfs}(p_i) \leq 2\mathrm{lfs}(p)$ since lfs is a 1-Lipschitz function and by taking $c_3\varepsilon < 1$. Hence $\|p - i_\phi\| \leq \frac{c_3}{2}\varepsilon\mathrm{lfs}(p)$ and $\|p - q\| \leq c_3\varepsilon\,\mathrm{lfs}(p)$.

3. Let $r_\phi = \sqrt{\|i_\phi - p_i\|^2 - \omega^2(p_i)}$. Since $i_\phi \in \mathrm{Vor}^\omega(\tau)$, the ball centered at i_ϕ with radius r_ϕ is orthogonal to the weighted vertices of τ. From Lemma 2, we have $r_\phi \geq R'_\tau$. Hence it suffices to show that $r_\phi \leq \rho_0 L_\tau$. Using $\|p_i - i_\phi\| \leq \frac{c_3}{4}\,\varepsilon\mathrm{lfs}(p_i)$ we get

$$r_\phi = \sqrt{\|i_\phi - p_i\| - \omega^2(p_i)} \leq \|i_\phi - p_i\| \leq \frac{c_3}{4}\,\varepsilon\mathrm{lfs}(p_i).$$

Let q be a vertex of a shortest edge of τ. We have, from part 2 of this lemma, $\|p_i - q\| \leq c_3\,\varepsilon\mathrm{lfs}(q)$. Therefore

$$R'_\tau \leq r_\phi \leq c_3\,\varepsilon\mathrm{lfs}(q) \leq c_3\,\left(\frac{\varepsilon}{\delta}\right)\delta \times \mathrm{lfs}(q) \leq c_3\,\eta_0\,L_\tau$$

From Lemma 2, we have $R_\tau \leq R'_\tau/\sqrt{1-4\omega_0^2}$. Therefore , $R_\tau/L_\tau \leq \rho_0 = c_3\eta_0/\sqrt{1-4\omega_0^2}$. $\square$

LEMMA 11. *Let τ be a j-simplex of $\mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^\omega(\mathcal{P})$, $j \leq k+1$. For any vertex p of τ, there exists a constant c_4 s.t. the distance between the orthocenter $o(\tau)$ of τ and $\mathrm{aff}(\tau_p)$ is at most $c_4\varepsilon\mathrm{lfs}(p)$.*

PROOF. 1. We first consider the case where $\tau \in \mathrm{In}^\omega(\mathcal{P})$. Then there exists an inconsistent configuration ϕ witnessed by points (p_i, p_j, p_l) s.t. τ is a j-dimensional subsimplex of ϕ. From Lemma 10, we have $\|i_\phi - p\| \leq \frac{c_3}{2}\,\varepsilon\mathrm{lfs}(p)$, $\|q - p\| \leq c_3\mathrm{lfs}(p)$ for all $q \in \phi$ and, since $\tau \in \mathrm{In}^\omega(\mathcal{P})$, $\rho(\tau) \leq \rho_0$. Using the above facts and Lemma 4, we get

$$\begin{aligned}\mathrm{dist}(\mathrm{o}_\tau, \mathrm{aff}(\tau_\mathrm{p})) &\leq \mathrm{dist}(\mathrm{i}_\phi, \mathrm{aff}(\tau_\mathrm{p})) \\ &\leq \frac{c_3}{2}\,\varepsilon\mathrm{lfs}(p) + (1+\rho_0+\omega_0)c_3\mathrm{lfs}(p) \\ &\leq (\frac{3}{2} + \rho_0 + \omega_0)c_3\varepsilon\mathrm{lfs}(p) = c_4\varepsilon\mathrm{lfs}(p),\end{aligned}$$

2. Consider now the case where $\tau \in \mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$. By definition, $\tau \in \mathrm{Del}^\omega_q(\mathcal{P})$ for some vertex q of τ. From Lemma 8, we have $\|q - c_q\| \leq c_1\varepsilon\mathrm{lfs}(q) \leq 2c_1\varepsilon\mathrm{lfs}(p)$ where $c_q = \mathrm{Vor}^\omega(\tau) \cap T_q$ and p is any vertex of τ. The last inequality follows from the facts that lfs is 1-Lipschitz, $\|p - q\| \leq c_2\varepsilon\mathrm{lfs}(p)$ (Lemma 9) and by taking $c_2\varepsilon \leq 1$. Therefore, using Lemma 4, we get

$$\mathrm{dist}(\mathrm{o}_\tau, \mathrm{aff}(\tau_\mathrm{p})) \leq \frac{2c_1\varepsilon\mathrm{lfs}(p)}{\sqrt{1-4\omega_0^2}} + (1+\rho'_0+\omega_0)c_2\,\varepsilon\mathrm{lfs}(p).$$

$\square$

The next lemma shows that inconsistent configurations are slivers provided that σ_0 is sufficiently large wrt ε.

LEMMA 12. *Let ϕ be an inconsistent configuration and*

$$\varepsilon < \mathrm{f}(\sigma_0) = \frac{(k+1)\,\sigma_0}{\rho_0\,(c_3\,+\,2\,\alpha(\sigma_0))}.$$

If the subfaces of ϕ are not slivers, then ϕ is a sliver.

PROOF. Let pq be the smallest edge of ϕ and let r be a vertex in $\phi \setminus \{p, q\}$. Since $\phi_r = \phi \setminus \{r\}$ is not a sliver (as all subfaces of ϕ are not slivers) and $\|r - x\| \leq c_3\mathrm{lfs}(x)$ (Lemma 10), we have $\sin\angle(\mathrm{aff}(\phi_r), T_x) \leq \alpha(\sigma_0)\,\varepsilon$ for all vertices x of ϕ_r by Lemma 7. From Lemma 3, $\sin\angle(pr, T_p) \leq \frac{c_3}{2}\,\varepsilon$. Therefore

$$\begin{aligned}D_\phi(r) &= \sin\angle(pr, \mathrm{aff}(\phi_r)) \times \|p - r\| \\ &\leq (\sin\angle(pr, T_p) + \sin\angle(\mathrm{aff}(\phi_r), T_p)) \times \|p - r\| \\ &\leq (\frac{c_3}{2} + \alpha(\sigma_0))\,\varepsilon\|p - r\|.\end{aligned}$$

Using the facts that $\mathrm{vol}(\phi_r) = \sigma^k(\phi_r)\,L^k_{\phi_r}$, $\sigma(\phi_r) \geq \sigma_0$, $\|p - r\| \leq 2R_\phi$, $\rho(\phi_r) \leq \rho_0$ and $\varepsilon < \mathrm{f}(\sigma_0)$, we get

$$\begin{aligned}\mathrm{vol}(\phi) &= D_\phi(r) \times \frac{\mathrm{vol}(\phi_r)}{k+1} \\ &\leq \left(\frac{c_3}{2} + \alpha(\sigma_0)\right)\,\varepsilon\|p - r\| \times \frac{\sigma^k(\phi_r)L^k_\phi}{k+1} \\ &\leq (c_3 + 2\alpha(\sigma_0))\,\varepsilon R_\phi \times \frac{\sigma_0^k L^k_\phi}{k+1} \\ &\leq \frac{\rho_0(c_3 + 2\alpha(\sigma_0))\varepsilon}{k+1} \times \sigma_0^k\,L^{k+1}_\phi \\ &< \sigma_0^{k+1}\,L^{k+1}_\phi\end{aligned}$$

$\square$

3.3 Number of local neighbors

Lemmas 9 and 10 show that, in order to construct star(p) and search for inconsistencies involving p, it is enough to consider the points of $\mathcal{P}$ that lie in ball $B_p = B(p, c_3\varepsilon\text{lfs}(p))$. Since ε and lfs(p) are not known in practice, we will consider instead the ball $B'_p = B(p, c_3\eta_0\, l_p)$ where

$$l_x = \min_{q\in\mathcal{P}, q\neq x} \|x - q\|.$$

It is easily seen that $l_x : \mathbb{M} \to \mathbb{R}$ is 1-Lipschitz and, by Lemma 3, we have $\delta\text{lfs}(p) \le l_p \le \frac{2\varepsilon}{1-\varepsilon}\,\text{lfs}(p)$. It follows that B'_p contains B_p if $\varepsilon/\delta \le \eta_0$. We call $LN_p = B'_p \cap \mathcal{P}$ the *local neighborhood* of p.

LEMMA 13. *The number of points of LN_p is less than $N < 2^{O(k)}$.*

PROOF. For convenience, write $\nu = 4c_3\eta_0\varepsilon$ and assume that $\nu \le 1/2$. We observe that $LN_p \subset B''_p = B(p, \nu\text{lfs}(p))$ since, by Lemma 3, $l_p \le \frac{2\varepsilon}{1-\varepsilon}\,\text{lfs}(p) \le 4\varepsilon\text{lfs}(p)$. We will count the number of points in $B''_p \cap \mathcal{P}$. Let x and y be two points of $B''_p \cap \mathcal{P}$. We have $\text{lfs}(x), \text{lfs}(y) \ge \text{lfs}(p)(1-\nu) \ge \text{lfs}(p)/2$, since lfs is 1-Lipschitz. The balls $B(x, l_x/2)$ and $B(y, l_y/2)$ are disjoint, and, since $l_x \ge \delta\,\text{lfs}(x) \ge \frac{\delta}{2}\,\text{lfs}(p)$ (and similarly for l_y), the balls $B_x = B(x, \frac{\delta}{4}\,\text{lfs}(p))$ and $B_y = B(y, \frac{\delta}{4}\,\text{lfs}(p))$ are also disjoint. Observe that both balls B_x and B_y are contained in $B^+_p = B(p, \mu\,\varepsilon\text{lfs}(p))$ where $\mu = \frac{\nu}{\varepsilon} + \frac{\delta}{4\varepsilon} \le 4c_3\eta_0 + \frac{1}{4}$.

A packing argument now allows to conclude. Specifically, by Lemma 15, we have that $\text{vol}(B_x \cap \mathbb{M}) > \phi_k \left(\frac{\delta\,\text{lfs}(p)}{8}\right)^k$ and $\text{vol}(B^+_p \cap \mathbb{M}) < \phi_k\ (\mu\varepsilon\text{lfs}(p))^k$, where ϕ_k is the volume of the k-dimensional unit ball. We conclude that the number of points of $\mathcal{P} \cap B''_p$ is less than $\left(\frac{8\mu\varepsilon}{\delta}\right)^k \le (32c_3\eta_0 + 2)^k\,\eta_0^k = 2^{O(k)}$. $\square$

4. MANIFOLD RECONSTRUCTION

In this section, we will show how to find a weight assignment for the point set $\mathcal{P}$ so that we can remove all inconsistent configurations. Once this is done, all the stars become coherent and the resulting weighted tangential Delaunay complex is a simplicial k-manifold.

4.1 Algorithm

The input to the algorithm is a point sample $\mathcal{P} = \{p_1, \ldots, p_n\}$ together with a bound on the relative density η_0 of $\mathcal{P}$. We assume in addition that the dimension k of $\mathbb{M}$ is given and that we know the tangent space T_p at each point $p \in \mathcal{P}$.

The algorithm fixes ω_0 in $(0, 1/2)$ and then computes a weight assignment $\omega \in [0, \omega_0]$ such that no inconsistent configuration remains in the weighted tangential Delaunay complex. More precisely, we will assign a weight to each point $p \in \mathcal{P}$ in turn so as to remove all j-dimensional slivers incident to p in $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \text{In}^\omega(\mathcal{P})$ for $j \in \{3, \ldots, k+1\}$. By Lemma 12, we know that, if σ_0 is large enough, removing inconsistent configurations from the tangential Delaunay complex reduces to removing slivers.

For a given j-simplex $\tau = [p_0, \ldots, p_j] \in \text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \text{In}^\omega(\mathcal{P})$ we have

$$\begin{aligned}\sigma^j(\tau) &= \frac{|\det(p_1 - p_0, \ldots, p_j - p_0)|}{j!\, L^j_\tau} \\ &\le \frac{\Pi^j_{i=1}\|p_i - p_0\|}{j!\, L^j_\tau} \le \frac{2^j\,\rho^j_0}{j!} \le 2\rho^j_0,\end{aligned}$$

the last inequalities follow from the facts that $L_\tau \le 2\,R_\tau$, and the radius-edge ratio of the simplices in $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \text{In}^\omega(\mathcal{P})$ is $\le \rho_0$ (Lemmas 9 and 10). This implies that $\sigma(\tau) \in (0,\, 2\,\rho_0]$. In the first step of our reconstruction algorithm, we pick a random value of σ_0 from $(0, 2\rho_0]$, and once σ_0 is selected, we try to remove all slivers from $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \text{In}^\omega(\mathcal{P})$. If we fail to remove slivers or if we still have inconsistencies, then we go back and select a new value of σ_0.

Manifold-Reconstruction($\mathcal{P} = \{p_1, \ldots, p_n\}$)

S0 Calculate $LN_{p_i}\ \forall p_i \in \mathcal{P}$ and ρ_1

S1 Select σ_0 at random from $(0, 2\,\rho_1]$

S2 **for** $i = 1$ **to** n **do**
  if weight(p_i, σ_0) = "FAIL" **then go-to** S1
  else update($\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}), p_i$)

S3 **if** $\text{In}^\omega(\mathcal{P}) \neq \emptyset$ **then go-to** S1
  else output $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$

Before we give the details of the function **weight**() we first define the subroutine **skyline**(p, S, σ_0) that will be used in the function. Let S be a set of simplices incident on p.

Let τ be a simplex in S whose subfaces are not σ_0-slivers. Such a simplex is called a *candidate simplex*. We associate to τ interval $W(\tau)$ that consists of all squared weights $\omega^2(p)$ for which τ appears as a simplex in $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \cup \text{In}^\omega(\mathcal{P})$. We define the skyline of p as the lower boundary of the union of all rectangles $R(\tau) = W(\tau) \times [\sigma(\tau), +\infty)$ for all candidate simplices over $W(p) = [0, \omega_0^2\, l_p^2]$, see Figure 5. For any $\omega^2(p)$, the minimum sliverity ratio of any *candidate simplex* incident to p is the height of **skyline**(p, S_p, σ_0) over $\omega^2(p) \in W(p)$. The best choice for $\omega(p)$ is the weight of p corresponding to the highest point on the **skyline**(p, S_p, σ_0), i.e **skyline**(p, S_p, σ_0) has the maximum height over $\omega^2(p)$.

Function **weight**(p, σ_0)

S1 $S_p \leftarrow$ **detect_simplices**(p, σ_0)

S2 **skyline**(p, S_p, σ_0) **begin**
  for $i = 1$ **to** n **do**
    if none of the subfaces of τ are σ_0-slivers **then**
      include $R(\tau) = W(\tau) \times [\sigma(\tau), +\infty)$
  end

S3 $\Theta \leftarrow$ highest point on **skyline**(p, S_p, σ_0)

S4 **if** $\Theta < \sigma_0$ **then return** "FAIL"
  else $\omega(p) \leftarrow$ weight of p corresponding to a highest point on **skyline**(p, S_p, σ_0)

Function **update**($\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P}), p$)

– Update the stars of p and of all points $x \in LN_p$ by modifying $\text{Del}^\omega_x(LN_x)$.

Function **detect_simplices**(p, σ_0)

1. We first detect all possible j-simplices for all $3 \le j \le k+1$ of $\text{Del}^\omega_{T\mathbb{M}}(\mathcal{P})$ incident on p for all possible $\omega(p)$. This is done in the following way: (1) we vary the weight of p from 0 to to $\omega_0 l_p$, keeping the weights of the other points constant; (2) for each new weight assignment to p, we modify the stars of the points in LN_p and detect from the stars the new j-simplices incident

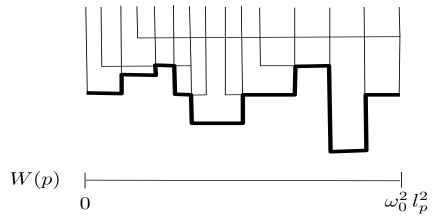

Figure 5: The above figure shows a skyline(p, S_p, σ_0) over $W(p) = [0, \omega^2\, l_p^2]$ for point p.

to p that have not been detected thus far. The weight of point p changes only in a finite number of instances $0 = P_0 < P_1 < \cdots < P_{n-1} < P_n = \omega_0 l_p$.

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.

2. Once we have detected all possible j-simplices, for all $j \in \{1, \ldots, k+1\}$, that can be incident on p in the weighted tangential Delaunay complex, we then detect all possible inconsistent configurations incident on p, by calling the function **detect_inconsistent-configuration**(p, σ_0).

Function **detect_inconsistent-configurations**(p, σ_0)

1. We vary the weight of p from 0 to $\omega_0 l_p$, keeping the weight of the rest of the points constant. Once we have assigned a new weight to p we modify the stars of the points in LN_p.

2. Detecting the inconsistent configurations incident to p is more complicated than detecting the simplices incident to p. We consider all points p_i in LN_p. Let τ be a k-simplex in the star of p_i, and let p_j be a vertex of τ such that τ is not in the star of p_j. We calculate the Voronoi diagram of the points in LN_p restricted to the line segment $[c_{p_i} c_{p_j}]$, where $c_{p_i} = T_{p_i} \cap \mathrm{Vor}^{\omega}(\tau)$ and $c_{p_j} = T_{p_j} \cap \mathrm{aff}(\mathrm{Vor}^{\omega}(\tau))$. From the restricted Voronoi diagram, we find a point p whose Voronoi cell intersects for the first time the line segment $[c_{p_i} c_{p_j}]$ oriented from c_{p_i} to c_{p_j}. If $p \in \phi = \tau \cup \{l\}$, then we report ϕ.

3. As in the **detect_simplices** function the weight of p is changed only a finite number of times. For each current inconsistent configuration ϕ incident to p, we keep in a priority queue the weight of p for which ϕ will disappear for the first time. The minimum weight in the priority queue gives the next weight assignment of p.

4.2 Analysis of the algorithm

DEFINITION 4 (SLIVERITY RANGE). *Let ω be a weight assignment of relative amplitude at most ω_0 we keep fixed except for $\omega(p)$. The sliverity range of a point $p \in \mathcal{P}$ is the measure of the set of all squared weights $\omega^2(p)$ for which p is incident to a sliver in $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$.*

LEMMA 14. *Under Hypothesis 1, the sliverity range of p is at most $\Sigma(p) = N^{k+2}(k+1)c_5\, \sigma_0\, l_p^2$ for some constant c_5.*

PROOF. Let τ be a simplex incident on p. We call the sliverity range of τ the measure of the set of squared weights for which τ is a sliver in $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$. If $\omega(p)$ is the weight of p, we write $H(\omega(p))$ for the signed distance of the orthocenter of τ to $\mathrm{aff}(\tau_p)$. From Lemma 11, we have $|H(\omega(p))| \le c_4 \varepsilon \mathrm{lfs}(p)$, for all $\tau \in \mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$. Moreover, using Lemma 5 and the fact that τ is a sliver, $H(\omega(p)) = H(0) - \frac{\omega^2(p)}{2D_p} \le H(0) - \frac{\omega^2(p)}{2j\sigma_0 L_\tau}$. It follows that the sliverity range of τ is at most $4j\sigma_0 L_\tau c_4 \varepsilon \mathrm{lfs}(p)$. Using the facts that $L_\tau \le c_3 \varepsilon \mathrm{lfs}(p)$ (from Lemmas 9 and 10), $\mathrm{lfs}(p) \le l_p/\delta$ and $\varepsilon/\delta \le \eta_0$, the sliverity range of τ is less than $4jc_3c_4\sigma_0\eta_0^2\, l_p^2 = j\, c_5 \sigma_0\, l_p^2$. By Lemma 13, the number of j-simplices that are incident to p is at most N^j. Hence, the sliverity range of p is less that

$$\Sigma(p) = \sum_{3}^{k+1} N^j\, jc_5\, \sigma_0\, l_p^2 \le N^{k+2}(k+1)c_5\, \sigma_0\, l_p^2.$$

□

THEOREM 1. *Under Hypothesis 1 and if*

$$\sigma_{\max} = \frac{\omega_0^2}{(k+1)\, c_5\, N^{k+2}} > \sigma_{\min} = \inf\{\sigma | \varepsilon < \mathrm{f}(\sigma)\},$$

then, for any $\sigma_0 \in [\sigma_{\min}, \sigma_{\max}]$, the above algorithm outputs $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ without any slivers nor inconsistent configuration.

PROOF. As in subroutine **skyline**(), let S_p denotes the set of all possible simplices that can be incident on p in the complex $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$ for all possible weight assignments ω of relative amplitude $\tilde{\omega} \le \omega_0$. By Lemma 14, the sliverity range of p is less than $N^{k+2}(k+1)c_5\, \sigma_0\, l_p^2$. If the sliverity range of p is less than $\omega_0^2 l_p^2$, the total range of all possible squared weights, or, equivalently, if

$$\sigma_0 < \sigma_{\max} = \frac{\omega_0^2}{(k+1)\, c_5\, N^{k+2}},$$

then we will be able to remove all slivers incident on p by selecting the highest point on the skyline.

If we select a value of σ_0 in the interval $(\sigma_{\min}, \sigma_{\max}]$, Lemma 14 ensures that removing all j-dimensional slivers for all $j \in \{3, \ldots, k+1\}$ in $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$ will also result in removing inconsistent configurations from $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$, i.e. $\mathrm{In}^{\omega}(\mathcal{P}) = \emptyset$. □

4.3 Time and space complexity

THEOREM 2. *The time complexity of the algorithm is*

$$O(d)\, |\mathcal{P}|^2 + \frac{1}{\lambda} \times 2^{O(k^2) + \log d}\, |\mathcal{P}|,$$

where $\lambda = (\sigma_{\max} - \sigma_{\min})/(2\rho_0)$, and its space complexity is

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

PROOF. We only sketch the complexity analysis. See [7] for a detailed discussion. Step S0 of **Manifold-Reconstruction** can easily be performed in $O(d)\, |\mathcal{P}|^2$ time.

We show now that the expected number of times Step **S1** of **ManifoldReconstruction** is repeated is $\frac{1}{\lambda}$. Indeed, for any simplex $\tau \in \mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \cup \mathrm{In}^{\omega}(\mathcal{P})$, we have $\sigma(\tau) \in (0, 2\rho_0]$. Hence, the probability that, for the selected value of σ_0, the algorithm removes all slivers and inconsistencies is at least $\lambda = (\sigma_{\max} - \sigma_{\min})/(2\rho_0)$. It follows that the expected number of times **S1** is performed is less than

$$\sum_{i=1}^{\infty} i(1-\lambda)^{i-1}\lambda = \frac{1}{\lambda}\,.$$

The time complexities of functions **update**$(\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}), p)$, **detect_simplices** (p, σ_0), and **detect_inconsistent-configurations**(p, σ_0) are $2^{O(k(k+\log d))}$. Indeed, we need to project $|LN_p| < N < 2^{O(k)}$ points onto T_p, which costs $O(kd) \times |LN_p| = O(d)\,2^{O(k)}$ time. The Delaunay simplices that are computed have their vertices in LN_p and are of dimension at most $k+1$. Hence their number is $2^{O(k^2)}$. The basic operations (mainly in-sphere predicates) amount to evaluating signs of determinants of $k \times k$ matrices. The cost of such a basic operation is $O(k^3)$. The total cost of both functions is thus $O(d)\,2^{O(k)} + O(k^3)\,2^{O(k^2)} = 2^{O(k^2)+\log d}$.

We deduce that the expected time complexity of **Manifold-Reconstruction** is

$$O(d)\,|\mathcal{P}|^2 + \frac{1}{\lambda} \times 2^{O(k^2)+\log d}\,|\mathcal{P}|.$$

We easily deduce from the above discussion that the total space complexity of the algorithm is

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

□

4.4 Topological and Geometric guarantees

We assume that the conditions of Theorem 1 are satisfied and denote by $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ the tangential complex output by our algorithm. Let π be the mapping that maps any point of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ to its closest point on $\mathbb{M}$. The proof of topological correctness (see [7]) uses ideas from [1, 8, 11, 14].

THEOREM 3. *Under the conditions of Theorem 1, our algorithm outputs* $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ *without any slivers and inconsistent configurations, and* $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ *has the following properties*

- **Bijection :** *The restriction of* π *to* $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ *is a bijection;*
- **Pointwise approximation :** $\forall x \in \mathbb{M}$, $\mathrm{dist}(\mathrm{x}, \pi^{-1}(\mathrm{x})) = O(\varepsilon^2 \mathrm{lfs}(x))$;
- **Normal approximation :** $\forall x \in \mathbb{M}$, $\angle N_x N_\tau = O(\varepsilon)$, *where* τ *is a* k*-simplex of* $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ *containing the point* $\pi^{-1}(x)$,
- **Topological correctness :** π *defines an isotopy between* $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ *and* $\mathbb{M}$.

PROOF SKETCH. By Lemma 7, one can show that the maximum distance from a point of a simplex $\tau \in \mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ to the closest point on $\mathbb{M}$ is $O(\varepsilon^2 \mathrm{lfs}(x))$. It follows that the projection π that maps every point of τ to its closest point on $\mathbb{M}$ is injective: if we extend an open segment of length $\mathrm{lfs}(y)$ from every manifold point y in all normal directions to $\mathbb{M}$, these segments do not intersect, and they can be used as the fibers of a tubular neighborhood $\hat{\mathbb{M}}$ of $\mathbb{M}$. Each point of such a segment has y as its unique closest neighbor on $\mathbb{M}$. For small enough ε, the simplex τ is contained in $\hat{\mathbb{M}}$. Thus, the mapping π defines an isotopy between τ and a corresponding manifold patch.

One can show that two k-simplices of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ that share a subface have normal spaces that differ by at most $O(\varepsilon)$, and that the mapping π extends continuously across the subfaces. It follows that the projection π restricted to the two adjacent k-simplices is a homeomorphism that is invertible locally. (In topological terms, $\pi\colon \mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}) \to \mathbb{M}$ is a covering map, if we can establish that it is surjective.) By assumption, on every component, there is at least one vertex of a simplex of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$. This ensures that $\pi(\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P}))$ contains that vertex, and since the mapping can be continued locally, it follows that every component of S is covered at least once. It is now still possible that some component is covered more than once by π. This would imply that some sample point $p \in P$ is covered more than once. However, one can show quite easily that no point p of $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ except p itself has p as its closest neighbor on $\mathbb{M}$.

It follows that the mapping π defines an ambient isotopy between $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ and $\mathbb{M}$. □

5. CONCLUSION

We have given the first algorithm that is able to reconstruct a smooth 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 very high dimension. This situation is quite common in practical applications in machine learning.

The algorithm is rather simple. The basic ingredients we need are data structures for constructing weighted Delaunay triangulations in k-flats. We will report experimental results in a forthcoming paper.

We have assumed that dimension of $\mathbb{M}$ is known. If not, we can use algorithms given in [22, 15] to estimate the dimension of $\mathbb{M}$ and the tangent space at each sample point. One interesting feature of our approach is that it is pretty robust and still works if we only have approximate tangent spaces at the sample points.

We have also assumed that we know an upper bound on the relative density η_0 of the input sample $\mathcal{P}$. Ideas from [21, 8] may be useful to convert a sample to a subsample with a bounded relative density.

We forsee 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].

Finally, let us mention that removing inconsistencies among stars that have been computed independently is a useful paradigm that has already been used for maintaining dynamic meshes [29] and generating anisotropic meshes [9]. We hope that this paper will motivate further applications.

6. ACKNOWLEDGEMENTS

The authors thank the anonymous referees for their insightful comments, and Mariette Yvinec for helpful discussions.

7. REFERENCES

[1] N. Amenta, S. Choi, T. K. Dey, , and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Intl. Journal of Computational Geometry and Application*, 12:125–141, 2002.

[2] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1), 1987.

[3] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984.

[4] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in Neural Information Processing Systems*, 2001.

[5] M. Belkin, J. Sun, and Y. Wang. Discrete laplace operator on meshed surfaces. In *Proc. ACM Symp. on Computational Geometry*, 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 and A. Ghosh. Manifold Reconstruction using Tangential Delaunay Complexes. *INRIA Research Report 7142*, 2009.

[8] J. D. Boissonnat, L. Guibas, and S. Y. Oudot. Manifold reconstruction in arbitary dimensions using witness complexes. In *Proc. ACM Symp. on Computational Geometry*, pages 193–204, 2007.

[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] 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.

[11] 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.

[12] F. Chazal and S. Y. Oudot. Towards Persistence-Based Reconstruction in Euclidean Spaces. In *Proc. ACM Symp. on Computational Geometry*, pages 232–241, 2008.

[13] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver Exudation. *Journal of ACM*, 47:883–904, 2000.

[14] 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.

[15] S.-W. Cheng, Y. Wang, and Z. Wu. Provable Dimension Detection using Principle Component Analysis. *Intl. Journal of Computational Geometry and Application*, 18(5):415–440, 2008.

[16] D. Cohen-Steiner and T. K. F. Da. A greedy Delaunay Based Surface Reconstruction Algorithm. *The Visual Computer*, 20:4–16, 2004.

[17] T. K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2006.

[18] D. L. Donohu and C. Grimes. Hessian eigenmaps: new locally linear embedding techniques for high dimensional data. *Proceedings of the Natural Academy of Sciences*, 100:5591–5596, 2003.

[19] 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.

[20] D. Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002.

[21] S. Funke and E. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pages 781–780, 2002.

[22] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. In *Proc. ACM Symp. on Computational Geometry*, pages 329–337, 2003.

[23] M. Gopi, S. Khrisnan, and C. T. Silva. Surface Reconstruction based on Lower Dimensional Localized Delaunay Triangulation. *Eurographics*, 19 (3), 2000.

[24] 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.

[25] 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.

[26] P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete and Computational Geometry*, 39(1-3), March 2008.

[27] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.

[28] H. S. Seung and D. D. lee. The manifold ways of perception. *Science*, 290:2268–2269, 2000.

[29] J. Shewchuk. Star splaying: An algorithm for repairing delaunay triangulations and convex hulls. In *Proc. ACM Symp. on Computational Geometry*, pages 237–246, 2005.

[30] J. B. Tenenbaum, V. deSilva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.

[31] 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.

APPENDIX

LEMMA 15 ([26]). *Let $A = B(p,r) \cap \mathbb{M}$ where $r \leq \varepsilon \mathrm{lfs}(\mathbb{M})$. Then, $\phi_k r^k \geq \mathrm{vol}(A) \phi_k r^k / 2^k$, where ϕ_k is volume of the k-dimensional unit ball.*