

Reconstruction Using Witness Complexes

Leonidas J. Guibas · Steve Y. Oudot

Received: 13 February 2007 / Revised: 23 April 2008 / Accepted: 23 April 2008 /
Published online: 28 June 2008
© Springer Science+Business Media, LLC 2008

Abstract We present a novel reconstruction algorithm that, given an input point set sampled from an object S , builds a one-parameter family of complexes that approximate S at different scales. At a high level, our method is very similar in spirit to Chew’s surface meshing algorithm, with one notable difference though: the restricted Delaunay triangulation is replaced by the witness complex, which makes our algorithm applicable in any metric space. To prove its correctness on curves and surfaces, we highlight the relationship between the witness complex and the restricted Delaunay triangulation in 2d and in 3d. Specifically, we prove that both complexes are equal in 2d and closely related in 3d, under some mild sampling assumptions.

Keywords Sampling · Reconstruction · Delaunay triangulation · Witness complex

1 Introduction

The problem of reconstructing a curve or a surface from scattered data points has received a lot of attention in the past. Although it is ill-posed by nature, since infinitely many shapes with different topological types can interpolate a given point cloud, a number of provably good methods have been proposed. The common denominator of these methods is the assumption that the input point set is densely sampled from a sufficiently regular shape: this assumption makes the reconstruction problem well posed, since all sufficiently regular shapes interpolating the point set have the

Communicated by Tamal Dey.

L.J. Guibas (✉) · S.Y. Oudot
Dept. Computer Science, Stanford University, Stanford, CA 94305, USA
e-mail: guibas@cs.stanford.edu

S.Y. Oudot
e-mail: steve.oudot@stanford.edu




Fig. 1 Helical curve drawn on a torus (left). From a uniform sampling of the curve, our algorithm builds a one-parameter family of complexes (right) and maintains their Betti numbers (left)

same topological type and are close to one another geometrically. It suffices then to approximate any of these shapes to get the right answer. The notion of ε -sample, introduced by Amenta and Bern [1], provides a sound mathematical framework for this kind of approach, the corresponding set of reconstructible shapes being the class of manifolds with *positive reach* [26]. A number of provably-good algorithms are based on the ε -sampling theory—see [9] for a survey, and several extensions have been proposed to reconstruct manifolds in higher-dimensional spaces [15] or from noisy point cloud data [24]. The theory itself has been recently extended to a larger class of shapes, known as the class of Lipschitz manifolds [8]. In all these methods, the Delaunay triangulation of the input point set plays a prominent role since the final reconstruction is extracted from it.

This approach to surface reconstruction is limited because it assumes implicitly that a point cloud should always represent a single class of shapes. Consider the example of a closed helical curve rolled around a torus in \mathbb{R}^3 —see Fig. 1. Take a very dense uniform point sample of the curve: what does this point set represent, the curve or the torus? Although both objects are well sampled according to Amenta and Bern’s sampling theory, classical reconstruction methods always choose a single class of shapes, here the curve or the torus, by restricting themselves either to a certain dimension or to a certain scale: for instance, the reconstruction method of [15] or the dimension detection algorithm of [23] will detect the curve but not the torus, since the point set is a sparse sample of the curve but not of the torus. Now, we claim that the result of the reconstruction should not be either the curve or the torus, but both of them. More generally, the result of the reconstruction should be a one-parameter family of complexes, whose elements approximate the original shape at different scales, as illustrated in Fig. 1. This point of view, inspired from recent results by Chazal et al. [10], stands in sharp contrast with previous work in the area

and is echoed in the literature on nonlinear dimensionality reduction [34, 36] and topological persistence [19, 25, 37].

This paper presents a novel reconstruction algorithm that, given an input point set W sampled from an object S , builds a one-parameter family of complexes that approximate S at different scales. At a high level, the method is very similar to Chew's surface meshing algorithm [7, 17]: it constructs a subset L of W iteratively, while maintaining a subcomplex of the Delaunay triangulation of L . The one-parameter family of complexes obtained from this iterative process is the result of the algorithm. The difference with Chew's approach is that, instead of maintaining the restricted Delaunay triangulation of L , we maintain its witness complex $\mathcal{C}^W(L)$ relative to W . The main advantages are that the underlying object S does not have to be known, and that the full-dimensional Delaunay triangulation $\mathcal{D}(L)$ does not have to be computed. Moreover, the algorithm can be used in any metric space, ultimately enabling new applications of the Delaunay-based reconstruction ideas.

The witness complex can be viewed as a weak version of the Delaunay triangulation, well defined and computable in any metric space—see Sect. 2 below. As such, it has played an important role in the context of topological data analysis [22]. It was first introduced by de Silva [20], who proved that $\mathcal{C}^W(L)$ is a subcomplex of $\mathcal{D}(L)$ whenever the points of L lie in general position in a Euclidean space. Moreover, if the set W of witnesses spans the whole ambient space, then $\mathcal{C}^W(L)$ is equal to $\mathcal{D}(L)$. Now the question is whether this property holds when the points of W are sampled from a subspace of the ambient space, such as for instance a submanifold: in [22], Carlsson and de Silva observed that $\mathcal{C}^W(L)$ is then closely related to the restricted Delaunay triangulation $\mathcal{D}^S(L)$, and they conjectured that both objects should coincide under some sampling assumptions on W and L . We prove that this conjecture is valid for a curve in the plane but not for a surface in 3d. In the latter case, we show how to relax the definition of the witness complex so that it contains $\mathcal{D}^S(L)$, and then how to extract a subcomplex that approximates $\mathcal{D}^S(L)$ (and hence S). This proves that our reconstruction algorithm is correct when applied to point samples of Lipschitz curves or surfaces. We are only aware of one related result: in [4], Attali et al. show that $\mathcal{C}^W(L)$ and $\mathcal{D}^S(L)$ coincide whenever the set W of witnesses spans an entire submanifold of \mathbb{R}^n of dimension one or two. This result differs from ours in two ways: our set W can be finite, which makes our result more practical, yet in return our set L has to be sparse compared to W , for $\mathcal{C}^W(L)$ to contain $\mathcal{D}^S(L)$. This sparseness condition is not an issue in practice, since the set L is constructed by the algorithm. Other noticeable differences are that our manifolds can have singularities and that our point samples can be noisy. Our assumption on the input point set W is fairly mild, since it amounts to saying that the Hausdorff distance between W and S is sufficiently small. In particular, there is no sparseness condition on W , and the amplitude of the noise can be as large as the sampling density. This noise model, introduced in [12] and used in subsequent work on reconstruction [10, 13, 29, 30], is less restrictive than its predecessors [16, 24, 28], and it makes our algorithm more practical.

The paper is organized as follows. In Sect. 2, we recall several concepts that will be used later on. In Sect. 3, we present our structural results. Specifically, we prove that the restricted Delaunay triangulation and the witness complex are equal in 2d

(Sect. 3.1) and closely related in 3d (Sect. 3.2), under some mild sampling assumptions. In Sect. 4, we introduce our reconstruction algorithm and present some experimental results.

2 Background and Definitions

Let S be a subset of \mathbb{R}^2 (resp. \mathbb{R}^3), L a finite set of points in \mathbb{R}^2 (resp. \mathbb{R}^3), and ε a positive number.

Definition 2.1

- L is an ε -noisy sample of S if no point of L is farther than ε from S .
- L is an ε -sample of S if no point of S is farther from L than ε .
- L is ε -sparse if the pairwise distances between the points of L are at least ε .

A 0-noisy sample is called a noise-free sample. When the first two conditions of the definition apply simultaneously, for a same ε , the Hausdorff distance between L and S is bounded by ε . We denote by $\mathcal{D}(L)$ the Delaunay triangulation of L .

Definition 2.2 The Delaunay triangulation of L restricted to S , or $\mathcal{D}^S(L)$ for short, is the subcomplex of $\mathcal{D}(L)$ made of the Delaunay faces whose dual Voronoi faces intersect S .

Let W be another set of points in \mathbb{R}^2 (resp. \mathbb{R}^3), finite or infinite.

Definition 2.3

- Given a point $w \in W$ and a simplex $\sigma = [p_0, \dots, p_l]$ with vertices in L , w witnesses σ if p_0, \dots, p_l belong to the $l + 1$ nearest neighbors of w , that is,

$$\forall i \in \{0, \dots, l\}, \forall q \in L \setminus \{p_0, \dots, p_l\}, \quad d(w, p_i) \leq d(w, q).$$

- The witness complex of L relative to W , or $\mathcal{C}^W(L)$ for short, is the maximum abstract simplicial complex with vertices in L , whose faces are witnessed by points of W .

The fact that $\mathcal{C}^W(L)$ is an abstract simplicial complex means that a simplex belongs to the complex only if all its faces do. By the so-called Weak Witness Theorem [20], we have $\mathcal{C}^W(L) \subseteq \mathcal{D}(L)$, which implies that $\mathcal{C}^W(L)$ is an embedded simplicial complex. In the sequel, L will be referred to as the set of landmarks, and W as the set of witnesses.

Lipschitz curves and surfaces Boissonnat and Oudot [8] introduced a new framework for the analysis of Delaunay-based sampling algorithms. This framework relies on a quantity, called the *Lipschitz radius*, which plays a role equivalent to the *local feature size* of Amenta and Bern [1], on a much larger class of shapes—the class of Lipschitz curves and surfaces.

Definition 2.4 Let S be the boundary of a bounded open subset \mathcal{O} of \mathbb{R}^2 (resp. \mathbb{R}^3). Given a point $p \in S$, the k -Lipschitz radius of S at p , or $\text{lr}_k(p)$ for short, is the maximum radius r such that $\mathcal{O} \cap B(p, r)$ is the intersection of $B(p, r)$ with the hypograph of some k -Lipschitz univariate (resp. bivariate) function. We call $\text{lr}_k(S)$ the infimum of lr_k over S .

Recall that the hypograph of a real-valued bivariate function f is the set of points $(x, y, z) \in \mathbb{R}^3$ such that $z < f(x, y)$. The function f is k -Lipschitz if: $\forall p, q \in \mathbb{R}^2, \frac{|f(p)-f(q)|}{\|p-q\|} \leq k$. It is proved in [8] that $\text{lr}_k(S) > 0$ whenever S is a k -Lipschitz curve in \mathbb{R}^2 or surface in \mathbb{R}^3 . In such a case, one can attach to each point $p \in S$ a so-called k -Lipschitz normal $\mathbf{n}_k(p)$ and a so-called k -Lipschitz support plane $T_k(p)$, which play a role similar to the usual normal vector and tangent plane in the Lipschitz setting:

Definition 2.5 Given $p \in S$, the k -Lipschitz normal of S at p , noted $\mathbf{n}_k(p)$, is the \mathbf{z} vector of an oriented orthonormal frame $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ in which $\mathcal{O}^- \cap B(p, \text{lr}_k(p))$ is the intersection of $B(p, \text{lr}_k(p))$ with the hypograph of a k -Lipschitz function $f(x, y)$. The k -Lipschitz support plane at p , noted $T_k(p)$, is the plane orthogonal to $\mathbf{n}_k(p)$ that passes through p .

The main result of [8] is the following:

Theorem 2.6 Let S be a k -Lipschitz surface in \mathbb{R}^3 and $L \subset S$ a finite point set such that:

- H1 L is an ε -sample of S with $\varepsilon < \frac{1}{7}\text{lr}_k(S)$,
- H2 the triangles of $\mathcal{D}^S(L)$ have radius-edge ratios of at most ϱ with $\varrho < \frac{\cos 2\theta}{2 \sin \theta}$ and $\theta = \arctan k$.

Then, $\mathcal{D}^S(L)$ is a 2-manifold isotopic to S , at Hausdorff distance at most ε from S , and whose oriented normals approximate the k -Lipschitz normals of S within an angle of $\arcsin(2\varrho \sin \theta)$.

Another useful result, proved in [31], is an equivalent of Proposition 13 of [5] for Lipschitz surfaces:

Lemma 2.7 Let S be a k -Lipschitz surface in \mathbb{R}^3 with $k < 1$. Then, $\forall p \in S, \forall r \leq \text{lr}_k(p), S \cap B(p, r)$ is a topological disk.

Similar results can be proved in the planar case (we recall their proofs for completeness):

Lemma 2.8 Let S be a k -Lipschitz curve in \mathbb{R}^2 with $k < 1$. Then, $\forall p \in S, \forall r \leq \text{lr}_k(p), S \cap B(p, r)$ is a topological arc. Moreover, the orthogonal projection of $S \cap B(p, r)$ onto $T_k(p)$ is a segment whose vertices are the orthogonal projections of the two endpoints of $S \cap B(p, r)$.

Proof Since S is an embedded curve without boundary, $S \cap B(p, r)$ is a collection of pairwise-disjoint simple arcs, whose endpoints lie on the circle $\partial B(p, r)$. Proving that $S \cap B(p, r)$ is a topological arc reduces then to showing that $|S \cap \partial B(p, r)| \leq 2$.

Since $r \leq \text{lr}_k(p)$, $S \cap B(p, r)$ is the graph of some k -Lipschitz univariate function f defined over $T_k(p)$ with $k < 1$. This implies that $S \cap B(p, r)$ lies in the double wedge K of apex p , of axis $T_k(p)$, and of half-angle θ . Now, at every point of $K \cap \partial B(p, r)$, the line tangent to $\partial B(p, r)$ makes an angle of at least $\frac{\pi}{2} - \theta$ with $T_k(c)$. Thus, for any two points u, v of $\partial B(p, r)$ lying in the same wedge of K , the line (u, v) makes an angle of at least $\frac{\pi}{2} - \theta > \theta$ with $T_k(c)$. It follows that $S \cap \partial B(p, r)$ cannot have more than one point per wedge of K , since otherwise the fact that $S \cap B(p, r)$ coincides with the graph of f would be contradicted. As a result, $S \cap B(p, r)$ is a topological arc, whose endpoints are precisely the two points u, v of $S \cap \partial B(p, r)$.

Since $S \cap B(p, r)$ coincides with the graph of f , the orthogonal projection of $S \cap B(p, r)$ is one-to-one, and thus it maps the arc $S \cap B(p, r)$ to a segment whose vertices are the orthogonal projections of the endpoints of $S \cap B(p, r)$. \square

Theorem 2.9 *If S is a k -Lipschitz curve in the plane with $k < 1$, and if L is an ε -sample of S with $\varepsilon < \text{lr}_k(S)$, then $\mathcal{D}^S(L)$ is a polygonal curve homeomorphic to S and at Hausdorff distance at most ε from S .*

Proof Since L is an ε -sample of S , no point of S is farther than ε from $L \subseteq \mathcal{D}^S(L)$. Conversely, every point of $\mathcal{D}^S(L)$ belongs to some Delaunay edge circumscribed by a Delaunay ball centered on S . This ball has a radius of at most ε , since L is an ε -sample of S . Thus, no point of $\mathcal{D}^S(L)$ is farther than ε from S . This proves that $d_{\mathcal{H}}(S, \mathcal{D}^S(L)) \leq \varepsilon$.

Let us now show that S and $\mathcal{D}^S(L)$ are homeomorphic. Let $[u, v]$ be an edge of $\mathcal{D}^S(L)$. Let $B(c, r)$ be a Delaunay ball centered on S that circumscribes $[u, v]$. Since L is an ε -sample of S , we have $r \leq \varepsilon < \text{lr}_k(S)$. Therefore, $S \cap B(c, r)$ is a topological arc whose endpoints are u and v , by Lemma 2.8. Moreover, S intersects $\partial B(c, r)$ only in u and v . Thus, u and v are consecutive points of L on S .

Assume conversely that u and v are two points of L consecutive on S . Let γ be the elementary¹ arc of S that joins u and v . Let l be the bisector line of $[u, v]$. Since l divides \mathbb{R}^2 into two regions, one of which contains u while the other contains v , γ intersects l . Let c be a point of intersection, and let $r = d(c, L)$. If the circle $\partial B(c, r)$ passes through u and v , then $B(c, r)$ is a Delaunay ball and hence $[u, v]$ belongs to $\mathcal{D}^S(L)$. Otherwise, $\partial B(c, r)$ passes through some other point $t \in L$. Since u and v are consecutive along S , t does not belong to γ . Hence, $S \cap B(c, r)$ contains a piece of γ (since $c \in \gamma$) and also a piece of $S \setminus \gamma$, and these pieces are disjoint, since $B(c, r)$ does not contain u nor v . As a result, $S \cap B(c, r)$ is disconnected, which implies that $r > \text{lr}_k(S)$, by Lemma 2.8. Now, since L is an ε -sample of S , we have $r \leq \varepsilon < \text{lr}_k(S)$, which contradicts the above statement. Therefore, $[u, v]$ belongs to $\mathcal{D}^S(L)$.

It follows from the last two paragraphs that there is a bijection between the elementary arcs of S and the edges of $\mathcal{D}^S(L)$. For any edge $e = [u, v]$ of $\mathcal{D}^S(L)$, let h_e be a homeomorphism from e to the elementary arc of S joining u and v , such that

¹Elementary means that γ does not contain any point of L other than u and v .

$h_e(u) = u$ and $h_e(v) = v$. Let h be the map defined by: $\forall p \in \mathcal{D}^S(L), h(p) = h_e(p)$, where e is any edge of $\mathcal{D}^S(L)$ containing p . Since the edges of $\mathcal{D}^S(L)$ have pairwise-disjoint relative interiors, and since the h_e coincide at vertices (they leave the vertices unchanged), $h(p)$ does not depend on the choice of the edge e of $\mathcal{D}^S(L)$ that contains p . It follows that h is a well-defined and continuous function from $\mathcal{D}^S(L)$ to S . Moreover, since the h_e are bijective, and since there is a bijection between the edges of $\mathcal{D}^S(L)$ and the elementary arcs of S , h is bijective. In addition, the inverse of h is continuous, since $\mathcal{D}^S(L)$ is compact. It follows that h is a homeomorphism from $\mathcal{D}^S(L)$ to S . □

3 Structural Results

In this section, we highlight the relationship between the witness complex and the restricted Delaunay triangulation in 2d and in 3d. Let S be a k -Lipschitz manifold, i.e., either a k -Lipschitz curve in the plane (Sect. 3.1) or a k -Lipschitz surface in 3d (Sect. 3.2), for some constant $k \geq 0$. For convenience, we define $\theta = \arctan k \in [0, \pi/2[$. Let W be a δ -noisy δ -sample of S and $L \subset W$ an ε -sparse ε -sample of W . The constants δ and ε will be explicited later on. Clearly, L is a $(\delta + \varepsilon)$ -sample of S .

For the sake of simplicity, we assume for now that the points of L belong to S , i.e., that L is a noise-free sample of S . We will see in Sect. 3.3 that this condition can actually be removed, since whenever L is a δ -noisy ε -sparse ε -sample of a k -Lipschitz surface S , there exists some nearby and isotopic k' -Lipschitz surface S' with $k' = O(k)$ and $\text{lr}_{k'}(S') = \Omega(\text{lr}_k(S))$ such that the points of L lie on S' (Theorem 3.16).

Our assumption that L is a sparse point set may look a bit restrictive at first sight. However, in our application the set of landmarks will not be given as input and will be made sparse by construction—see Sect. 4.

3.1 The Planar Case

Theorem 3.1 *Assume that $\theta < \arcsin \frac{1}{8} \approx 7.2$ deg and that $\delta < \min\{\frac{1-8\sin\theta}{12}, \frac{3\cos\theta-2}{4(9\cos\theta+6)}\}\text{lr}_k(S)$. If ε satisfies $\max\{\frac{12\sin\theta}{1-8\sin\theta}, \frac{6}{3\cos\theta-2}\}\delta < \varepsilon < \frac{1}{8}\text{lr}_k(S) - \frac{3}{2}\delta$, then $C^W(L)$ coincides with $\mathcal{D}^S(L)$.*

The lower bound on ε means that the set W of witnesses must be sufficiently dense² compared to the set L of landmarks, for the simplices of $\mathcal{D}^S(L)$ to be witnessed. An illustration is given in Fig. 2, which shows that $C^W(L)$ contains $\mathcal{D}^S(L)$ when L is sparse (left picture), whereas when $L = W$ ($\varepsilon = 0$), $C^W(L)$ coincides with the nearest neighbor graph of L , which has nothing to do with $\mathcal{D}^S(L)$ (right picture). The upper bound on ε ensures that the set L of landmarks is sufficiently dense, so that the nice properties of $\mathcal{D}^S(L)$ stated in Sect. 2 hold and that classical local arguments can be applied to show that $C^W(L)$ is included in $\mathcal{D}^S(L)$.

To prove the theorem, we will need the following technical result:

²In particular, we have $\varepsilon > 6\delta$.




Fig. 2 (Color online) A set of witnesses sampling a smooth closed curve with noise, and two different subsets of landmarks (in green) together with their witness complexes

Lemma 3.2 Assume that $\theta < \frac{\pi}{6}$. Then, every Voronoi vertex lies at least $\min\{\frac{1}{2}lr_k(S), \frac{\epsilon}{2\sin\theta}\}$ away from L .

Proof Let c be a Voronoi vertex, and let (u, v, w) be its dual Delaunay triangle. Since $L \subset S$, the points u, v, w belong to S . Let $r = d(c, u)$. Assume for a contradiction that $r < \min\{\frac{1}{2}lr_k(S), \frac{\epsilon}{2\sin\theta}\}$. Then, (u, v, w) is included in the ball $B(u, 2r) \subseteq B(u, lr_k(S))$, and thus lines (u, v) , (u, w) , and (v, w) form angles of at most θ with $T_k(u)$. It follows that each inner angle of triangle (u, v, w) is smaller than 2θ or larger than $\pi - 2\theta$. Since $\theta < \frac{\pi}{6}$, the larger inner angle of the triangle is at least $\pi - 2\theta$, hence the radius-edge ratio of (u, v, w) is at least $\frac{1}{2\sin\theta}$. As a consequence, we have $r \geq \frac{\epsilon}{2\sin\theta}$, since the lengths of the edges of (u, v, w) are at least ϵ (recall that L is ϵ -sparse). This contradicts the assumption that $r < \min\{\frac{1}{2}lr_k(S), \frac{\epsilon}{2\sin\theta}\}$ and thereby concludes the proof of the lemma. \square

We will now prove Theorem 3.1 in two steps: first, we will show that $\mathcal{D}^S(L)$ is included in $\mathcal{C}^W(L)$ (Lemma 3.3), then we will show that $\mathcal{C}^W(L)$ is included in $\mathcal{D}^S(L)$ (Lemma 3.4).

Lemma 3.3 Assume that $\theta < \frac{\pi}{6}$ and that $\delta < \min\{\frac{1-2\sin\theta}{2}, \frac{3\cos\theta-2}{2(3\cos\theta+4)}\}lr_k(S)$. If ϵ satisfies $\max\{\frac{2\sin\theta}{1-2\sin\theta}, \frac{6}{3\cos\theta-2}\}\delta < \epsilon < \frac{1}{2}lr_k(S) - \delta$, then $\mathcal{D}^S(L)$ is included in $\mathcal{C}^W(L)$ and has no simplex of dimension two or more.

Proof Let $e = [u, v]$ be an edge of $\mathcal{D}^S(L)$. By the definition of $\mathcal{D}^S(L)$, the dual Voronoi edge of e intersects S at some point c . Let $r = d(c, u) = d(c, v) = d(c, L)$, which is at most $\epsilon + \delta$, since L is a $(\delta + \epsilon)$ -sample of S . Since W is a δ -sample of S , there is some $w \in W$ at distance at most δ from c . Then, u and v are both included in $B(w, \epsilon + 2\delta)$.

Let p be any point of $L \setminus \{u, v\}$. We will prove that $p \notin B(w, \epsilon + 2\delta)$, which means that w witnesses e . Consider the portion of S that lies in $B(c, r)$. Since $r \leq \epsilon + \delta < lr_k(S)$, we know from Lemma 2.8 that $S \cap B(c, r)$ is a topological arc whose endpoints are u, v and whose orthogonal projection onto $T_k(c)$ is the line segment $[\bar{u}, \bar{v}]$, where \bar{u}, \bar{v} are the orthogonal projections of u, v . If p does not belong to

$B(c, \text{lr}_k(S))$, then it does not belong to $B(w, \varepsilon + 2\delta)$ either, since $B(w, \varepsilon + 2\delta) \subset B(c, \varepsilon + 3\delta)$, which by hypothesis is included in $B(c, \text{lr}_k(S))$. Otherwise, since $p \in L \setminus \{u, v\}$, $S \cap B(c, r)$ contains c but not p . Therefore, \bar{p} (the orthogonal projection of p onto $T_k(c)$) does not belong to $[\bar{u}, \bar{v}]$, because the projection from $S \cap B(c, \text{lr}_k(S))$ to $T_k(c)$ is one-to-one. As a consequence, $d(c, p)$ is at least $d(c, \bar{p}) \geq d(c, \{\bar{u}, \bar{v}\}) + d(\bar{p}, \{\bar{u}, \bar{v}\})$. Since L is ε -sparse, we have $d(p, u) \geq \varepsilon$, $d(p, v) \geq \varepsilon$, and $d(c, u) = d(c, v) \geq \frac{d(u, v)}{2} \geq \frac{\varepsilon}{2}$. Moreover, since $S \cap B(c, \text{lr}_k(S))$ is the graph of a k -Lipschitz univariate function defined over $T_k(c)$, we have $d(\bar{p}, \bar{u}) \geq d(p, u) \cos \theta$, $d(\bar{p}, \bar{v}) \geq d(p, v) \cos \theta$, $d(c, \bar{u}) \geq d(c, u) \cos \theta$, and $d(c, \bar{v}) \geq d(c, v) \cos \theta$. As a result,

$$d(c, p) \geq d(c, \bar{p}) \geq d(c, \{\bar{u}, \bar{v}\}) + d(\bar{p}, \{\bar{u}, \bar{v}\}) \geq \frac{\varepsilon}{2} \cos \theta + \varepsilon \cos \theta = \frac{3}{2} \varepsilon \cos \theta.$$

This expression is greater than $\varepsilon + 3\delta$, since $\varepsilon > \frac{6\delta}{3 \cos \theta - 2}$, by hypothesis. It follows that p is farther than $\varepsilon + 3\delta$ from c and hence farther than $\varepsilon + 2\delta$ from w . Thus, w witnesses $[u, v]$. Similarly, every other edge of $\mathcal{D}^S(L)$ is witnessed by some point of W . Since $L \subseteq W$, the vertices of $\mathcal{D}^S(L)$ witness themselves, hence the 1-skeleton of $\mathcal{D}^S(L)$ is included in $C^W(L)$.

Let us now focus on higher-dimensional simplices. Since L is a $(\delta + \varepsilon)$ -sample of S with $\delta + \varepsilon < \min\{\frac{1}{2}\text{lr}_k(S), \frac{\varepsilon}{2 \sin \theta}\}$ by hypothesis, all the points of S lie within distance $\min\{\frac{1}{2}\text{lr}_k(S), \frac{\varepsilon}{2 \sin \theta}\}$ of L . Therefore, S cannot contain any Voronoi vertex, by Lemma 3.2. This means that $\mathcal{D}^S(L)$ has no triangle. As a consequence, $\mathcal{D}^S(L)$ has no simplex of dimension two or more, and therefore it is equal to its 1-skeleton, which is included in $C^W(L)$. □

Lemma 3.4 *Assume that $\theta < \arcsin \frac{1}{8}$ and that $\delta < \frac{1-8 \sin \theta}{12} \text{lr}_k(S)$. If ε satisfies $\frac{12 \sin \theta}{1-8 \sin \theta} \delta < \varepsilon < \frac{1}{8} \text{lr}_k(S) - \frac{3}{2} \delta$, then $C^W(L)$ is included in $\mathcal{D}^S(L)$ and has no simplex of dimension two or more.*

Proof Let v be a vertex of $C^W(L)$. By assumption, we have $v \in L \subset S$, therefore the Voronoi cell of v intersects S , which implies that v is a vertex of $\mathcal{D}^S(L)$. Let now $[u, v]$ be an edge of $C^W(L)$, and let $w \in W$ be a witness of $[u, v]$. We assume without loss of generality that $d(w, u) \leq d(w, v)$.

Claim 3.4.1 $d(w, u) \leq d(w, v) \leq 2\varepsilon + 3\delta$.

Proof Let \bar{w} be a point of S closest to w . Since W is a δ -noisy sample of S , we have $d(w, \bar{w}) \leq \delta$. It follows that $d(\bar{w}, u) \leq \delta + \varepsilon < \text{lr}_k(S)$. Since $S \cap B(\bar{w}, \text{lr}_k(S))$ is the graph of a k -Lipschitz univariate function defined over $T_k(\bar{w})$, u belongs to the double wedge K of apex \bar{w} , of axis $T_k(\bar{w})$, and of half-angle θ . Let K_u be the wedge of K that contains u (if $u = \bar{w}$, then take either wedge of K for K_u). Since $\theta < \frac{\pi}{4}$, for any point $p \in K \setminus K_u$, the point of K_u closest to p is \bar{w} . In particular, given any $r \in]\delta + \varepsilon, \text{lr}_k(S)[$, every point p of $S \cap \partial B(\bar{w}, r) \setminus K_u$ satisfies $d(p, u) \geq d(p, \bar{w}) = r$. If v does not belong to $B(\bar{w}, 2r)$, then $d(p, L) \geq r > \delta + \varepsilon$, which contradicts the fact that L is a $(\delta + \varepsilon)$ -sample of S . Since this is true for all $r \in]\delta + \varepsilon, \text{lr}_k(S)[$, v belongs to the ball $B(\bar{w}, 2(\delta + \varepsilon))$, which implies that $d(w, v) \leq 3\delta + 2\varepsilon$, since $d(w, \bar{w}) \leq \delta$. □

Consider the set of balls centered on the line segment $[w, v]$ and whose bounding spheres pass through v . Since $B(w, d(w, v))$ contains u , one of these balls (called B) is such that ∂B contains both u and v . Clearly, $B \subseteq B(w, d(w, v))$. Since w witnesses $[u, v]$, the interior of $B(w, d(w, v))$ cannot contain any point of L other than u , thus no point of L lies in the interior of B . Therefore, B is a Delaunay ball. Let c be its center. It lies on the Voronoi edge e dual to $[u, v]$. Moreover, its distance to L is at most $d(w, v)$. By Claim 3.4.1, we have $d(w, v) \leq 2\varepsilon + 3\delta$, which, according to the hypothesis of the lemma, is less than $\min\{\frac{1}{4}\text{lr}_k(S), \frac{\varepsilon}{4\sin\theta}\}$. Hence, $d(c, L) < \min\{\frac{1}{4}\text{lr}_k(S), \frac{\varepsilon}{4\sin\theta}\}$.

Claim 3.4.2 If a Voronoi edge contains a point closer to L than $\min\{\frac{1}{4}\text{lr}_k(S), \frac{\varepsilon}{4\sin\theta}\}$, then this edge intersects S .

Proof Let e be a Voronoi edge containing a point c closer to L than $r = \min\{\frac{1}{4}\text{lr}_k(S), \frac{\varepsilon}{4\sin\theta}\}$. Let $[u, v]$ be the Delaunay edge dual to e . We have $d(c, u) = d(c, v) = d(c, L) < r$. It follows that $d(u, v) \leq 2r$. Since $2r < \text{lr}_k(S)$, $S \cap B(u, 2r)$ is the graph of a k -Lipschitz univariate function defined over $T_k(u)$. As a consequence, $S \cap B(u, 2r)$ lies outside the double wedge $K(u)$ of apex u , of axis aligned with $\mathbf{n}_k(u)$, and of half-angle $\frac{\pi}{2} - \theta$. Moreover, by Lemma 2.8, $S \cap B(u, 2r)$ is a topological arc whose endpoints lie on the circle $\partial B(u, 2r)$. Since c lies in $B(u, r)$, it can be proved that $S \cap B(u, 2r)$ intersects any line³ included in the double wedge $K(c)$ of apex c , of axis aligned with $\mathbf{n}_k(u)$, and of half-angle θ : a proof of this fact is given in Claim 6.6.3 of [8] for the 3d case, and it holds the same in 2d. Now, since v lies outside $K(u)$, the angle between $\mathbf{n}_k(u)$ and the line l supporting e is at most θ , which means that l is included in $K(c)$ and thus intersects S inside $B(u, 2r)$. Finally, we assumed that e intersects $B(u, 2r)$, while Lemma 3.2 states that the vertices of e lie outside of $B(u, 2r)$. Hence, $l \cap B(u, 2r) = e \cap B(u, 2r)$, and therefore e intersects S . □

It follows from this claim that e intersects S . Therefore, $[u, v]$ is an edge of $\mathcal{D}^S(L)$. As a result, the 1-skeleton of $\mathcal{C}^W(L)$ is included in $\mathcal{D}^S(L)$. Let us now show that $\mathcal{C}^W(L)$ has no simplex of dimension two or more. Recall that a simplex σ belongs to $\mathcal{C}^W(L)$ only if all its faces belong to $\mathcal{C}^W(L)$. Therefore, it suffices to prove that $\mathcal{C}^W(L)$ has no triangle.

Claim 3.4.3 For any triangle t of $\mathcal{D}(L)$, at least one edge of t does not belong to $\mathcal{D}^S(L)$.

Proof Let (u, v, w) be a triangle of $\mathcal{D}(L)$. Assume for a contradiction that edges $[u, v]$, $[v, w]$, and $[u, w]$ belong to $\mathcal{D}^S(L)$. Since L is a $(\delta + \varepsilon)$ -sample of S , the three edges are circumscribed by Delaunay balls of radius at most $\delta + \varepsilon$. It follows that the triangle (u, v, w) is included in $B(v, 2(\delta + \varepsilon))$. Let c be the circumcenter of this triangle. Since c is a Voronoi vertex, Lemma 3.2 states that $d(c, L) \geq \min\{\frac{1}{2}\text{lr}_k(S), \frac{\varepsilon}{2\sin\theta}\}$,

³Any such line must pass through c .




Fig. 3 Order-two Voronoi diagrams of two sets of points in the plane

which by assumption is greater than $2(\epsilon + \delta)$. Therefore, c does not belong to $B(v, 2(\epsilon + \delta))$, and hence it does not belong to triangle (u, v, w) either. This means that c is not a local maximum of the distance to L . As a consequence, there is a Voronoi edge e incident to c such that $d(x, L)$ increases as point x moves away from c along e . Since the local maxima of the distance to L are Voronoi vertices, $d(x, L)$ keeps increasing until x reaches the other vertex of e . Thus, every point of e is at least $d(c, L)$ away from L with $d(c, L) > 2(\delta + \epsilon)$. This implies that e cannot intersect S , because otherwise some point of e would be at most $\delta + \epsilon$ away from L (recall that L is a $(\delta + \epsilon)$ -sample of S). Now, the Delaunay edge dual to e is an edge of (u, v, w) . Therefore, at least one edge of (u, v, w) does not belong to $\mathcal{D}^S(L)$, which contradicts our assumption. \square

Since $\mathcal{C}^W(L)$ is included in $\mathcal{D}(L)$, and since its edges belong to $\mathcal{D}^S(L)$, any triangle of $\mathcal{C}^W(L)$ is a Delaunay triangle whose edges belong to $\mathcal{D}^S(L)$. Hence, by Claim 3.4.3, $\mathcal{C}^W(L)$ has no triangle, which concludes the proof of the lemma. \square

3.2 The 3d Case

Unlike in the planar case, the witness complex and the restricted Delaunay triangulation of points sampled from a surface in 3d may not always coincide, even in situations where the sets of witnesses and landmarks satisfy strong sampling conditions. The reason is that, when a tetrahedron t of $\mathcal{D}(L)$ has almost cocircular vertices, the chance for any of the diagonal edges of t to be witnessed by a point of W is small—such a tetrahedron is called a *sliver* in the literature [35]. In order to give an intuition of this fact, let us assume for simplicity that the surface is flat and that the vertices of t are cocircular, as in Fig. 3 (left). The order-two Voronoi diagram of the vertices is then degenerate, the Voronoi cells of the diagonal edges being reduced to a single point p that lies at the intersection of the edges of the diagram. Therefore, any diagonal edge can be witnessed only by p , which means that the probability for any triangle of the quadrangle to be witnessed when W is finite is zero. As a result, holes appear with probability one in the witness complex, as illustrated in Fig. 4 (left).




Fig. 4 Witness complex, v -witness complex, and extracted manifold

When the vertices of tetrahedron t are almost-cocircular, as in Fig. 3 (right), the order-two Voronoi cell of one diagonal edge is empty, while the cell of the other diagonal edge is arbitrarily small. Thus, the probability for any triangle of the quadrangle to be witnessed when W is finite is also arbitrarily small. Although it is always possible to perturbate the point set L so that the points are in general position, guaranteeing that the order-two Voronoi cells of the edges of $\mathcal{D}^S(L)$ are sufficiently large requires large perturbations, which are not tractable in practice, since the underlying surface S is unknown.

3.2.1 The v -witness Complex

Our approach for dealing with the above issue consists in relaxing the definition of the witness complex, so that the latter includes the restricted Delaunay triangulation. This requires to modify the concept of witness:

Definition 3.5 Given an integer m , a point $w \in W$, and a simplex $\sigma = [p_0, \dots, p_l]$ with vertices in L , w m -witnesses σ if all the $d(w, p_i)$ are among the m smallest values of the set $\{d(w, q), q \in L\}$.

Observe that, in the case where $m \leq l$, some vertices of σ must be equidistant to w for w to m -witness σ . In particular, if $m = 1$, then all the points of σ must be equidistant to w , which means that w is a strong witness of σ , or equivalently, that σ is a Delaunay simplex. If $m = 0$, then no point $w \in W$ can m -witness σ . In [22], the authors use m -witnesses only for edges. More generally, we use them for simplices of all dimensions:

Definition 3.6 Given a countable sequence ν of integers, the ν -witness complex of L , or $\mathcal{C}_\nu^W(L)$ for short, is the maximum abstract simplicial complex with vertices in L such that each i -face is ν_i -witnessed by some point of W .

Since the simplices of $\mathcal{C}_\nu^W(L)$ have their vertices in L , their dimension is at most $|L| - 1$. Hence, in the sequence of integers, only ν_0 through $\nu_{|L|-1}$ are used. There is a strong relationship between $\mathcal{C}^W(L)$ and $\mathcal{C}_\nu^W(L)$: if $\nu_i \geq i + 1 \forall i$, then $\mathcal{C}_\nu^W(L)$ contains $\mathcal{C}^W(L)$; in contrast, if $\nu_i \leq i + 1 \forall i$, then $\mathcal{C}_\nu^W(L)$ is included in $\mathcal{C}^W(L)$;

thus, $C^W(L) = C_v^W(L)$ whenever $v_i = i + 1 \forall i$. In addition, if $v_i = 0$ for some i , then the i -skeleton of $C_v^W(L)$ is empty, since a simplex cannot be 0-witnessed, and the dimension of $C_v^W(L)$ is then at most $i - 1$.

Theorem 3.7 Assume that $\theta < \arccos(2 \sin \frac{\pi}{7}) \approx 29.7$ deg and that $\delta < \frac{\cos \theta - 2 \sin \frac{\pi}{7}}{3 \cos \theta + 2 \sin \frac{\pi}{7}} \cdot \text{Ir}_k(S)$. If ε satisfies $\frac{8 \sin \frac{\pi}{7}}{\cos \theta - 2 \sin \frac{\pi}{7}} \delta < \varepsilon < \text{Ir}_k(S) - 3\delta$, then, for any sequence v of integers such that $v_0 \geq 1, v_1 \geq 6$, and $v_2 \geq 6$, $\mathcal{D}^S(L)$ is included in $C_v^W(L)$.

Proof Since the vertices of $\mathcal{D}^S(L)$ belong to L , which is included in W , they witness themselves and thus belong to $C_v^W(L)$. In addition, $\mathcal{D}^S(L)$ contains no tetrahedron generically.

Let σ be a simplex (edge or triangle) of $\mathcal{D}^S(L)$, and let $B(c, r)$ be a Delaunay ball centered on S and circumscribing σ . Since W is a δ -sample of S , there is a point $w \in W$ at distance at most δ from c . Then, σ lies in the ball $B(w, r + \delta)$, which is included in $B(c, r + 2\delta)$.

Claim 3.7.1 The ball $B(c, r + 2\delta)$ contains at most six points of L .

Proof Since L is a $(\delta + \varepsilon)$ -sample of S , the radius r of the surface Delaunay ball is at most $\delta + \varepsilon$. Therefore, $r + 2\delta \leq \varepsilon + 3\delta$, which is less than $\text{Ir}_k(S)$ according to the hypothesis of Theorem 3.7. It follows that $S \cap B(c, r + 2\delta)$ is the graph of a k -Lipschitz bivariate function defined over the plane $T_k(c)$.

Let p_1, \dots, p_l be the points of $L \cap B(c, r + 2\delta)$. We call $\bar{p}_1, \dots, \bar{p}_l$ their orthogonal projections onto $T_k(c)$. Since L is ε -sparse, the p_i are at least ε away from one another. And since they belong to $S \cap B(c, \text{Ir}_k(S))$, which is the graph of a k -Lipschitz bivariate function defined over the plane $T_k(c)$, their projections \bar{p}_i are at least $\varepsilon \cos \theta$ away from one another. Moreover, since $B(c, r)$ is a Delaunay ball, the p_i are at least r away from c , hence the \bar{p}_i are at least $r \cos \theta$ away from c .

The rest of the proof depends on whether $r \geq \varepsilon$ or $r < \varepsilon$. In fact, the overall ideas are the same, but some technical details differ.

- If $r \geq \varepsilon$, then, inside $T_k(c)$, c and the \bar{p}_i are centers of pairwise-disjoint open disks of radius $\frac{\varepsilon}{2} \cos \theta$. Let D_c, D_1, \dots, D_l denote these disks. Since the p_i belong to $B(c, r + 2\delta)$, the \bar{p}_i belong to the disk $D(c, r + 2\delta)$. Therefore, D_c, D_1, \dots, D_l form a congruent packing of the disk $D(c, r + 2\delta + \frac{\varepsilon}{2} \cos \theta)$. Now, according to the hypotheses of the theorem, we have:

$$r + 2\delta + \frac{\varepsilon}{2} \cos \theta \leq \varepsilon + 3\delta + \frac{\varepsilon}{2} \cos \theta < \frac{\varepsilon}{2} \cos \theta \left(\frac{1}{\sin \frac{\pi}{7}} + 1 \right).$$

Hence, by a classical result on congruent packings of disks [27, 33], there are at most seven disks of radius $\frac{\varepsilon}{2} \cos \theta$ packed in $D(c, r + 2\delta + \frac{\varepsilon}{2} \cos \theta)$. The fact that D_c is one of them implies that $l \leq 6$, which proves the claim in the case where $r \geq \varepsilon$.

- If $r < \varepsilon$, then c and the \bar{p}_i are centers of pairwise-disjoint open disks of radius $\frac{r}{2} \cos \theta$. Let D_c, D_1, \dots, D_l denote these disks. Since the p_i belong to $B(c, r + 2\delta)$,

D_c, D_1, \dots, D_l are included in $D(c, r + 2\delta + \frac{r}{2} \cos \theta)$. Now, $B(c, r)$ is a Delaunay ball, hence its bounding sphere contains at least two points of L , which implies that $r \geq \frac{\varepsilon}{2}$. Moreover, the hypotheses of the theorem state that $\delta < \frac{\varepsilon}{4} (\frac{\cos \theta}{2 \sin \frac{\pi}{7}} - 1)$, which is at most $\frac{r}{2} (\frac{\cos \theta}{2 \sin \frac{\pi}{7}} - 1)$. Therefore, D_c, D_1, \dots, D_l form a congruent packing of a disk of radius

$$r + 2\delta + \frac{r}{2} \cos \theta < r + r \left(\frac{\cos \theta}{2 \sin \frac{\pi}{7}} - 1 \right) + \frac{r}{2} \cos \theta = \frac{r}{2} \cos \theta \left(\frac{1}{\sin \frac{\pi}{7}} + 1 \right).$$

It follows, by the same result as above on congruent packings of disks, that $l \leq 6$, which proves the claim in the case where $r < \varepsilon$. □

The claim implies that the vertices of σ are among the six nearest neighbors of w . Since this is true for any edge or triangle of $\mathcal{D}^S(L)$, and since the vertices of $\mathcal{D}^S(L)$ belong to $C_v^W(L)$, $C_v^W(L)$ contains all the edges and triangles of $\mathcal{D}^S(L)$. This ends the proof of the theorem. □

The next theorem guarantees that the simplices of $C_v^W(L)$ are not too large as far as the v_i remain bounded. It follows that the size of $C_v^W(L)$ is linear in $|L|$, since L is sparse. This property can be generalized to higher dimensions, at the price of an exponential growth of the constant factor. This motivates the use of the witness complex instead of the Delaunay triangulation.

Theorem 3.8 *Assume that δ, ε satisfy $\delta + \varepsilon < \frac{\cos \theta}{\sqrt{6}} \text{Irr}_k(S)$. Then, for any point $w \in W$, the distance between w and its sixth nearest neighbor among the points of L is at most $\delta + (\frac{\sqrt{6}}{\cos \theta} + 1)(\delta + \varepsilon)$. As a consequence, for any sequence v of integers such that $v_1 \leq 6$, the total number of simplices of $C_v^W(L)$ is at most $2^{O((\frac{\delta + \varepsilon}{\varepsilon \cos \theta})^3)} |L|$, which is linear with respect to $|L|$ as far as θ is fixed and δ is within a constant factor of ε .*

Proof Let $w \in W$, and let \bar{w} be a point of S closest to w . Since W is a δ -noisy sample of S , we have $d(w, \bar{w}) \leq \delta$. We call p_1, \dots, p_l the points of L that lie in $B(\bar{w}, (\frac{\sqrt{6}}{\cos \theta} + 1)(\delta + \varepsilon))$, and $\bar{p}_1, \dots, \bar{p}_l$ their orthogonal projections onto the plane $T_k(\bar{w})$. We will prove that $l \geq 6$.

Since L is a $(\delta + \varepsilon)$ -sample of S , the balls $B_i = B(p_i, \delta + \varepsilon)$ cover $S \cap B(\bar{w}, \frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta})$ (observe that, among the balls of radius $\delta + \varepsilon$ centered at the points of L , only the B_i intersect $B(\bar{w}, \frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta})$). It follows that, inside $T_k(\bar{w})$, the disks $D_i = D(\bar{p}_i, \delta + \varepsilon)$ cover the orthogonal projection of $S \cap B(\bar{w}, \frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta})$. Now, according to the hypothesis of the lemma, we have $\frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta} < \text{Irr}_k(S)$. Thus, by Lemma 2.7, $S \cap B(\bar{w}, \frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta})$ is a topological disk whose orthogonal projection onto $T_k(\bar{w})$ contains the projection D of the intersection of $B(\bar{w}, \frac{(\delta + \varepsilon)\sqrt{6}}{\cos \theta})$ with the cone of apex \bar{w} , of axis aligned with $\mathbf{n}_k(\bar{w})$, and of half-angle $\frac{\pi}{2} - \theta$. Therefore, the D_i cover D , which is a disk of center \bar{w} and radius $(\delta + \varepsilon)\sqrt{6}$. Thus, the number of

disks D_i is at least

$$\frac{\text{Area}(D)}{\text{Area}(D_i)} = \frac{6\pi(\delta + \varepsilon)^2}{\pi(\delta + \varepsilon)^2} = 6.$$

It follows that the number of points of L that lie in $B(\bar{w}, (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon))$ is at least 6. As a result, the distance from w to its sixth nearest landmark is at most $d(w, \bar{w}) + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon) \leq \delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon)$.

Let us now bound the size of $\mathcal{C}_v^W(L)$. Let p be a point of L . From the above paragraph we deduce that the edges of $\mathcal{C}_v^W(L)$ incident to p are included in balls of radii at most $\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon)$. Hence, all edges belong to a common ball of center p and radius $r \leq 2\delta + 2(\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon)$, which is equal to $2\varepsilon(1 + \frac{\sqrt{6}}{\cos\theta} + (\frac{\sqrt{6}}{\cos\theta} + 2)\frac{\delta}{\varepsilon})$. The neighboring vertices q_1, \dots, q_l of p in $\mathcal{C}_v^W(L)$ belong to $B(p, r)$ as well. Now, since the points of L are farther than ε from one another, the q_i are centers of pairwise-disjoint balls of radius $\frac{\varepsilon}{2}$, hence their number l is at most $\frac{\text{Vol}(B(p, r))}{\text{Vol}(B(p, \frac{\varepsilon}{2}))} \leq 64(1 + \frac{\sqrt{6}}{\cos\theta} + (\frac{\sqrt{6}}{\cos\theta} + 2)\frac{\delta}{\varepsilon})^3 = O(\frac{1}{\cos^3\theta}(1 + \frac{\delta}{\varepsilon})^3)$. Since every simplex of $\mathcal{C}_v^W(L)$ incident to p is uniquely defined as a subset of $\{q_1, \dots, q_l\}$, the number of simplices of $\mathcal{C}_v^W(L)$ incident to p is at most 2^l , which gives the result. \square

3.2.2 Manifold Extraction

It follows from Theorem 3.7 that $\mathcal{C}_v^W(L)$ contains $\mathcal{D}^S(L)$, but Fig. 4 (center)⁴ shows that $\mathcal{C}_v^W(L)$ is not restricted to $\mathcal{D}^S(L)$ and contains additional simplices that are small enough to be ν -witnessed. Nevertheless, it is possible to extract from $\mathcal{C}_v^W(L)$ a simplicial surface \hat{S} isotopic to S and at Hausdorff distance $O(\varepsilon + \delta)$ of S . The extraction procedure takes a number ϱ as parameter and proceeds as follows:

1. Since the goal is to extract a 2-manifold, only the 2-skeleton of $\mathcal{C}_v^W(L)$ is considered. Since it may not be an embedded complex, we intersect it with $\mathcal{D}(L)$. The result is a pure 2-dimensional subcomplex \mathcal{C} of $\mathcal{D}(L)$.
2. To guarantee that the output simplicial surface has no skinny triangle, we delete from \mathcal{C} all the triangles whose radius-edge ratios are greater than ϱ .
3. We greedily remove from \mathcal{C} all the triangles incident to *sharp edges*. An edge is sharp if all its incident triangles in \mathcal{C} lie in a small wedge of angle at most $\frac{\pi}{2}$. This definition applies in particular to edges that are incident to one single triangle.
4. By a depth-first walk in the dual graph of the remaining part of \mathcal{C} , we extract the outer boundary of \mathcal{C} .

Observe that steps 3 and 4 correspond to the manifold extraction procedure of [1–3]. As argued in these papers, the outcome is a simplicial complex \hat{S} whose dihedral angles are greater than $\frac{\pi}{2}$. Moreover, thanks to step 2, the radius-edge ratios of the facets of \hat{S} are at most ϱ . However, two issues arise: first, by greedily removing non-Delaunay triangles or triangles with sharp edges or large radius-edge ratios

⁴Thanks to a bug in Geomview, we can see some hidden triangles in the vicinity of slivers, such as in green areas.

from $C_v^W(L)$, steps 1 through 3 might end up with an empty complex \mathcal{C} . As a result, \hat{S} may be empty. Second, the outer boundary of \mathcal{C} might not be an embedded surface, since it may contain multiple vertices or edges. By proceeding with a depth-first search on the dual graph of \mathcal{C} , step 4 duplicates multiple vertices and edges, so that the resulting complex \hat{S} is a simplicial surface whose immersion in \mathbb{R}^3 coincides with the outer boundary of \mathcal{C} .

Theorem 3.9 *Let $\varrho = 1 + \frac{1-2\sin\frac{\pi}{7}}{8\sin\frac{\pi}{7}} \approx 1.038$. Assume that $\theta < \arctan \frac{\sqrt{3}}{1+4\varrho} \approx 18.6$ deg and that $\delta < \frac{\cos^4\theta(\cos\theta-2\sin\frac{\pi}{7})}{(2\cos^2\theta+\sqrt{6}\cos\theta+4\cos\theta\sin\frac{\pi}{7}-2\sqrt{6}\sin\frac{\pi}{7})(\cos\theta+\sqrt{6})(4+3\cos^2\theta)\varrho\sqrt{3}} \text{lr}_k(S)$. If ε satisfies $\frac{8\sin\frac{\pi}{7}}{\cos\theta-2\sin\frac{\pi}{7}}\delta < \varepsilon < \frac{\cos^3\theta}{(\cos\theta+\sqrt{6})(4+3\cos^2\theta)\varrho\sqrt{3}} \text{lr}_k(S) - \frac{2\cos\theta+\sqrt{6}}{\cos\theta}\delta$, then, for any sequence v of integers such that $v_0 \geq 1$ and $v_1 = v_2 = 6$, the simplicial complex \hat{S} extracted from $C_v^W(L)$ with parameter ϱ is an embedded surface isotopic to S and at Hausdorff distance at most $(\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon)) \frac{\varrho\sqrt{3}}{\cos^2\theta}$ from S .*

The proof of the theorem follows roughly the same scheme as in [8] and it uses material that comes mainly from that paper. Here is a short overview:

- First, we show that the triangles of $\mathcal{D}^S(L)$ are not skinny and make large dihedral angles. This fact, combined with Theorem 3.7, implies that complex \mathcal{C} contains $\mathcal{D}^S(L)$ after step 3 above. We deduce that \hat{S} is not empty, since $\mathcal{D}^S(L)$ is a manifold without boundary, by Theorem 2.6.
- Second, we use Theorem 3.3(ii) of [8] to show that \hat{S} is a Lipschitz surface, which implies in particular that it is an embedded surface.
- Third, we use Proposition 6.4 of [8] to bound the Hausdorff distance $d_{\mathcal{H}}(\hat{S}, S)$ between \hat{S} and S . We show that $d_{\mathcal{H}}(\hat{S}, S)$ is small compared to the Lipschitz radii of \hat{S} and S .
- Finally, we apply Theorem 6.2 of [11] to show that \hat{S} and S are isotopic.

The rest of Sect. 3.2 is devoted to the details of the proof of Theorem 3.9 and it can therefore be skipped in a first reading. The following results from [8] will be useful:

Lemma 3.10 (Triangle Normal) *Let S be a k -Lipschitz surface, and $\theta = \arctan k$. For any $p \in S$ and for any triangle $f = (u, v, w)$ such that $u, v, w \in S \cap B(p, \text{lr}_k(S))$ and that the radius-edge ratio of f is at most ϱ , the angle between $\mathbf{n}_k(p)$ and the line orthogonal to the plane $\text{aff}(u, v, w)$ is at most $\arcsin(2\varrho \sin\theta)$.*

Lemma 3.11 (Normal Variation) *Let S be a k -Lipschitz surface, and $\theta = \arctan k$. $\forall p, q \in S$ s.t. $d(p, q) < \text{lr}_k(p)$, $(\mathbf{n}_k(p), \mathbf{n}_k(q)) \leq 2\theta$.*

Moreover, the hypotheses of Theorem 3.9 imply that the following conditions are satisfied:

$$C1 \quad \delta < \frac{\cos\theta-2\sin\frac{\pi}{7}}{8\sin\frac{\pi}{7}}\varepsilon.$$

$$C2 \quad \delta + \left(\frac{\sqrt{6}}{\cos\theta} + 1\right)(\delta + \varepsilon) < \frac{\cos^2\theta}{(4+3\cos^2\theta)\varrho\sqrt{3}} \text{Ir}_k(S).$$

\hat{S} is not Empty

Lemma 3.12 $\mathcal{D}^S(L)$ is a 2-manifold with obtuse dihedral angles, isotopic to S , at Hausdorff distance $\delta + \varepsilon$ from S , and whose triangles have radius-edge ratios of at most ϱ .

Proof Recall that L is a $(\delta + \varepsilon)$ -sample of S with $\delta + \varepsilon < \frac{1}{7}\text{Ir}_k(S)$, by C2. Therefore, hypothesis H1 of Theorem 2.6 is satisfied with ε replaced by $\delta + \varepsilon$. Let now f be a triangle of $\mathcal{D}^S(L)$, and let $B(c, r)$ be a surface Delaunay ball that circumscribes f . Since L is a $(\delta + \varepsilon)$ -sample of S , we have $r \leq \delta + \varepsilon$, which implies that the circumradius of f is bounded by $\delta + \varepsilon$. Moreover, since L is ε -sparse, the length of the shortest edge of f is at least ε . Therefore, the radius-edge ratio of f is at most $1 + \frac{\delta}{\varepsilon}$, which by C1 is less than $1 + \frac{\cos\theta - 2\sin\frac{\pi}{7}}{8\sin\frac{\pi}{7}} \leq \varrho$. This property holds for every triangle f of $\mathcal{D}^S(L)$. Since in addition we assumed that $\theta < \arctan\frac{\sqrt{3}}{1+4\varrho}$, we have $2\varrho \sin\theta < \cos(\theta + \frac{\pi}{6}) < \cos(2\theta)$, and thus hypothesis H2 of Theorem 2.6 is satisfied as well. It follows then from Theorem 2.6 that $\mathcal{D}^S(L)$ is a 2-manifold isotopic to S and at Hausdorff distance $\delta + \varepsilon$ of S . Moreover, the oriented normals of $\mathcal{D}^S(L)$ approximate the k -Lipschitz normals of S within an angle of $\arcsin(2\varrho \sin\theta)$. This means that, for any triangles f_1 and f_2 of $\mathcal{D}^S(L)$ that share an edge, for any common vertex v , the normals $\mathbf{n}(f_1)$ and $\mathbf{n}(f_2)$ make angles of at most $\arcsin(2\varrho \sin\theta)$ with the k -Lipschitz normal $\mathbf{n}_k(v)$. Since $\theta < \arctan\frac{\sqrt{3}}{1+4\varrho}$, $\arcsin(2\varrho \sin\theta)$ is less than $\frac{\pi}{4}$, hence the dihedral angle between f_1 and f_2 is greater than $\frac{\pi}{2}$. \square

Let us now prove that the closed 2-manifold $\mathcal{D}^S(L)$ is included in complex \mathcal{C} upon termination of step 3 of the extraction process. This will imply that \hat{S} is not empty, since it coincides with the outer boundary of \mathcal{C} , which encloses $\mathcal{D}^S(L)$. According to Theorem 3.7, $\mathcal{D}^S(L)$ is included in $\mathcal{C}_v^W(L)$ initially. By definition, $\mathcal{D}^S(L)$ is a subcomplex of $\mathcal{D}(L)$, and by Lemma 3.12 it is a 2-manifold (and thus a pure 2-dimensional complex). Thus, $\mathcal{D}^S(L)$ is included in \mathcal{C} upon termination of step 1 of the extraction process. Still by Lemma 3.12, the radius-edge ratios of the triangles of $\mathcal{D}^S(L)$ are at most ϱ , and $\mathcal{D}^S(L)$ has obtuse dihedral angles. As a consequence, $\mathcal{D}^S(L)$ remains in \mathcal{C} throughout steps 2 and 3.

\hat{S} is a Lipschitz Surface

As argued in Sect. 3.2.2, the subcomplex \hat{S} extracted from $\mathcal{C}_v^W(L)$ is a combinatorial surface, but its immersion in \mathbb{R}^3 may not be an embedding. However, it coincides with the outer boundary of complex \mathcal{C} , thus, by moving the vertices of \hat{S} slightly in \mathbb{R}^3 , one can easily turn the immersion into an embedding. Therefore, up to an infinitesimal perturbation of the vertices of \hat{S} , we consider from now on that \hat{S} is an embedded surface. This is needed for the sake of our proofs, but we will see after Lemma 3.15 that actually no perturbation is needed to embed \hat{S} .

C2 and Theorem 3.8 imply that any triangle f of $\mathcal{C}_v^W(L)$ is included in a ball of radius $\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon)$. Since the radius-edge ratio of f is at most ϱ , a quick computation shows that the circumradius of f is at most $(\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon))\varrho\sqrt{3}$, which by C2 is less than $\frac{\cos^2\theta}{4+3\cos^2\theta}\text{lr}_k(S) < \frac{1}{7}\text{lr}_k(S)$.

Orientation Convention 3.13 *We orient every triangle f of \hat{S} such that $\mathbf{n}(f) \cdot \mathbf{n}_k(v) \geq 0$ for each vertex v of f .*

The existence of such an orientation follows from C2 and Lemmas 3.10 and 3.11. We will now prove that this convention induces a valid orientation of \hat{S} .

Lemma 3.14 *The Orientation Convention 3.13 induces a valid orientation of \hat{S} .*

Proof It suffices to prove that the orientation of the triangles of \hat{S} is consistent across the edges of \hat{S} , i.e., that the orientations of any two adjacent triangles of \hat{S} induce opposite orientations of their common edge. Let f and f' be two triangles of \hat{S} with a common edge e . Let v be a vertex of e . We call f_v and f'_v the respective orthogonal projections of f and f' onto the plane $T_k(v)$. Inside $T_k(v)$, we endow triangles f_v and f'_v with the same orientations as f and f' respectively. The Orientation Convention 3.13 states that $\mathbf{n}(f) \cdot \mathbf{n}_k(v)$ and $\mathbf{n}(f') \cdot \mathbf{n}_k(v)$ are both positive, hence so are $\mathbf{n}(f_v) \cdot \mathbf{n}_k(v)$ and $\mathbf{n}(f'_v) \cdot \mathbf{n}_k(v)$. Now, by the triangle normal Lemma 3.10, the planes $\text{aff}(f)$ and $\text{aff}(f')$ make angles of at most $\arcsin(2\varrho \sin\theta) < \frac{\pi}{4}$ with $T_k(v)$. Since the dihedral angle between f and f' is at least $\frac{\pi}{2}$, f_v and f'_v lie on either sides of the projection of e onto $T_k(v)$. It follows that the orientations of f_v and f'_v induce opposite orientations of their common edge. Similarly, the orientations of f and f' induce opposite orientations of e . Since this is true for any two triangles of \hat{S} with a common edge, \hat{S} is consistently oriented by the orientation convention. \square

Lemma 3.15 *\hat{S} is a k' -Lipschitz surface with $\text{lr}_{k'}(\hat{S}) \geq \text{lr}_k(S) - 3\varrho\sqrt{3}(\delta + (\frac{\sqrt{6}}{\cos\theta} + 1) \times (\delta + \varepsilon))$.*

Proof Let $d = (\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon))\varrho\sqrt{3}$. For any point $p \in \hat{S}$ and any triangle f of \hat{S} intersecting $B(p, \text{lr}_k(S) - 3d)$, the circumradius of f is at most d , hence f is included in $B(p, \text{lr}_k(S) - d)$. Let f_p be a triangle of \hat{S} containing p , and let v be a vertex of f_p closest to p . We have $d(p, v) \leq d$, hence f is included in $B(v, \text{lr}_k(S))$. Since the radius-edge ratio of f is at most ϱ , the triangle normal Lemma 3.10 states that $(\mathbf{n}(f), \mathbf{n}_k(v)) \leq \arcsin(2\varrho \sin\theta)$. Since this is true for any point $p \in \hat{S}$ and any triangle $f \in \hat{S}$ intersecting $B(p, \text{lr}_k(S) - 3d)$, Theorem 3.3(ii) of [8] states that \hat{S} is a k' -Lipschitz surface with $k' = \tan(\arcsin(2\varrho \sin\theta))$ and $\text{lr}_{k'}(\hat{S}) \geq \text{lr}_k(S) - 3d$. \square

The fact that \hat{S} is a Lipschitz surface implies that any infinitesimal perturbation of its vertices yields another Lipschitz (and thus embedded) surface. As a consequence, no perturbation of the vertices of \hat{S} is in fact needed to embed \hat{S} .

Hausdorff Distance $d_{\mathcal{H}}(\hat{S}, S)$

Let $d = (\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon))\varrho\sqrt{3}$. To apply Proposition 6.4 of [8] and bound $d_{\mathcal{H}}(S, \hat{S})$, we need to check that the following conditions are satisfied:

- The vertices of \hat{S} lie on S : this is true, since \hat{S} is a subcomplex of $\mathcal{C}_v^W(L)$ and since $L \subset S$, by assumption.
- The circumradii of the triangles of \hat{S} are at most d with $d < \frac{1}{7}\text{lr}_k(S)$: this is ensured by C2 and Theorem 3.8.
- For any triangle f of \hat{S} and any vertex v of f , the angle between $\mathbf{n}(f)$ and $\mathbf{n}_k(v)$ is less than $\frac{\pi}{3} - \theta$: this is ensured by the triangle normal Lemma 3.10 and by the fact that $\theta < \arctan \frac{\sqrt{3}}{1+4\varrho}$.
- \hat{S} has vertices on all the connected components of S : this is true because \hat{S} encloses complex \mathcal{C} , which contains the closed surface $\mathcal{D}^S(L)$ that has vertices on all the connected components of S (by Lemma 3.12), and because the triangles of \mathcal{C} are too small to connect distinct connected components of S .

It follows then from Proposition 6.4 of [8] that

$$d_{\mathcal{H}}(S, \hat{S}) \leq \frac{d}{\cos^2\theta} = \left(\delta + \left(\frac{\sqrt{6}}{\cos\theta} + 1 \right) (\delta + \varepsilon) \right) \frac{\varrho\sqrt{3}}{\cos^2\theta}. \tag{1}$$

\hat{S} is isotopic to S

By (1), we have $d_{\mathcal{H}}(\hat{S}, S) \leq \frac{d}{\cos^2\theta}$, where $d = (\delta + (\frac{\sqrt{6}}{\cos\theta} + 1)(\delta + \varepsilon))\varrho\sqrt{3}$. By C2, this quantity is less than $\frac{1}{4}(\text{lr}_k(S) - 3d)$, which, by Lemma 3.15, is at most $\frac{1}{4} \min\{\text{lr}_k(S), \text{lr}_k(\hat{S})\}$. As a consequence, $d_{\mathcal{H}}(\hat{S}, S)$ is less than half the so-called *weak feature sizes* of $\mathbb{R}^3 \setminus \hat{S}$ and of $\mathbb{R}^3 \setminus S$, by Theorem 3.8 of [8]. This implies that \hat{S} and S are isotopic, by Theorem 6.2 of [11].

3.3 Dealing with Noisy Data

Our previous results hold as far as the set L of landmarks lies on the surface S . This property can be ensured by assuming that the set W of witnesses, which contains L , is noise-free. However, Theorem 3.16 below shows that the property can still be ensured when W, L are noisy samples of S . In the sequel, $\lambda_0 \approx 0.078$ denotes the smallest positive root of the polynomial $64\lambda^6 + 832\lambda^5 + 1008\lambda^4 - 160\lambda^3 - 4\lambda^2 - 12\lambda + 1$. Moreover, given any $\theta \in [0, \frac{\pi}{2}]$, $\lambda(\theta)$ denotes the smallest positive root of $16(4\sin^2\theta - 1)\lambda^6 + 32\lambda^5 - 12(2 + 3\sin^2\theta)\lambda^4 + 8\lambda^3 + (4\sin^2\theta + 63)\lambda^2 + 64\lambda - 16$.

Theorem 3.16 *Let S be a k -Lipschitz surface and L a δ -noisy ε -sparse $(\delta + \varepsilon)$ -sample of S . Assume that δ, ε satisfy the following conditions, where $\theta = \arctan k$:*

- N1 $\delta < \min\{\frac{1}{4}, \frac{\sqrt{2}-4\sin\theta}{2(\sqrt{2}+4\sin\theta)}, \frac{\cos(2\theta)-2\sin\theta}{2(\cos(2\theta)+2\sin\theta)}, \lambda_0, \lambda(\theta)\}\varepsilon$.
- N2 $\delta < \min\{\frac{1}{14}\text{lr}_k(S) - \frac{1}{2}\varepsilon, \frac{1}{6}\text{lr}_k(S) - \frac{7}{12}\varepsilon\}$.

Then, there exists a k' -Lipschitz surface S' , passing through the points of L , isotopic to S , and at Hausdorff distance of S at most $\varepsilon + 3\delta$, such that:

$$\begin{cases} k' = \tan\left(\arcsin\left(\frac{4((1+2\sin\theta)\frac{\delta}{\varepsilon} + \sin\theta)}{2\sin\left(\arcsin\frac{1-2\frac{\delta}{\varepsilon}}{2(1+2\frac{\delta}{\varepsilon})} - 2\arcsin\frac{2\frac{\delta}{\varepsilon}}{1-2\frac{\delta}{\varepsilon}}\right)}\right)\right) = O\left(k + \frac{\delta}{\varepsilon}\right), \\ \text{lr}_{k'}(S') \geq \text{lr}_k(S) - (2\varepsilon + 7\delta) > \frac{1}{2}\text{lr}_k(S). \end{cases}$$

If the set W of witnesses is a δ -noisy δ -sample of a k -Lipschitz surface S for some sufficiently small δ (as compared to $\text{lr}_k(S)$), then Theorem 3.16 ensures that there exists an interval of values of ε such that any ε -sparse ε -sample L of W lies on a k' -Lipschitz surface S' with $k' = O(k + \frac{\delta}{\varepsilon})$ and $\text{lr}_{k'}(S') = \Omega(\text{lr}_k(S))$. Theorem 3.9 can then be applied to S', W, L , provided that k, ε , and $\frac{\delta}{\varepsilon}$ are small enough to meet its hypothesis. Then, because S' is isotopic to S and close to it for the Hausdorff distance, the conclusion of Theorem 3.9 holds for S, W, L as well, with a slightly worse bound on the Hausdorff distance. There also exists a version of Theorem 3.16 for Lipschitz curves, which can be combined with the structural results of Sect. 3.1. The proof is roughly the same.

The rest of Sect. 3.3 is devoted to the proof of Theorem 3.16, which builds an isotopy $\phi : [0, 1] \times S \rightarrow \mathbb{R}^3$ such that $S' = \phi(1, S)$ is a k' -Lipschitz surface passing through the points of L with $k' = O(k + \frac{\delta}{\varepsilon})$ and $\text{lr}_{k'}(S') = \Omega(\text{lr}_k(S))$. Intuitively, since the points of L lie ε away from one another, with ε large compared to the amplitude δ of the noise, the surface S can be snapped onto the points of L without changing its normals too much. This can be easily seen on simple examples, such as for instance when S is the x -axis in \mathbb{R}^2 (in this case, the snapped curve is the polygonal chain connecting the points of L in the order of their abscissae).

To build the isotopy ϕ , we first map the points of L to their nearest neighbors on S . If some point $p \in L$ has two or more nearest neighbors on S , then we choose either of them. We get a new point set $\tilde{L} \subset S$, together with a surjective map $\mu : L \rightarrow \tilde{L}$. The restricted Delaunay triangulation of $\tilde{L}, \mathcal{D}^S(\tilde{L})$, will play a prominent role in our construction.

Lemma 3.17 *Under the hypotheses of Theorem 3.16, for any distinct points p, q of L , $d(\mu(p), \mu(q)) \geq \varepsilon - 2\delta > 0$.*

Proof Since L is δ -noisy, we have $d(p, \mu(p)) = d(p, S) \leq \delta$ and $d(q, \mu(q)) = d(q, S) \leq \delta$. Moreover, since L is ε -sparse, we have $d(p, q) \geq \varepsilon$. It follows by the triangle inequality that $d(\mu(p), \mu(q)) \geq \varepsilon - 2\delta$, which is positive, by N1. \square

It follows from this lemma that μ is one-to-one and hence bijective. Therefore, we can define a function $\tilde{\phi} : [0, 1] \times \tilde{L} \rightarrow \mathbb{R}^3$ as follows:

$$\forall t \in [0, 1], \forall p \in \tilde{L}, \quad \tilde{\phi}(t, p) = (1 - t)p + t\mu^{-1}(p).$$

We can then extend $\tilde{\phi}$ to $[0, 1] \times \mathcal{D}^S(\tilde{L})$ by linearity:

$$\forall t \in [0, 1], \forall p \in \mathcal{D}^S(\tilde{L}), \quad \tilde{\phi}(t, p) = \lambda_u \tilde{\phi}(t, u) + \lambda_v \tilde{\phi}(t, v) + \lambda_w \tilde{\phi}(t, w),$$

where $\lambda_u, \lambda_v, \lambda_w$ are the barycentric coordinates of p in any triangle (u, v, w) of $\mathcal{D}^S(\tilde{L})$ containing p . If p belongs to several triangles, then $\tilde{\phi}(t, p)$ does not depend on the choice of the triangle, since $\mathcal{D}^S(\tilde{L})$ is an embedded simplicial complex. Moreover, $\tilde{\phi}$ is a continuous function in variables (t, p) .

Since L is a δ -noisy $(\varepsilon + \delta)$ -sample of S , \tilde{L} is a noise-free $(\varepsilon + 2\delta)$ -sample of S with $\varepsilon + 2\delta < \frac{1}{7}\text{lr}_{\tilde{k}}(S)$, by N2. Therefore, hypothesis H1 of Theorem 2.6 is satisfied. Moreover, by Lemma 3.17, \tilde{L} is $(\varepsilon - 2\delta)$ -sparse, hence the radius-edge ratios of the triangles of $\mathcal{D}^S(\tilde{L})$ are at most $\tilde{\varrho} = \frac{\varepsilon + 2\delta}{\varepsilon - 2\delta}$, which by N1 is less than $\frac{\cos(2\theta)}{2\sin\theta}$. It follows that hypothesis H2 is also satisfied, which implies that $\mathcal{D}^S(\tilde{L})$ is isotopic to S , by Theorem 2.6. This means that there is an isotopy $I : [0, 1] \times S \rightarrow \mathbb{R}^3$ such that $I(1, S) = \mathcal{D}^S(\tilde{L})$. We define ϕ by applying I and $\tilde{\phi}$ successively, at double speed:

$$\begin{cases} \forall t \in [0, \frac{1}{2}], \forall p \in S, & \phi(t, p) = I(2t, p), \\ \forall t \in [\frac{1}{2}, 1], \forall p \in S, & \phi(t, p) = \tilde{\phi}(2t - 1, I(1, p)). \end{cases}$$

This is the standard way of composing isotopies. In particular, ϕ is well defined, since $I(1, p) \in \mathcal{D}^S(\tilde{L})$ for all $p \in S$. To prove that ϕ is an isotopy, it suffices to show that $\tilde{\phi}$ itself is an isotopy, which boils down to showing that $\tilde{\phi}(t, \cdot)$ is one-to-one as a function of p for all $t \in [0, 1]$, since we already know that $\tilde{\phi}$ is continuous and since $\mathcal{D}^S(\tilde{L})$ is compact. We will use an intermediate result:

Lemma 3.18 *Under the hypotheses of Theorem 3.16, the following assertions hold:*

- (i) *For any two vertices u, v of $\mathcal{D}^S(\tilde{L})$, $d(u, v) \geq \varepsilon - 2\delta$.*
- (ii) *For any vertex v and any non-incident simplex (edge or triangle) σ of $\mathcal{D}^S(\tilde{L})$, the distance between v and σ is at least $\varepsilon \frac{(1 - 2\frac{\delta}{\varepsilon})^2}{2(1 + 2\frac{\delta}{\varepsilon})}$.*
- (iii) *For any two disjoint edges e, e' of $\mathcal{D}^S(\tilde{L})$, the distance between e and e' is bounded from below by $\varepsilon \frac{(1 - 2\frac{\delta}{\varepsilon})^2}{2(1 + 2\frac{\delta}{\varepsilon})} \sqrt{1 - 4 \frac{(1 + 2\frac{\delta}{\varepsilon})^2}{(1 - 2\frac{\delta}{\varepsilon})^2} \sin^2 \theta}$.*

Proof (i) is a direct consequence of the fact that \tilde{L} is $(\varepsilon - 2\delta)$ -sparse.

Concerning (ii), it is proved in [8, Lemma 5.8] that, since assumptions H1–H2 of Theorem 2.6 are satisfied, $\mathcal{D}^S(\tilde{L})$ is a \tilde{k} -Lipschitz surface with $\tilde{k} = \tan(\arcsin(2\tilde{\varrho} \sin\theta))$ and $\text{lr}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})) \geq \text{lr}_k(S) - 3(\varepsilon + 2\delta)$. Let v be a vertex of $\mathcal{D}^S(\tilde{L})$, and σ a simplex (edge or triangle) of $\mathcal{D}^S(\tilde{L})$ that is not incident to v . If σ does not intersect $B(v, \text{lr}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})))$, then its distance to v is at least $\text{lr}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})) \geq \text{lr}_k(S) - 3(\varepsilon + 2\delta)$, which by N2 is greater than $\frac{\varepsilon}{2} \geq \varepsilon \frac{(1 - 2\frac{\delta}{\varepsilon})^2}{2(1 + 2\frac{\delta}{\varepsilon})}$. Otherwise, the point $p \in \sigma$ closest to v lies inside the ball $B(v, \text{lr}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})))$. Since $\mathcal{D}^S(\tilde{L}) \cap B(v, \text{lr}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})))$ coincides with the graph of a bivariate function defined over the \tilde{k} -Lipschitz support plane of $\mathcal{D}^S(\tilde{L})$ at v (called $T_{\tilde{k}}(v)$, for short), the orthogonal projection \bar{p} of p onto $T_{\tilde{k}}(v)$ lies outside $\bar{U}(v)$, the orthogonal projection of the umbrella $U(v)$ of triangles of $\mathcal{D}^S(\tilde{L})$ incident to v . Since v lies in the interior of $\bar{U}(v)$, the line segment $[v, \bar{p}]$ intersects the boundary of $\bar{U}(v)$. Let \bar{q} be a point of intersection, and q the point of the boundary of $U(v)$ that projects onto \bar{q} . Inside the

plane that contains v , p and \bar{p} (and hence also q and \bar{q}), $\mathcal{D}^S(\tilde{L}) \cap B(v, \text{Ir}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})))$ coincides with the graph of a \tilde{k} -Lipschitz univariate function defined over the line (v, \bar{p}) . Since⁵ $\tilde{k} < 1$, the fact that $\bar{q} \in [v, \bar{p}]$ implies that $d(v, q) \leq d(v, p)$. Hence, $d(v, p)$ is at least the distance between v and the boundary of $U(v)$, which we must now bound from below. For any triangle (u, v, w) of $\mathcal{D}^S(\tilde{L})$ incident to v , the distance h from v to line (u, w) , also known as the height of triangle (u, v, w) , is given by $d(v, u) \sin \widehat{vu}w$, and it is a lower bound on the distance between v and edge $[u, w]$. According to (i), we have $d(v, u) \geq \varepsilon - 2\delta$. Moreover, since the radius-edge ratios of the triangles are at most $\bar{\varrho} = \frac{\varepsilon+2\delta}{\varepsilon-2\delta}$, every inner angle of (u, v, w) is at least $\arcsin \frac{\varepsilon-2\delta}{2(\varepsilon+2\delta)}$. Therefore, we have $\widehat{vu}w \geq \arcsin \frac{\varepsilon-2\delta}{2(\varepsilon+2\delta)}$ and $h \geq \frac{(\varepsilon-2\delta)^2}{2(\varepsilon+2\delta)} = \varepsilon \frac{(1-2\frac{\delta}{\varepsilon})^2}{2(1+2\frac{\delta}{\varepsilon})}$. Since this is true for any triangle of $\mathcal{D}^S(\tilde{L})$ incident to v , the distance between v and the boundary of $U(v)$ is at least $\varepsilon \frac{(1-2\frac{\delta}{\varepsilon})^2}{2(1+2\frac{\delta}{\varepsilon})}$, which proves (ii).

By the same argument, for any disjoint edges e, e' of $\mathcal{D}^S(L)$, and for any point $p \in e$, the distance from p to e' is at least the distance from p to the boundary of $U(e) = U(u) \cup U(v)$, where u, v are the vertices of e . Therefore, the proof of (iii) reduces to finding a lower bound on the distance between e and the points of $\partial U(e)$. Let q be the point of $\partial U(e)$ whose orthogonal projection \bar{q} onto the \tilde{k} -Lipschitz support plane $T_{\tilde{k}}(v)$ minimizes the distance to the orthogonal projection $\bar{e} = [\bar{u}, v]$ of e . The distance from any point $p \in \partial U(e)$ to e is at least the distance from its orthogonal projection \bar{p} to \bar{e} , which is bounded from below by $d(\bar{q}, \bar{e})$. Moreover, we have $d(\bar{q}, \bar{e}) \geq d(q, e) \cos(\arctan \tilde{k})$, since $U(e) \subset \mathcal{D}^S(\tilde{L}) \cap B(v, \text{Ir}_{\tilde{k}}(\mathcal{D}^S(\tilde{L})))$, which coincides with the graph of a \tilde{k} -Lipschitz bivariate function defined over $T_{\tilde{k}}(v)$. Therefore, the distance of any point of $\partial U(e)$ to e is at least $d(q, e) \cos(\arctan \tilde{k}) = d(q, e) \sqrt{1 - 4 \frac{(1+2\frac{\delta}{\varepsilon})^2}{(1-2\frac{\delta}{\varepsilon})^2} \sin^2 \theta}$. We must find a lower bound on $d(q, e)$. Let $\Sigma \subset T_{\tilde{k}}(v)$ be the strip whose boundaries are the lines orthogonal to \bar{e} passing through \bar{u} and v . For any point $\bar{p} \notin \Sigma$, the nearest neighbor of \bar{p} on \bar{e} is either \bar{u} or v . Hence, if $\bar{q} \notin \Sigma$, then $d(q, e) \geq \varepsilon \frac{(1-2\frac{\delta}{\varepsilon})^2}{2(1+2\frac{\delta}{\varepsilon})}$, by (ii). For any point $\bar{p} \in \Sigma$, the nearest neighbor of \bar{p} on \bar{e} is also its nearest neighbor on the line $\text{aff}(\bar{e})$. Hence, if $\bar{q} \in \Sigma$, then \bar{q} is either a vertex of $\partial U(e)$ (meaning that q is a vertex of $\partial U(e)$) or a point of $\partial \Sigma$ (meaning that the nearest neighbor of \bar{q} on \bar{e} is either \bar{u} or v). In both cases, we have $d(q, e) \geq \varepsilon \frac{(1-2\frac{\delta}{\varepsilon})^2}{2(1+2\frac{\delta}{\varepsilon})}$, by (ii). This proves (iii). \square

Let $t \in [0, 1]$. Observe that, since L is a δ -noisy sample of S , $\tilde{\varphi}(t, \cdot)$ moves the points of \tilde{L} by at most δ . It follows, by convexity,⁶ that every point of $\mathcal{D}^S(\tilde{L})$ is moved by at most δ . Using Lemma 3.18, we can prove that $\tilde{\varphi}(t, \cdot)$ is one-to-one:

- For any simplex (vertex, edge, or triangle) σ of $\mathcal{D}^S(\tilde{L})$, the restriction of $\tilde{\varphi}(t, \cdot)$ to σ is one-to-one. Indeed, since $\tilde{\varphi}(t, \cdot)$ is linear on σ , its restriction to σ is one-to-one, provided that $\tilde{\varphi}(t, \sigma)$ has the same dimension as σ . This is clearly the case

⁵By N1, we have $\frac{\varepsilon+2\delta}{\varepsilon-2\delta} < \frac{\sqrt{2}}{4 \sin \theta}$, which implies that $\arcsin(2\bar{\varrho} \sin \theta) = \arcsin(2 \frac{\varepsilon+2\delta}{\varepsilon-2\delta} \sin \theta) < \frac{\pi}{4}$.

⁶ $\forall p \in \mathcal{D}^S(\tilde{L})$, $(\tilde{\varphi}(t, p) - p)$ is indeed a convex combination of the $\{(\tilde{\varphi}(t, q) - q), q \in \tilde{L}\}$.

if σ is a vertex. If σ is an edge $[u, v]$, then, since by N1 we have $2\delta < \varepsilon - 2\delta$, $\tilde{\phi}(t, u)$ and $\tilde{\phi}(t, v)$ are distinct, by Lemma 3.18(i). If σ is a triangle (u, v, w) , then we must show that $\tilde{\phi}(t, u)$, $\tilde{\phi}(t, v)$, and $\tilde{\phi}(t, w)$ are not collinear. Since every vertex of the triangle moves by at most δ , vertices u and v move by at most 2δ relative to w . Thus, each of them reduces the angle \widehat{uwv} by at most $\arcsin \frac{2\delta}{\varepsilon - 2\delta}$, since the vertices are at least $\varepsilon - 2\delta$ away from one another initially. Recall that the radius-edge ratio of (u, v, w) is at most $\tilde{\varrho} = \frac{\varepsilon + 2\delta}{\varepsilon - 2\delta}$, hence angle \widehat{uwv} is at least $\arcsin \frac{1}{2\tilde{\varrho}} = \arcsin \frac{\varepsilon - 2\delta}{2(\varepsilon + 2\delta)}$ initially. It follows that \widehat{uwv} does not vanish as u, v, w are moved by $\tilde{\phi}$, since by N1⁷ we have $2 \arcsin \frac{2\delta}{\varepsilon - 2\delta} < \arcsin \frac{\varepsilon - 2\delta}{2(\varepsilon + 2\delta)}$. Since this is true for any inner angle of (u, v, w) , $\tilde{\phi}(t, (u, v, w))$ is a nondegenerate triangle.

- Since the restriction of $\tilde{\phi}(t, \cdot)$ to any simplex of $\mathcal{D}^S(\tilde{L})$ is one-to-one, $\tilde{\phi}(t, \cdot)$ itself is one-to-one provided that the images of any two disjoint simplices of $\mathcal{D}^S(\tilde{L})$ (which are simplices of same dimensions) are disjoint, and that the images of any two non-disjoint simplices of $\mathcal{D}^S(\tilde{L})$ have disjoint relative interiors. Assume for a contradiction that two disjoint (resp. non-disjoint) simplices σ_1, σ_2 of $\mathcal{D}^S(\tilde{L})$ have non-disjoint images (resp. images with non-disjoint relative interiors) through $\tilde{\phi}(t, \cdot)$. Since $\sigma_1 \cap \sigma_2 = \emptyset$ (resp. $\text{int}(\sigma_1) \cap \text{int}(\sigma_2) = \emptyset$), and since $\tilde{\phi}$ moves the points of $\mathcal{D}^S(\tilde{L})$ continuously at some time $t' \in [0, t]$, either $\tilde{\phi}(t', v_1) \in \tilde{\phi}(t', \sigma_2)$ for some vertex v_1 of σ_1 non-incident to σ_2 , or $\tilde{\phi}(t', v_2) \in \tilde{\phi}(t', \sigma_1)$ for some vertex v_2 of σ_2 non-incident to σ_1 , or $\tilde{\phi}(t', e_1) \cap \tilde{\phi}(t', e_2) \neq \emptyset$ for two disjoint edges $e_1 \in \sigma_1$ and $e_2 \in \sigma_2$. In all three cases, the two simplices involved (namely, v_1 and σ_2 , or v_2 and σ_1 , or e_1 and e_2) lie at most 2δ away from each other in $\mathcal{D}^S(\tilde{L})$, since $\tilde{\phi}(t, \cdot)$ does not move the points of $\mathcal{D}^S(\tilde{L})$ by more than δ . This contradicts

Lemma 3.18, since by N1⁸ we have $\delta < \frac{\varepsilon}{2} \frac{(1 - 2\frac{\delta}{\varepsilon})^2}{(1 + 2\frac{\delta}{\varepsilon})^2} \sqrt{1 - 4 \frac{(1 + 2\frac{\delta}{\varepsilon})^2}{(1 - 2\frac{\delta}{\varepsilon})^2} \sin^2 \theta}$. Therefore,

disjoint simplices of $\mathcal{D}^S(\tilde{L})$ have disjoint images through $\tilde{\phi}(t, \cdot)$, and non-disjoint simplices of $\mathcal{D}^S(\tilde{L})$ have images through $\tilde{\phi}(t, \cdot)$ with disjoint relative interiors, which implies that $\tilde{\phi}(t, \cdot)$ is one-to-one.

It follows that $\tilde{\phi}$ is an isotopy from $\mathcal{D}^S(\tilde{L})$ to $S' = \tilde{\phi}(1, \mathcal{D}^S(\tilde{L}))$. As a consequence, ϕ is an isotopy from S to S' . In addition, since \tilde{L} is a noise-free $(\varepsilon + 2\delta)$ -sample of S , we have $d_{\mathcal{H}}(S, \mathcal{D}^S(\tilde{L})) \leq \varepsilon + 2\delta$, by Theorem 5.4 of [8]. Moreover, since $\tilde{\phi}(1, \cdot)$ moves the points of $\mathcal{D}^S(\tilde{L})$ by at most δ , we have $d_{\mathcal{H}}(\mathcal{D}^S(\tilde{L}), S') \leq \delta$. As a result, the Hausdorff distance between S and S' is at most $\varepsilon + 3\delta$. For each oriented simplex $\sigma = [v_0, \dots, v_l]$ of $\mathcal{D}^S(\tilde{L})$, we set the orientation of simplex $\tilde{\phi}(1, \sigma)$ to $[\tilde{\phi}(1, v_0), \dots, \tilde{\phi}(1, v_l)]$. Since $\mathcal{D}^S(\tilde{L})$ is an oriented surface and $\tilde{\phi}$ is an isotopy from $\mathcal{D}^S(\tilde{L})$ to S' , we get a valid orientation of S' . We conclude the proof of Theorem 3.16 with the following lemma:

Lemma 3.19 *Under the hypotheses of Theorem 3.16, S' is a k' -Lipschitz surface, with $k' = \tan(\arcsin(\frac{4((1+2\sin\theta)\frac{\delta}{\varepsilon} + \sin\theta)}{2\sin(\arcsin \frac{1-2\frac{\delta}{\varepsilon}}{2(1+2\frac{\delta}{\varepsilon})} - 2\arcsin \frac{2\frac{\delta}{\varepsilon}}{1-2\frac{\delta}{\varepsilon}})))$ and $\text{lr}_{k'}(S') \geq \text{lr}_k(S) - (2\varepsilon + 7\delta)$.*

⁷More precisely, by the fact that $\delta < \lambda_0\varepsilon$.

⁸More precisely, by the fact that $\delta < \lambda(\theta)\varepsilon$.

Proof The proof is roughly the same as for Lemma 5.8 of [8]. Let f' be a triangle of S' , and \tilde{f} the triangle of $\mathcal{D}^S(\tilde{L})$ such that $f' = \tilde{\phi}(1, \tilde{f})$. Since \tilde{L} is an $(\varepsilon + 2\delta)$ -sample of S , \tilde{f} is circumscribed by a ball of radius at most $\varepsilon + 2\delta$. And since $\tilde{\phi}(1, \cdot)$ does not move the vertices of f by more than δ , f' is included in a ball of radius at most $\varepsilon + 3\delta$, although its circumradius might be much larger. Moreover, as shown above in the proof of isotopy, the radius-edge ratio of f' is at most $q' = \frac{1}{2 \sin(\arcsin \frac{\varepsilon - 2\delta}{2(\varepsilon + 2\delta)} - 2 \arcsin \frac{2\delta}{\varepsilon - 2\delta})}$.

Let p' be a point of S' , and p a point of S closest to p' . Since $d_{\mathcal{H}}(S, \mathcal{D}^S(\tilde{L})) \leq \varepsilon + 2\delta$ and $d_{\mathcal{H}}(\mathcal{D}^S(\tilde{L}), S') \leq \delta$, we have $d(p, p') \leq \varepsilon + 3\delta$. Let $r = \text{lr}_k(S) - (2\varepsilon + 7\delta)$. For any triangle $f' = (u', v', w')$ of S' that intersects $B(p', r)$, f' intersects $B(p, r + \varepsilon + 3\delta)$ and is therefore included in $B(p, r + 2\varepsilon + 6\delta)$. Let $\tilde{f} = (\tilde{u}, \tilde{v}, \tilde{w})$ be the triangle of $\mathcal{D}^S(\tilde{L})$ such that $f' = \tilde{\phi}(1, \tilde{f})$. Since $\tilde{\phi}(1, \cdot)$ does not move the vertices of \tilde{f} by more than δ , \tilde{f} is included in $B(p, r + 2\varepsilon + 7\delta) = B(p, \text{lr}_k(S))$. Therefore, lines (\tilde{u}, \tilde{v}) , (\tilde{v}, \tilde{w}) , and (\tilde{w}, \tilde{u}) make angles of at least $\frac{\pi}{2} - \theta$ with the k -Lipschitz normal $\mathbf{n}_k(p)$ of S at p . As a consequence,

$$\begin{aligned} |(\mathbf{w}' - \mathbf{v}') \cdot \mathbf{n}_k(p)| &\leq |(\mathbf{w}' - \tilde{\mathbf{w}}) \cdot \mathbf{n}_k(p)| + |(\tilde{\mathbf{w}} - \tilde{\mathbf{v}}) \cdot \mathbf{n}_k(p)| + |(\tilde{\mathbf{v}} - \mathbf{v}') \cdot \mathbf{n}_k(p)| \\ &\leq \delta + \|\tilde{\mathbf{w}} - \tilde{\mathbf{v}}\| \sin \theta + \delta. \end{aligned}$$

Since \tilde{L} is an $(\varepsilon + 2\delta)$ -sample of S , we have $\|\tilde{\mathbf{w}} - \tilde{\mathbf{v}}\| \leq 2\varepsilon + 4\delta$. Moreover, since L is ε -sparse, we have $\|\mathbf{w}' - \mathbf{v}'\| \geq \varepsilon$. It follows that the angle between $\mathbf{n}_k(p)$ and line (v', w') is:

$$\begin{aligned} \arccos \frac{|(\mathbf{w}' - \mathbf{v}') \cdot \mathbf{n}_k(p)|}{\|\mathbf{w}' - \mathbf{v}'\|} &\geq \arccos \frac{2\delta + (2\varepsilon + 4\delta) \sin \theta}{\varepsilon} \\ &= \frac{\pi}{2} - \arcsin \left(2(1 + 2 \sin \theta) \frac{\delta}{\varepsilon} + 2 \sin \theta \right). \end{aligned}$$

The same relation holds for lines (u', v') and (u', w') as well. Since the radius-edge ratio of f' is at most q' , the triangle normal Lemma 3.10 states that the angle between $\mathbf{n}_k(p)$ and the oriented normal $\mathbf{n}(f')$ of f' is at most $\arcsin(4q'((1 + 2 \sin \theta) \frac{\delta}{\varepsilon} + \sin \theta))$. Since this is true for any triangle of S' that intersects $B(p, r)$, Theorem 3.3(ii) of [8] states that S' is a k' -Lipschitz surface with $k' = \tan(\arcsin(4q'((1 + 2 \sin \theta) \frac{\delta}{\varepsilon} + \sin \theta)))$ and $\text{lr}_{k'}(S') \geq r = \text{lr}_k(S) - (2\varepsilon + 7\delta)$. \square

4 Application to Reconstruction from Point Samples

4.1 Algorithm

The algorithm works in any metric space. It takes as input a finite point set W , identified as the set of witnesses, and an optional countable sequence ν of integers, whose default value is $\nu_i = i + 1 \forall i$ (which corresponds to $\mathcal{C}_\nu^W = \mathcal{C}^W$). The algorithm constructs a set $L \subseteq W$ of landmarks iteratively, starting with $L = \emptyset$, and in the meantime it maintains $\mathcal{C}_\nu^W(L)$. At each iteration, the witness lying furthest away from L

is inserted in L , and $\mathcal{C}_v^W(L)$ is updated as described below. The process stops when $L = W$. The output of the algorithm is either the one-parameter family of complexes $\mathcal{C}_v^W(L)$ built throughout the process, or simply the diagram of their Betti numbers computed on the fly using the persistence algorithm⁹ of [37]. With this diagram at hand, the user can determine the scale at which to process the data: it is then easy to generate the corresponding subset of landmarks (the points of W have been sorted according to their order of insertion in L) and to rebuild its witness complex.

4.2 Update of $\mathcal{C}_v^W(L)$

Our strategy to update $\mathcal{C}_v^W(L)$ relies on the following observation: when a witness p is inserted in L , every simplex that appears in $\mathcal{C}_v^W(L)$ is incident to p , whereas every simplex that disappears from $\mathcal{C}_v^W(L)$ has a face that is no longer v -witnessed because of the insertion of p in the set of landmarks. It follows that the v -witnesses of all these simplices belong to the *reverse κ -nearest witnesses*¹⁰ of p , where $\kappa = \min\{|L|, \max_i v_i\}$. Hence, $\mathcal{C}_v^W(L)$ can be updated, first by searching for the reverse κ -nearest witnesses of p , and then, for each witness w in the outcome, by searching for the κ nearest landmarks of w , to determine which simplices to insert or delete from $\mathcal{C}_v^W(L)$. A number of dynamic data structures exist that can perform these queries efficiently—see [18] for a survey. Note however that κ can be as large as $|W|$, a case in which the above queries take linear time. Moreover, when $v_i \geq |L| \forall i$, $\mathcal{C}_v^W(L)$ coincides with the complete hypergraph of L and hence has an exponential size. Nevertheless, in Euclidean space \mathbb{R}^n , κ is more likely to be a constant depending (exponentially) on n , which reduces the size of $\mathcal{C}_v^W(L)$ to $O(|L|)$, by Theorem 3.8. The total time spent by the algorithm to maintain $\mathcal{C}_v^W(L)$ is then $O(|W|^2)$, since any newly-created landmark has $\Theta(|W|)$ reverse κ -nearest witnesses (these can be detected naively by an exhaustive search on the set W), each of which witnesses a constant number of simplices (these can be found by maintaining the list of κ nearest landmarks of each witness). We conjecture that it should be possible to reduce the time complexity to $O(|W| \log |W|)$, under some sparseness condition on W .

4.3 2d and 3d Cases

We take $v = (1, 2, 3)$ in 2d and $v = (1, 6, 6, 4)$ in 3d, as prescribed by the theory. Moreover, we replace $\mathcal{C}_v^W(L)$ by its intersection with $\mathcal{D}(L)$, noted $(\mathcal{C}_v^W \cap \mathcal{D})(L)$. This makes sense because, $\mathcal{D}^S(L)$ being a subset of $\mathcal{D}(L)$, Theorems 3.1 and 3.9 hold the same if $\mathcal{C}_v^W(L)$ is replaced by $(\mathcal{C}_v^W \cap \mathcal{D})(L)$. The advantage is that $(\mathcal{C}_v^W \cap \mathcal{D})(L)$ can be stored as a subcomplex of $\mathcal{D}(L)$, which allows to speed-up the (reverse) nearest neighbors queries in practice. Another thing in 3d is that we also maintain the subcomplex \hat{S} extracted from $(\mathcal{C}_v^W \cap \mathcal{D})(L)$ by the procedure of Sect. 3.2.2.

⁹The filtration used in [37] is rebuilt at each iteration, since some simplices are deleted from our complex $\mathcal{C}_v^W(L)$.

¹⁰These are the witnesses that have p among their κ nearest landmarks.

4.4 Theoretical Guarantees

Let $L(i)$ denote the set L at the end of the i th iteration of the algorithm. We call $\varepsilon(i)$ the minimum number such that $L(i)$ is an $\varepsilon(i)$ -sample of W . Since $L(i)$ grows with i , $\varepsilon(i)$ is a decreasing function of i . Moreover, we have the following:

Lemma 4.1 *At any iteration i , $L(i)$ is an $\varepsilon(i)$ -sparse $\varepsilon(i)$ -sample of W .*

Proof At each iteration $j \leq i$ of the algorithm, the witness $p(j)$ farthest from $L(j-1)$ is inserted in $L(j-1)$. Right before this insertion, $L(j-1)$ is an $\varepsilon(j-1)$ -sample of W . This means that the distance from $p(j)$ to $L(j-1)$ is $\varepsilon(j-1)$. Since L keeps growing during the process, we have $\varepsilon(j) \geq \varepsilon(j+1) \forall j \leq i$. Therefore, every point inserted in L before or at iteration i is at least $\varepsilon(i)$ away from L at the time of its insertion. This implies that $L(i)$ is an $\varepsilon(i)$ -sparse point sample. \square

Combining this lemma with the structural results of Sects. 2 and 3, we obtain a guarantee on the stability of the topological type of the witness complex during the course of the algorithm:

Theorem 4.2 *Assume that the input W is a δ -noisy δ -sample of some k -Lipschitz curve (resp. surface) S in \mathbb{R}^2 (resp. \mathbb{R}^3). Then, $(\mathcal{C}_v^W \cap \mathcal{D})(L(i))$ (resp. $\hat{S}(i)$) is isotopic to S and close to S for the Hausdorff distance whenever $L(i)$ satisfies the hypotheses of Theorems 2.9, 3.1, and 3.16 (resp. 3.9 and 3.16), which eventually happens if δ and k are small enough.*

Under mild assumptions on the input, the theorem guarantees that the topological type of $(\mathcal{C}_v^W \cap \mathcal{D})(L(i))$ (resp. $\hat{S}(i)$) stabilizes for some time during the course of the algorithm and therefore that a plateau appears in the diagram of the Betti numbers of $(\mathcal{C}_v^W \cap \mathcal{D})(L(i))$ (resp. $\hat{S}(i)$), showing the correct values. In practice, the user does not know exactly where the plateau should be located in the diagram, because δ , k , and $\text{lr}_k(S)$ are unknown. Nevertheless, our structural results give a lower bound on the length of the plateau, which is of order of $(a \text{lr}_k(S) - b\delta)$, where a, b are two constants. As a consequence, for sufficiently small δ , the plateau is long enough to be detected by the user or any statistical analysis performed in a post-processing step. In the case where W samples several manifolds (such as in the example of Fig. 1), Theorem 4.2 guarantees¹¹ that several plateaus appear in the diagram of Betti numbers, each one showing a plausible reconstruction, depending on the scale at which the input data set W is processed. Once again, these guarantees hold as long as the input point set W is sufficiently dense compared to $\text{lr}_k(S)$. In particular, no sparseness condition on W is required, although S can be non-smooth and W can be noisy. This makes our reconstruction method more practical than previous work in this context [8, 24].

¹¹In the example of Fig. 1, the two plausible reconstructions are the helical curve and the torus in \mathbb{R}^3 . Strictly speaking, our theoretical statements do not consider curves in \mathbb{R}^n with $n \geq 3$. However, as in the smooth setting, our proofs can be extended to arbitrary values of n , provided that a suitable definition of Lipschitz curve is used.




Fig. 5 Diagram of Betti numbers of \hat{S} for the Tanglecube point set

4.5 Experimental Results and Discussion

Figures 5 and 6 show the evolution of the Betti numbers of $\hat{S}(i)$ throughout the course of the algorithm on two different input point sets. The abscissa represents $\frac{1}{\epsilon(i)}$, which is an increasing function of i and gives an idea of the scale at which the input data is being processed. The first point set was generated using Chew’s algorithm on an algebraic surface commonly known as the tanglecube. The second point set is courtesy of the Stanford shape repository and was obtained from a 3d scanner.

Figure 5 shows three phases in the evolution of the Betti numbers. At first, their behavior is erratic, which means that the topology of $\hat{S}(i)$ keeps changing because $\epsilon(i)$ is not yet small enough compared to $lr_k(S)$. The top picture in the right column shows $\hat{S}(i)$ at some point during this phase. Then, $\epsilon(i)$ satisfies the hypothesis of Theorem 3.9 for a while, which implies that the Betti numbers stabilize to the right value and a plateau appears in the diagram—here, the surface is connected, of genus five, and without boundary, hence $\beta_0 = \beta_2 = 1$ and $\beta_1 = 10$. The size of the plateau depends on the value of δ compared to $lr_k(S)$. The two middle pictures in the right column show $\hat{S}(i)$ at two different times during this phase. Finally, $\epsilon(i)$ becomes too small, and $(C_v^W \cap \mathcal{D})(L(i))$ does not contain $\mathcal{D}^S(L(i))$ any longer. As a result, some holes appear in the complex, and $\hat{S}(i)$ is simply the boundary of a myriad of randomly agglomerated tetrahedra. An example is shown in the bottom-right corner of the image.

The same kind of phenomenon occurs in Fig. 6, where the plateau is smaller due to the fairly high value of δ . The top-right picture shows $\hat{S}(i)$ at some point on the plateau. Interestingly enough, the value of β_1 on the plateau ($\beta_1 = 12$) means that the genus of $\hat{S}(i)$ is six, which coincides with the genus of the scanned physical model

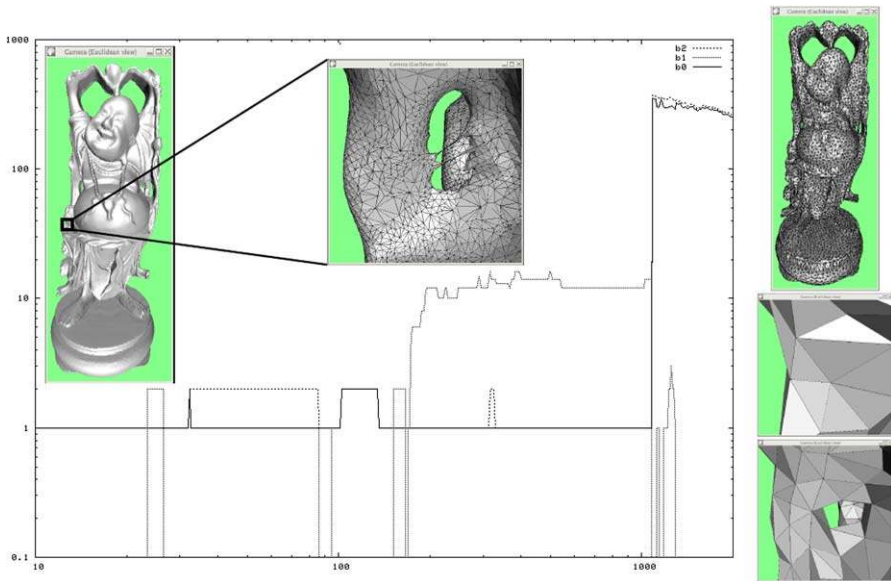


Fig. 6 Diagram of Betti numbers of \hat{S} for the Happy Buddha point set

but not with that of the mesh provided by the repository, which contains extra handles such as the one shown at the center of Fig. 6. The reason is that, due to the presence of noise and holes in the input point cloud, standard reconstruction algorithms fail at retrieving the topology of the original object because they consider the data at full scale only. In contrast, our algorithm performs reconstructions at various scales and generates plateaus whenever the topological type of the reconstruction is stable, which corresponds to plausible topological types of the original object. As an illustration, observe that a new plateau starts on the right-hand side of the main plateau, right before the data structure becomes unstable. This new plateau shows $\beta_1 = 14$, which indicates that a new handle has been detected, as illustrated in the bottom-right corner of Fig. 6. Another example is shown in Fig. 1, where the diagram has two well-separated plateaus corresponding to two plausible reconstructions: a torus, and a simple closed curve drawn on that torus. To handle the change in dimension (shown on the right-hand side of the figure), we maintained $\mathcal{C}_\nu^W(L(i))$ for both $\nu = (1, 2, 3, 4)$ and $\nu = (1, 6, 6, 4)$ simultaneously and determined at each step the complex to keep according to their Betti number β_2 .

The definitive choice of the scale(s) at which to deal with the data is highly dependent on the application, and therefore it should be made by the user. We expect that statistical analysis on our diagrams would allow our algorithm to detect the plateaus and to suggest reasonable choices. Note however that not all the plateaus are guaranteed to be meaningful. Indeed, the fact that the homotopy type of the complex remains stable over a period of time does not imply that the complex itself is stable combinatorially. In the example of Fig. 7 for instance, the insertion of a new landmark creates a loop but at the same time destroys another loop in the complex, thereby keeping its Betti numbers unchanged, whereas the sets of generators of the homology groups of




Fig. 7 (Color online) A set W of witnesses (in black) sampling a smooth curve in the plane, and a sparse subset L of landmarks. The witness complex of L relative to W is shown in green, while the edges of the order-two Voronoi diagram of L are shown in grey. Note how the insertion of $a \in W$ in L creates a new loop but at the same time removes edge $[b, c]$ from the complex, thereby destroying the former loop. The homotopy type of the complex remains unchanged

the complex do differ. Such events leading to *fake* plateaus in the diagram are likely to be rare and short, since they correspond to highly unstable phases of the course of the algorithm, where the sampling conditions are not satisfied. Nevertheless, according to the example of Fig. 7, they can occur, and the question of whether they can be detected and discarded automatically is still open.

5 Conclusion

We have introduced a new reconstruction method based on the use of the witness complex. This method uses inter-sample distances alone and can therefore be used in any metric space. Moreover, it stands in sharp contrast with previous work in the area, since it is multiscale and gives some insights on the various plausible topological types of the original object. We believe that this approach to manifold reconstruction is highly practical and has a number of potential applications, such as for instance topological noise removal or mesh compression. As a side product, in order to prove our algorithm correct on Lipschitz curves and surfaces, we have worked out interesting relationships between the witness complex and the restricted Delaunay triangulation in 2d and 3d. Our structural results assume that the set L of landmarks is an ε -sparse ε -sample of the set W of witnesses. This hypothesis is stringent but not unrealistic in practice, since L is generated by the algorithm. In contrast, our assumption on the input W is fairly mild.

Let us mention that our structural results still hold in the somewhat more general setting where W is an *adaptive* δ -sample of S in the sense of [1] and where L is an (ε, κ) -sample of W in the sense of [24]. The proofs are roughly the same, with an additional twist which slightly degrades the constants.

A number of theoretical questions are left open by this work, including:

- Is it possible to devise automatic procedures for discriminating the *fake* plateaus from the meaningful plateaus in the diagram of Betti numbers? One possible approach, suggested by the example of Fig. 7, would be to keep track of the generators of the homology groups of the complex throughout the course of the algorithm,

similarly to what is done in persistent homology [37]. The difficulty here is that our family of complexes is not a filtration, the complex at iteration i being neither a subset nor a superset of the complex at iteration $i + 1$.

- Can the time complexity of the algorithm be reduced to $O(|W| \log |W|)$? Reverse nearest neighbors queries can take linear time in the worst case, because a landmark can be among the nearest landmarks of a linear number of witnesses. As a consequence, reducing the time complexity requires to control the sizes of the lists of reverse nearest witnesses, which can probably be done under some sparseness condition on W . However, can it be done under weaker conditions?
- The behavior of our reconstruction algorithm depends highly on the order in which the witnesses are inserted in L —in some sense, the algorithm is *parameterized* by this order. In the present paper, we use one possible ordering of the witnesses, prescribed by the furthest-point criterion. However, many other orderings could be considered, some of which might be relevant in different contexts. One appealing example is the problem of reconstructing curves and surfaces in the presence of background noise, where the furthest-witness strategy applied in this paper fails for obvious reasons.
- Can our structural results be extended to manifolds of higher dimensions? Since the publication of the conference version of this paper, significant steps have been made in this direction: de Silva has extended the Weak Witness Theorem to the case of point clouds lying on Riemannian manifolds of constant sectional curvature [21]. Unfortunately, focusing on the Euclidean case, we have seen that the Weak Witness Theorem does not hold for general smooth submanifolds of \mathbb{R}^n [32]. Nevertheless, Cheng, Dey, and Ramos had proved that restricted Delaunay triangulations of dense point samples enjoy similar properties in higher dimensions as they do in \mathbb{R}^3 , provided that the standard Delaunay test is replaced by a weighted Delaunay test [15]. Building on their work, we have extended the results of this paper to the general case of smooth submanifolds of \mathbb{R}^n , with the Euclidean metric replaced by a weighted metric [6]. Combined with recent advances on the front of topological persistence, these results have led to a new class of reconstruction algorithms, inspired from this work and based on the use of the witness complex, whose space and time complexities scale up with the intrinsic dimension of the point cloud data, as opposed to the dimension of the ambient space [14].

Acknowledgements The authors wish to thank V. de Silva, G. Carlsson, and the anonymous referees for their insightful comments. They also thank A. Zomorodian for providing them with his implementation of the persistence algorithm. This work was supported by DARPA grant 32905—Topological Data Analysis, NSF grants FRG-0354543 and ITR-0205671, and NIH grant GM-072970.

References

1. Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* **22**(4), 481–504 (1999)
2. Amenta, N., Choi, S.: One-pass Delaunay filtering for homeomorphic 3D surface reconstruction. *Manuscript* (1999)
3. Amenta, N., Choi, S., Dey, T.K., Leekha, N.: A simple algorithm for homeomorphic surface reconstruction. *Int. J. Comput. Geom. Appl.* **12**, 125–141 (2002)

4. Attali, D., Edelsbrunner, H., Mileyko, Y.: Weak witnesses for Delaunay triangulations of submanifolds. In: Proc. ACM Symp. on Solid and Physical Modeling, pp. 143–150 (2007)
5. Boissonnat, J.-D., Cazals, F.: Natural neighbor coordinates of points on a surface. *Comput. Geom.: Theory Appl.* **19**, 155–173 (2001)
6. Boissonnat, J.-D., Guibas, L.J., Oudot, S.Y.: Manifold reconstruction in arbitrary dimensions using witness complexes. In: Proc. 23rd Symp. on Comput. Geom., pp. 194–203 (2007)
7. Boissonnat, J.-D., Oudot, S.: Provably good sampling and meshing of surfaces. *Graph. Models* **67**(5), 405–451 (2005)
8. Boissonnat, J.-D., Oudot, S.: Provably good sampling and meshing of Lipschitz surfaces. In: Proc. 22nd Annu. Symp. Comput. Geom., pp. 337–346 (2006)
9. Cazals, F., Giesen, J.: Delaunay triangulation based surface reconstruction. In: Boissonnat, J.D., Teilaud, M. (eds.) *Effective Computational Geometry for Curves and Surfaces*, pp. 231–273. Springer, Berlin (2006)
10. Chazal, F., Cohen-Steiner, D., Lieutier, A.: A sampling theory for compact sets in Euclidean space. In: Proc. 22nd Annu. Symp. Comput. Geom., pp. 319–326 (2006)
11. Chazal, F., Lieutier, A.: Weak feature size and persistent homology: Computing homology of solids in \mathbb{R}^n from noisy data samples. Technical Report 378, Institut de Mathématiques de Bourgogne (2004). Partially published in Proc. 21st Annu. ACM Symp. on Comput. Geom., pp. 255–262 (2005)
12. Chazal, F., Lieutier, A.: The λ -medial axis. *Graph. Models* **67**(4), 304–331 (2005)
13. Chazal, F., Lieutier, A.: Topology guaranteeing manifold reconstruction using distance function to noisy data. In: Proc. 22nd Annu. Symp. on Comput. Geom., pp. 112–118 (2006)
14. Chazal, F., Oudot, S.Y.: Towards persistence-based reconstruction in Euclidean spaces. In: Proc. 24th ACM Symp. Comput. Geom., pp. 232–241 (2008)
15. Cheng, S.-W., Dey, T.K., Ramos, E.A.: Manifold reconstruction from point samples. In: Proc. 16th Symp. Discrete Algorithms, pp. 1018–1027 (2005)
16. Cheng, S.-W., Funke, S., Golin, M., Kumar, P., Poon, S.-H., Ramos, E.: Curve reconstruction from noisy samples. *Comput. Geom.: Theory Appl.* **31**(1–2), 63–100 (2005)
17. Chew, L.P.: Guaranteed-quality mesh generation for curved surfaces. In: Proc. 9th Annu. ACM Symp. Comput. Geom., pp. 274–280 (1993)
18. Clarkson, K.L.: Nearest-neighbor searching and metric space dimensions. In: Shakhnarovich, G., Darrell, T., Indyk, P. (eds.) *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, pp. 15–59. MIT Press, Cambridge (2006)
19. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. In: Proc. 21st ACM Symp. Comput. Geom., pp. 263–271 (2005)
20. de Silva, V.: A weak definition of Delaunay triangulation. Technical report, Stanford University, October 2003
21. de Silva, V.: A weak characterisation of the Delaunay triangulation. *Geom. Dedicata* (2008). doi:[10.1007/s10711-008-9261-1](https://doi.org/10.1007/s10711-008-9261-1). Online First
22. de Silva, V., Carlsson, G.: Topological estimation using witness complexes. In: Proc. Symp. Point-Based Graphics, pp. 157–166 (2004)
23. Dey, T.K., Giesen, J., Goswami, S., Zhao, W.: Shape dimension and approximation from samples. *Discrete Comput. Geom.* **29**, 419–434 (2003)
24. Dey, T.K., Goswami, S.: Provable surface reconstruction from noisy samples. *Comput. Geom.: Theory Appl.* **35**(1–2), 124–141 (2006)
25. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. In: Proc. 41st Annu. IEEE Symp. Found. Comput. Sci., pp. 454–463 (2000)
26. Federer, H.: Curvature measures. *Trans. Am. Math. Soc.* **93**, 418–491 (1959)
27. Graham, R.L.: Sets of points with given minimum separation (solution to Problem E1921). *Am. Math. Mon.* **75**, 192–193 (1968)
28. Kolluri, R.: Provably good moving least squares. In: Proc. 16th ACM-SIAM Symposium on Discrete Algorithms, pp. 1008–1017 (2005)
29. Mederos, B., Amenta, N., Velho, L., de Figueiredo, L.H.: Surface reconstruction for noisy point clouds. In: Proc. 3rd Symp. on Geometry Processing, pp. 53–62 (2005)
30. Niyogi, P., Smale, S., Weinberger, S.: Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.* **39**(1), 419–441 (2008)
31. Oudot, S., Rineau, L., Yvinec, M.: Meshing volumes with curved boundaries. *Eng. Comput.* (2008, to appear). Special issue IMR 2005

32. Oudot, S.Y.: On the topology of the restricted Delaunay triangulation and witness complex in higher dimensions. Technical report, Stanford University, November 2006. LANL [arXiv:0803.1296v1](https://arxiv.org/abs/0803.1296) [cs.CG], <http://arxiv.org/abs/0803.1296>
33. Pirl, U.: Der Mindestabstand von n in der Einheitskreisscheibe gelegenen Punkten. *Math. Nachr.* **40**, 111–124 (1969)
34. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**(5500), 2323–2326 (2000)
35. Shewchuk, J.R.: What is a good linear element? Interpolation, conditioning, and quality measures. In: *Proc. 11th Int. Meshing Roundtable*, pp. 115–126 (2002)
36. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000)
37. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 249–274 (2005)