Delaunay-like Triangulation of Smooth Orientable Submanifolds by ℓ_1-Norm Minimization

Dominique Attali[1*†] and André Lieutier[2†]

[1]Univ. Grenoble Alpes, CNRS, GIPSA-lab, Grenoble, France.
[2]Aix-en-Provence, France.

*Corresponding author(s). E-mail(s): Dominique.Attali@grenoble-inp.fr;
Contributing authors: andre.lieutier@gmail.com;
†These authors contributed equally to this work.

Abstract

In this paper, we study the shape reconstruction problem, when the shape we wish to reconstruct is an orientable smooth d-dimensional submanifold of the Euclidean space. Assuming we have as input a simplicial complex K that approximates the submanifold (such as the Čech complex or the Rips complex), we recast the problem of reconstucting the submanifold from K as an ℓ_1-norm minimization problem in which the optimization variable is a d-chain of K over the field $\mathbb{R}$. Providing that K satisfies certain reasonable conditions, we prove that the considered minimization problem has a unique solution which triangulates the submanifold and coincides with the flat Delaunay complex introduced and studied in a companion paper [1]. Since the objective is a weighted ℓ_1-norm and the constraints are linear, the triangulation process can thus be implemented by linear programming.

Keywords: manifold reconstruction, Delaunay complex, triangulation, simplicial complex, ℓ_1-norm minimization, chain

1 Introduction

In many practical situations, the shape of interest is only known through a finite set of sample points. Given as input these points, a natural question is how to construct a *triangulation* of the shape, that is, a set of simplices whose union is homeomorphic to the shape. This problem, known as *shape reconstruction*, has been widely studied [2–9]. In the paper, the shape is assumed to be a smooth orientable d-dimensional submanifold of the Euclidean space. We show that, under appropriate conditions, a triangulation of that submanifold can be expressed as the solution of a weighted ℓ_1-norm minimization problem under linear constraints.

Overview of the method in the particular case of planar curves.

We first give an informal description of our variational formulation in the easy case of the triangulation of a closed, connected smooth curve in the plane. Assume that we are given a set of points P that sample an unknown curve $\mathcal{C}$ (example in Figure 1a). Consider the graph K whose vertices are the points of P and whose edges connect pair of points that are within a certain given distance; see Figure 1b. Our goal is to compute a triangulation of the curve $\mathcal{C}$ (*i.e.* a closed polygonal line) whose vertices and edges are in K and which follows "nicely" $\mathcal{C}$.

Let us orient (arbitrarily) edges in K; see Figure 1b. A 1-chain γ on K with real coefficients is the assignation of a real number $\gamma(e)$ to each oriented edge e in K. A 1-cycle is a 1-chain γ such that $\partial\gamma = 0$, which means that, at each vertex v of K, the sum of the coefficients of the edges entering v equals the sum of coefficients of edges leaving v; see Figure 1d.

We now define informally what we mean by a normalized 1-cycle. Consider a tubular neighborhood of $\mathcal{C}$ sufficiently small, so that its complement consists only of two connected components, one bounded component called the inside region and one unbounded component called the outside region. Suppose furthermore that K is contained in this small tubular neighborhood; see Figure 1c. The segment $[a, b]$ shown on Figure 1c has one endpoint a in the outside region, one endpoint b in the inside region and it cuts both the curve $\mathcal{C}$ and the graph K transversally. Let us say that an oriented edge $[p, q]$ of K intersects $[a, b]$ in a *forward direction* (*resp. backward direction*) if a lies on the left (*resp. on the right*) of the directed line through p and q. Given a 1-chain γ, we then define the *flux* through $[a, b]$ of γ as

$$\mathrm{flux}_{[a,b]}(\gamma) = \sum_{e^+} \gamma(e^+) - \sum_{e^-} \gamma(e^-),$$

where the first sum is over all edges e^+ of K that cross $[a, b]$ in a forward direction and the second sum is over all edges e^- of K that cross $[a, b]$ in a backward direction. For example, the flux of γ through $[a, b]$ is 2 on Figure 1d. Note that the flux through $[a, b]$ is a linear form on the vector space of 1-cycles. Moreover, the flux of a 1-cycle γ through $[a, b]$ does not depend upon the location of $[a, b]$, as long as a remains in the outside region and b remains in the inside region. Indeed, the expression of the flux changes only when the edge $[a, b]$ passes through a vertex of K, at which time one can check that the value of the flux remains constant if γ is a 1-cycle. We say that a

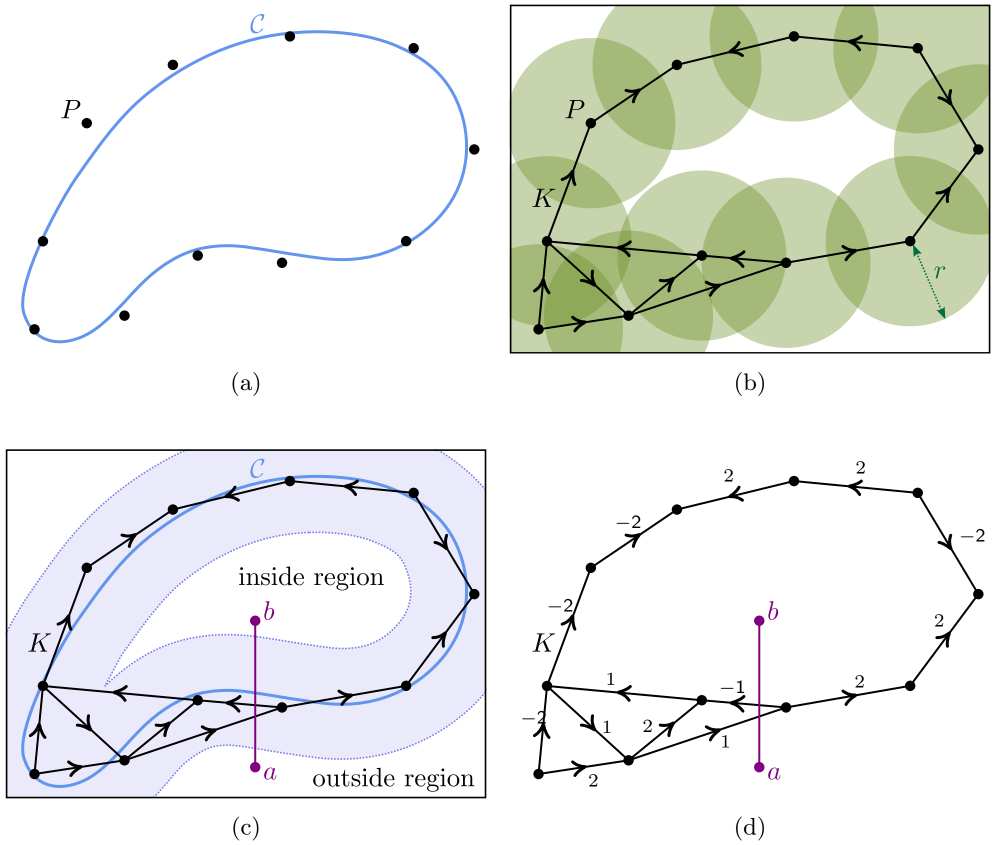

Fig. 1: (a) A finite set of (black) points P that sample a (blue) curve $\mathcal{C}$. (b) The graph K is a proximity graph constructed from P by connecting every pair of points that are less than $2r$ apart. Each edge in K is arbitrarily oriented. (c) The segment $[a, b]$ intersects both $\mathcal{C}$ and K transversally with b in the inside region and a in the outside region. (d) A 1-cycle of K whose flux through $[a, b]$ equals 2.

1-cycle is *normalized* if its flux through $[a, b]$ is equal to 1. Figure 2a depicts such a normalized cycle.

Our reconstruction method consists in computing the minimal cycle among all normalized cycles. By minimal cycle, we mean here a cycle that minimizes a weighted ℓ^1-norm. Perhaps the most natural weighted ℓ^1-norm for geometric 1-cycles could be the length:

$$\text{length}(\gamma) = \sum_e \text{length}(e)\, |\gamma(e)|,$$

where the sum is over all edges e of K. The normalized cycle minimizing the length is depicted in Figure 2b, where one can observe that it is indeed a triangulation of the curve $\mathcal{C}$. However, in order to minimize the length, the resulting minimal cycle

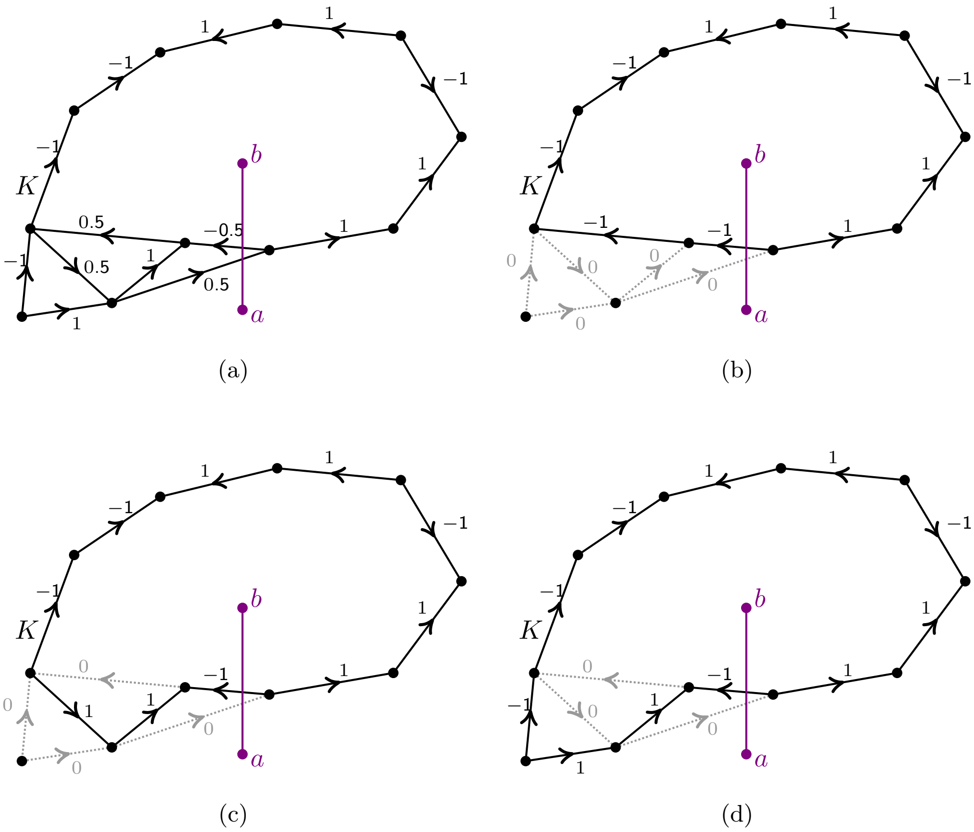

Fig. 2: (a) A 1-cycle is said to be normalized if its flux through $[a, b]$ is equal to 1. (b) A normalized cycle that minimizes the length. (c) A normalized cycle that minimizes the sum of the edge lengths squared. (d) A normalized cycle that minimizes the sum of the edge lengths cubed, also called the Delaunay energy (up to a multiplicative constant). Our reconstruction method returns the support of that cycle.

favors long edges and skips intermediate sample points whenever possible, so that small features of the curve are ignored.

In contrast, thanks to Pythagorean Theorem, weighting edges by the square of the length would make the minimum cycle follow intermediate points, as long as there exist at those intermediate points an incoming edge and an outcoming edge forming an angle larger than $\frac{\pi}{2}$; see Figure 2c.

Instead, we consider minimizing the Delaunay energy, which, in the particular case of one dimensional simplices (*i.e.* edges), consists in weighting edges (up to a constant factor) by the cube of the length:

$$\mathrm{E}_{\mathrm{del}}(\gamma) = \frac{1}{6} \sum_{e} \mathrm{length}(e)^3 \, |\gamma(e)|$$

The cycle minimizing this Delaunay energy is depicted on Figure 2d. One can again check (and we shall actually prove in the paper) that this is indeed a triangulation of $\mathcal{C}$. Our method for reconstructing one dimensional curves can then be expressed as solving the following optimization problem over 1-chains of K:

$$\begin{aligned} \underset{\gamma}{\text{minimize}} \quad & \mathrm{E}_{\mathrm{del}}(\gamma) \\ \text{subject to} \quad & \partial\gamma = 0, \\ & \mathrm{flux}_{[a,b]}(\gamma) = 1 \end{aligned}$$

We aim at generalizing the above problem beyond dimension one manifolds and at identifying sufficient conditions for the support of its solution to be a triangulation. But for that we need first to make a detour to Delaunay complexes.

Variational formulation for Delaunay complexes.

The starting point of the work presented in this paper is the observation that when we consider a point cloud P in $\mathbb{R}^d$, its Delaunay complex can be expressed as the solution of a particular ℓ_p-norm minimization problem. This fact is best explained by lifting the point set P vertically onto the paraboloid $\mathscr{P} \subseteq \mathbb{R}^{d+1}$ whose equation is $x_{d+1} = \sum_{i=1}^{d} x_i^2$. Denoting by $\hat{P}$ the lifted points, it is well-known that the Delaunay complex of P is isomorphic to the boundary complex of the lower convex hull of $\hat{P}$.

Starting from this equivalence, Chen has observed in [10] that the Delaunay complex of P minimizes over all triangulations T of P the ℓ_p-norm of the difference between the following two functions: the first function maps each point x to its vertical projection onto the lifted triangulation $\hat{T}$ and the second function maps each point x to the lifted point $\hat{x} \in \mathscr{P}$. This variational formulation has been successfully exploited in [11–13] for the generation of *Optimal Delaunay Triangulations*. When $p = 1$, the ℓ_p-norm associated to T is what we call in this paper the Delaunay energy of T and, can be interpreted as the $(d+1)$-volume enclosed between the lifted triangulation $\hat{T}$ and the paraboloid $\mathscr{P}$. Given a d-simplex σ, we call the $(d+1)$-volume enclosed between the convex hull of $\hat{\sigma}$ and $\mathscr{P}$ the *Delaunay weight* of σ and denote it as $\omega(\sigma)$. The Delaunay energy of T can then be simply expressed as the sum of the Delaunay weights of its d-simplices.

Contributions.

We present a variational formulation to submanifold reconstruction in Euclidean space, that both extends our variational approach for curve reconstruction in the plane and the variational approach for generating Delaunay complexes in $\mathbb{R}^d$. Consider a finite set of points P that sample an unknown d-dimensional submanifold $\mathcal{M} \subseteq \mathbb{R}^N$ and suppose that we have at hand a simplicial complex K whose vertex set is P (such as for instance the Čech complex of P or the Vietoris-Rips complex of P). Given as input K, our goal is to find a triangulation of $\mathcal{M}$ contained in K.

A crucial ingredient in our variational formulation is to embed the triangulations contained in K inside the vector space formed by simplicial d-cycles[1] of K over the field $\mathbb{R}$. In spirit, this is similar to what is done in the theory of minimal surfaces, when oriented surfaces are identified to particular elements of a much larger set, namely the space of currents [14], which enjoys the nice property of being a vector space. Our method for reconstructing submanifolds consists in solving the following optimization problem over d-chains of K:

$$\begin{array}{ll} \underset{\gamma}{\text{minimize}} & \mathrm{E}_{\mathrm{del}}(\gamma) \\ \text{subject to} & \partial\gamma = 0, \\ & \mathrm{load}_A(\gamma) = 1 \end{array}$$

The objective function is the Delaunay energy, whose definition we adapt to the d-chains γ of K by setting

$$\mathrm{E}_{\mathrm{del}}(\gamma) = \sum_{\sigma} \omega(\sigma)\,|\gamma(\sigma)|,$$

where the sum is over all d-simplices σ of K and $\omega(\sigma)$ is the Delaunay weight of σ. The first constraint expresses the fact that we are searching for d-cycles. The second constraint is a linear equation which may be interpreted as a kind of normalization of γ that extends the condition $\mathrm{flux}_{[a,b]}(\gamma) = 1$ beyond dimension one. The letter A designates a set of parameters that specifies where the load is computed. Thus, our optimization problem is an ℓ_1-norm minimization problem under linear constraints. As such, it can be turned into a linear optimization problem in the standard form through slack variables as explained in Appendix A, and can be addressed by standard linear programming techniques such as the simplex algorithm.

One important point is that the objective function (*i.e.* the Delaunay energy) is a weighted ℓ_1-norm. Put it simply, we are searching for weighted ℓ_1-minima. The celebrated sparsity of ℓ_1-minima manifests itself in our context by the fact that the support of such minima is sparse, in other words it is non-zero only on a small subset of simplices of K. Our main result is a set of conditions under which our optimization problem has a unique solution whose support is a triangulation of $\mathcal{M}$ (either with theoretical or practical normalization).

Proof technique.

The proof requires us to introduce an elaborate construction, the *Delloc complex* of P, as a tool to describe the solution. The d-simplices of that complex possess exactly the property that we need for our analysis. In a companion paper [1] we show that the Delloc complex indeed provides a triangulation of the manifold, assuming the set of sample points P to be sufficiently dense, safe, and not too noisy. Incidently, the Delloc complex coincides with the *flat Delaunay complex* introduced in our companion

[1] Or relative d-cycles when the considered domain has a boundary.

paper [1] and is akin to the *tangential Delaunay complex* introduced and studied in [15, 16]. When the manifold is sufficiently densely sampled by the data points, all three constructions are locally isomorphic to a (weighted) Delaunay triangulation computed in a local tangent space to the manifold. Intuitively, this indicates that the Delaunay energy should locally reach a minimum for all three constructions and, therefore ought to be also a global minimum. Actually, turning this intuitive reasoning into a correct proof turns out to be more tricky than it appears and is the main purpose of the present paper. In particular, we need to globally compare the Delaunay energy of the cycle carrying the Delloc complex with that of alternate d-cycles, and this requires us to carefully distribute the Delaunay energy along barycentric coordinates.

For the purpose of the proof, it is convenient to first consider a rather artificial problem (Problem ($\star$)) where, besides the sample P, the manifold $\mathcal{M}$ is known. At the end of the paper, we show how to turn this problem into a more realistic one (Problem ($\star\star$)) that takes as input only the sample of the unknown manifold, and is correct assuming that reasonable sampling conditions are met.

An overview of the proof is provided in Section 6.2.

Related work.

A closely related problem is the computation of ℓ_1-minimum homology representative cycles. Several authors, with computational topology or topological data analysis motivations, have considered the computation of such cycles, generally for integers or integers modulo p coefficients [17–21].

A combinatorial counterpart of Delaunay energy, called *lexicographic order*, has been considered, using the field $2\mathbb{Z}/\mathbb{Z}$ of integer modulo 2 instead of $\mathbb{R}$ as chain coefficients. Given a point cloud P in Euclidean space, the lexicographic minimal chain, among chains with vertices in P and whose boundary support is the boundary of the convex hull of P has the Delaunay triangulation as its support [22]. In [23], the authors consider lexicographic minimal chains for practical applications to surface reconstruction in $\mathbb{R}^3$.

Outline.

Section 2 introduces the necessary terminology. Section 3 reviews Delaunay complexes and characterizes them as the triangulations with smallest Delaunay energy. Section 4 defines Delaunay weights and expresses the Delaunay energy as a sum of Delaunay weights. Section 5 tackles the problem of reconstructing a submanifold $\mathcal{M}$ from a simplicial complex K whose vertices sample $\mathcal{M}$. The section presents a convex optimization problem on the d-chains of K whose objective function is the Delaunay energy and whose constraints are linear. We then state our main result, which are conditions under which a solution to that optimization problem provides a triangulation of $\mathcal{M}$. Section 6 is dedicated to proving our main result. Section 7 discusses practical aspects.

2 Preliminaries

In this section, we review the necessary background and explain some of our terms.

2.1 Subsets and submanifolds

Given a set $A \subseteq \mathbb{R}^N$ and a point $x \in \mathbb{R}^N$, we say that x is an *affine combination* of points in A if we can find $a_1, \ldots, a_p$ in A and real numbers $\lambda_1, \ldots, \lambda_p$ summing up to 1 such that $x = \sum_{i=1}^{p} \lambda_i a_i$. The set of all affine combinations of points in A is called the *affine space* spanned by A and is denoted as aff A. Similarly, we say that x is a *convex combination* of points in A if we can find $a_1, \ldots, a_p$ in A and $\lambda_1, \ldots, \lambda_p$ such that $x = \sum_{i=1}^{p} \lambda_i a_i$, where $\sum_{i=1}^{p} \lambda_i = 1$ and $\lambda_i \geq 0$ for all $1 \leq i \leq p$. The set of all convex combinations of points in A form the *convex hull* of A and is denoted as conv A. The *relative interior* of A, denoted as relint(A), represents the interior of A within aff A. For any $x \in \mathbb{R}^N$ and any $r \in \mathbb{R}$, we denote the closed ball with center x and radius r by $B(x, r)$. We shall say that $A \subseteq \mathbb{R}^N$ is r*-small* if it can be enclosed in a ball of radius r. The r*-tubular neighborhood* of A is the set of points $A^{\oplus r} = \bigcup_{a \in A} B(a, r)$. The *medial axis* of A, denoted as axis(A), is the set of points in $\mathbb{R}^N$ that have at least two closest points in A. The *reach* of A is the infimum of distances between A and its medial axis, and is denoted as reach A. Furthermore, we define the *projection map* $\pi_A : \mathbb{R}^N \setminus \text{axis}(A) \to A$, which associates to each point x its unique closest point in A. This projection map is well-defined on every subset of $\mathbb{R}^N$ that does not intersect the medial axis of A. In particular, it is well-defined on every r-tubular neighborhood of A with $r < \text{reach} A$. Recall that the *angle* between two vector spaces V_0 and V_1 is defined as $\angle(V_0, V_1) = \max_{v_0 \in V_0} \min_{v_1 \in V_1} \angle v_0, v_1$. The definition is not symmetric in V_0 and V_1, unless the two vector spaces V_0 and V_1 share the same dimension. The angle between two affine spaces A_0 and A_1 whose corresponding vector spaces are V_0 and V_1 is $\angle(A_0, A_1) = \angle(V_0, V_1)$ [24].

2.2 Simplicial complexes

In this section, we review some background notation on algebraic topology and refer the reader to [25] for a detailed introduction to the topic.

All simplices and simplicial complexes that we consider in the paper are abstract. We recall that an *abstract simplicial complex* is a collection K of finite non-empty sets with the property that if σ belongs to K, so does every non-empty subset of σ. Each element σ of K is called an *abstract simplex* and its *dimension* is one less than its cardinality, $\dim \sigma = \text{card}\, \sigma - 1$. A simplex of dimension i is called an i-simplex. If τ and σ are two simplices such that $\tau \subseteq \sigma$, then τ is called a *face* of σ, and σ is called a *coface* of τ. The elements of σ are also referred to as the *vertices* of σ and the *vertex set* of K is the set of vertices of all simplices in K, $\text{Vert}\, K = \bigcup_{\sigma \in K} \sigma$. When an abstract simplex $\sigma \subseteq \mathbb{R}^N$ has its vertices in $\mathbb{R}^N$, it is naturally associated to the geometric simplex defined as conv σ. The dimension of conv σ, which is the dimension of the affine space aff σ, cannot be larger than the dimension of the abstract simplex σ. When the dimension of the geometric simplex conv σ coincides with that of the abstract simplex σ, we say that σ is *non-degenerate*. For a simplicial complex K with vertices in $\mathbb{R}^N$, we say that K is *geometrically realized* (or *embedded*) if (1) $\dim(\sigma) = \dim(\text{aff}\, \sigma)$ for all $\sigma \in K$, and (2) $\text{conv}(\alpha \cap \beta) = \text{conv}\, \alpha \cap \text{conv}\, \beta$ for all $\alpha, \beta \in K$.

Given a set of abstract simplices Σ with vertices in $\mathbb{R}^N$ (not necessarily forming a simplicial complex), we let $\Sigma^{[i]}$ designate the set of i-simplices of Σ. We define the *shadow* of Σ as the subset of $\mathbb{R}^N$ covered by the relative interior of the geometric simplices associated to the abstract simplices in Σ, $|\Sigma| = \bigcup_{\sigma\in\Sigma} \mathrm{relint}(\mathrm{conv}\,\sigma)$. The closure of Σ is the smallest simplicial complex that contains Σ.

2.3 Barycentric coordinates

Consider an abstract simplex $\alpha \subseteq \mathbb{R}^N$ and note that α is non-degenerate if and only if its vertices are affinely independent. Suppose that α is non-degenerate and consider the affine combination $x = \sum_{a\in\alpha} \lambda_a a$ with $\sum_{a\in\alpha} \lambda_a = 1$. Then, the λ_a are uniquely determined by x and are called the *(normalized) barycentric coordinates* of x with respect to α. In the paper, we shall denote each λ_a as $\mathrm{BarycentricCoord}_a^\alpha(x)$.

2.4 Chains and weighted norms

Chains play an important role in this work as they provide a tool to embed the discrete set of candidate solutions (triangulations of $\mathcal{M}$ in some simplicial complex K) into a larger continuous space (the d-chains of K). In this section, we recall some standard definitions concerning chains from [25]. Given an abstract simplex σ, two orderings of the vertices of σ are said to be *equivalent* if they differ from one another by an even permutation. The orderings of the vertices of σ fall into equivalent classes: two classes if $\dim\sigma > 0$ and one class if $\dim\sigma = 0$. Each of these classes is called an *orientation* of σ. An *oriented simplex* is a simplex σ together with an orientation of σ. We denote as $[v_0,\dots,v_d]$ the oriented d-simplex consisting of the d-simplex $\{v_0,\dots,v_d\}$ together with the equivalent class of the particular ordering $(v_0,\dots,v_d)$. Consider an abstract simplicial complex K and assume that each simplex σ in K is given an arbitrary orientation. A *d-chain* of K with coefficients in $\mathbb{R}$ is a formal sum $\gamma = \sum_\sigma \gamma(\sigma)\sigma$, where σ ranges over all d-simplices of K and $\gamma(\sigma) \in \mathbb{R}$ is the value (or the coordinate) assigned to the d-simplex σ with the rule that if σ and σ' are the same simplex but have two different orientations, then $\sigma = -\sigma'$. The set of such d-chains is a vector space denoted by $C_d(K,\mathbb{R})$. Recall that the ℓ_1-norm of γ is defined by $\|\gamma\|_1 = \sum_\sigma |\gamma(\sigma)|$. Let W be a weight function which assigns a non-negative weight $W(\sigma)$ to each d-simplex σ of K. The W-weighted ℓ_1-norm of γ is expressed as $\|\gamma\|_{1,W} = \sum_\sigma W(\sigma)|\gamma(\sigma)|$. We shall say that a chain γ is *carried by* a subcomplex D of K if γ has value 0 on every simplex that is not in D. The *support* of γ is the set of simplices on which γ has a non-zero value. It is denoted by $\mathrm{Supp}\,\gamma$. The *boundary operator* is a homomorphism $\partial : C_d(K,\mathbb{R}) \to C_{d-1}(K,\mathbb{R})$ that associates to each oriented d-simplex $\sigma = [v_0,\dots,v_d]$ the $(d-1)$-chain:

$$\partial\sigma = \sum_{i=0}^{d} (-1)^i \, [v_0,\dots,\hat{v}_i,\dots,v_d],$$

where the symbol $\hat{v}_i$ means that the vertex v_i has been deleted from the sequence of vertices forming σ. A d-chain $\gamma \in C_d(K,\mathbb{R})$ whose boundary vanishes, $\partial\gamma = 0$, is called a *d-cycle*.

3 Background on Delaunay complexes

In this section, we recall basic facts about Delaunay complexes (Section 3.1). We then give a variational characterization of Delaunay complexes (Section 3.2). Throughout the section, P designates a finite point set of $\mathbb{R}^N$.

3.1 Definitions and basic property

Definition 1 (Delaunay simplex) A *Delaunay simplex* of P is an abstract simplex $\sigma \subseteq P$ for which there exists a ball B whose boundary circumscribes σ and whose interior does not contain any point of P.

Definition 2 (Delaunay complex) The set of Delaunay simplices form an abstract simplicial complex called the *Delaunay complex* of P and is denoted as $\mathrm{Del}(P)$.

We now state a classical result on Delaunay complexes, for which we need two extra definitions.

Definition 3 (General position) Letting $d = \dim(\mathrm{aff}\, P)$, we say that $P \subseteq \mathbb{R}^N$ is in *general position* if no $d+2$ points of P lie on a common $(d-1)$-dimensional sphere and no $k+2$ points of P lie on the same k-dimensional flat for $k < d$.

Definition 4 (Triangulation) A *triangulation* of P is an abstract simplicial complex whose vertex set is P, whose shadow is $\mathrm{conv}\, P$, and which is geometrically realized.

Theorem 1 *When P is in general position,* $\mathrm{Del}(P)$ *is a triangulation of P.*

3.2 A variational characterization

The Delaunay complex of P optimizes many functionals over the set of triangulations of P [26–28], one of them being the Delaunay energy that we shall now define [29].

In preparation for this, we recall a famous result which says that building a Delaunay complex in $\mathbb{R}^N$ is topologically equivalent to building a lower convex hull in $\mathbb{R}^{N+1}$. For simplicity, we shall identify each point $x \in \mathbb{R}^N$ with the point $(x, 0)$ in $\mathbb{R}^{N+1}$. Consider the paraboloid $\mathscr{P} \subseteq \mathbb{R}^{N+1}$ defined as the graph of the function $\|\cdot\|^2 : \mathbb{R}^N \to \mathbb{R}$, $x \mapsto \|x\|^2$, where $\|\cdot\|$ designates the Euclidean norm; see Figure 3, left. For each point $x \in \mathbb{R}^N$, its vertical projection onto $\mathscr{P}$ is the point $\hat{x} = (x, \|x\|^2) \in \mathbb{R}^{N+1}$, which we call the *lifted image* of x. Similarly, the lifted image of $P \subseteq \mathbb{R}^N$ is $\hat{P} = \{\hat{p} \mid p \in P\}$. Recall that the lower convex hull of $\hat{P}$ is the portion of $\mathrm{conv}\, \hat{P}$ visible to a viewer standing at $x_{d+1} = -\infty$. A classical result says that for all $\sigma \subseteq P$, the following equivalence holds: σ is a Delaunay simplex of P if and only if $\mathrm{conv}\, \hat{\sigma}$ is contained in the lower convex hull of $\hat{P}$ [30].

We are now ready to define the Delaunay energy of any triangulation T of P. Let $d = \dim(\mathrm{aff}\, P)$. Given a triangulation T of P, the *Delaunay energy* $\mathrm{E}_{\mathrm{del}}(T)$ of T is defined as the $(d+1)$-volume between the d-manifold $|\hat{T}| = \bigcup_{\sigma \in T} \mathrm{conv}\, \hat{\sigma}$ and the

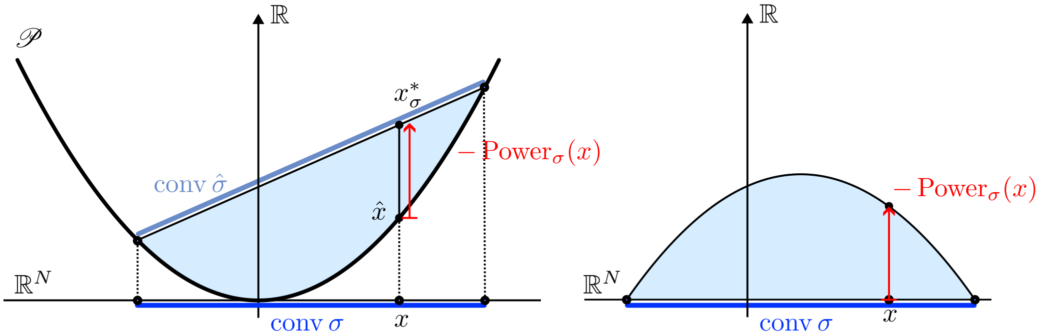

Fig. 3: Left: the Delaunay weight of σ can be depicted as the $(d+1)$-volume of the blue region between the lifted geometric simplex $\operatorname{conv}\hat{\sigma}$ and the paraboloid $\mathscr{P}$ (see Lemma 4). Right: the Delaunay weight of σ is also the $(d+1)$-volume of the blue region lying below the graph of $-\operatorname{Power}_\sigma$ and above $\operatorname{conv}\sigma$.

paraboloid $\mathscr{P}$. Let us give a formula for this energy. Consider a point $x \in \operatorname{conv} P$. By construction, x belongs to at least one geometric d-simplex $\operatorname{conv}\sigma$ for some $\sigma \in T$. Erect an infinite vertical half-line going up from x. This half-line intersects the paraboloid $\mathscr{P}$ at point $\hat{x}$ and $\operatorname{conv}\hat{\sigma}$ at point $x^\star_\sigma$; see Figure 3, left. We have

$$\mathrm{E}_{\mathrm{del}}(T) = \sum_{\sigma \in T^{[d]}} \int_{x \in \operatorname{conv}\sigma} \|\hat{x} - x^\star_\sigma\| \, dx.$$

Let us recall a well-known result [10, 28], that is a direct consequence of the lifting construction:

Theorem 2 *When P is in general position, the triangulation of P that minimizes the Delaunay energy is unique and is the Delaunay complex of P.*

4 Delaunay weight

In this section, we define the Delaunay weight of a simplex with vertices in $\mathbb{R}^N$ (Section 4.1). We show that the Delaunay energy defined in the previous section can be expressed as a sum of Delaunay weights (Section 4.2). We then provide an intrinsec expression for Delaunay weights (Section 4.3) that motivates extending the Delaunay energy to collection of d-simplices in $\mathbb{R}^N$ and more generally to d-chains of simplicial complexes with vertices in $\mathbb{R}^N$, as will be done in the next section.

4.1 Definition

Let $\sigma \subseteq \mathbb{R}^N$ be an abstract simplex. If σ is non-degenerate, we define $S(\sigma)$ as the smallest $(N-1)$-sphere that circumscribes σ. We also let $Z(\sigma)$ and $R(\sigma)$ denote the center and radius of $S(\sigma)$, respectively. Finally, we introduce the map

$$x \mapsto \mathrm{Power}_\sigma(x) = \|x - Z(\sigma)\|^2 - R(\sigma)^2,$$

which associates each point $x \in \mathbb{R}^N$ with the power distance of x from $S(\sigma)$.

Definition 5 (Delaunay weight) The *Delaunay weight* of an abstract d-simplex $\sigma \subseteq \mathbb{R}^N$ is:

$$\omega(\sigma) = \begin{cases} -\int_{x \in \mathrm{conv}\,\sigma} \mathrm{Power}_\sigma(x)\,dx & \text{if } \sigma \text{ is non-degenerate,} \\ 0 & \text{otherwise.} \end{cases}$$

Figure 3, right depicts graphically the Delaunay weight of a simplex $\sigma \subseteq \mathbb{R}^N$. It is worth noting that the Delaunay weight is defined for any abstract d-simplex σ, irrespective of whether σ is the Delaunay simplex of some point set P or not. The reason for calling it a Delaunay weight will become clear in the next section.

4.2 Delaunay energy as a sum of Delaunay weights

Before showing that the Delaunay energy can be expressed as a sum of Delaunay weights in Lemma 4, we first provide a useful expression of $\mathrm{Power}_\sigma(x)$ when x is an affine combination of the vertices of σ (Lemma 3) and deduce an alternative expression of the Delaunay weight in Lemma 4.

Lemma 3 *Let $\sigma \subseteq \mathbb{R}^N$ be a non-degenerate simplex and let x be an affine combination of the vertices of σ. For every $z \in \mathbb{R}^N$*

$$\mathrm{Power}_\sigma(x) = \|x - z\|^2 - \sum_{a \in \sigma} \mathrm{BarycentricCoord}_a^\sigma(x)\,\|a - z\|^2.$$

Proof Recall that $\mathrm{Power}_\sigma(x) = \|x - Z(\sigma)\|^2 - R(\sigma)^2$. On one hand, we have

$$\|x - Z(\sigma)\|^2 \;=\; \|x - z\|^2 + 2(x - z)\cdot(z - Z(\sigma)) + \|z - Z(\sigma)\|^2.$$

On the other hand, writing $\lambda_a = \mathrm{BarycentricCoord}_a^\sigma(x)$ for short, we have

$$\begin{aligned} R(\sigma)^2 \;&=\; \sum_{a \in \sigma} \lambda_a \|Z(\sigma) - a\|^2 \\ &=\; \sum_{a \in \sigma} \lambda_a \left[\|Z(\sigma) - z\|^2 + 2(Z(\sigma) - z)\cdot(z - a) + \|z - a\|^2\right] \\ &=\; \|Z(\sigma) - z\|^2 + 2(Z(\sigma) - z)\cdot(z - x) + \sum_{a \in \sigma} \lambda_a \|z - a\|^2. \end{aligned}$$

Substracting the above expressions of $\|x - Z(\sigma)\|^2$ and $R(\sigma)^2$ yields the result. □

Lemma 4 *For any non-degenerate abstract d-simplex $\sigma \subseteq \mathbb{R}^N$, its Delaunay weight represents the $(d+1)$-volume between the lifted geometric simplex* $\mathrm{conv}\,\hat{\sigma}$ *and the paraboloid $\mathscr{P}$,* i.e.

$$\omega(\sigma) = \int_{x \in \mathrm{conv}\,\sigma} \|\hat{x} - x_\sigma^\star\|\,dx,$$

where $\hat{x} = (x, \|x\|^2)$ *and* $x^\star_\sigma = \sum_{a\in\sigma} \mathrm{BarycentricCoord}^\sigma_a(x)\hat{a}$*. If* P *designates a finite point set of* $\mathbb{R}^N$ *in general position and* $d = \dim(\mathrm{aff}\, P)$*, then the Delaunay energy of any triangulation* T *of* P *can be expressed as*

$$\mathrm{E}_{\mathrm{del}}(T) = \sum_{\sigma\in T^{[d]}} \omega(\sigma).$$

Proof Letting $x \in \mathrm{conv}\,\sigma$, we show that $-\mathrm{Power}_\sigma(x) = \|\hat{x} - x^\star_\sigma\|$. Applying Lemma 3 with $z = 0$ and writing $\lambda_a = \mathrm{BarycentricCoord}^\sigma_a(x)$ for short, we get that

$$\begin{aligned} -\mathrm{Power}_\sigma(x) &= \sum_{a\in\sigma} \lambda_a \|a\|^2 - \|x\|^2 \\ &= \left\| \sum_{a\in\sigma} \lambda_a (a, \|a\|^2) - (x, \|x\|^2) \right\| \\ &= \left\| \sum_{a\in\sigma} \lambda_a \hat{a} - \hat{x} \right\| \\ &= \|x^\star_\sigma - \hat{x}\|. \end{aligned}$$

The expressions of both $\omega(\sigma)$ and $\mathrm{E}_{\mathrm{del}}(T)$ follow immediately. □

The above lemma suggests the following interpretation of the Delaunay energy. Given a triangulation T of P, consider the map $w_T : \mathrm{conv}\, P \to \mathbb{R}$ whose restriction to any d-simplex σ of T is defined by $w_T(x) = -\mathrm{Power}_\sigma(x)$. The graph of w_T is a d-dimensional piecewise parabolic manifold $\mathcal{W}_T$ which has been depicted for two different triangulations in Figure 4. The Delaunay energy can then be interpreted as the $(d + 1)$-volume of the region lying below $\mathcal{W}_T$ and above $\mathrm{conv}\, P$.

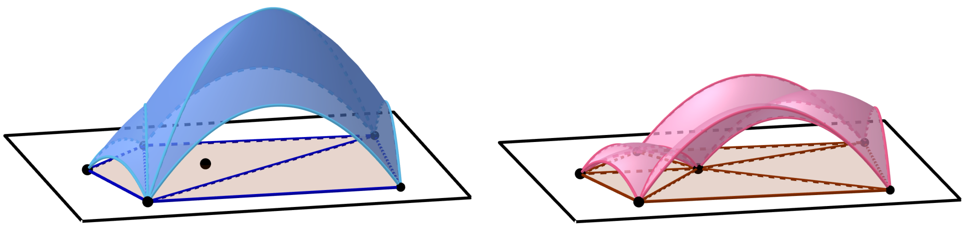

Fig. 4: Two triangulations of six points (black dots) in the plane. For each triangulation T, the Delaunay energy is the volume between the convex hull of the points and the piecewise parabolic surface $\mathcal{W}_T$. The surface $\mathcal{W}_T$ is lowest and therefore the Delaunay energy of T is smallest when T is the Delaunay complex as is the case on the right.

4.3 Intrinsec closed expression

Below, we give a closed expression for the Delaunay weight due to Chen and Holst in [12]. For completeness, we provide a proof. Writing $\mathrm{vol}(\sigma)$ for the d-dimensional volume of $\mathrm{conv}\,\sigma$, we have:

Lemma 5 ([12]) *The weight of the abstract d-simplex $\sigma = \{a_0, \ldots, a_d\}$ is*

$$\omega(\sigma) = \frac{1}{(d+1)(d+2)} \mathrm{vol}(\sigma) \left[\sum_{0 \le i < j \le d} \|a_i - a_j\|^2 \right].$$

Proof Let $\sigma = \{a_0, a_1, \ldots, a_d\} \subseteq \mathbb{R}^N$. If σ is degenerate, then $\mathrm{vol}(\sigma) = 0$ and the result is clear. Suppose that σ is non-degenerate and recall that the standard simplex is

$$\Delta_d = \{\lambda \in \mathbb{R}^d \mid \sum_{i=1}^{d} \lambda_i \le 1; \lambda_i \ge 0, i = 1, 2, \ldots, d\}.$$

We introduce the map $\psi : \mathbb{R}^d \to \mathbb{R}^d$, defined by $\psi(\lambda) = a_0 + \sum_{i=1}^d \lambda_i (a_i - a_0)$, which establishes a one-to-one correspondence between the points λ of the standard simplex Δ_d and the points $x = \psi(\lambda)$ of $\mathrm{conv}\,\sigma$. Making the change of variable $x = \psi(\lambda) \to \lambda$, we get that:

$$w(\sigma) = \int_{\lambda \in \Delta_d} -\mathrm{Power}_\sigma(\psi(\lambda)) \cdot |\det(\mathrm{D}\psi)(\lambda)|\, d\lambda.$$

Noting that $\mathrm{D}\psi(\lambda)$ is the $d \times d$ matrix whose ith column is the vector $a_i - a_0$, we deduce that $|\det(\mathrm{D}\psi)(\lambda)| = d!\,\mathrm{vol}(\sigma)$. Observing that $\psi(\lambda)$ has (normalized) barycentric coordinates $(1 - \sum_{i=1}^d \lambda_i, \lambda_1, \lambda_2, \ldots, \lambda_d)$ and applying Lemma 3 with $z = a_0$, we can write:

$$\mathrm{Power}_\sigma(\psi(\lambda)) = -\left(\sum_{i=1}^{d} \lambda_i \|a_i - a_0\|^2 \right) + \|\psi(\lambda) - a_0\|^2,$$

and thus obtain (after plugging in the expression of $\psi(\lambda)$)

$$w(\sigma) = d!\,\mathrm{vol}(\sigma) \int_{\lambda \in \Delta_d} \left[\sum_{i=1}^{d} \lambda_i \|a_i - a_0\|^2 - \left\| \sum_{i=1}^{d} \lambda_i (a_i - a_0) \right\|^2 \right] d\lambda.$$

We then use a formula for integrating a homogeneous polynomial on the standard simplex that may be found in [31]:

$$\int_{\lambda \in \Delta_d} \lambda_1^{\eta_1} \ldots \lambda_d^{\eta_d}\, d\lambda = \frac{\eta_1! \ldots \eta_d!}{(d + \sum_i \eta_i)!}.$$

We obtain that

$$w(\sigma) = \frac{1}{(d+1)(d+2)} \mathrm{vol}(\sigma) \left[d \sum_{i=1}^{d} \|a_i - a_0\|^2 - 2 \sum_{1 \le i < j \le d} (a_i - a_0) \cdot (a_j - a_0) \right].$$

Observing that $\|a_i - a_0\|^2 + \|a_j - a_0\|^2 - 2(a_i - a_0) \cdot (a_j - a_0) = \|a_i - a_j\|^2$, we can further rearrange the above formula to get the result. □

It follows from Definition 5 but also from the expression of the Delaunay weight given in Lemma 5 that two isometric simplices have the same Delaunay weight. Hence, a Delaunay energy can be straightforwardly associated to any collection Σ of d-simplices living in $\mathbb{R}^N$ by setting $E(\Sigma) = \sum_{\sigma \in \Sigma} \omega(\sigma)$. It is then tempting to ask what would happen if one minimizes this energy over all collections Σ of d-simplices whose vertices sample a d-dimensional submanifold $\mathcal{M}$ and whose union is homeomorphic to that submanifold. As is, the problem is non-convex. We shall transform it into a convex problem in the next section.

5 Variational formulation for submanifold reconstruction

Afterwards, we assume that the shape $\mathcal{M}$ we wish to reconstruct is a compact orientable C^2 d-dimensional submanifold of $\mathbb{R}^N$ for some $d < N$. We let P be a finite point set that samples $\mathcal{M}$ and suppose furthermore that we have at our disposal a simplicial complex K whose vertices are the points of P. The complex K can be thought of as some rough approximation of $\mathcal{M}$ as illustrated in Figure 5.

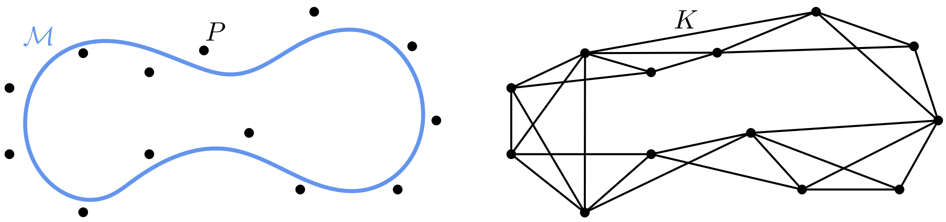

Fig. 5: Left: a d-dimensional submanifold $\mathcal{M}$ (for $d = 1$) and a noisy sample P of $\mathcal{M}$. Right: a simplicial complex K whose vertex set is P.

Details on how to derive K from P are given at the end of the section. In this section, we describe a convex optimization problem on the d-chains of K and state conditions under which the solution to that problem is unique and provides a faithful reconstruction of $\mathcal{M}$. The concept of faithful reconstruction encapsulates what we mean by a "desirable" reconstruction of $\mathcal{M}$:

Definition 6 (Faithful reconstruction) Consider a subset $\mathcal{M} \subseteq \mathbb{R}^N$ whose reach is positive, and a simplicial complex D with vertex set in $\mathbb{R}^N$. We say that D *reconstructs* $\mathcal{M}$ *faithfully* (or is a *faithful reconstruction* of $\mathcal{M}$) if the following three conditions hold:

Embedding: D is geometrically realized;

Closeness: $|D| \subseteq \mathcal{M}^{\oplus r}$ for some $0 \leq r < \operatorname{reach} \mathcal{M}$;

Homeomorphism: the projection map $\pi_{\mathcal{M}} : |D| \to \mathcal{M}$ is a homeomorphism.

We note that when D reconstructs $\mathcal{M}$ faithfully, $|D|$ and $\mathcal{M}$ are homeomorphic and D is a triangulation of $\mathcal{M}$.

The rest of the section is organized as follows. Section 5.1 presents our convex optimization problem on the d-chains of K. Section 5.2 shows that the feasible set of that problem contains all faithful reconstructions of $\mathcal{M}$ in K. In Section 5.3, we introduce the necessary definitions to state our main theorem in Section 5.4.

Because we have assumed $\mathcal{M}$ to be a compact C^2 submanifold of $\mathbb{R}^N$, the reach of $\mathcal{M}$ is positive and finite [32]. Afterwards, we denote it as $\mathcal{R} = \operatorname{reach} \mathcal{M}$. Given $m \in \mathcal{M}$, we denote the vector tangent space to $\mathcal{M}$ at m as $T_m\mathcal{M}$ and the affine tangent space to $\mathcal{M}$ at m as $\mathbf{T}_m\mathcal{M}$. Clearly, $\mathbf{T}_m\mathcal{M} = x + T_x\mathcal{M}$. In the rest of the section, we assume that $\mathcal{M}$ together with all d-simplices of K have received an arbitrary orientation. We also assume that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $0 \leq r < \operatorname{reach} \mathcal{M}$ and that none of the d-simplices of K are orthogonal to $\mathcal{M}$. Precisely, defining the *angular deviation* of a simplex σ relatively to $\mathcal{M}$ as

$$\operatorname{angularDeviation}_{\mathcal{M}}(\sigma) = \max_{m \in \pi_{\mathcal{M}}(\operatorname{conv} \sigma)} \angle(\operatorname{aff} \sigma, \mathbf{T}_m\mathcal{M}),$$

we assume that each d-simplex $\sigma \in K$ is such that $\operatorname{angularDeviation}_{\mathcal{M}}(\sigma) < \frac{\pi}{2}$. This allows us to assign to each d-simplex $\sigma \in K$ a sign with respect to $\mathcal{M}$ as follows:

$$\operatorname{sign}_{\mathcal{M}}(\sigma) = \begin{cases} 1 & \text{if the orientation of } \sigma \text{ is consistent with that of } \mathcal{M}, \\ -1 & \text{otherwise.} \end{cases}$$

We refer the reader to Appendix F in [33] for a formal definition of consistency and more details.

5.1 Least ℓ_1-norm problem

We need notation to describe the convex optimization problem that we are considering. Let ω be the weight function which assigns to each d-simplex σ of K its Delaunay weight $\omega(\sigma)$ introduced in Section 2. We define the *Delaunay energy* of the chain $\gamma \in C_d(K, \mathbb{R})$ to be its ω-weighted ℓ_1-norm:

$$\mathrm{E}_{\mathrm{del}}(\gamma) = \|\gamma\|_{1,\omega} = \sum_{\sigma} \omega(\sigma) \cdot |\gamma(\sigma)| = \sum_{\sigma} \left(\int_{x \in \operatorname{conv} \sigma} - \operatorname{Power}_{\sigma}(x)\, dx \right) \cdot |\gamma(\sigma)|,$$

where σ ranges over all d-simplices of K. The Delaunay energy is the objective function of our optimization problem. To describe the constraint functions, let $\mathbf{1}_X : \mathbb{R}^N \to \{0, 1\}$ denote the indicator function of a subset $X \subseteq \mathbb{R}^N$. Suppose that $|K| \subseteq \mathcal{M}^{\oplus r}$ for some $0 \leq r < \operatorname{reach} \mathcal{M}$ and that each d-simplex $\sigma \in K$ satisfies $\operatorname{angularDeviation}_{\mathcal{M}}(\sigma) < \frac{\pi}{2}$. Given $m_0 \in \mathcal{M}$, we assign to each d-chain γ of K the real number:

$$\operatorname{load}_{m_0, \mathcal{M}}(\gamma) = \sum_{\sigma \in K^{[d]}} \gamma(\sigma) \operatorname{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\pi_{\mathcal{M}}(\operatorname{conv} \sigma)}(m_0)$$

and call it the *load* of $\gamma \in C_d(K, \mathbb{R})$ on $\mathcal{M}$ at m_0. Roughly, it measures the "flux" of the chain γ above point $m_0 \in \mathcal{M}$. We illustrate its evaluation in Figure 6.

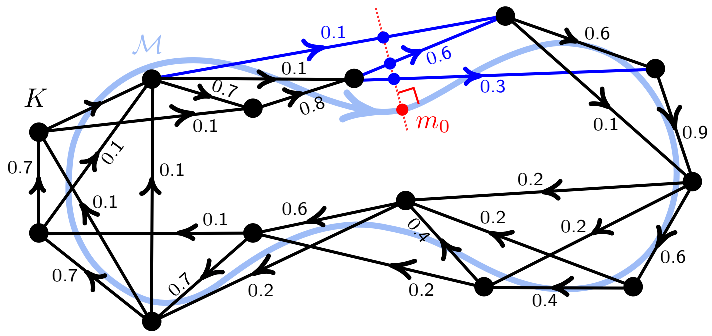

Fig. 6: A simplicial complex K whose edges have an orientation consistent with that of $\mathcal{M}$ and a 1-cycle γ of K. To evaluate the load of $\gamma \in C_1(K, \mathbb{R})$ on the curve $\mathcal{M}$ at m_0, one has first to select edges (depicted in blue) for which there is a point (blue dot) whose projection onto $\mathcal{M}$ is m_0 and sum up the coefficients of γ on these edges. In this example, $\text{load}_{m_0,\mathcal{M}}(\gamma) = 1$ and therefore γ satisfies the constraints of Problem ($\star$).

Letting m_0 be a generic[2] point on $\mathcal{M}$, we are interested in the following optimization problem over the set of chains in $C_d(K, \mathbb{R})$:

$$\begin{aligned} \underset{\gamma}{\text{minimize}} \quad & \text{E}_{\text{del}}(\gamma) \\ \text{subject to} \quad & \partial\gamma = 0, \\ & \text{load}_{m_0,\mathcal{M}}(\gamma) = 1 \end{aligned} \qquad (\star)$$

Problem ($\star$) is a least-norm problem whose constraint functions ∂ and $\text{load}_{m_0,\mathcal{M},K}$ are clearly linear. It is therefore a convex optimization problem. The first constraint $\partial\gamma = 0$ expresses the fact that we are searching for d-cycles. The second constraint $\text{load}_{m_0,\mathcal{M}}(\gamma) = 1$ is a normalization of γ and forbids the zero chain to belong to the feasible set. Two chains that satisfy the constraints are depicted in Figures 6 and 7. We shall see that, under the assumptions of our main theorem, the solution to Problem ($\star$) takes its coordinate values in $\{0, +1, -1\}$ and is furthermore the code of a faithful reconstruction of $\mathcal{M}$.

In Problem ($\star$), besides the simplicial complex K that we shall see how to build from P, the knowledge of the manifold $\mathcal{M}$ seems to be required as well for expressing the normalization constraint. In Section 7.1, we discuss how to transform Problem ($\star$) into an equivalent problem that does not refer to $\mathcal{M}$ anymore.

[2] Generic in the sense that it is not in the projection on $\mathcal{M}$ of the convex hull of any $(d-1)$-simplex of K.

5.2 Faithful reconstructions are encoded in the feasible set

Given a subcomplex $D \subseteq K$, we associate to D the d-chain code_D of K whose coordinate on the d-simplex σ is:

$$\mathrm{code}_D(\sigma) = \begin{cases} \mathrm{sign}_{\mathcal{M}}(\sigma) & \text{if } \sigma \in D^{[d]}, \\ 0 & \text{otherwise.} \end{cases}$$

We note that whenever D is a faithful reconstruction of $\mathcal{M}$, then code_D provides a way of encoding D as a d-chain, since D can be recovered straightforwardly from code_D by taking the closure of the support of code_D. The code of a faithful reconstruction is depicted in Figure 7.

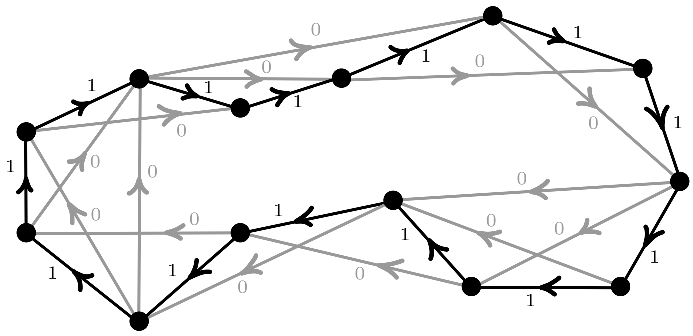

Fig. 7: A 1-chain of K that encodes a faithful reconstruction of curve $\mathcal{M}$. Its coefficients are either 0 (on grey edges) or 1 (on black edges), assuming the orientation of edges is consistent with that of $\mathcal{M}$. This 1-chain satisfies the constraints of Problem ($\star$).

In this section, we show that, under weak conditions on K, if D is a faithful reconstruction of $\mathcal{M}$, then code_D satisfies the constraints of Problem ($\star$). Indeed, $|D|$ being homeomorphic to $\mathcal{M}$, the generic point $m_0 \in \mathcal{M}$ is covered by the projection of the convex hull of a unique d-simplex $\sigma \in D$ and

$$\mathrm{load}_{m_0,\mathcal{M}}(\mathrm{code}_D) = \mathrm{code}_D(\sigma)\,\mathrm{sign}_{\mathcal{M}}(\sigma) = 1.$$

The next lemma states conditions under which the constraint $\partial\,\mathrm{code}_D = 0$ is also satisfied.

Lemma 6 *Let $r, \rho \geq 0$ such that $\rho < \frac{\sqrt{2}}{4}(\mathcal{R} - r)$. Let K be a simplicial complex such that $|K| \subseteq \mathcal{M}^{\oplus r}$ and whose d-simplices are ρ-small and have an angular deviation smaller than $\frac{\pi}{4}$ relatively to $\mathcal{M}$. If the subcomplex $D \subseteq K$ is a faithful reconstruction of $\mathcal{M}$, then code_D is a cycle.*

Proof We first prove that for all simplices $\sigma \in D$ and all points $m \in \pi_{\mathcal{M}}(\operatorname{conv}\sigma)$, we have that

$$\pi_{\mathcal{M}}(\operatorname{conv}\sigma) \subseteq B\left(m, \sin\left(\frac{\pi}{4}\right)\mathcal{R}\right)^{\circ}. \tag{1}$$

Indeed, consider $x, x' \in \operatorname{conv}\sigma$. Suppose that $m = \pi_{\mathcal{M}}(x)$ and let $m' = \pi_{\mathcal{M}}(x')$. We know from [32, page 435] that for $0 \leq r < \operatorname{reach}\mathcal{M}$, the projection map $\pi_{\mathcal{M}}$ is $\left(\frac{\mathcal{R}}{\mathcal{R}-r}\right)$-Lipschitz for points at distance less than r from $\mathcal{M}$. It follows that

$$\|m - m'\| \leq \frac{\mathcal{R}}{\mathcal{R}-r}\|x - x'\| \leq \frac{2\rho}{\mathcal{R}-r}\mathcal{R} < \frac{\sqrt{2}}{2}\mathcal{R}$$

and Inclusion (1) follows.

Given a simplicial complex L and a point $x \in \mathbb{R}^N$, we define the *star* of x in L as the set of simplices $\operatorname{St}(x, L) = \{\sigma \in L \mid x \in \operatorname{conv}\sigma\}$. Since D is a faithful reconstruction of $\mathcal{M}$, $|D|$ is a d-dimensional submanifold. Hence, each $(d-1)$-simplex $\tau \in D$ has exactly two d-cofaces σ_1 and σ_2; see Figure 8. Consider a point x in the relative interior of τ and its projection $m = \pi_{\mathcal{M}}(x)$ onto $\mathcal{M}$. The star of x in D consists of the two d-simplices σ_1 and σ_2 and the common $(d-1)$-face τ. It follows that the set $\pi_{\mathbf{T}_m\mathcal{M}}(\operatorname{St}(x, D))$ possesses exactly two d-simplices $\sigma_1' = \pi_{\mathbf{T}_m\mathcal{M}}(\sigma_1)$ and $\sigma_2' = \pi_{\mathbf{T}_m\mathcal{M}}(\sigma_2)$, and one $(d-1)$-simplex $\tau' = \pi_{\mathbf{T}_m\mathcal{M}}(\tau)$. As we project the d-simplex $\sigma_i = [u_0, \ldots, u_d]$, let us preserve the ordering of the vertices, that is, let $\sigma_i' = [\pi_{\mathbf{T}_m\mathcal{M}}(u_0), \pi_{\mathbf{T}_m\mathcal{M}}(u_1), \ldots, \pi_{\mathbf{T}_m\mathcal{M}}(u_d)]$. Let us give to $\mathbf{T}_m\mathcal{M}$ an orientation that is consistent with that of $\mathcal{M}$. Inclusion (1) allows us to apply Lemma **??** in [33, Appendix F]: each d-simplex σ_i' has the same orientation with respect to $\mathbf{T}_m\mathcal{M}$ than that of σ_i with respect to $\mathcal{M}$. Let $s_i = \operatorname{sign}_{\mathbf{T}_m\mathcal{M}}(\sigma_i') = \operatorname{sign}_{\mathcal{M}}(\sigma_i)$.

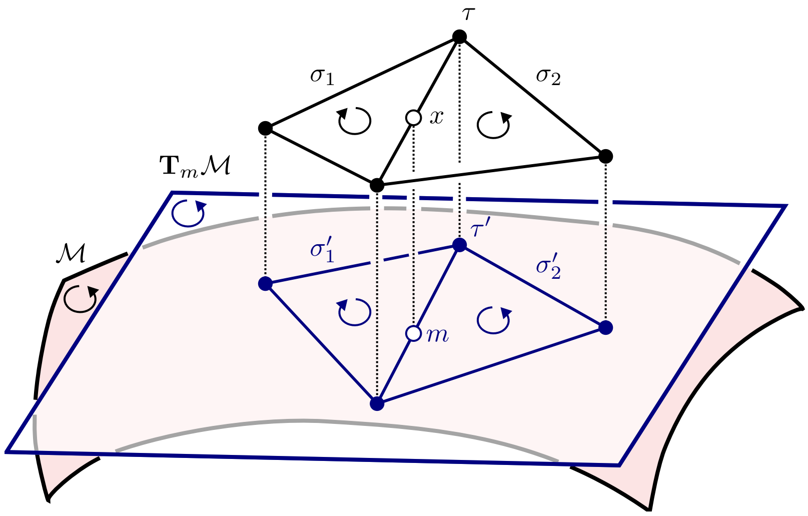

Fig. 8: Notation for the proof of Lemma 6. In this example, $s_1 = -1$ and $s_2 = +1$.

We claim that the two geometric d-cofaces $\operatorname{conv}\sigma_1'$ and $\operatorname{conv}\sigma_2'$ of $\operatorname{conv}\tau'$ have disjoint relative interiors. Indeed, let us denote by U_x an open neighborhood of x in $\mathbb{R}^N$ and suppose that U_x is sufficiently small so that its restriction to $|K|$ is contained in $|\operatorname{St}(x, D)|$. For

$i = 1, 2$, let $U_x^i = U_x \cap \operatorname{aff} \sigma_i$. Note that the map $\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^i}$ is differentiable and the map $\pi_{\mathbf{T}_m\mathcal{M}}|_{U_x^i}$ is affine. Both maps have equal differential maps at x, that is:

$$\mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^i} \right) = \mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}}|_{U_x^i} \right). \tag{2}$$

Let T_i^+ denote the set of all vectors parallel to $\operatorname{aff} \sigma_i$ and pointing inside $\operatorname{conv} \sigma_i$ after translation at x. This set forms a closed half-space in the vector tangent space to $\operatorname{aff} \sigma_i$. Since $\pi_{\mathbf{T}_m\mathcal{M}}|_{U_x^i}$ is affine, it coincides, up to a constant, with its differential at x and using Equation (2), we get that

$$\operatorname{conv} \sigma_i' \;\subseteq\; x + \mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}}|_{U_x^i} \right) (T_i^+) \;=\; x + \mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^i} \right) (T_i^+). \tag{3}$$

Observe that the map $\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^i}$, being the composition of two injective functions, is injective. It follows that $\mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^1} \right) (T_1^+)$ and $\mathrm{D}_x \left(\pi_{\mathbf{T}_m\mathcal{M}} \circ \pi_{\mathcal{M}}|_{U_x^2} \right) (T_2^+)$ are two half-spaces in the vector space $T_m\mathcal{M}$ with disjoint interiors. Using Equation (3), we obtain that $\operatorname{conv} \sigma_1'$ and $\operatorname{conv} \sigma_2'$ also have disjoint interiors, as claimed.

It follows that $\partial(s_1\sigma_1' + s_2\sigma_2')$ is 0 on τ', and consequently $\partial(s_1\sigma_1 + s_2\sigma_2)$ is 0 on τ. We have shown that $\partial \operatorname{code}_D = 0$. □

5.3 Geometric conditions on the sample

Recall that our goal is to give conditions under which a solution to Problem ($\star$) provides a faithful reconstruction of $\mathcal{M}$. To express the conditions that we need, let us introduce some definitions and notations.

Definition 7 (Dense sample) We say that P is an ε-*dense* sample of $\mathcal{M}$ if for every point $m \in \mathcal{M}$, there is a point $p \in P$ with $\|p - m\| \leq \varepsilon$ or, equivalently, if $\mathcal{M} \subseteq P^{\oplus\varepsilon}$.

Definition 8 (Accurate sample) We say that P is a δ-*accurate* sample of $\mathcal{M}$ if for every point $p \in P$, there is a point $m \in \mathcal{M}$ with $\|p - m\| \leq \delta$ or, equivalently, if $P \subseteq \mathcal{M}^{\oplus\delta}$.

The *separation* of a point set P is

$$\operatorname{separation}(P) = \min_{p \neq q \in P} \|p - q\|.$$

We recall that the *height* of a simplex σ is

$$\operatorname{height}(\sigma) = \min_{v \in \sigma} d(v, \operatorname{aff}(\sigma \setminus \{v\})).$$

The height of σ vanishes if and only if σ is degenerate. The *protection* of a simplex σ relatively to a point set Q is

$$\operatorname{protection}(\sigma, Q) = \min_{q \in \pi_{\operatorname{aff} \sigma}(Q \setminus \sigma)} d(q, S(\sigma)).$$

We stress that our definition of a simplex protection differs slightly from the one in [16, 34]. We now associate to a finite point set P and a scale ρ three quantities that describe the quality of the pair (P, ρ) at dimension d:

$$\begin{aligned} \text{height}(P, \rho) &= \min_{\sigma} \text{height}(\sigma), \\ \text{angularDeviation}_{\mathcal{M}}(P, \rho) &= \max_{\sigma} \text{angularDeviation}_{\mathcal{M}}(\sigma), \\ \text{protection}(P, \rho) &= \min_{\sigma} \text{protection}(\sigma, P \cap B(c_\sigma, \rho)), \end{aligned}$$

where the two minima and the maximum are over all ρ-small d-simplices $\sigma \subseteq P$. Observe that assuming $\text{height}(P, \rho) > 0$ is equivalent to assuming that all ρ-small d-simplices of P are non-degenerate.

Definition 9 (Safety condition) Let ε, δ, and ρ be non-negative real numbers. The *safety condition* on (P, ε, δ) at scale ρ is the existence of a real number $\theta \in \left[0, \frac{\pi}{6}\right]$ such that:

$$\begin{aligned} \text{angularDeviation}_{\mathcal{M}}(P, \rho) &\leq \frac{\theta}{2} - \arcsin\left(\frac{\rho + \delta}{\mathcal{R}}\right), \\ \text{separation}(P) &> 8(\delta\theta + \rho\theta^2) + 6\delta + \frac{2\rho^2}{\mathcal{R}}, \\ \text{protection}(P, 3\rho) &> 8(\delta\theta + \rho\theta^2)\left(1 + \frac{4d\varepsilon}{\text{height}(P, \rho)}\right). \end{aligned}$$

Roughly speaking, assuming the safety condition on (P, ε, δ) at scale ρ enforces ρ-small d-simplices of P to make a sufficiently small angle relatively to $\mathcal{M}$. It also enforces P to be both sufficiently separated and protected at scale 3ρ. As explained in the companion paper [1], the safety condition on (P, ε, δ) can be met by considering a $\left(\frac{20\varepsilon}{21}\right)$-dense $\left(\frac{\delta}{2}\right)$-accurate point set P and perturbing it as described in [1].

5.4 Main theorem

In the statement of our main theorem, there is a constant $\Omega(\Delta_d)$ that depends only upon the dimension d and whose definition is given in the proof of Lemma 16. Let $\mathcal{C}(P, r)$ denote the set of simplices of P that are r-small, also known as the *Čech complex* of P at scale r.

Theorem 7 (Faithful reconstruction by a variational approach) *Let $\mathcal{M}$ be a compact orientable C^2 d-dimensional submanifold of $\mathbb{R}^N$ for some $d < N$. Let ε, δ, and ρ be non-negative real numbers such that $\delta \leq \varepsilon$ and $16\varepsilon \leq \rho < \frac{\mathcal{R}}{4}$. Let $\Theta = \text{angularDeviation}_{\mathcal{M}}(P, \rho)$ and assume that $\Theta \leq \frac{\pi}{6}$. Set*

$$J = \frac{(\mathcal{R} + \rho)^d}{(\mathcal{R} - \rho)^d \ (\cos \Theta)^{\min\{d, N-d\}}} - 1.$$

Let P be a δ-accurate ε-dense sample of $\mathcal{M}$ such that $\operatorname{height}(P,\rho) > 0$. *Suppose that the safety condition on (P,ε,δ) is satisfied at scale ρ. Suppose furthermore that*

$$\operatorname{protection}(P,3\rho)^2 + \operatorname{protection}(P,3\rho)\operatorname{separation}(P) > \max\left\{10\rho\Theta(\varepsilon+\rho\Theta), \frac{4J(1+J)}{(d+2)(d-1)!\,\Omega(\Delta_d)}\rho^2\right\}. \tag{4}$$

Consider a simplicial complex K such that

$$\operatorname{Del}(P)\cap\mathcal{C}(P,\varepsilon)\subseteq K\subseteq\mathcal{C}(P,\rho). \tag{5}$$

Then Problem ($\star$) has a unique solution and the closure of the support of that solution is a faithful reconstruction of $\mathcal{M}$.

Observe that our main theorem does not require K to be geometrically realized nor to retain the homotopy type of $\mathcal{M}$. One may ask about the feasability of realizing the assumptions of Theorem 7. In Section 7.2, we explain how to apply Moser Tardos Algorithm ([35] and [16, Section 5.3.4]) as a perturbation scheme for enforcing both the safety condition and Condition (4) required by Theorem 7.

Choosing the simplicial complex K.

Recall that the Čech complex of P at scale r, denoted as $\mathcal{C}(P,r)$, is the set of simplices of P that are r-small. The Rips complex of P at scale r, denoted as $\mathcal{R}(P,r)$, consists of all simplices of P with diameter at most $2r$. It is a more easily-computed version of the Čech complex. We stress that our main theorem applies to any simplicial complex K such that $\operatorname{Del}(P)\cap\mathcal{C}(P,\varepsilon)\subseteq K\subseteq\mathcal{C}(P,\rho)$. Since $\mathcal{C}(P,r)\subseteq\mathcal{R}(P,r)\subseteq\mathcal{C}(P,\sqrt{2}r)$ for all $r\geq 0$, it applies in particular to any $K=\mathcal{R}(P,r)$ with $\varepsilon\leq r\leq\frac{\rho}{\sqrt{2}}$. This choice of K is well-suited for applications in high dimensional spaces, while choosing $K=\operatorname{Del}(P)\cap\mathcal{C}(P,r)$ for any $\varepsilon\leq r\leq\rho$ may be more suited for applications in low dimensional spaces.

6 Proving the main theorem

6.1 Technical lemma

The proof of Theorem 7 relies on a technical lemma which we now state and prove.

Lemma 8 *Let $\mathcal{D}$ be an orientable d-dimensional submanifold (with or without boundary) of $\mathbb{R}^N$ and let K be a simplicial complex with vertices in $\mathbb{R}^N$. Assume that there is a continuous function $\varphi:|K|\to\mathcal{D}$. Suppose that for each d-simplex $\sigma\in K$, we have two positive weights $W(\sigma)\geq W_{\min}(\sigma)$ and that there exists an integrable function $f:\mathcal{D}\to\mathbb{R}^+$ such that $W_{\min}(\sigma)=\int_{\varphi(\operatorname{conv}\sigma)}f$. Consider the d-chain $\gamma_{\min}$ on K defined by*

$$\gamma_{\min}(\sigma)=\begin{cases}\operatorname{sign}_{\mathcal{D}}(\sigma) & \text{if } W_{\min}(\sigma)=W(\sigma),\\ 0 & \text{otherwise.}\end{cases}$$

Suppose that $\sum_{\sigma\in K^{[d]}}\gamma_{\min}(\sigma)\operatorname{sign}_{\mathcal{D}}(\sigma)\mathbf{1}_{\varphi(\operatorname{conv}\sigma)}(x)=1$, for almost all $x\in\mathcal{D}$. Then, $\gamma_{\min}$ is the unique solution to the following optimization problem over the set of chains in $C_d(K,\mathbb{R})$:

$$\begin{aligned} &\underset{\gamma}{\text{minimize}} \quad \|\gamma\|_{1,W} \\ &\text{subject to} \quad \sum_{\sigma \in K^{[d]}} \gamma(\sigma)\, \mathrm{sign}_{\mathcal{D}}(\sigma) \mathbf{1}_{\varphi(\mathrm{conv}\,\sigma)}(x) = 1, \ \textit{for almost all } x \in \mathcal{D} \end{aligned}$$

Proof We note that the problem is invariant under change of orientation of d-simplices in K and thus we may assume that every d-simplex σ in K has an orientation that is consistant with that of $\mathcal{D}$, that is, $\mathrm{sign}_{\mathcal{D}}(\sigma) = 1$ for all $\sigma \in K^{[d]}$. With this assumption, the lemma simply asserts the following. Consider the d-chain γ_{min} on K defined by

$$\gamma_{\mathrm{min}}(\sigma) = \begin{cases} 1 & \text{if } W_{\mathrm{min}}(\sigma) = W(\sigma), \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that $\sum_{\sigma \in K^{[d]}} \gamma_{\mathrm{min}}(\sigma) \mathbf{1}_{\varphi(\mathrm{conv}\,\sigma)}(x) = 1$, for almost all $x \in \mathcal{D}$. Then the ℓ_1-like norm $\|\gamma\|_{1,W}$ attains its minimum over all d-chains γ such that

$$\sum_{\sigma \in K^{[d]}} \gamma(\sigma) \mathbf{1}_{\varphi(\mathrm{conv}\,\sigma)}(x) = 1, \quad \text{for almost all } x \in \mathcal{D} \tag{6}$$

if and only if $\gamma = \gamma_{\mathrm{min}}$.

We write $\tilde{\sigma} = \varphi(\mathrm{conv}\,\sigma)$ throughout the proof for a shorter notation. We prove the lemma by showing that for all d-chains γ on K that satisfy constraint (6), we have:

$$\|\gamma\|_{1,W} \ \geq \ \|\gamma\|_{1,W_{\mathrm{min}}} \ \geq \ \int_{\mathcal{D}} f \ = \ \|\gamma_{\mathrm{min}}\|_{1,W_{\mathrm{min}}} \ = \ \|\gamma_{\mathrm{min}}\|_{1,W}, \tag{7}$$

with the first inequality being an equality if and only if $\gamma = \gamma_{\mathrm{min}}$. Clearly, $\|\gamma\|_{1,W} \geq \|\gamma\|_{1,W_{\mathrm{min}}}$ because $W(\sigma) \geq W_{\mathrm{min}}(\sigma)$. To obtain the second inequality, recall that we have assumed $\sum_\sigma \gamma(\sigma)\mathbf{1}_{\tilde{\sigma}}(x) = 1$ almost everywhere in $\mathcal{D}$. We use this to write that:

$$\|\gamma\|_{1,W_{\mathrm{min}}} \ \geq \ \sum_\sigma \gamma(\sigma) \int_{\tilde{\sigma}} f \ = \ \sum_\sigma \gamma(\sigma) \int_{\mathcal{D}} f \mathbf{1}_{\tilde{\sigma}} \ = \ \int_{\mathcal{D}} f \sum_\sigma \gamma(\sigma) \mathbf{1}_{\tilde{\sigma}} \ = \ \int_{\mathcal{D}} f, \tag{8}$$

where sums are over all d-simplices σ in K. Setting $\gamma = \gamma_{\mathrm{min}}$ in (8), we observe that the inequality in (8) becomes an equality because none of the coefficients of γ_{min} are negative by construction. It follows that $\int_{\mathcal{D}} f = \|\gamma_{\mathrm{min}}\|_{1,W_{\mathrm{min}}}$. Finally, $\|\gamma_{\mathrm{min}}\|_{1,W_{\mathrm{min}}} = \|\gamma_{\mathrm{min}}\|_{1,W}$ because γ_{min} has been defined so that for all simplices σ in its support, $W_{\mathrm{min}}(\sigma) = W(\sigma)$. We have thus established (7). Suppose now that $\gamma \neq \gamma_{\mathrm{min}}$ and let us prove that $\|\gamma\|_{1,W} > \|\gamma\|_{1,W_{\mathrm{min}}}$, or equivalently that

$$\sum_{\sigma \in \mathrm{Supp}\,\gamma} |\gamma(\sigma)| \left(W(\sigma) - W_{\mathrm{min}}(\sigma)\right) > 0.$$

Since none of the terms in the above sum are negative, it suffices to show that there exists at least one simplex $\sigma \in \mathrm{Supp}\,\gamma$ for which $W(\sigma) > W_{\mathrm{min}}(\sigma)$. By contradiction, assume that for all $\sigma \in \mathrm{Supp}\,\gamma$, $W(\sigma) = W_{\mathrm{min}}(\sigma)$. By construction, we thus have the implication: $\gamma(\sigma) \neq 0 \implies \gamma_{\mathrm{min}}(\sigma) = 1$, and therefore $\mathrm{Supp}\,\gamma \subseteq \mathrm{Supp}\,\gamma_{\mathrm{min}}$. But, since $\sum_\sigma \gamma_{\mathrm{min}}(\sigma)\mathbf{1}_{\tilde{\sigma}}(x) = 1$ for almost all $x \in \mathcal{D}$ and coefficients of γ_{min} are either 0 or 1, it follows that for almost all $x \in \mathcal{D}$, point x is covered by a unique d-simplex in the support of γ_{min}. Hence, the simplices in $\mathrm{Supp}\,\gamma_{\mathrm{min}}$ have pairwise disjoint interiors while their union covers $\mathcal{D}$. Since $\sum_\sigma \gamma(\sigma)\mathbf{1}_{\tilde{\sigma}}(x) = 1$ for almost all $x \in \mathcal{D}$, the simplices in $\mathrm{Supp}\,\gamma$ must also cover $\mathcal{D}$ while using only a subset of simplices in $\mathrm{Supp}\,\gamma_{\mathrm{min}}$. The only possibility is that $\gamma = \gamma_{\mathrm{min}}$, yielding a contradiction. □

6.2 Overview of the proof

We first illustrate the use of the technical lemma by establishing a simple variant of Theorem 7 in which $\mathcal{M}$ has no curvature. We then pinpoint what has to be modified in the proof to establish Theorem 7.

Euclidean setting.

Theorem 9 *Let P be a finite point set of $\mathbb{R}^N$ in general position and let K be a simplicial complex with vertex set P such that $\mathrm{Del}(P) \subseteq K$. Let $\mathcal{M} = \mathrm{conv}\, P$ and let $d = \dim(\mathrm{aff}\, P)$. Then, $\mathrm{code}_{\mathrm{Del}(P)}$ is the unique solution to the following optimization problem over the set of chains in $C_d(K, \mathbb{R})$:*

$$\begin{aligned} &\underset{\gamma}{\text{minimize}} && \mathrm{E}_{\mathrm{del}}(\gamma) \\ &\text{subject to} && \sum_{\sigma \in K^{[d]}} \gamma(\sigma)\, \mathrm{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\mathrm{conv}\,\sigma}(x) = 1, \ \ \textit{for almost all } x \in \mathcal{M} \end{aligned}$$

Proof The proof consists in applying the technical lemma (Lemma 8). In preparation for this, we make the following definitions. Let $\mathcal{D} = \mathcal{M} = \mathrm{conv}\, P = |K|$. Clearly, $\mathcal{D}$ is an orientable d-dimensional submanifold (with boundary). Let φ be the identity map of $\mathcal{D}$ and let $f : \mathcal{D} \to \mathbb{R}^+$ be the map defined by:

$$f(x) = \min_{\sigma} \left(-\mathrm{Power}_{\sigma}(x) \right),$$

where the minimum is taken over all d-simplices $\sigma \in K$ such that $x \in \mathrm{conv}\, \sigma$. Finally, for any $\sigma \in K$, we let $W(\sigma) = \omega(\sigma)$ be the Delaunay weight of σ and define the weight:

$$W_{\min}(\sigma) = \int_{x \in \mathrm{conv}\,\sigma} f(x)\, dx. \tag{9}$$

By construction, $W(\sigma) \geq W_{\min}(\sigma)$. Because P is in general position, all d-simplices of K are non-degenerate and $W_{\min}(\sigma) > 0$. Figure 9 depicts the two weights associated to a simplex β. Consider the d-chain $\gamma_{\min}$ on K:

$$\gamma_{\min}(\sigma) = \begin{cases} \mathrm{sign}_{\mathcal{D}}(\sigma) & \text{if } W_{\min}(\sigma) = W(\sigma), \\ 0 & \text{otherwise.} \end{cases}$$

Before applying the technical lemma, we make three observations, one per step.

Step 1: For any d-simplices $\alpha \in \mathrm{Del}(P)$ and $\beta \in K \setminus \mathrm{Del}(P)$ and for any $x \in \mathrm{conv}\, \alpha \cap \mathrm{conv}\, \beta$, we have that $-\mathrm{Power}_{\alpha}(x) \leq -\mathrm{Power}_{\beta}(x)$ and the inequality is strict whenever $x \notin \mathrm{conv}(\alpha \cap \beta)$ as illustrated in Figure 9.

Step 2: For every point $x \in \mathcal{D}$, we thus have that $f(x) = -\mathrm{Power}_{\alpha}(x)$, where α is any d-simplex of $\mathrm{Del}(P)$ whose convex hull contains x.

Step 3: For all simplices $\sigma \in K$, the following property holds: $W_{\min}(\sigma) = W(\sigma)$ if and only if σ is a Delaunay d-simplex of P.

Hence, $\gamma_{\min} = \mathrm{code}_{\mathrm{Del}(P)}$ and applying Lemma 8 yields the result. □

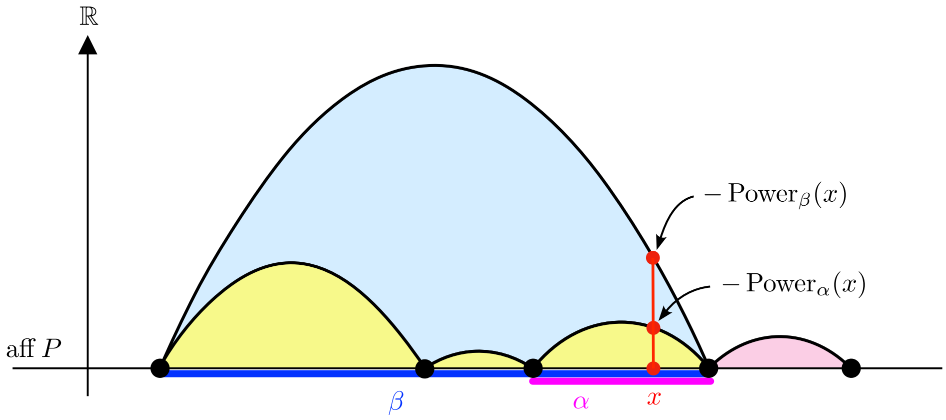

Fig. 9: The two weights $W(\beta)$ and $W_{\min}(\beta)$ defined in the proof of Theorem 9 can be depicted as the volume of the blue region and the volume of the the yellow region, respectively. The blue region is the part of the subgraph of $-\,\mathrm{Power}_\beta$ lying above $\mathrm{conv}\,\beta$ and the yellow region is the part of the subgraph of f lying above $\mathrm{conv}\,\beta$.

Adapting the proof to the submanifold setting.

In the submanifold setting, we let $\mathcal{M}$ be an orientable C^2 d-dimensional submanifold of $\mathbb{R}^N$ and let P be a finite δ-accurate ε-dense sample of $\mathcal{M}$. We also consider a simplicial complex K with vertex set P and let $\rho \geq 0$ be a scale parameter. Assuming that $\mathcal{M}$, P ε, δ and ρ satisfy the assumptions of Theorem 7, we now give an overview of our proof of Theorem 7.

The proof consists in applying the technical lemma, following for this the same steps as in the proof of Theorem 9. However, the different steps are now more involved as is the definition of the various objects required to apply the technical lemma. First of all, we introduce in Section 6.3 a simplicial complex, called the *Delloc complex* of P at scale ρ and denoted as $\mathrm{Delloc}_d(P, \rho)$. This complex is going to act as the counterpart of the Delaunay complex of P in the Euclidean setting. In particular, we aim at showing that the unique solution to Problem ($\star$) is $\mathrm{code}_{\mathrm{Delloc}_d(P,\rho)}$ (instead of $\mathrm{code}_{\mathrm{Del}(P)}$ in the Euclidean setting).

We first show in Section 6.3 that, under the assumptions of Theorem 7, $|\,\mathrm{Delloc}_d(P, \rho)|$ is a faithful reconstruction of $\mathcal{M}$ (Theorem 10) and $\mathrm{Delloc}_d(P, \rho) \subseteq K$ (Remark 2). The set $\mathcal{D} = |\,\mathrm{Delloc}_d(P, \rho)|$ is thus an orientable d-dimensional submanifold of $\mathbb{R}^N$, which is going to play the role of $\mathrm{conv}\,P$ in the Euclidean setting. Next we define $\varphi : |K| \to \mathcal{D}$ which maps any $y \in |K|$ to the point $x \in \mathcal{D}$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. We then define the map $f : \mathcal{D} \to \mathbb{R}^+$ by

$$f(x) = \min_{\sigma} \left(-\,\mathrm{Power}_\sigma(y)\right),$$

where the minimum is taken over all d-simplices $\sigma \in K$ and all points $y \in \mathrm{conv}\,\sigma$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. Finally, for any $\sigma \in K$, we let $W(\sigma) = \omega(\sigma)$ be the Delaunay weight of σ and define the weight:

$$W_{\min}(\sigma) = \int_{x \in \varphi(\operatorname{conv} \sigma)} f(x)\, dx. \tag{10}$$

By construction, $W(\sigma) \geq W_{\min}(\sigma)$ and because every d-simplex $\sigma \in K$ is non-degenerate, $W_{\min}(\sigma) > 0$. Figure 10 depicts the two weights associated to a simplex β.

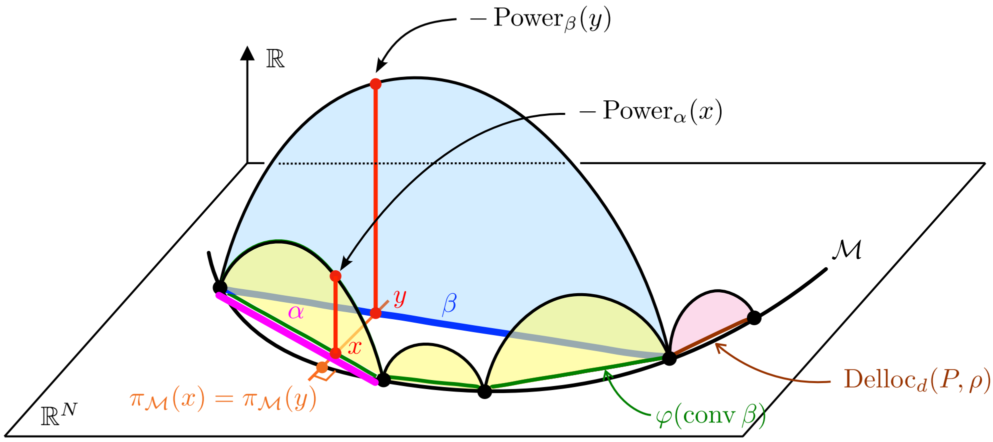

Fig. 10: The two weights $W(\beta)$ and $W_{\min}(\beta)$ used in the proof of Theorem 7 can be depicted as the volume of the blue region and the volume of the yellow region, respectively. The complex $\operatorname{Delloc}_d(P, \rho)$ is formed of the five black dots, the brown segment and the three green segments. The green polygonal chain represents $\varphi(\operatorname{conv} \beta)$.

The first step in the proof of Theorem 9 shows that $-\operatorname{Power}_\alpha(x) \leq -\operatorname{Power}_\beta(x)$ for any d-simplices $\alpha \in \operatorname{Del}(P)$ and $\beta \subseteq P$ and for any $x \in \operatorname{conv} \alpha \cap \operatorname{conv} \beta$. One difficulty in the submanifold setting is that now d-simplices of K do not lie anymore in the same d-dimensional affine space as illustrated in Figure 10. Nonetheless, we can still compare $\operatorname{Power}_\alpha(x)$ and $\operatorname{Power}_\beta(y)$ for two d-simplices $\alpha \in \operatorname{Delloc}_d(P, \rho)$ and $\beta \subseteq P$, and for two points $x \in \operatorname{conv} \alpha$ and $y \in \operatorname{conv} \beta$, assuming that they share the same projection onto $\mathcal{M}$, *i.e.* assuming that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. This is the goal of Section 6.4.

In a second step, we establish that for every point $x \in \mathcal{D}$, we have that $f(x) = -\operatorname{Power}_\alpha(x)$, where α is any d-simplex of $\operatorname{Delloc}_d(P, \rho)$ whose convex hull contains x (Lemma 15 in Section 6.5).

In a third step, we check that we have defined the two weight functions W and $W_{\min}$ in such a way that $W(\sigma) = W_{\min}(\sigma)$ if and only if σ belongs to the Delloc complex of P. This is done in Section 6.5 thanks to Lemmas 15 and 16. The tricky part consists in showing that $W_{\min}(\beta) < W(\beta)$ whenever β is not in the Delloc complex of P. Indeed, to compare the two quantities, we make a change of variable which possibly can jeopardize the strict inequality but which we are able to counterbalance, thanks to the conditions we are assuming on P (and in particular a sufficient protection of P).

Finally, all the elements are put together in Section 6.5 where one can find the proof of the main theorem.

6.3 Delloc complexes

In this section, we define the Delloc complex. We then recall a key result established in the companion paper [1]: when the Delloc complex is computed over a finite point set P that samples some d-dimensional submanifold of $\mathbb{R}^N$, it provides a faithful reconstruction of that submanifold. Incidentally, under the right assumptions, the Delloc complex coincides with the flat Delaunay complex [1] and the tangential Delaunay complex [15, 16]. Since all the results in this paper are based on the property for a simplex to belong to the Delloc complex, we find it more enlightening to formulate the results of this paper using the Delloc complex.

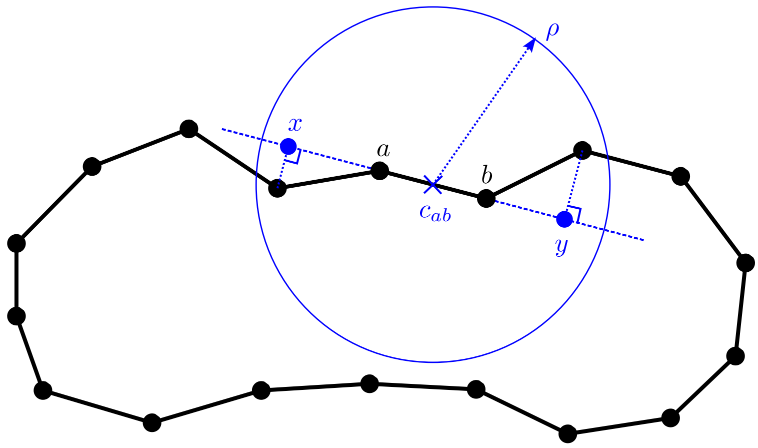

Fig. 11: A set of black dots and its 1-dimensional Delloc complex at scale ρ. The edge ab belongs to the Delloc complex because ab is a Delaunay edge of the four points $\{a, b, x, y\}$ obtained by projecting the four black dots inside $B(c_{ab}, \rho)$ onto the line passing through a and b.

Definition.

Afterwards, P designates a finite set of points in $\mathbb{R}^N$, d designates an integer in $[0, N)$ and $\rho \geq 0$ designates a scale parameter.

Definition 10 (Delloc complex) We say that a simplex σ is *delloc* in P at scale ρ if

$$\sigma \in \mathrm{Del}(\pi_{\mathrm{aff}\,\sigma}(P \cap B(c_\sigma, \rho))),$$

where c_σ denotes the center of the smallest N-ball enclosing σ. The d-dimensional *Delloc complex* of P at scale ρ, denoted by $\mathrm{Delloc}_d(P, \rho)$, is the set of d-simplices that are delloc in P at scale ρ together with all their faces; see Figure 11.

Remark 1 It is easy to see that if $2R(\sigma) \leq \rho$, then the smallest circumsphere $S(\sigma)$ of σ is contained in $B(c_\sigma, \rho)$. It follows that a delloc simplex σ in P at scale ρ is also a *Gabriel simplex* of P, by which we mean that $S(\sigma)$ does not enclose any point of P in its interior. In particular, when $2R(\sigma) \leq \rho$, then σ is a Delaunay simplex of P.

Key result.

We now recall a key result established in the companion paper [1] and which gives condition under which the Delloc complex of P is a faithful reconstruction of $\mathcal{M}$.

Theorem 10 (Faithful reconstruction by a geometric approach) *Let ε, δ, and ρ be non-negative real numbers such that $\delta \leq \varepsilon$ and $16\varepsilon \leq \rho < \frac{\mathcal{R}}{4}$. Let P be a δ-accurate ε-dense sample of $\mathcal{M}$ such that* $\text{height}(P, \rho) > 0$*. Suppose that the safety condition on (P, ε, δ) is satisfied at scale ρ. Then,* $\text{Delloc}_d(P, \rho)$ *is a faithful reconstruction of $\mathcal{M}$. Furthermore, for all d-simplices $\sigma \in \text{Delloc}_d(P, \rho)$, we have $R(\sigma) \leq \varepsilon$.*

Remark 2 Under the assumptions of Theorem 10, Remark 1 implies that the Delloc complex of P at scale ρ is a subset of the Delaunay complex of P, that is, $\text{Delloc}_d(P, \rho) \subseteq \text{Del}(P)$. It follows that under the assumptions of Theorem 10:

$$\text{Delloc}_d(P, \rho) \subseteq \text{Del}(P) \cap \mathcal{C}(P, \varepsilon)$$

and therefore any simplicial complex K that satisfies the assumptions of Theorem 7 contains $\text{Delloc}_d(P, \rho)$.

6.4 Comparing power distances

The goal of this section is to relate the two maps $\text{Power}_\alpha(x)$ and $\text{Power}_\beta(y)$ for two d-simplices $\alpha \in \text{Delloc}_d(P, \rho)$ and $\beta \subseteq P$, and for two points $x \in \text{conv}\,\alpha$ and $y \in \text{conv}\,\beta$, such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. The main result of the section is stated in the following lemma and proved at the end of the section. We recall that given a non-degenerate simplex α and a point $x \in \text{aff}\,\sigma$, the (normalized) barycentric coordinates of x relatively to the simplex α are real numbers $\{\lambda_a\}_{a \in \alpha}$ such that $x = \sum_{a \in \alpha} \lambda_a a$ and $\sum_{a \in \alpha} \lambda_a = 1$. We write

$$\text{BarycentricCoord}_a^\alpha(x) = \lambda_a$$

Lemma 11 *Let $\varepsilon, \delta, \rho \geq 0$ such that $0 \leq 2\varepsilon \leq \rho$, and $16\delta \leq \rho \leq \frac{\mathcal{R}}{3}$. Suppose that $P \subseteq \mathcal{M}^{\oplus \delta}$. Let $\mathsf{p} = \text{protection}(P, 3\rho)$, $\mathsf{s} = \text{separation}(P)$, and $\Theta = \text{angularDeviation}_{\mathcal{M}}(P, \rho)$. Assume that $\Theta \leq \frac{\pi}{6}$ and*

$$10\rho\,\Theta\,(\varepsilon + \rho\,\Theta) < \mathsf{p}^2 + \mathsf{ps}.$$

Then, for every non-degenerate ε-small d-simplex $\alpha \in \text{Delloc}_d(P, \rho)$, every non-degenerate ρ-small d-simplex $\beta \subseteq P$, every $x \in \text{conv}\,\alpha$, and every $y \in \text{conv}\,\beta$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$:

$$-\text{Power}_\beta(y) + \text{Power}_\alpha(x) \geq \frac{1}{2}\left(\mathsf{p}^2 + \mathsf{ps}\right) \sum_{b \in \beta \setminus \alpha} \text{BarycentricCoord}_b^\beta(y).$$

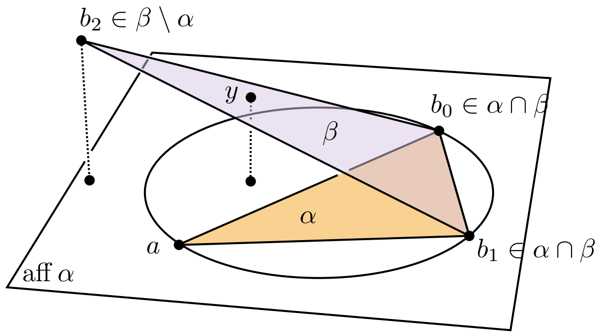

Fig. 12: Notation for the proof of Lemma 12.

Lemma 12 *Let α and β be two non-degenerate abstract d-simplices in $\mathbb{R}^N$ such that $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\alpha}(\alpha \cup \beta))$. Let $\mathsf{p} = \mathrm{protection}(\alpha, \beta)$. Then for every $y \in \mathrm{conv}\,\beta$, we have*

$$\mathrm{Power}_\beta(y) \leq \mathrm{Power}_\alpha(\pi_{\mathrm{aff}\,\alpha}(y)) - (\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b \in \beta \setminus \alpha} \mathrm{BarycentricCoord}_b^\beta(y).$$

Proof See Figure 12. Let $Z(\alpha)$ be the radius of the $(d-1)$-dimensional circumsphere of α. Clearly, $\|a - Z(\alpha)\| = R(\alpha)$ for all $a \in \alpha$. Since $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\alpha}(\alpha \cup \beta))$ and $\mathsf{p} = \mathrm{protection}(\alpha, \beta)$, we get:

$$\begin{aligned} (R(\alpha) + \mathsf{p})^2 &\leq \|\pi_{\mathrm{aff}\,\alpha}(b) - Z(\alpha)\|^2, \qquad \text{for all } b \in \beta \setminus \alpha, \\ R(\alpha)^2 &= \|\pi_{\mathrm{aff}\,\alpha}(b) - Z(\alpha)\|^2, \qquad \text{for all } b \in \beta \cap \alpha. \end{aligned}$$

Let $\mu_b = \mathrm{BarycentricCoord}_b^\beta(y)$ and note that $\mu_b \geq 0$. Multiplying both sides of each equation above by μ_b and summing over all $b \in \beta$, we obtain:

$$R(\alpha)^2 + (\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b \in \beta \setminus \alpha} \mu_b \leq \sum_{b \in \beta} \mu_b \|\pi_{\mathrm{aff}\,\alpha}(b) - Z(\alpha)\|^2. \tag{11}$$

For short, write $y' = \pi_{\mathrm{aff}\,\alpha}(y)$ and $\beta' = \pi_{\mathrm{aff}\,\alpha}(\beta)$. Noting that $y' = \sum_{b \in \beta} \mu_b b'$ and applying Lemma 3 with $z = Z(\alpha)$, we get that

$$\mathrm{Power}_{\beta'}(y') = \|y' - Z(\alpha)\|^2 - \sum_{b \in \beta} \mu_b \|\pi_{\mathrm{aff}\,\alpha}(b) - Z(\alpha)\|^2.$$

Substracting $\|y' - Z(\alpha)\|^2$ from both sides of (11) and using the above expression, we obtain

$$-\mathrm{Power}_\alpha(y') + (\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b \in \beta \setminus \alpha} \mu_b \leq -\mathrm{Power}_{\beta'}(y').$$

Applying Lemma 3 again, with $Z = y'$ and $Z = y$ respectively, we get that:

$$-\mathrm{Power}_{\beta'}(y') = \sum_{b \in \beta} \mu_b \|\pi_{\mathrm{aff}\,\alpha}(b) - \pi_{\mathrm{aff}\,\alpha}(y)\|^2 \leq \sum_{b \in \beta} \mu_b \|b - y\|^2 = -\mathrm{Power}_\beta(y),$$

which concludes the proof. □

Lemma 13 *Let α and β be two non-degenerate abstract d-simplices in $\mathbb{R}^N$ such that $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\alpha}(\alpha \cup \beta))$. Let $\mathsf{p} = \mathrm{protection}(\alpha, \beta)$ and $\Theta = \mathrm{angularDeviation}_{\mathcal{M}}(\alpha)$. Suppose that both $\mathrm{conv}\,\alpha$ and $\mathrm{conv}\,\beta$ are contained in the $\left(\frac{\rho}{4}\right)$-tubular neighborhood of $\mathcal{M}$ with $\rho < 4\mathcal{R}$. Suppose furthermore that β is ρ-small. If $\Theta \leq \frac{\pi}{6}$ and*

$$5\rho \sin(\Theta) \left(2R(\alpha) + \frac{\rho}{2}\sin(\Theta)\right) < \mathsf{p}^2 + 2\mathsf{p}R(\alpha),$$

then for every $x \in \mathrm{conv}\,\alpha$ and every $y \in \mathrm{conv}\,\beta$ with $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$, we have

$$\mathrm{Power}_\beta(y) \leq \mathrm{Power}_\alpha(x) - \frac{1}{2}(\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b \in \beta\setminus\alpha} \mathrm{BarycentricCoord}_b^\beta(y).$$

Proof Consider a point $x \in \mathrm{conv}\,\alpha$ and a point $y \in \mathrm{conv}\,\beta$ with $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. We distinguish two cases depending on whether y belongs to $\mathrm{conv}(\alpha \cap \beta)$ or not.

First, assume that $y \in \mathrm{conv}(\alpha \cap \beta)$. In that case, we claim that the only possibility is that $x = y$. Indeed, assume for a contradiction that this is not the case. Then, we would have two distinct points $x \neq y$ of $\mathrm{conv}\,\alpha$ that share the same projection onto $\mathcal{M}$, showing that $\angle(\mathrm{aff}\,\alpha, \mathbf{T}_{\pi_{\mathcal{M}}(x)}\mathcal{M}) = \frac{\pi}{2}$ for some $x \in \mathrm{conv}\,\alpha$ and contradicting our assumption that $\Theta < \frac{\pi}{6}$. Hence, $x = y \in \mathrm{conv}(\alpha \cap \beta)$. We claim that furthermore $\mathrm{Power}_\alpha(x) = \mathrm{Power}_\beta(y)$. Indeed, Lemma 3 implies that when x is an affine combination of points in α, that is, when $x = \sum_{a\in\alpha} \lambda_a a$ with $\sum_a \lambda_a = 1$, then $\mathrm{Power}_\alpha(x) = -\sum_{a\in\alpha} \lambda_a \|x - a\|^2$. In particular, if x belongs to the convex hull of a face of α, the expression of the power distance depends only upon the vertices of that face. It follows that

$$\mathrm{Power}_\alpha(x) = \mathrm{Power}_{\alpha\cap\beta}(x) = \mathrm{Power}_{\alpha\cap\beta}(y) = \mathrm{Power}_\beta(y).$$

Since $y \in \mathrm{conv}(\alpha \cap \beta)$, we have $\sum_{b\in\beta\setminus\alpha} \mathrm{BarycentricCoord}_b^\beta(y) = 0$ and combining this with the above equality, we get the desired inequality.

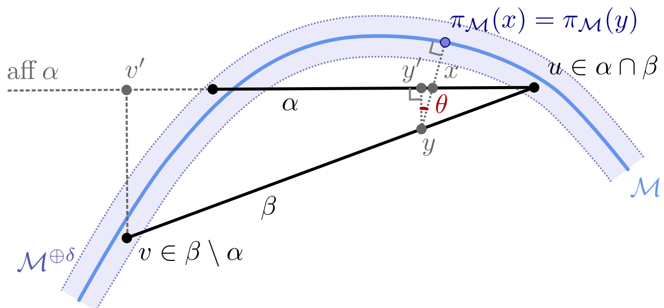

Fig. 13: Right: Notation for the proof of Lemma 13. We know by Lemma 14 that if $\alpha, \beta \subseteq \mathcal{M}^{\oplus\delta}$ with $\delta \leq \frac{\rho}{16}$, then $\mathrm{conv}\,\alpha, \mathrm{conv}\,\beta \subseteq \mathcal{M}^{\oplus\frac{\rho}{4}}$ as assumed in Lemma 13.

Second, assume that $y \in \mathrm{conv}\,\beta \setminus \mathrm{conv}(\alpha\cap\beta)$; see Figure 13. Write $\mu_b = \mathrm{BarycentricCoord}_b^\beta(y)$ and note that $\mu_b \geq 0$. Letting $y' = \pi_{\mathrm{aff}\,\alpha}(y)$, we know by Lemma 12 that:

$$\mathrm{Power}_\beta(y) \leq \mathrm{Power}_\alpha(y') - (\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b\in\beta\setminus\alpha} \mu_b. \tag{12}$$

Because $y \notin \operatorname{conv}(\alpha \cap \beta)$, we have $\sum_{b\in\alpha\cap\beta} \mu_b \neq 1$ and therefore $\sum_{b\in\beta\setminus\alpha} \mu_b \neq 0$. First, suppose that $y = x$. In that case, $y' = x$ and the result follows immediately. Second, suppose that $y \neq x$. We claim that in that case we also have $y \neq y'$. Indeed, if we were to have that $y = y'$, then both x and y would belong to $\operatorname{aff} \alpha$ and since $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$, this would mean that $\angle(\operatorname{aff} \alpha, \mathbf{T}_{\pi_{\mathcal{M}}(x)}\mathcal{M}) = \frac{\pi}{2}$ for $x \in \operatorname{conv} \alpha$, contradicting our assumption that $\Theta < \frac{\pi}{6}$. Thus, $x \neq y$ and $y \neq y'$, and we can define the angle $\theta = \angle xyy'$. Noting that $\theta \leq \angle(\operatorname{aff} \alpha, \mathbf{T}_{\pi_{\mathcal{M}}(x)}\mathcal{M}) \leq \Theta$ and $\|x - y'\| = \|x - y\| \sin\theta$, and recalling that $Z(\alpha)$ is the radius of the $(d-1)$-dimensional circumsphere of α, we have:

$$\begin{aligned}
\operatorname{Power}_\alpha(y') - \operatorname{Power}_\alpha(x) &= \|y' - Z(\alpha)\|^2 - \|x - Z(\alpha)\|^2 = (y' - x)\cdot(x + y' - 2Z(\alpha)) \\
&\leq \|x - y'\| \cdot \big(\|x - Z(\alpha)\| + \|y' - Z(\alpha)\|\big) \\
&\leq \|x - y'\| \cdot \big(2\|x - Z(\alpha)\| + \|x - y'\|\big) \\
&\leq \|x - y\| \sin(\theta)\,(2R(\alpha) + \|x - y\|\sin(\theta))\,. \qquad (13)
\end{aligned}$$

Writing $m = \pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$, we have $\|x - y\| \leq \|x - m\| + \|m - y\| \leq \frac{\rho}{2}$. Summing up Inequalities (12) and (13), we get

$$\operatorname{Power}_\beta(y) - \operatorname{Power}_\alpha(x) \leq -\underbrace{(\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b\in\beta\setminus\alpha} \mu_b}_{A} + \underbrace{\|x - y\| \sin(\Theta) \left(2R(\alpha) + \frac{\rho}{2}\sin(\Theta)\right)}_{B}.$$

To establish the lemma in the second case, it suffices to show that $2B < A$, that is,

$$2\|x - y\| \sin\Theta \cdot \left(2R(\alpha) + \frac{\rho}{2}\sin\Theta\right) < (\mathsf{p}^2 + 2\mathsf{p}R(\alpha)) \sum_{b\in\beta\setminus\alpha} \mu_b. \qquad (14)$$

We consider two subcases:

Subcase 1: $\alpha \cap \beta = \emptyset$. In that case, $\sum_{b\in\beta\setminus\alpha} \mu_b = 1$, and because $2\|x - y\| \leq 4\rho \leq 5\rho$, one can see that (14) follows from our assumptions.

Subcase 2: $\alpha \cap \beta \neq \emptyset$. In that case, we know that there exists a point $u \in \operatorname{conv}(\beta \cap \alpha)$ and a point $v \in \operatorname{conv}(\beta \setminus \alpha)$ such that $y = \sum_{b\in\beta\cap\alpha} \mu_b u + \sum_{b\in\beta\setminus\alpha} \mu_b v$; see Figure 13. Furthermore, letting $v' = \pi_{\operatorname{aff} \alpha}(v)$ we have

$$\sum_{b\in\beta\setminus\alpha} \mu_b = \frac{\|y - u\|}{\|v - u\|} \geq \frac{\|y - y'\|}{\|v - u\|} \geq \frac{\|x - y\|\cos\theta}{\operatorname{Diam}(\beta)} \geq \frac{\|x - y\|\cos\theta}{2\rho} \geq \frac{\sqrt{3}}{4\rho} \cdot \|x - y\|.$$

Again, (14) follows from our assumptions. □

The next lemma says that if a subset $\sigma \subseteq \mathbb{R}^N$ is sufficiently small and sufficiently close to a subset $A \subseteq \mathbb{R}^N$ compare to the reach of A, then the convex hull of σ is not too far away from A.

Lemma 14 *Let $16\delta \leq \rho \leq \frac{\operatorname{reach} A}{3}$. If the subset $\sigma \subseteq A^{\oplus\delta}$ is ρ-small, then $\operatorname{conv}\sigma \subseteq A^{\oplus\frac{\rho}{4}}$.*

Proof Let $\mathcal{R} = \operatorname{reach} A$. Applying Lemma 14 in [36], we get that $\operatorname{conv}\sigma \subseteq A^{\oplus r}$ for $r = \mathcal{R} - \sqrt{(\mathcal{R} - \delta)^2 - \rho^2}$. Since $\delta \leq \frac{\rho}{16}$, we deduce that $\frac{r}{\mathcal{R}} \leq 1 - \sqrt{\left(1 - \frac{\rho}{16\mathcal{R}}\right)^2 - \frac{\rho}{\mathcal{R}}^2}$ and since for all $0 \leq t \leq \frac{1}{3}$ we have $1 - \sqrt{\left(1 - \frac{t}{16}\right)^2 - t^2} \leq \frac{t}{4}$, we obtain the result. □

We are now ready to prove Lemma 11.

Proof of Lemma 11 Let α be a non-degenerate ε-small d-simplex of $\mathrm{Delloc}_d(P, \rho)$. Because $\alpha \in \mathrm{Delloc}_d(P, \rho)$, we have that $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\sigma}(P \cap B(c_\sigma, \rho)))$, and because α is ε-small, we have that $B(Z(\alpha), R(\alpha)) \subseteq B(c_\sigma, 2\varepsilon) \subseteq B(c_\sigma, \rho)$ and consequently $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\sigma}(P \cap B(c_\sigma, 3\rho)))$.

Let β be a non-degenerate ρ-small d-simplex of P. Assume that $\pi_{\mathcal{M}}(\alpha) \cap \pi_{\mathcal{M}}(\beta) \neq \emptyset$ and let us show that $\beta \subseteq P \cap B(c_\sigma, 3\rho)$. Suppose that $x \in \mathrm{conv}\,\alpha$ and $y \in \mathrm{conv}\,\beta$ share the same projection m onto $\mathcal{M}$, that is, $m = \pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. Since both α and β are ρ-small, Lemma 14 implies that both $\mathrm{conv}\,\alpha$ and $\mathrm{conv}\,\beta$ are contained in the $\left(\frac{\rho}{4}\right)$-tubular neighborhood of $\mathcal{M}$ and in particular $\|x - y\| \leq \|x - m\| + \|m - y\| \leq \frac{\rho}{4} + \frac{\rho}{4} \leq \frac{\rho}{2}$. For all vertices $b \in \beta$, we thus have

$$\|c_\alpha - b\| \leq \|c_\alpha - x\| + \|x - y\| + \|y - b\| \leq \varepsilon + \frac{\rho}{2} + 2\rho \leq 3\rho,$$

showing that $\beta \subseteq P \cap B(c_\sigma, 3\rho)$. Hence, we get that $\alpha \in \mathrm{Del}(\pi_{\mathrm{aff}\,\sigma}(\alpha \cup \beta))$ and can easily see that $\mathsf{p} = \mathrm{protection}(P, 3\rho) \leq \tilde{\mathsf{p}} = \mathrm{protection}(\alpha, \beta)$. Let $\tilde{\Theta} = \mathrm{angularDeviation}_{\mathcal{M}}(\alpha) \leq \Theta$. To apply Lemma 13, we need to verify that

$$5\rho \sin(\tilde{\Theta}) \left(2R(\alpha) + \frac{\rho}{2} \sin(\tilde{\Theta})\right) < \tilde{\mathsf{p}}^2 + 2\tilde{\mathsf{p}}R(\alpha).$$

Since $\frac{\mathsf{s}}{2} \leq R(\alpha) \leq \varepsilon$ and $\sin t \leq t$ for all $t \geq 0$, this follows from:

$$10\rho\Theta\left(\varepsilon + \rho\Theta\right) < \mathsf{p}^2 + \mathsf{ps},$$

which is a consequence of our hypotheses. □

6.5 Final

Suppose that K is a simplicial complex with vertex set P. Write $D = \mathrm{Delloc}_d(P, \rho)$, $\mathcal{D} = |D|$ and $\mathcal{K} = |K|$ for short. In this section, we prove our main theorem by applying Lemma 8. This requires us to define two maps $\varphi : \mathcal{K} \to \mathcal{D}$ and $f : \mathcal{D} \to \mathbb{R}^+$, two weights $W(\alpha)$ and $W_{\min}(\alpha)$ for each d-simplex $\alpha \in K$, and to check that these maps and weights satisfy the requirements of Lemma 8. For each $\alpha \in K$, let $W(\alpha) = \omega(\alpha)$ be the Delaunay weight of α. To be able to define φ, f, and $W_{\min}$, we assume that the following conditions are met:

- D is a faithful reconstruction of $\mathcal{M}$;
- For every d-simplex $\sigma \subseteq K$, the map $\pi_{\mathcal{M}}\big|_{\mathrm{conv}\,\sigma}$ is well-defined and injective.

These conditions are easily derived from the assumptions of the main theorem. We are now ready to introduce additional notation. Consider a subset $X \subseteq \mathbb{R}^N$ and suppose that the map $\pi_{\mathcal{M}}\big|_X$ is well-defined and injective. Then it is possible to define a bijective map $\pi_{X \to \mathcal{M}} : X \to \pi_{\mathcal{M}}(X)$. Because D is a faithful reconstruction of $\mathcal{M}$, the map $\pi_{\mathcal{D} \to \mathcal{M}}$ is well-defined and bijective. Similarly, for every d-simplex $\sigma \in K$, the map $\pi_{\mathrm{conv}\,\sigma \to \mathcal{M}}$ is well-defined and bijective. We now introduce the map $\varphi : \mathcal{K} \to \mathcal{D}$ defined by $\varphi = [\pi_{\mathcal{D} \to \mathcal{M}}]^{-1} \circ \pi_{\mathcal{M}}$ and let $f : \mathcal{D} \to \mathbb{R}^+$ be the map defined by:

$$f(x) = \min_{\sigma} \left(-\,\mathrm{Power}_\sigma([\pi_{\mathrm{conv}\,\sigma \to \mathcal{M}}]^{-1} \circ \pi_{\mathcal{M}}(x))\right), \tag{15}$$

where the minimum is taken over all d-simplices $\sigma \in K$ such that $x \in \varphi(\operatorname{conv} \sigma)$. Note that $f(x)$ can be defined equivalently as the minimum of $-\operatorname{Power}_\beta(y)$ over all d-simplices $\beta \in K$ and all points $y \in \operatorname{conv} \beta$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. Given a d-simplex $\sigma \in K$, we associate to σ the weight:

$$W_{\min}(\sigma) = \int_{x \in \varphi(\operatorname{conv} \sigma)} f(x)\, dx. \tag{16}$$

Lemma 15 *Under the assumptions of Theorem 7:*

- *For every d-simplex $\alpha \in D$ and every point $x \in \operatorname{conv} \alpha$, we have $f(x) = -\operatorname{Power}_\alpha(x)$.*
- *For every d-simplex $\alpha \in D$, we have $W_{\min}(\alpha) = W(\alpha)$.*

Proof Consider a d-simplex $\alpha \in D$, a d-simplex $\beta \in K$, $x \in \operatorname{conv} \alpha$ and $y \in \operatorname{conv} \beta$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. Applying Lemma 11, we obtain that $\operatorname{Power}_\beta(y) \leq \operatorname{Power}_\alpha(x)$ or equivalently $\operatorname{Power}_\beta([\pi_{\operatorname{conv} \beta \to \mathcal{M}}]^{-1} \circ \pi_{\mathcal{M}}(x)) \leq \operatorname{Power}_\alpha(x)$ and therefore $f(x) = -\operatorname{Power}_\alpha(x)$. To establish the second item of the lemma, notice that for all $\alpha \in D$, the restriction of φ to $\operatorname{conv} \alpha$ is the identity function, $\varphi_{|\operatorname{conv} \alpha} = \operatorname{Id}$ and therefore $\varphi(\operatorname{conv} \alpha) = \operatorname{conv} \alpha$. Since we have just established that $f(x) = -\operatorname{Power}_\alpha(x)$, we get that

$$W_{\min}(\alpha) = \int_{x \in \varphi(\operatorname{conv} \alpha)} f(x)\, dx = \int_{x \in \operatorname{conv} \alpha} -\operatorname{Power}_\alpha(x)\, dx = \omega(\alpha) = W(\alpha),$$

which concludes the proof. □

Lemma 16 *Under the assumptions of Theorem 7, for every d-simplex $\beta \in K \setminus D$, we have $W_{\min}(\beta) < W(\beta)$.*

Proof We need some notation. Given α and β in K, we write $\operatorname{conv}_{|\alpha} \beta$ for the set of points $y \in \operatorname{conv} \beta$ for which there exists a point $x \in \operatorname{conv} \alpha$ such that $\pi_{\mathcal{M}}(x) = \pi_{\mathcal{M}}(y)$. We define the map $\varphi_{\beta \to \alpha} : \operatorname{conv}_{|\alpha} \beta \to \operatorname{conv}_{|\beta} \alpha$ as $\varphi_{\beta \to \alpha}(y) = [\pi_{\operatorname{conv} \alpha \to \mathcal{M}}]^{-1} \circ \pi_{\operatorname{conv} \beta \to \mathcal{M}}(y)$. Note that $\varphi_{\beta \to \alpha}$ is invertible and its inverse is $\varphi_{\alpha \to \beta}$. Also, note that J in Theorem 7 has been chosen precisely so that one can apply Lemma **??** in [33, Appendix E] and guarantee that $|\det(\mathrm{D}\varphi_{\beta \to \alpha})(y)| \in [\frac{1}{1+J}, 1+J]$ for all $\alpha, \beta \in K$ and all $y \in \operatorname{conv}_{|\alpha} \beta$. Consider a d-simplex $\beta \in K \setminus D$. By Lemma 15, $f(x) = -\operatorname{Power}_\alpha(x)$ and therefore:

$$W_{\min}(\beta) \;=\; \sum_{\alpha \in D^{[d]}} \int_{x \in \operatorname{conv}_{|\beta} \alpha} -\operatorname{Power}_\alpha(x)\, dx.$$

For any convex combination y of points in β, let $\{\lambda_b^\beta(y)\}_{b \in \beta}$ designate the family of non-negative real numbers summing up to 1 such that $y = \sum_{b \in \beta} \lambda_b^\beta(y) b$. Plugging in the upper bound on $-\operatorname{Power}_\alpha(x)$ provided by Lemma 11, letting

$$c = \frac{1}{2}\left(\mathsf{p}^2 + \mathsf{ps}\right),$$

and making the change of variable $x = \varphi_{\beta \to \alpha}(y)$, we upper bound $W_{\min}(\beta)$ as follows:

$$W_{\min}(\beta) \;\leq\; \sum_{\alpha \in D^{[d]}} \int_{x \in \operatorname{conv}_{|\beta} \alpha} \left[-\operatorname{Power}_\beta(\varphi_{\alpha \to \beta}(x)) - c \sum_{b \in \beta \setminus \alpha} \lambda_\beta^b(\varphi_{\alpha \to \beta}(x)) \right] dx$$

$$= \sum_{\alpha \in D^{[d]}} \int_{y \in \text{conv}_{|\alpha} \beta} \left[-\text{Power}_\beta(y) - c \sum_{b \in \beta \setminus \alpha} \lambda_\beta^b(y) \right] |\det(\text{D}\varphi_{\beta \to \alpha})(y)| \, dy$$

$$\leq (1+J)W(\beta) - (1+J)^{-1} c \sum_{\alpha \in D^{[d]}} \int_{y \in \text{conv}_{|\alpha} \beta} \sum_{b \in \beta \setminus \alpha} \lambda_\beta^b(y) \, dy.$$

A key observation is that, because $\beta \neq \alpha$, then $\beta \setminus \alpha \neq \emptyset$. Therefore the sum $\sum_{b \in \beta \setminus \alpha} \lambda_\beta^b(y)$ does not vanish and is always lower bounded by $\inf_{b \in \beta} \lambda_\beta^b(y)$. Associating the quantity

$$\Omega(\beta) = \int_{y \in \text{conv}\, \beta} \inf_{b \in \beta} \lambda_\beta^b(y) \, dy,$$

to β we thus obtain that $W_{\min}(\beta) \leq (1+J)W(\beta) - (1+J)^{-1} c\, \Omega(\beta)$. Hence, $W_{\min}(\beta) < W(\beta)$ as long as

$$JW(\beta) < (1+J)^{-1} c\, \Omega(\beta). \tag{17}$$

Using a change of variable, it is not too difficult to show that $\Omega(\beta) = d! \,\text{vol}(\beta) \Omega(\Delta_d)$, where $\Delta_d = \{\lambda \in \mathbb{R}^d \mid \sum_{i=1}^d \lambda_i \leq 1; \lambda_i \geq 0, i = 1, 2, \ldots, d\}$ represents the standard d-simplex. Remark that $\Omega(\Delta_d)$ is a constant that depends only upon the dimension d and is thus universal. Plugging in $\Omega(\beta) = d! \,\text{vol}(\beta) \Omega(\Delta_d)$ on the right side of (17), and the expression of $W(\beta) = \omega(\beta)$ given by Lemma 5 on the left side of (17), and recalling that β is ρ-small, we find that condition (17) is implied by the following condition:

$$J\rho^2 < (1+J)^{-1} \frac{(d+2)(d-1)!}{4} \left(\mathsf{p}^2 + \mathsf{ps} \right) \Omega(\Delta_d),$$

which we have assumed to hold. □

Proof of Theorem 7 Let $D = \text{Delloc}_d(P, \rho)$, $\mathcal{D} = |D|$ and $\mathcal{K} = |K|$. Theorem 10 ensures that $\mathcal{D}$ is a d-manifold and $\pi_{\mathcal{M}} : \mathcal{D} \to \mathcal{M}$ is a homeomorphism. Give to $\mathcal{D}$ an orientation that is consistent with that of $\mathcal{M}$, that is, $\text{sign}_{\mathcal{D}}(\sigma) = \text{sign}_{\mathcal{M}}(\sigma)$ for all $\sigma \in K^{[d]}$. Define $\varphi : \mathcal{K} \to \mathcal{D}$, $f : \mathcal{D} \to \mathbb{R}^+$, W, and $W_{\min}$ as explained at the beginning of the section. Consider the d-chain $\gamma_{\min}$ on K:

$$\gamma_{\min}(\sigma) = \begin{cases} \text{sign}_{\mathcal{M}}(\sigma) & \text{if } W_{\min}(\sigma) = W(\sigma), \\ 0 & \text{otherwise.} \end{cases}$$

By Lemma 15 and Lemma 16, the following property holds: for all $\sigma \in K$, $W_{\min}(\sigma) = W(\sigma)$ if and only if σ is a d-simplex of D. It follows that $\gamma_{\min} = \text{code}_D$. Furthermore, we have $\sum_{\sigma \in K^{[d]}} \gamma_{\min}(\sigma) \,\text{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\varphi(\text{conv}\,\sigma)}(x) = \sum_{\sigma \in D^{[d]}} \mathbf{1}_{\text{conv}\,\sigma}(x) = 1$ for almost all $x \in \mathcal{D}$. Recalling that $W = \omega$ and therefore $\|\gamma\|_{1,W} = \text{E}_{\text{del}}(\gamma)$, and applying Lemma 8, we deduce that $\gamma_{\min} = \text{code}_D$ is the unique solution to the following optimization problem over the set of chains in $C_d(K, \mathbb{R})$:

$$\begin{aligned} &\underset{\gamma}{\text{minimize}} && \text{E}_{\text{del}}(\gamma) \\ &\text{subject to} && \sum_{\sigma \in K^{[d]}} \gamma(\sigma) \,\text{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\varphi(\text{conv}\,\sigma)}(x) = 1, \quad \text{for almost all } x \in \mathcal{D} \end{aligned} \tag{$\circledast$}$$

We now claim that the feasible set of Problem ($\circledast$) contains the feasible set of Problem ($\star$). Indeed, consider a d-chain γ that satisfies the constaints of Problem ($\star$), that is, such that

$$\begin{cases} \partial\gamma = 0, \\ \sum_{\sigma \in K^{[d]}} \gamma(\sigma) \,\text{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\pi_{\mathcal{M}}(\text{conv}\,\sigma)}(m_0) = 1. \end{cases}$$

Then, by Lemma 33 in [33, Appendix G], we obtain that γ also satisfies the following constraint:

$$\sum_{\sigma \in K^{[d]}} \gamma(\sigma)\, \mathrm{sign}_{\mathcal{M}}(\sigma) \mathbf{1}_{\pi_{\mathcal{M}}(\mathrm{conv}\, \sigma)}(m) = 1, \quad \text{for almost all } m \in \mathcal{M},$$

which is equivalent to the constaint of Problem (⊛). Since the unique solution to Problem (⊛) is code_D, Theorem 10 guarantees that $D = \mathrm{Delloc}_d(P, \rho)$ is a faithful reconstruction of $\mathcal{M}$. By Lemma 6, code_D is thus a cycle. Hence, the unique solution code_D to Problem (⊛) also satisfies the containts of Problem ($\star$) and, because the feasible set of Problem (⊛) contains the feasible set of Problem ($\star$), code_D is also the unique solution to Problem ($\star$). □

7 Practical aspects

In this section, we discuss practical aspects.

7.1 Transforming the problem into a realistic algorithm

Besides the complex K that one can build from P, Problem ($\star$) seems to require the knowledge of $\mathcal{M}$ for expressing the normalization constraint $\mathrm{load}_{m_0,\mathcal{M}}(\gamma) = 1$. What we call a *realistic* algorithm is an algorithm that takes only the point set P as input. In this section, we explain how to transform Problem ($\star$) into an equivalent problem that does not refer to $\mathcal{M}$ anymore, thus providing a realistic algorithm. Roughly, we simply replace the constraint $\mathrm{load}_{m_0,\mathcal{M}}(\gamma) = 1$ by a constraint of the form $\mathrm{load}_{x_0,\Pi,\Sigma}(\gamma) = 1$, where x_0 is a point "close" to $\mathcal{M}$, Π is a d-flat that "roughly approximates" $\mathcal{M}$ near x_0 and Σ are simplices of K "close" to x_0. To make this idea precise, we use the following localized version of the load:

$$\mathrm{load}_{x_0,\Pi,\Sigma}(\gamma) = \sum_{\sigma \in \Sigma^{[d]}} \gamma(\sigma)\, \mathrm{sign}_{\Pi}(\sigma) \mathbf{1}_{\pi_{\Pi}(\mathrm{conv}\, \sigma)}(x_0)$$

and state conditions in Lemma 17 (see below) under which Problem ($\star$) is equivalent to the problem obtained by replacing the constraint $\mathrm{load}_{m_0,\mathcal{M}}(\gamma) = 1$ with the constraint $\mathrm{load}_{x_0,\Pi,\Sigma}(\gamma) = 1$. Given a point $x \in \mathbb{R}^N$ and $r \geq 0$, let us introduce the subset of K:

$$K[x, r] = \{\sigma \in K \mid \mathrm{conv}\, \sigma \cap B(x, r) \neq \emptyset\}.$$

Note that $K[x, r]$ is not necessarily a simplicial complex.

Lemma 17 *Suppose $0 \leq \rho \leq \frac{\mathcal{R}}{25}$. Consider a point $x_0 \in \mathcal{M}^{\oplus \rho}$ and a d-dimensional affine space Π passing through x_0. Suppose that $\angle(\Pi, \mathbf{T}_{\pi_{\mathcal{M}}(x_0)}\mathcal{M}) \leq \frac{\pi}{8}$ and that the orientation of Π is consistent with that of $\mathbf{T}_{\pi_{\mathcal{M}}(x_0)}\mathcal{M}$. Then, Problem ($\star$) is equivalent to the following problem*

$$\begin{aligned} &\underset{\gamma}{\text{minimize}} && \mathrm{E}_{\mathrm{del}}(\gamma) \\ &\text{subject to} && \partial\gamma = 0, \\ & && \mathrm{load}_{x_0,\Pi,K[x_0,4\rho]}(\gamma) = 1 \end{aligned} \qquad (\star\star)$$

Proof This is a direct consequence of Lemma 34 in [33, Appendix G]. □

Observe that the conditions on the d-flat Π in the above lemma are rather mild. Indeed, we only require Π to pass through a point x_0 such that $d(x_0, \mathcal{M}) \leq \frac{\mathcal{R}}{25}$ and $\angle(\Pi, \mathbf{T}_{\pi_{\mathcal{M}}(x_0)}\mathcal{M}) \leq \frac{\pi}{8}$. Hence, Π only needs to be what we could call a *rough approximation* of $\mathcal{M}$ near x_0. In practice, we may take for x_0 any point $p_0 \in P$ and for Π the d-dimensional affine space $T_{p_0}(P,\rho)$ passing through p_0 and parallel to the d-dimensional vector space $V_{p_0}(P,\rho)$ defined as follows: $V_{p_0}(P,\rho)$ is spanned by the eigenvectors associated to the d largest eigenvalues of the inertia tensor of $(P \cap B(p_0,\rho)) - c$, where c is the center of mass of $P \cap B(p_0,\rho)$. By Lemma 35 in [33, Appendix H], for $\frac{\rho}{\mathcal{R}}$ small enough and $\varepsilon < \frac{\rho}{16}$, we have

$$\angle(T_{p_0}(P,\rho), T_{\pi_{\mathcal{M}}(p_0)}\mathcal{M}) \leq \frac{\pi}{8}.$$

See Appendix H in [33] for more details. Hence, the assumptions of the above lemma hold for $x_0 = p_0$ and $\Pi = T_{p_0}(P,\rho)$. This shows that the normalization constraint in Problem ($\star$) can be replaced by a constraint whose definition depends only upon the point set P, thus providing a realistic algorithm.

7.2 Perturbing the data set for ensuring the safety conditions

In this section, we assume that P_0 is a δ_0-accurate ε_0-dense sample of $\mathcal{M}$ and perturbe it to obtain a point set P that satisfies the assumptions of our main theorem. For this, we use the Moser Tardos Algorithm [35] as a perturbation scheme in the spirit of what is done in [16, Section 5.3.4].

The perturbation scheme is parametrized with real numbers $\rho \geq 0$, $r_{\text{pert.}} \geq 0$, $\text{Height}_{\min} > 0$, and $\text{Prot}_{\min} > 0$. To describe it, we need some notations and terminology. Let $T^*_{p_0} = T_{p_0}(P_0, 3\rho)$ be the d-dimensional affine space passing through p_0 and parallel to the d-dimensional vector space $V_{p_0}(P_0, 3\rho)$. To each point $p_0 \in P_0$, we associate a perturbed point $p \in P$, computed by applying a sequence of elementary operations called reset. Precisely, given a point $p \in P$ associated to the point $p_0 \in P_0$, the *reset* of p is the operation that consists in drawing a point q uniformely at random in $T^*_{p_0} \cap B(p_0, r_{\text{pert.}})$ and assigning q to p. Finally, we call any of the two situations below a *bad event*:

Violation of the height condition: There exists a ρ-small d-simplex $\sigma \subseteq P$ such that $\text{height}(\sigma) < \text{Height}_{\min}$;

Violation of the protection condition: There exists a pair (p, σ) made of a point $p \in P$ and a d-simplex $\sigma \subseteq P \setminus \{p\}$ such that $p \in B(c_\sigma, 3\rho)$ and $\text{protection}(\sigma, \{p\}) \leq \text{Prot}_{\min}$.

In both situations, we associate to the bad event E a set of points called the points *correlated* to E. In the first situation, the points correlated to E are the $d+1$ vertices of σ and in the second situation, they are the $d+2$ points of $\{p\} \cup \sigma$.

```
Moser-Tardos Algorithm:
1. For each p_0 ∈ P_0, compute the d-dimensional affine space T*_{p_0}
2. For each point p ∈ P, reset p
3. WHILE (some bad event E occurs):
-------------- For each point p correlated to E, reset p
----- END WHILE
4. Return P
```

Roughly speaking, in our context, the Moser Tardos Algorithm reassigns new coordinates to any point $p \in P$ that is correlated to a bad event as long as a bad event occurs. A beautiful result from [35] tells us that if bad events are mostly independent from one another and have each a sufficiently small probability to occur, then the Moser-Tardos Algorithm terminates and does so in a number of steps that is expected to be linear in the size of P_0. Precisely, suppose that each bad event is independent of all but at most Γ of the other bad events and the probability of a bad event is at most ϖ. Then, the result in [35] tells us that the Moser-Tardos algorithm terminates with expected time $O(\sharp P_0)$ whenever

$$\varpi \leq \frac{1}{e(\Gamma+1)}, \tag{18}$$

where e is the base of natural logarithms. Using this result, one can establish the following lemma, the proof of which is beyond the scope of this paper:

Lemma 18 *Let $\varepsilon_0 \geq 0$, $\eta_0 > 0$, and $\rho = C_{\mathrm{ste}}\varepsilon_0$, where $C_{\mathrm{ste}} \geq 32$. Let $\delta_0 = \frac{\rho^2}{\mathcal{R}}$, $r_{\mathrm{pert.}} = \frac{\eta_0\varepsilon_0}{20}$, $\varepsilon = \frac{21}{20}\varepsilon_0$, and $\delta = 2\delta_0$. There are positive constants c_0, c_1, and c_2 that depend only upon η_0, C_{ste}, and d such that if $\frac{\rho}{\mathcal{R}} < c_0$ then, given a point set P_0 such that $\mathcal{M} \subseteq (P_0)^{\oplus\varepsilon_0}$, $P_0 \subseteq \mathcal{M}^{\oplus\delta_0}$, and* separation$(P_0) > \eta_0\varepsilon_0$*, the point set P obtained after resetting each of its points satisfies $\mathcal{M} \subseteq P^{\oplus\varepsilon}$, $P \subseteq \mathcal{M}^{\oplus\delta}$, and* separation$(P) > \frac{9}{10}\eta_0\varepsilon_0$*. Moreover, whenever we apply the Moser-Tardos Algorithm with* $\mathrm{Height}_{\min} = c_1\left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}}\rho$ *and* $\mathrm{Prot}_{\min} = c_2\left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}}\rho$*, the algorithm terminates with expected time $O(\sharp P_0)$ and returns a point set P that is a δ-accurate ε-dense sample of $\mathcal{M}$ and that satisfies the assumptions of Theorem 7.*

We only sketch the proof of Lemma 18 below.

Sketch of proof The proof consists in applying the Moser Tardos theorem [35]. In other words, we show that Condition (18) holds, for a well-chosen upper bound ϖ on the probability of each bad event and a well-chosen upper bound Γ on the number of bad events to which each bad event is dependent upon. Upper bounds ϖ and Γ are obtained by adapting the proof of a similar simpler result presented in the appendix of [1]. The intuition is that thanks to Lemma 35 in [33, Appendix H], one can compute from the sample P_0 a local approximation $T_{p_0}(P_0, 3\rho)$ of a local tangent space with accuracy $\mathcal{O}\left(\frac{\rho}{\mathcal{R}}\right)$. It follows that, if $\frac{\rho}{\mathcal{R}}$ is small enough, the volume, in $\Pi_{p_0 \in P_0} T_{p_0}(P_0, 3\rho)$, for which a height or protection condition is violated, can be made arbitrary small, and Condition (18) required for Moser-Tardos algorithm to terminate will be met.

When the Moser-Tardos algorithm terminates, we thus have two positive constants c_1 and c_2 such that

$$\text{height}(P, \rho) > c_1 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}} \rho, \tag{19}$$

$$\text{protection}(P, 3\rho) > c_2 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}} \rho. \tag{20}$$

For short, write

$$\begin{aligned} \mathsf{p} &= \text{protection}(P, 3\rho), \\ \mathsf{s} &= \text{separation}(P), \\ \Theta &= \text{angularDeviation}_{\mathcal{M}}(P, \rho). \end{aligned}$$

Let us check that the safety assumptions of Theorem 7 are then satisfied. For this, we need to show that one can find $\theta \in \left[0, \frac{\pi}{6}\right]$ such that:

$$\Theta \leq \frac{\theta}{2} - \arcsin\left(\frac{\rho + \delta}{\mathcal{R}}\right), \tag{21}$$

$$\mathsf{s} > 8(\delta\theta + \rho\theta^2) + 6\delta + \frac{2\rho^2}{\mathcal{R}}, \tag{22}$$

$$\mathsf{p} > 8(\delta\theta + \rho\theta^2)\left(1 + \frac{4d\varepsilon}{\text{height}(P, \rho)}\right), \tag{23}$$

$$\mathsf{p}^2 + \mathsf{ps} > \max\left\{10\rho\Theta(\varepsilon + \rho\Theta), \frac{4J(1+J)}{(d+2)(d-1)!\,\Omega(\Delta_d)}\rho^2\right\}, \tag{24}$$

where

$$J = \frac{(\mathcal{R} + \rho)^d}{(\mathcal{R} - \rho)^d\ (\cos\Theta)^{\min\{d, N-d\}}} - 1.$$

By Lemma 28 in [33, Appendix 33], we obtain that

$$\Theta \leq \arcsin\left(\frac{2d}{\text{height}(P, \rho)}\left(\frac{3\rho^2}{\mathcal{R}} + \delta\right)\right).$$

Hence, since there exists a positive constant c_1 such that $\text{height}(P, \rho) > c_1 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}} \rho$, we deduce that there exists a positive constant c_3 such that for $\frac{\rho}{\mathcal{R}}$ small enough we have

$$\Theta \leq c_3 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{2}{3}}.$$

Let $\theta = 3\Theta$ and observe that for $\frac{\rho}{\mathcal{R}}$ small enough, $\theta \in \left[0, \frac{\pi}{6}\right]$. With this choice of θ and using $\mathsf{s} > \frac{\eta_0}{C_{\text{ste}}}\rho$, $\mathsf{p} > c_2 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}} \rho$, $\varepsilon = \frac{21}{20 C_{\text{ste}}}\rho$, $\delta = \frac{2\rho^2}{\mathcal{R}}$, and $\text{height}(P, \rho) > c_1 \left(\frac{\rho}{\mathcal{R}}\right)^{\frac{1}{3}} \rho$, it is easy to check that for $\frac{\rho}{\mathcal{R}}$ small enough, Inequalities (21), (22), (23), and (24) hold and therefore the safety assumptions of Theorem 7 are met. □

8 Conclusion

We have shown that the submanifold reconstruction problem can be recast as a weighted ℓ_1-norm minimization problem under linear constraints and as such is solvable by linear programming.

In the future, it would be interesting to study variants of this minimization problem. For instance, one could imagine constraining the solution to be a homology representative d-cycle (instead of a normalized d-cycle). Indeed, when $\mathcal{M}$ is orientable and connected, its d-homology group with real coefficients is one-dimensional, and the homology class of a normalized generator of it is called the *manifold fundamental class*. Furthermore, it is known that if K is either the Čech complex of P or the Vietoris-Rips complex of P, then K and $\mathcal{M}$ are homotopy equivalent, assuming that P samples sufficiently densely and accurately the manifold $\mathcal{M}$ and for a careful choice of the scale parameter of these complexes [37–41]. Hence, it follows that the d-homology group of K is also one-dimensional. One representative of a generator of the fundamental class of $\mathcal{M}$ can then be obtained, up to a multiplicative constant, by taking any non-boundary d-cycle γ_0 of K (performing for this standard linear algebra operations on the boundary operators ∂_d and ∂_{d+1} of K). Thus, a variant to our problem can be expressed as follows: among all the d-chains of K homologous to γ_0, search for the one with smallest Delaunay energy. We believe that it would be possible to adapt the proof presented in the paper and establish conditions under which the solution to this variant is also a d-chain which carries a triangulation of $\mathcal{M}$ (namely, the Delloc complex of P).

References

[1] Attali, D., Lieutier, A.: Flat Delaunay complexes for homeomorphic manifold reconstruction (2022)

[2] Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. In: Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH), pp. 71–78 (1992)

[3] Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. Discrete and Computational Geometry **22**(4), 481–504 (1999)

[4] Amenta, N., Choi, S., Kolluri, R.K.: The power crust. In: Anderson, D.C., Lee, K. (eds.) Sixth ACM Symposium on Solid Modeling and Applications, pp. 249–266. ACM, Sheraton Inn, Ann Arbor, Michigan, USA (2001). https://doi.org/10.1145/376957.376986 . https://doi.org/10.1145/376957.376986

[5] Edelsbrunner, H.: Surface reconstruction by wrapping finite sets in space. Discrete and computational geometry: the Goodman-Pollack Festschrift, 379–404 (2003)

[6] Giesen, J., John, M.: Surface reconstruction based on a dynamical system. In: Computer Graphics Forum, vol. 21, pp. 363–371 (2002). Wiley Online Library

[7] Cheng, S.-W., Dey, T.K., Ramos, E.A.: Manifold reconstruction from point samples. In: SODA, vol. 5, pp. 1018–1027 (2005)

[8] Dey, T.K.: Curve and Surface Reconstruction: Algorithms with Mathematical Analysis. Cambridge University Press, New York, USA (2006)

[9] Bauer, U., Roll, F.: Wrapping cycles in delaunay complexes: Bridging persistent homology and discrete morse theory. In: Mulzer, W., Phillips, J.M. (eds.) 40th International Symposium on Computational Geometry (SoCG 2024). Leibniz International Proceedings in Informatics (LIPIcs), vol. 293, pp. 15–11516. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2024)

[10] Chen, L., Xu, J.-c.: Optimal delaunay triangulations. Journal of Computational Mathematics, 299–308 (2004)

[11] Alliez, P., Cohen-Steiner, D., Yvinec, M., Desbrun, M.: Variational tetrahedral meshing. ACM Transactions on Graphics (TOG) **24**(3), 617–625 (2005)

[12] Chen, L., Holst, M.: Efficient mesh optimization schemes based on optimal delaunay triangulations. Computer Methods in Applied Mechanics and Engineering **200**(9), 967–984 (2011)

[13] Chen, Z., Wang, W., Lévy, B., Liu, L., Sun, F.: Revisiting optimal delaunay triangulation for 3d graded mesh generation. SIAM Journal on Scientific Computing **36**(3), 930–954 (2014)

[14] Morgan, F.: Geometric Measure Theory: a Beginner's Guide. Academic press, Cambridge, Massachussets, United States (2016)

[15] Boissonnat, J.-D., Ghosh, A.: Manifold reconstruction using tangential delaunay complexes. Discrete & Computational Geometry **51**(1), 221–267 (2014)

[16] Boissonnat, J.-D., Chazal, F., Yvinec, M.: Geometric and Topological Inference vol. 57. Cambridge University Press, Cambridge, United Kingdom (2018)

[17] Chen, C., Freedman, D.: Hardness Results for Homology Localization. Discrete & Computational Geometry **45**(3), 425–448 (2011) https://doi.org/10.1007/s00454-010-9322-8

[18] Borradaile, G., Maxwell, W., Nayyeri, A.: Minimum Bounded Chains and Minimum Homologous Chains in Embedded Simplicial Complexes. In: Cabello, S., Chen, D.Z. (eds.) 36th International Symposium on Computational Geometry (SoCG 2020). Leibniz International Proceedings in Informatics (LIPIcs), vol. 164, pp. 21–12115. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2020). https://doi.org/10.4230/LIPIcs.SoCG.2020.21

[19] Dey, T.K., Hirani, A.N., Krishnamoorthy, B.: Optimal Homologous Cycles, Total Unimodularity, and Linear Programming. SIAM Journal on Computing **40**(4), 1026–1044 (2011) https://doi.org/10.1137/100800245

[20] Chambers, E.W., Erickson, J., Nayyeri, A.: Minimum cuts and shortest homologous cycles. In: Proceedings of the 25th Annual Symposium on Computational Geometry - SCG '09, p. 377. ACM Press, Aarhus, Denmark (2009). https:

//doi.org/10.1145/1542362.1542426

[21] Dey, T.K., Hou, T., Mandal, S.: Computing minimal persistent cycles: Polynomial and hard cases. In: Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms. SODA'20, pp. 2587–2606. Society for Industrial and Applied Mathematics, USA (2020)

[22] Cohen-Steiner, D., Lieutier, A., Vuillamy, J.: Delaunay and regular triangulations as lexicographic optimal chains. Discrete & Computational Geometry **70**(1), 1–50 (2023)

[23] Cohen-Steiner, D., Lieutier, A., Vuillamy, J.: Lexicographic optimal homologous chains and applications to point cloud triangulations. Discrete & Computational Geometry **68**(4), 1155–1174 (2022)

[24] Jordan, C.: Essai sur la géométrie à n dimensions. Bulletin de la Société Mathématique de France **3**, 103–174 (1875) https://doi.org/10.24033/bsmf.90

[25] Munkres, J.R.: Elements of Algebraic Topology. Addison-Wesley Publishing Company, Inc., MA, USA (1993)

[26] Boissonnat, J.-D., Yvinec, M.: Algorithmic Geometry. Cambridge university press, Cambridge, United Kingdom (1998)

[27] Rippa, S.: Minimal roughness property of the delaunay triangulation. Computer Aided Geometric Design **7**(6), 489–497 (1990)

[28] Musin, O.R.: Properties of the delaunay triangulation. In: Proceedings of the Thirteenth Annual Symposium on Computational Geometry, pp. 424–426 (1997)

[29] Chen, L.: Mesh smoothing schemes based on optimal delaunay triangulations. In: Proceedings of the 13th International Meshing Roundtable, IMR 2004, Williamsburg, Virginia, USA, September 19-22, 2004, pp. 109–120 (2004). http://imr.sandia.gov/papers/abstracts/Ch317.html

[30] Edelsbrunner, H., Shah, N.R.: Incremental topological flipping works for regular triangulations. Algorithmica **15**(3), 223–241 (1996)

[31] Zeller, S., Thaler, M.: Almost sure escape from the unit interval under the logistic map. American Mathematical Monthly, 155–158 (2001)

[32] Federer, H.: Curvature measures. Trans. Amer. Math. Soc **93**, 418–491 (1959)

[33] Attali, D., Lieutier, A.: Delaunay-like Triangulation of Smooth Orientable Submanifolds by ℓ_1-Norm Minimization (2022). https://arxiv.org/abs/2203.06008

[34] Boissonnat, J.-D., Dyer, R., Ghosh, A.: The stability of delaunay triangulations. International Journal of Computational Geometry & Applications **23**(04n05),

303–333 (2013) https://doi.org/10.1142/s0218195913600078

[35] Moser, R.A., Tardos, G.: A constructive proof of the general lovász local lemma. Journal of the ACM (JACM) **57**(2), 1–15 (2010)

[36] Attali, D., Lieutier, A.: Reconstructing shapes with guarantees by unions of convex sets. In: Proc. 26th Ann. Sympos. Comput. Geom., Snowbird, Utah, pp. 344–353 (2010)

[37] Chazal, F., Lieutier, A.: Smooth Manifold Reconstruction from Noisy and Non Uniform Approximation with Guarantees. Computational Geometry: Theory and Applications **40**, 156–170 (2008)

[38] Chazal, F., Cohen-Steiner, D., Lieutier, A.: A sampling theory for compact sets in euclidean space. Discrete and Computational Geometry **41**(3), 461–479 (2009)

[39] Attali, D., Lieutier, A., Salinas, D.: Vietoris–rips complexes also provide topologically correct reconstructions of sampled shapes. Computational Geometry: Theory and Applications **46**(4), 448–465 (2013)

[40] Niyogi, P., Smale, S., Weinberger, S.: Finding the Homology of Submanifolds with High Confidence from Random Samples. Discrete Computational Geometry **39**(1-3), 419–441 (2008)

[41] Kim, J., Shin, J., Chazal, F., Rinaldo, A., Wasserman, L.: Homotopy reconstruction via the cech complex and the vietoris-rips complex. arXiv preprint arXiv:1903.06955 (2019)

A A linear programming formulation

For the sake of completeness, we recall in this appendix how minimizing a weighted ℓ^1-norm under linear constraints can be expressed as a linear programming problem through slack variables. Consider the following minimization problem:

$$\begin{aligned} \text{minimize} \quad & \sum_i \omega_i |x_i| \\ \text{subject to} \quad & Ax = b. \end{aligned}$$

Here $x \in \mathbb{R}^n$ is the variable, and $\omega_1, \omega_2, \ldots, \omega_n \in \mathbb{R}$, $A \in \mathbb{R}^{k \times n}$, $b \in \mathbb{R}^k$ are parameters. This problem is equivalent to the linear programming problem:

$$\begin{aligned} \text{minimize} \quad & \sum_i \omega_i s_i \\ \text{subject to} \quad & Ax = b, \\ & x_i \leq s_i, \qquad i = 1, \ldots, n, \\ & x_i \geq -s_i, \quad i = 1, \ldots, n, \end{aligned}$$

where the variables are $x \in \mathbb{R}^n$ and $s \in \mathbb{R}^n$. Each s_i is called a *slack variable.*