

Towards Persistence-Based Reconstruction in Euclidean Spaces

Frédéric Chazal
INRIA, Geometrica group
4, rue Jacques Monod
91893 ORSAY Cedex, France
frederic.chazal@inria.fr

Steve Y. Oudot
INRIA, Geometrica group
4, rue Jacques Monod
91893 ORSAY Cedex, France
steve.oudot@inria.fr

ABSTRACT

Manifold reconstruction has been extensively studied for the last decade or so, especially in two and three dimensions. Recent advances in higher dimensions have led to new methods to reconstruct large classes of compact subsets of \mathbb{R}^d . However, the complexities of these methods scale up exponentially with d , making them impractical in medium or high dimensions, even on data sets of low intrinsic dimensionality.

In this paper, we introduce a novel approach that stands in-between classical reconstruction and topological estimation, and whose complexity scales up with the intrinsic dimension of the data. Our algorithm combines two paradigms: greedy refinement, and topological persistence. Given a point cloud in \mathbb{R}^d , we build a set of landmarks iteratively, while maintaining a nested pair of abstract complexes, whose images in \mathbb{R}^d lie close to the data, and whose persistent homology eventually coincides with the homology of the underlying shape. When the data points are densely sampled from a smooth m -submanifold X of \mathbb{R}^d , our method retrieves the homology of X in time at most $c(m)n^5$, where n is the size of the input and $c(m)$ is a constant depending solely on m .

To prove the correctness of our algorithm, we investigate on Čech, Rips, and witness complex filtrations in Euclidean spaces. More precisely, we show how previous results on unions of balls can be transposed to Čech filtrations, and from there to Rips and witness complex filtrations. Finally, investigating further on witness complexes, we quantify a conjecture of Carlsson and de Silva, which states that witness complex filtrations should have cleaner persistence barcodes than Čech or Rips filtrations, at least on smooth submanifolds of Euclidean spaces.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems — Geometrical problems and computations, Computations on discrete structures.

General Terms:




Figure 1: Topological analysis of a synthetic 1000-dimensional data set. The 50,000 data points have been sampled uniformly at random from a helical curve drawn on the 2d Clifford torus, embedded into \mathbb{R}^{1000} via a quadratic mapping. The image shows the persistence barcode of the Rips filtration built over a carefully-chosen subset of 2000 landmarks.

groups, whose images are isomorphic to the homology groups of X . Combined with the structure theorem of [33], which states that the persistent homology of the sequence $\{L^\alpha\}_{\alpha \geq 0}$ is fully described by a finite set of intervals, called a *persistence barcode* or a *persistence diagram* — see Figure 1, the above result means that the homology of X can be deduced from this barcode, simply by removing the intervals of length less than 2ε , which are therefore viewed as topological noise.

From an algorithmic point of view, the persistent homology of a nested sequence of simplicial complexes (called a *filtration*) can be efficiently computed using the persistence algorithm [21, 33]. Among the many filtrations that can be built on top of a point set L , the α -shape enables to reliably recover the homology of the underlying space X , since it is known to be a deformation retract of L^α [20]. However, this property is useless in high dimensions, since computing the α -shape requires to build the full-dimensional Delaunay triangulation. It is therefore appealing to consider other filtrations that are easy to compute in arbitrary dimensions, such as the Rips and witness complex filtrations. In this paper, we produce an equivalent of the result of [12, 15] for these filtrations, and more generally for any filtration that is intertwined with the Čech filtration. Recall that, for all $\alpha > 0$, the Čech complex $C^\alpha(L)$ is the nerve of the union of the open balls of same radius α about the points of L . It is known to be homotopy equivalent to L^α . However, combining this fact with the result of [12, 15] is not enough to prove that the persistent homology of $C^\alpha(L) \hookrightarrow C^{\alpha+2\varepsilon}(L)$ coincides with the homology of X , because it is unclear whether the homotopy equivalences $C^\alpha(L) \rightarrow L^\alpha$ and $C^{\alpha+2\varepsilon}(L) \rightarrow L^{\alpha+2\varepsilon}$ commute with the canonical inclusions $C^\alpha(L) \hookrightarrow C^{\alpha+2\varepsilon}(L)$ and $L^\alpha \hookrightarrow L^{\alpha+2\varepsilon}$. In the paper, we show that there exist homotopy equivalences that commute with canonical inclusions, at least at homology and homotopy levels. This enables us to extend the result of [12, 15] to the Čech filtration, and from there to the Rips and witness complex filtrations.

Another common concern in topological data analysis is the size of the vertex set on top of which the filtration is built. Indeed, in practical situations where the input data

set W samples the underlying shape very finely, it makes sense to build the filtration on top of a small subset L of *landmarks* to avoid a waste of computational resources. However, downsampling the vertex set may result in a significant degradation in the quality of the persistence barcode. This is true in particular with the Čech and Rips filtrations, whose barcodes can have topological noise of amplitude depending directly on the density of the landmark set L . The introduction of the witness complex filtration appeared as an elegant way of solving this issue [17]. The witness complex of L relative to W , or $C_W(L)$ for short, can be viewed as a relaxed version of the Delaunay triangulation of L , in which the points of $W \setminus L$ are used to drive the construction of the complex [16]. Due to its special nature, which takes advantage of the points of $W \setminus L$, the witness complex filtration is likely to give persistence barcodes whose topological noise depends on the density of W rather than on the one of L , as conjectured in [17]. We prove that this statement is only true to some extent, namely: whenever the points of W are sufficiently densely sampled from some smooth submanifold of \mathbb{R}^d , the topological noise in the barcode can be arbitrarily small compared to the density of L . Nevertheless, it cannot depend solely on the density of W . This shows that the witness complex filtration does provide cleaner persistence barcodes, though maybe not as clean as expected.

Taking advantage of the above results, we propose a novel approach to reconstruction that stands in-between the classical reconstruction and topological estimation paradigms. Our algorithm is a variant of the method of [6, 26] that combines greedy refinement and topological persistence. Given an input point cloud W , the algorithm builds a subset L of landmarks iteratively, and in the meantime it maintains a nested pair of simplicial complexes (Rips or witness complexes) and computes its persistent Betti numbers. The outcome of the algorithm is the diagram showing the evolution of these persistent Betti numbers. Using this diagram, a user or software agent can determine a relevant scale at which to process the data. It is then easy to rebuild the corresponding set of landmarks, as well as its nested pair of complexes. Although our method does not really compute an embedded complex that is close to X topologically and geometrically, it comes with theoretical guarantees, it is easily implementable, and it has reasonable complexity. Indeed, in the case where the input point cloud W is densely sampled from a smooth submanifold X of \mathbb{R}^d , we show that the complexity of our algorithm is bounded by $c(m)n^5$, where $c(m)$ is a quantity depending solely on the intrinsic dimension m of X , and n is the size of W . To the best of our knowledge, this is the first provably-good topological estimation or reconstruction method whose complexity scales up with the intrinsic dimension of the data. When X is a more general compact set in \mathbb{R}^d , our complexity bound becomes $c(d)n^5$.

The paper is organized as follows: after introducing the Čech, Rips, and witness complex filtrations in Section 2, we prove our structural results in Sections 3 and 4, focusing on compact subsets of \mathbb{R}^d in Section 3, and on the particular case of smooth submanifolds in Section 4. Finally, we present our algorithm and its analysis in Section 5.

2. VARIOUS RELATED FILTRATIONS

The definitions, results and proofs of this section hold in any arbitrary metric space. However, for the sake of consistency with the rest of the paper, we state them in the

particular case of \mathbb{R}^d , endowed with the Euclidean norm $\|\cdot\|$. Although our bounds can be proved to be tight in the general metric case, it is possible to work out somewhat tighter bounds in the Euclidean case, at the price of a loss of simplicity in the statements.

For any compact set $X \subset \mathbb{R}^d$, we call $\text{diam}(X)$ the diameter of X , and $\text{diam}_{\text{CC}}(X)$ the *component-wise diameter* of X , defined by: $\text{diam}_{\text{CC}}(X) = \inf_i \text{diam}(X_i)$, where the X_i are the path-connected components of X . Finally, given two compact sets X, Y in \mathbb{R}^d , we call $d_{\mathcal{H}}(X, Y)$ their Hausdorff distance. Given a finite set L of points of \mathbb{R}^d and a pos-

relationship with the topology of the *offsets* $X^\alpha = \{x \in \mathbb{R}^d : d_X(x) \leq \alpha\}$, stressed in the following result from [25]:

LEMMA 3.1. *If $0 < \alpha < \alpha'$ are such that there is no critical value of d_X in the closed interval $[\alpha, \alpha']$, then $X^{\alpha'}$ deformation retracts onto X^α .*

In particular, the hypothesis of the lemma is satisfied when $0 < \alpha_1 < \alpha_2 < \text{wfs}(X)$. Therefore, all the offsets of X have the same homotopy type in the interval $(0, \text{wfs}(X))$. In the sequel, we repeatedly make use of the following standard result of linear algebra:

LEMMA 3.2. *Given a sequence $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F$ of homomorphisms between finite-dimensional vector spaces, if $\text{rank}(A \rightarrow F) = \text{rank}(C \rightarrow D)$, then this quantity also equals the rank of $B \rightarrow E$. Similarly, if $A \rightarrow B \rightarrow C \rightarrow E \rightarrow F$ is a sequence of homomorphisms such that $\text{rank}(A \rightarrow F) = \dim C$, then $\text{rank}(B \rightarrow E) = \dim C$.*

3.1 Čech filtration

Since the Čech complex is the nerve of a union of balls, its homotopy type is closely related to the one of this union. We will use the following extension of Theorem 4.7 of [12]:

LEMMA 3.3. *Let X be a compact set and L a finite set in \mathbb{R}^d , such that $d_{\mathcal{H}}(X, L) < \varepsilon$ for some $\varepsilon < \frac{1}{4} \text{wfs}(X)$. Then, for all $\alpha, \alpha' \in [\varepsilon, \text{wfs}(X) - \varepsilon]$ such that $\alpha' - \alpha \geq 2\varepsilon$, and for all $\lambda \in (0, \text{wfs}(X))$, we have: $\forall k \in \mathbb{N}$, $H_k(X^\lambda) \cong \text{im } i_*$, where $i_* : H_k(L^\alpha) \rightarrow H_k(L^{\alpha'})$ is the homomorphism between homology groups induced by the canonical inclusion $i : L^\alpha \hookrightarrow L^{\alpha'}$. Given an arbitrary point $x_0 \in X$, the same conclusion holds for homotopy groups with base-point x_0 .*

PROOF. We can assume without loss of generality that $\varepsilon < \alpha < \alpha' - 2\varepsilon < \text{wfs}(X) - 3\varepsilon$, since otherwise we can replace ε by any $\varepsilon' \in (d_H(X, L), \varepsilon)$. From the hypothesis we deduce the following sequence of inclusions:

$$X^{\alpha-\varepsilon} \hookrightarrow L^\alpha \hookrightarrow X^{\alpha+\varepsilon} \hookrightarrow L^{\alpha'} \hookrightarrow X^{\alpha'+\varepsilon} \quad (1)$$

By the Isotopy Lemma 3.1, for all $0 < \beta < \beta' < \text{wfs}(X)$, the canonical inclusion $X^\beta \hookrightarrow X^{\beta'}$ is a homotopy equivalence. As a consequence, Eq. (1) induces a sequence of homomorphisms between homology groups, such that all homomorphisms between homology groups of $X^{\alpha-\varepsilon}, X^{\alpha+\varepsilon}, X^{\alpha'+\varepsilon}$ are isomorphisms. It follows then from Lemma 3.2 that $i_* : H_k(L^\alpha) \rightarrow H_k(L^{\alpha'})$ has same rank as these isomorphisms. Now, this rank is equal to the dimension of $H_k(X^\lambda)$, since the X^β are homotopy equivalent to X^λ for all $0 < \beta < \text{wfs}(X)$. It follows that $\text{im } i_* \cong H_k(X^\lambda)$, since our ring of coefficients is a field.

The case of homotopy groups is a little trickier, since the above rank argument cannot be used. However, we can use the same proof as in Theorem 4.7 of [12] to conclude. \square

Observe that Lemma 3.3 does not guarantee the retrieval of the homology of X . Instead, it deals with sufficiently small offsets of X , which are homotopy equivalent to one another but possibly not to X itself [12, Fig. 4]. In the special case where X is a smooth submanifold of \mathbb{R}^d however, X^λ and X are homotopy equivalent, and therefore the theorem guarantees the retrieval of the homology of X .

Consider now the Čech complex $\mathcal{C}^\alpha(L)$, for any value $\alpha > 0$. By definition, $\mathcal{C}^\alpha(L)$ is the nerve of the open cover $\{L^\alpha\}$ of

L^α . Since the elements of $\{L^\alpha\}$ are convex, they form a *good* open cover of L^α , i.e. their intersections are either empty or contractible. It follows from the *nerve theorem* [27, Corollary 4G.3] that L^α and its nerve $\mathcal{C}^\alpha(L)$ are homotopy equivalent. We thus get the following diagram, where horizontal arrows are canonical inclusions and vertical arrows are homotopy equivalences provided by the nerve theorem:

$$\begin{array}{ccc} L^\alpha & \hookrightarrow & L^{\alpha'} \\ \uparrow & & \uparrow \\ \mathcal{C}^\alpha(L) & \hookrightarrow & \mathcal{C}^{\alpha'}(L) \end{array} \quad (2)$$

Unfortunately, the nerve theorem does not guarantee that this diagram commutes. However, standard arguments of algebraic topology imply the following result, where $\mathcal{N}\mathcal{U}$ (resp. $\mathcal{N}\mathcal{U}'$) stands for the nerve of the open cover \mathcal{U} (resp. \mathcal{U}'):

LEMMA 3.4. *Let $X \subseteq X'$ be two paracompact spaces, and let $\mathcal{U} = \{U_\alpha\}_{\alpha \in A}$ and $\mathcal{U}' = \{U'_\alpha\}_{\alpha \in A}$ be good open covers of X and X' respectively, based on a same finite parameter set A , such that $U_\alpha \subseteq U'_\alpha$ for all $\alpha \in A$. Then, there exist homotopy equivalences $\mathcal{N}\mathcal{U} \rightarrow X$ and $\mathcal{N}\mathcal{U}' \rightarrow X'$ that commute with the canonical inclusions $X \hookrightarrow X'$ and $\mathcal{N}\mathcal{U} \hookrightarrow \mathcal{N}\mathcal{U}'$ at homology and homotopy levels.*

Letting $X = L^\alpha$, $X' = L^{\alpha'}$, $\mathcal{U} = \{L^\alpha\}$, and $\mathcal{U}' = \{L^{\alpha'}\}$, we get from Lemma 3.4 that there exist homotopy equivalences $\mathcal{C}^\alpha(L) \rightarrow L^\alpha$ and $\mathcal{C}^{\alpha'}(L) \rightarrow L^{\alpha'}$ that make the diagram of Eq. (2) commute at homology and homotopy levels. Combined with Lemma 3.3, this fact implies the following result:

THEOREM 3.5. *Let X be a compact set and L a finite set in \mathbb{R}^d , such that $d_{\mathcal{H}}(X, L) < \varepsilon$ for some $\varepsilon < \frac{1}{4} \text{wfs}(X)$. Then, for all $\alpha, \alpha' \in [\varepsilon, \text{wfs}(X) - \varepsilon]$ such that $\alpha' - \alpha > 2\varepsilon$, and for all $\lambda \in (0, \text{wfs}(X))$, we have: $\forall k \in \mathbb{N}$, $H_k(X^\lambda) \cong \text{im } j_*$, where $j_* : H_k(\mathcal{C}^\alpha(L)) \rightarrow H_k(\mathcal{C}^{\alpha'}(L))$ is the homomorphism between homology groups induced by the canonical inclusion $j : \mathcal{C}^\alpha(L) \hookrightarrow \mathcal{C}^{\alpha'}(L)$. Given a point $x_0 \in X$, the same result holds for homotopy groups with base-point x_0 .*

Using the terminology of [33], this theorem guarantees that the homology of X^λ is obtained from the persistence barcode of the filtration $\{\mathcal{C}^\alpha(L)\}_{\alpha \geq 0}$ by removing the intervals of persistence less than 2ε .

We now give our proof¹ of Lemma 3.4, which consists in a generalization to our context of the main arguments of the proof of the nerve theorem provided in Section 4G of [27]:

Proof of Lemma 3.4. Recall that \mathcal{U} is a good open cover of X , namely: $\forall k \in \mathbb{N}$, $\forall \alpha_0, \dots, \alpha_k \in L$, $\bigcap_{l=0}^k U_{\alpha_l}$ is either empty or contractible. From this cover we construct a topological space ΔX as follows: let Δ^n denote the standard n -simplex, where $n = \#A - 1$. To each non-empty subset S of A we associate the face $[S]$ of Δ^n spanned by the elements of S , as well as the subspace $U_S = \bigcap_{s \in S} U_s$ of X . ΔX is then the subspace of $X \times \Delta^n$ defined by:

$$\Delta X = \bigcup_{\emptyset \neq S \subseteq A} U_S \times [S].$$

The subspace $\Delta X' \subseteq X' \times \Delta^n$ is built similarly. Note that we have $\Delta X \subseteq \Delta X'$, since the hypothesis of the lemma implies $U_S \subseteq U'_S$ for all $S \subseteq A$. Furthermore, the product

¹Another proof of Lemma 3.4 is provided in [5], for the special case where $X = L^\alpha$, $X' = L^{\alpha'}$, $\mathcal{U} = \{L^\alpha\}$, and $\mathcal{U}' = \{L^{\alpha'}\}$ lie in \mathbb{R}^d . Our proof is simpler, and it holds in a more general setting.

structures of ΔX and $\Delta X'$ imply the existence of canonical projections $p : \Delta X \rightarrow X$ and $p' : \Delta X' \rightarrow X'$. These projections commute with the canonical inclusions $\Delta X \hookrightarrow \Delta X'$ and $X \hookrightarrow X'$, therefore the following diagram:

$$\begin{array}{ccc} X & \hookrightarrow & X' \\ \uparrow p & & \uparrow p' \\ \Delta X & \hookrightarrow & \Delta X' \end{array} \quad (3)$$

induces commutative diagrams at homology and homotopy levels. Moreover, since \mathcal{U} is an open cover of X , which is paracompact, p is a homotopy equivalence [27, Prop. 4G.2]. The same holds for p' , and therefore p and p' induce isomorphisms at homology and homotopy levels.

We now show that, similarly, there exist homotopy equivalences $\Delta X \rightarrow \mathcal{N}\mathcal{U}$ and $\Delta X' \rightarrow \mathcal{N}\mathcal{U}'$ that commute with the canonical inclusions $\Delta X \hookrightarrow \Delta X'$ and $\mathcal{N}\mathcal{U} \hookrightarrow \mathcal{N}\mathcal{U}'$. This follows in fact from the proof of Corollary 4G.3 of [27]. Indeed, using the notion of *complex of spaces* introduced in [27, Section 4G], it can be shown that ΔX is the realization of the complex of spaces associated with the cover \mathcal{U} — see the proof of [27, Prop. 4G.2]. Its base is the barycentric subdivision Γ of $\mathcal{N}\mathcal{U}$, where each vertex corresponds to a non-empty finite intersection U_S for some set $S \subseteq A$, and where each edge connecting two vertices $S \subset S'$ corresponds to the canonical inclusion $U_{S'} \hookrightarrow U_S$. In the same way, $\Delta X'$ is the realization of a complex of spaces built over the barycentric subdivision Γ' of $\mathcal{N}\mathcal{U}'$. Now, since the non-empty finite intersections U_S (resp. U'_S) are contractible, the map $q : \Delta X \rightarrow \Gamma$ (resp. $q' : \Delta X' \rightarrow \Gamma'$) induced by sending each open set U_S (resp. U'_S) to a point is a homotopy equivalence [27, Prop. 4G.1 and Corol. 4G.3]. Furthermore, by construction, q is the restriction of q' to ΔX . Therefore,

$$\begin{array}{ccc} \Delta X & \hookrightarrow & \Delta X' \\ \downarrow q & & \downarrow q' \\ \Gamma & \hookrightarrow & \Gamma' \end{array} \quad (4)$$

is a commutative diagram where vertical arrows are homotopy equivalences. Now, it is well-known that Γ and Γ' are homeomorphic to $\mathcal{N}\mathcal{U}$ and $\mathcal{N}\mathcal{U}'$ respectively, and that the homeomorphisms commute with the inclusion. Combined with (3) and (4), this fact proves Lemma 3.4. \square

3.2 Intertwined filtrations

3.2.1 Results on homology

Using Lemma 2.1 and Theorem 3.5, we get the following guarantees on the Rips filtration:

THEOREM 3.6. *Let $X \subset \mathbb{R}^d$ be a compact set and $L \subset \mathbb{R}^d$ a finite point set such that $d_{\mathcal{H}}(X, L) < \varepsilon$ for some $\varepsilon < \frac{1}{9} \text{wfs}(X)$. Then, for all $\alpha \in [2\varepsilon, \frac{1}{4}(\text{wfs}(X) - \varepsilon)]$ and all $\lambda \in (0, \text{wfs}(X))$, we have: $\forall k \in \mathbb{N}$, $H_k(X^\lambda) \cong \text{im } j_*$, where j_* is the homomorphism between homology groups induced by the canonical inclusion $j : \mathcal{R}^\alpha(L) \hookrightarrow \mathcal{R}^{4\alpha}(L)$.*

PROOF. Lemma 2.1 provides the following sequence:

$$C^{\frac{\alpha}{2}}(L) \hookrightarrow \mathcal{R}^\alpha(L) \hookrightarrow C^\alpha(L) \hookrightarrow C^{2\alpha}(L) \hookrightarrow \mathcal{R}^{4\alpha}(L) \hookrightarrow C^{4\alpha}(L)$$

Since $\alpha \geq 2\varepsilon$, Theorem 3.5 implies that this sequence of inclusions induces a sequence of homomorphisms between homology groups, such that $H_k(C^{\frac{\alpha}{2}}(L)) \rightarrow H_k(C^{4\alpha}(L))$ and $H_k(C^\alpha(L)) \rightarrow H_k(C^{2\alpha}(L))$ have ranks equal to $\dim H_k(X^\lambda)$. Hence, by Lemma 3.2, $\text{rank } j_*$ is also equal to $\dim H_k(X^\lambda)$. It follows that $\text{im } j_* \cong H_k(X^\lambda)$. \square

Similarly, Corollary 2.3 provides the following sequence:

$$C^{\frac{\alpha}{4}}(L) \hookrightarrow C_W^\alpha(L) \hookrightarrow C^{8\alpha}(L) \hookrightarrow C^{9\alpha}(L) \hookrightarrow C_W^{36\alpha}(L) \hookrightarrow C^{288\alpha}(L),$$

from which follows a result similar to Theorem 3.6 on the witness complex, by the same proof:

THEOREM 3.7. *Let X be a compact set in \mathbb{R}^d , and let $L \subseteq W \subseteq \mathbb{R}^d$ be such that L is finite. Assume that $d_{\mathcal{H}}(X, W) \leq \delta$ and that $d_{\mathcal{H}}(W, L) \leq \varepsilon$, with $\delta \leq \varepsilon < \min\{\frac{1}{8} \text{diam}_{\text{CC}}(X), \frac{1}{1153} \text{wfs}(X)\}$. Then, for all $\alpha \in [4\varepsilon, \frac{1}{288}(\text{wfs}(X) - \varepsilon)]$ and all $\lambda \in (0, \text{wfs}(X))$, we have: $\forall k \in \mathbb{N}$, $H_k(X^\lambda) \cong \text{im } j_*$, where j_* is the homomorphism between homology groups induced by the canonical inclusion $j : C_W^\alpha(L) \hookrightarrow C_W^{36\alpha}(L)$.*

More generally, the above arguments show that the homology of X^λ can be recovered from the persistence barcode of any filtration $\{F_\alpha\}_{\alpha \geq 0}$ that is intertwined with the Čech filtration in the sense of Lemmas 2.1 and 2.2. Note however that Theorems 3.6 and 3.7 suggest a different behavior of the barcode in this case, since its topological noise might scale up with α (specifically, it might be up to linear in α), whereas it is uniformly bounded by a constant in the case of the Čech filtration. This difference of behavior is easily explained by the way $\{F_\alpha\}_{\alpha \geq 0}$ is intertwined with the Čech filtration. A trick to get a uniformly-bounded noise is to represent the barcode of $\{F_\alpha\}_{\alpha \geq 0}$ on a logarithmic scale, that is, with $\log_2 \alpha$ instead of α in abscissa.

3.2.2 Results on homotopy

The results on homology obtained in Section 3.2.1 follow from simple algebraic arguments. Using a more geometric approach, we can get similar results on homotopy. From now on, $x_0 \in X$ is a fixed point and all the homotopy groups $\pi_k(X) = \pi_k(X, x_0)$ are assumed to be with base-point x_0 . Theorems 3.6 and 3.7 extend to homotopy as follows:

THEOREM 3.8.

- *Under the hypotheses of Theorem 3.6, we have: $\forall k \in \mathbb{N}$, $\pi_k(X^\lambda) \cong \text{im } j_*$, where $j_* : \pi_k(\mathcal{R}^\alpha(L)) \rightarrow \pi_k(\mathcal{R}^{4\alpha}(L))$ is the homomorphism between homotopy groups induced by the inclusion $\mathcal{R}^\alpha(L) \hookrightarrow \mathcal{R}^{4\alpha}(L)$.*

- *Under the hypotheses of Theorem 3.7, we have: $\forall k \in \mathbb{N}$, $\pi_k(X^\lambda) \cong \text{im } j_*$, where $j_* : \pi_k(C_W^\alpha(L)) \rightarrow \pi_k(C_W^{36\alpha}(L))$ is the homomorphism between homotopy groups induced by the inclusion $C_W^\alpha(L) \hookrightarrow C_W^{36\alpha}(L)$.*

The proof of the theorem relies on the following result, which is an immediate generalization of Proposition 4.1 of [12]:

LEMMA 3.9. *Let X be a compact set and L a finite set in \mathbb{R}^d , such that $d_{\mathcal{H}}(X, L) < \varepsilon$ for some $\varepsilon < \frac{1}{4} \text{wfs}(X)$. Let $\alpha, \alpha' \in [\varepsilon, \text{wfs}(X) - \varepsilon]$ be such that $\alpha' - \alpha \geq 2\varepsilon$. Given $k \in \mathbb{N}$, two k -loops $\sigma_1, \sigma_2 : \mathbb{S}^k \rightarrow (L^\alpha, x_0)$ in L^α are homotopic in $X^{\alpha'+\varepsilon}$ if and only if they are homotopic in $L^{\alpha'}$.*

Proof of Theorem 3.8. As mentioned at the beginning of the proof of Lemma 3.3, we can assume without loss of generality that $2\varepsilon < \alpha < \frac{1}{4}(\text{wfs}(X) - \varepsilon)$. Consider the sequence of inclusions introduced in the proof of Theorem 3.6. We use the homotopy equivalences $h_\beta : L^\beta \rightarrow C^\beta(L)$ provided by Lemma 3.4 for all values $\beta > 0$, which commute with inclusions at homotopy level. Note that, for any element σ of $\pi_k(C^\beta(L))$, there exists a k -loop in L^β that is mapped through h_β to a k -loop representing the homotopy

class σ . In the following, we denote by σ_g such a k -loop. Let E, F and G be the images of $\pi_k(\mathcal{C}^{\frac{\alpha}{2}}(L))$ in $\pi_k(\mathcal{C}^\alpha(L))$, $\pi_k(\mathcal{C}^{2\alpha}(L))$ and $\pi_k(\mathcal{C}^{4\alpha}(L))$ respectively, through the homomorphisms induced by inclusion. We thus have a sequence of surjective homomorphisms: $\pi_k(\mathcal{C}^{\frac{\alpha}{2}}(L)) \rightarrow E \rightarrow F \rightarrow G$.

Note that, by Theorem 3.5, F and G are isomorphic to $\pi_k(X^\lambda)$. Let $\sigma \in F$ be a homotopy class. Since F is the image of $\pi_k(\mathcal{C}^{\frac{\alpha}{2}}(L))$, we can assume without loss of generality that $\sigma_g \subset L^{\frac{\alpha}{2}}$. Assume that the image of σ in G is zero. Then, σ_g is null-homotopic in $L^{4\alpha}$ and, since $L^{4\alpha} \subset X^{4\alpha+\varepsilon}$, σ_g is also null-homotopic in $X^{4\alpha+\varepsilon}$. But $\sigma_g \subset L^{\frac{\alpha}{2}} \subset X^{\frac{\alpha}{2}+\varepsilon}$, and $X^{2\alpha+\varepsilon}$ deformation retracts onto $X^{\frac{\alpha}{2}+\varepsilon}$, by the Isotopy Lemma 3.1. Therefore, σ_g is null-homotopic in $X^{\frac{\alpha}{2}+\varepsilon}$, which is contained in $L^{2\alpha}$ since $\frac{\alpha}{2} + 2\varepsilon < 2\alpha$. Hence, σ_g is null-homotopic in $L^{2\alpha}$, *i.e.* $\sigma = 0$ in F . So, the homomorphism $F \rightarrow G$ is injective, and therefore it is an isomorphism. Thus, $F \rightarrow \pi_k(\mathcal{R}^{4\alpha}(L))$ is injective, and it is now enough to prove that the image of the homomorphism $\phi_* : \pi_k(\mathcal{R}^\alpha(L)) \rightarrow \pi_k(\mathcal{C}^{2\alpha}(L))$ induced by inclusion is F .

Obviously, F is contained in the image of ϕ_* . Now, let $\sigma \in \pi_k(\mathcal{R}^\alpha(L))$ and let $\phi_*(\sigma)_g$ be a k -loop in $L^{2\alpha}$ that is mapped through $h_{2\alpha}$ to a k -loop representing the homotopy class $\phi_*(\sigma)$. Since $\phi_*(\sigma)$ is in the image of ϕ_* , and since $\mathcal{R}^\alpha(L) \subset \mathcal{C}^\alpha(L)$, we can assume that $\phi_*(\sigma)_g$ belongs to L^α . Let $\tilde{\sigma}_g$ be the image of $\phi_*(\sigma)_g$ through a deformation retraction of $X^{2\alpha+\varepsilon}$ onto X^{α_0} , where $0 < \alpha_0 < \frac{\alpha}{2}$ is such that $\frac{\alpha}{2} - \alpha_0 > \varepsilon$. Obviously, $\tilde{\sigma}_g$ and $\phi_*(\sigma)_g$ are homotopic in $X^{2\alpha+\varepsilon}$, and it follows from Lemma 3.9 that $\tilde{\sigma}_g$ and $\phi_*(\sigma)_g$ are also homotopic in $L^{2\alpha}$. And since $\tilde{\sigma}_g$ is contained in $X^{\alpha_0} \subset L^{\frac{\alpha}{2}}$, the equivalence class of $h_{\frac{\alpha}{2}}(\tilde{\sigma}_g)$ in $\pi_k(\mathcal{C}^{\frac{\alpha}{2}}(L))$ is mapped to $\phi_*(\sigma) \in \pi_k(\mathcal{C}^{2\alpha}(L))$ through the homomorphism induced by $\mathcal{C}^{\frac{\alpha}{2}}(L) \hookrightarrow \mathcal{C}^{2\alpha}(L)$, which commutes with the homotopy equivalences. As a result, $\phi_*(\sigma)$ belongs to F , which is thus equal to $\text{im } \phi_*$. This proves the first part of the theorem. The proof of the second part is mostly the same. \square

4. SMOOTH SUBMANIFOLDS OF \mathbb{R}^D

In this section, we consider the case of submanifolds X of \mathbb{R}^d that have positive *reach*. Recall that the reach of X , or $\text{rch}(X)$ for short, is the minimum distance between the points of X and the points of its medial axis [1]. A point cloud $L \subset X$ is an ε -*sample* of X if every point of X lies within distance ε of L . In addition, L is ε -*sparse* if its points lie at least ε away from one another.

Theorem 4.1 below is a first attempt at quantifying a conjecture of Carlsson and de Silva [17], according to which the witness complex filtration should have *cleaner* persistence barcodes than the Čech and Rips filtrations, at least on smooth submanifolds of \mathbb{R}^d . By *cleaner* is meant that the amplitude of the topological noise in the barcodes should be smaller, and also that the long intervals should appear earlier. We prove this latter statement correct to some extent:

THEOREM 4.1. *There exist a constant $\varrho > 0$ and a continuous, non-decreasing map $\bar{\omega} : [0, \varrho] \rightarrow [0, \frac{1}{2}]$, with $\bar{\omega}(0) = 0$, such that, for any submanifold X of \mathbb{R}^d , for all ε, δ satisfying $0 < \delta \leq \varepsilon < \varrho \text{rch}(X)$, for any δ -sample W of X and any ε -sparse ε -sample L of W , $\mathcal{C}_W^\alpha(L)$ contains a subcomplex \mathcal{D} homeomorphic to X and such that the canonical inclusion $\mathcal{D} \hookrightarrow \mathcal{C}_W^\alpha(L)$ induces an injective homomorphism between homology groups, provided that α satisfies: $\frac{8}{3}(\delta + \bar{\omega}(\frac{\varepsilon}{\text{rch}(X)}))^2 \varepsilon \leq \alpha < \frac{1}{2} \text{rch}(X) - (3 + \frac{\sqrt{2}}{2})(\varepsilon + \delta)$.*

This theorem guarantees that, for values of α ranging from $O(\delta + \bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})^2 \varepsilon)$ to $O(\text{rch}(X))$, the topology of X is captured by a subcomplex \mathcal{D} that injects itself suitably in $\mathcal{C}_W^\alpha(L)$. As a result, long intervals showing the homology of X appear around $\alpha = O(\delta + \bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})^2 \varepsilon)$ in the persistence barcode of the witness complex filtration. This can be much sooner than the time $\alpha = 2\varepsilon$ prescribed by Theorem 3.7, since $\bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})$ can be arbitrarily small. Specifically, the denser the landmark set L , the smaller the ratio $\frac{\varepsilon}{\text{rch}(X)}$, and therefore the smaller $\frac{8}{3}(\delta + \bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})^2 \varepsilon)$ compared to 2ε .

Our proof of Theorem 4.1 stresses the close relationship that exists between the α -witness complex and the so-called *weighted restricted Delaunay triangulation* $\mathcal{D}_\omega^X(L)$. Given a submanifold X of \mathbb{R}^d , a finite landmark set $L \subset \mathbb{R}^d$, and an assignment of non-negative weights to the landmarks, specified through a map $\omega : L \rightarrow [0, \infty)$, $\mathcal{D}_\omega^X(L)$ is the nerve of the restriction to X of the *power diagram*² of the weighted set L . By a result of Cheng *et al.* (see Theorem 4.2 below), $\mathcal{D}_\omega^X(L)$ is homeomorphic to X under a sufficient landmark density and under a suitable choice of weights – bounded from above by $\bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})$. The main point of our proof is then to show that $\mathcal{C}_W^\alpha(L)$ contains $\mathcal{D}_\omega^X(L)$ and that the latter injects itself *nice*ly into the former.

In the special case where X is a smooth curve or surface, all weights can be taken to be zero, since the unweighted restricted Delaunay triangulation is known to be homeomorphic to X [1, 2]. As a result, function $\bar{\omega}$ is zero, and the long intervals showing the homology of X in the barcode of the witness complex filtration appear already at time $\alpha = O(\delta)$.

In the general case however, the upper bound on the appearance time of long bars cannot depend solely on δ , since otherwise, in the limit case where $\delta = 0$ (*i.e.* $W = X$), we would get that the homology groups of X can be injected into the ones of $\mathcal{C}_X(L)$, which is known to be true for curves and surfaces [3], but not for 3-manifolds [30]. Now, whether $O(\delta + \bar{\omega}(\frac{\varepsilon}{\text{rch}(X)})^2 \varepsilon)$ is a tight upper bound or not is open.

The rest of Section 4 is devoted to the proof of Theorem 4.1. After introducing the weighted restricted Delaunay triangulation formally in Section 4.1, we stress its relationship with the α -witness complex in Section 4.2, and then we detail the proof of Theorem 4.1 in Section 4.3.

4.1 Weighted restricted Delaunay triangulation

Given a finite point set $L \subset \mathbb{R}^d$, an *assignment of weights over L* is a non-negative real-valued function $\omega : L \rightarrow [0, \infty)$. The quantity $\max_{u \in L, v \in L \setminus \{u\}} \frac{\omega(u)}{\|u-v\|}$ is called the *relative amplitude* of ω . Given $p \in \mathbb{R}^d$, the *weighted distance* from p to some weighted point $v \in L$ is $\|p-v\|^2 - \omega(v)^2$. This is actually not a metric, since it is not symmetric. Given a finite point set L and an assignment of weights ω over L , we denote by $\mathcal{V}_\omega(L)$ the power diagram of the weighted set L , and by $\mathcal{D}_\omega(L)$ its nerve, also known as the weighted Delaunay triangulation. If the relative amplitude of ω is at most $\frac{1}{2}$, then the points of L have non-empty cells in $\mathcal{V}_\omega(L)$, and in fact each point of L belongs to its own cell [13]. For any simplex σ of $\mathcal{D}_\omega(L)$, $\mathcal{V}_\omega(\sigma)$ denotes its dual face in $\mathcal{V}_\omega(L)$.

Given a subset X of \mathbb{R}^d , we call $\mathcal{V}_\omega^X(L)$ the restriction of $\mathcal{V}_\omega(L)$ to X , and we denote by $\mathcal{D}_\omega^X(L)$ its nerve, also known as the weighted Delaunay triangulation of L restricted to

²More on power diagrams and on restricted Delaunay triangulations can be found in [4] and [22] respectively.

Now that we have built a map $\tilde{h}_0 : C_W^\alpha(L) \rightarrow \mathbb{R}^d$, our next step is to turn it into a map $C_W^\alpha(L) \rightarrow X$. To do so, we compose it with the projection p_X that maps every point of \mathbb{R}^d to its nearest neighbor on X , if the latter is unique. This projection is known to be well-defined and continuous over $\mathbb{R}^d \setminus M$, where M denotes the medial axis of X [23].

LEMMA 4.5. *Let $X, W, L, \delta, \varepsilon$ satisfy the hypotheses of Theorem 4.1. Then, $\tilde{h}_0(C_W^\alpha(L)) \cap M = \emptyset$ provided that $\alpha < \frac{1}{2} \text{rch}(X) - \left(3 + \frac{\sqrt{2}}{2}\right)(\varepsilon + \delta)$.*

Since by Lemma 4.5 we have $\tilde{h}_0(C_W^\alpha(L)) \cap M = \emptyset$, the map $p_X \circ \tilde{h}_0 : C_W^\alpha(L) \rightarrow X$ is well-defined and continuous. Our final step is to compose it with h_0^{-1} , to get a continuous map $h = h_0^{-1} \circ p_X \circ \tilde{h}_0 : C_W^\alpha(L) \rightarrow \mathcal{D}_\omega^X(L)$. The restriction of h to $\mathcal{D}_\omega^X(L)$ is simply $h_0^{-1} \circ p_X \circ h_0$, which coincides with $h_0^{-1} \circ h_0 = \text{id}$ since $h_0(\mathcal{D}_\omega^X(L)) = X$. It follows that $h \circ j$ is the identity in $\mathcal{D}_\omega^X(L)$, and therefore that the induced map $h_* \circ j_*$ is also the identity. This implies that $j_* : H_k(\mathcal{D}_\omega^X(L)) \rightarrow H_k(C_W^\alpha(L))$ is injective, which concludes the proof of Theorem 4.1. \square

We end the section by providing the proof of Lemma 4.5:

Proof of Lemma 4.5. First, we claim that the image through \tilde{h}_0 of any simplex of $C_W^\alpha(L)$ is included in the Euclidean convex hull of the restricted Voronoi cells of its simplices, that is:

$$\forall \sigma \in C_W^\alpha(L), \tilde{h}_0(\sigma) \subseteq \text{CH} \left(\bigcup_{v \text{ vertex of } \sigma} V_\omega(v) \cap X \right).$$

This is clearly true if σ belongs to $\mathcal{D}_\omega^X(L)$, since in this case we have $\tilde{h}_0(\sigma) = h_0(\sigma) \subseteq \bigcup_v \text{vertex of } \sigma V_\omega(v) \cap X$, as mentioned after Theorem 4.2. Now, if the property holds for all the proper faces of a simplex $\sigma \in C_W^\alpha(L)$, then by induction it also holds for the simplex itself. Indeed, for each proper face $\tau \subset \sigma$, we have $\tilde{h}_0(\tau) \subseteq \text{CH} \left(\bigcup_v \text{vertex of } \tau V_\omega(v) \cap X \right)$, which is included in $\text{CH} \left(\bigcup_v \text{vertex of } \sigma V_\omega(v) \cap X \right)$. Therefore, $\text{CH} \left(\bigcup_v \text{vertex of } \sigma V_\omega(v) \cap X \right)$ contains $\text{CH} \left(\tilde{h}_0(\partial\sigma) \right)$, which, by Dugundji's extension theorem, contains $\tilde{h}_0(\sigma)$. Thus, the property holds for every simplex of $C_W^\alpha(L)$.

We can now prove that the image through \tilde{h}_0 of any arbitrary simplex σ of $C_W^\alpha(L)$ does not intersect the medial axis of X . This is clearly true if σ is a simplex of $\mathcal{D}_\omega^X(L)$, since in this case $\tilde{h}_0(\sigma) = h_0(\sigma)$ is included in X . Assume now that $\sigma \notin \mathcal{D}_\omega^X(L)$. In particular, σ is not a vertex. Let v be an arbitrary vertex of σ . Consider any other vertex u of σ . Edge $[u, v]$ is α -witnessed by some point $w_{uv} \in W$. We then have $\|v - u\| \leq \|v - w_{uv}\| + \|w_{uv} - u\| \leq 2d_2(w_{uv}) + 2\alpha$, where $d_2(w_{uv})$ stands for the Euclidean distance from w_{uv} to its second nearest landmark. According to Lemma 3.4 of [6], we have $d_2(w) \leq 3(\varepsilon + \delta)$, since L is an $(\varepsilon + \delta)$ -sample of X . Thus, all the vertices of σ are included in the Euclidean ball $B(v, 2\alpha + 6(\varepsilon + \delta))$. Moreover, for any vertex u of σ and any point $p \in V_\omega(u) \cap X$, we have $\|p - u'\| \leq \varepsilon + \delta$, where u' is a landmark closest to p in the Euclidean metric. Combined with the fact that $\|p - u\|^2 - \omega(u)^2 \leq \|p - u'\|^2 - \omega(u')^2$, we get: $\|p - u\|^2 \leq \|p - u'\|^2 + \omega(u)^2 \leq 2(\varepsilon + \delta)^2$, since by Lemma 3.3 of [6] we have $\omega(u) \leq 2\bar{\omega} \left(\frac{\varepsilon}{\text{rch}(X)} \right) (\varepsilon + \delta) \leq \varepsilon + \delta$. Hence, $V_\omega(u) \cap X$ is included in $B(u, \sqrt{2}(\varepsilon + \delta)) \subset B(v, 2\alpha + (6 + \sqrt{2})(\varepsilon + \delta))$. Since this is true for every vertex

Input: W , distances $\{l(w, w'), w, w' \in W\}$.

Init: Let $L := \emptyset$, $\varepsilon := +\infty$;

While $L \subsetneq W$ **do**

 Let $p := \text{argmax}_{w \in W} \min_{v \in L} l(w, v)$;

 // p is chosen arbitrarily in W if $L = \emptyset$

$L := L \cup \{p\}$;

$\varepsilon := \max_{w \in W} \min_{v \in L} l(w, v)$;

 Update $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$;

 Compute persistence of $\mathcal{R}^{4\varepsilon}(L) \hookrightarrow \mathcal{R}^{16\varepsilon}(L)$;

End_while

Output: diagram showing the evolution of persistent Betti numbers versus ε .

Figure 2: Pseudo-code of the algorithm.

u of σ , we get: $\tilde{h}_0(\sigma) \subseteq \text{CH} \left(\bigcup_u \text{vertex of } \sigma V_\omega(u) \cap X \right) \subseteq B(v, 2\alpha + (6 + \sqrt{2})(\varepsilon + \delta))$. Now, v belongs to $L \subseteq W \subseteq X$, and by assumption we have $2\alpha + (6 + \sqrt{2})(\varepsilon + \delta) < \text{rch}(X)$, therefore $\tilde{h}_0(\sigma)$ does not intersect the medial axis of X . \square

5. APPLICATION TO RECONSTRUCTION

Taking advantage of the structural results of Section 3, we devise a very simple yet provably-good algorithm for constructing nested pairs of complexes that capture the homology of a large class of compact subsets of \mathbb{R}^d . This algorithm is a variant of the greedy refinement technique of [26], which builds a set L of landmarks iteratively and in the meantime maintains a suitable data structure. In our case, the data structure is composed of a nested pair of simplicial complexes, which can be either $\mathcal{R}^\alpha(L) \hookrightarrow \mathcal{R}^{\alpha'}(L)$ or $C_W^\alpha(L) \hookrightarrow C_W^{\alpha'}(L)$, for specific values $\alpha < \alpha'$. Both variants of the algorithm enjoy similar theoretical guarantees, but the variant using witness complexes is likely to bevgugmax(

the algorithm. Since $L(i)$ keeps growing with i , $\varepsilon(i)$ is a decreasing function of i . In addition, $L(i)$ is an $\varepsilon(i)$ -sample of W , by definition of $\varepsilon(i)$. Hence, if W lies at Hausdorff distance δ of some compact set $X \subset \mathbb{R}^d$, then we have $d_{\mathcal{H}}(L(i), X) \leq \delta + \varepsilon(i)$. Therefore, Theorem 3.6 provides us with the following theoretical guarantee:

THEOREM 5.1. *If the input W lies at Hausdorff distance δ of some compact set $X \subset \mathbb{R}^d$, with $\delta < \frac{1}{18} \text{wfs}(X)$, then, at each iteration i such that $\delta < \varepsilon(i) < \frac{1}{18} \text{wfs}(X)$, the persistent homology groups of $\mathcal{R}^{4\varepsilon(i)}(L(i)) \hookrightarrow \mathcal{R}^{16\varepsilon(i)}(L(i))$ are isomorphic to the homology groups of X^λ , $\forall \lambda \in (0, \text{wfs}(X))$.*

Under mild conditions on the input, this theorem guarantees the existence of a plateau showing the homology of X^λ , of length at least $(\frac{1}{18} \text{wfs}(X) - \delta)$, in the diagram of persistent Betti numbers. When δ is small enough compared to $\text{wfs}(X)$, the plateau is large enough to be detected. In cases where W samples several compact sets with different weak feature sizes, the theorem ensures that several plateaus appear in the diagram, showing plausible reconstructions at various scales – see Figure 3. Once a relevant scale has been selected, the corresponding landmark set and nested complexes are easily rebuilt. Differently from the algorithm of [26], this outcome is not a single embedded simplicial complex but a nested pair of abstract complexes, whose images in \mathbb{R}^d lie at Hausdorff distance⁶ $O(\varepsilon)$ of X , and whose persistent homology gives the homology of X^λ .

Update of $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$.

We now describe how to maintain $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$. In fact, we settle for describing how to rebuild $\mathcal{R}^{16\varepsilon}(L)$ completely at each iteration, which is sufficient for achieving our complexity bounds, although it is clearly much preferable in practice to use more local rules to update the simplicial complexes. Consider the one-skeleton graph G of $\mathcal{R}^{16\varepsilon}(L)$. By definition, a simplex that is not a vertex belongs to $\mathcal{R}^{16\varepsilon}(L)$ if and only if all its edges are in G . Therefore, the simplices of $\mathcal{R}^{16\varepsilon}(L)$ are precisely the cliques of G . The simplicial complex can then be built as follows: (1.) build graph G , (2.) find all maximal cliques in G , and (3.) report the maximal cliques and all their subcliques. We perform Step 1. naively in $O(|L|^2)$ time, where $|L|$ denotes the size of L . For Step 2., we use the output-sensitive algorithm of [32], which finds all the maximal cliques of G in $O(k|L|^3)$ time, where k is the size of the answer. Finally, we report all the subcliques of the maximal cliques in a time that is linear in the total number of cliques, which is also the size of $\mathcal{R}^{16\varepsilon}(L)$. Therefore, at each iteration of the algorithm, $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$ are rebuilt within $O(|\mathcal{R}^{16\varepsilon}(L)| |L|^3)$ time, where $|\mathcal{R}^{16\varepsilon}(L)|$ denotes the size of $\mathcal{R}^{16\varepsilon}(L)$.

Running time of the algorithm.

Let $|W|$ denote the size of W . At each iteration of the algorithm, point p and parameter ε are computed naively by iterating over the points of W , and for each such point, by reviewing its distances to all the landmarks. This procedure takes $O(|W||L|)$ time. Once $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$ have been updated, the persistence algorithm runs in $O(|\mathcal{R}^{16\varepsilon}(L)|^3)$ time [21, 33]. Hence,

⁶Indeed, every simplex of $\mathcal{R}^{16\varepsilon}(L)$ has all its vertices in $X^{\varepsilon+\delta} \subset X^{2\varepsilon}$, and the lengths of its edges are at most 16ε .




Figure 3: Output of our algorithm when applied blindly to the 1000-dimensional data set of Figure 1.

LEMMA 5.2. *The time complexity of every iteration of the algorithm is $O(|W||L| + |\mathcal{R}^{16\varepsilon}(L)||L|^3 + |\mathcal{R}^{16\varepsilon}(L)|^3)$.*

In addition, standard packing arguments (omitted in this extended abstract) provide the following tight upper bounds on the size of $\mathcal{R}^{16\varepsilon}(L)$ in Euclidean spaces:

LEMMA 5.3. *If L is a finite ε -sparse point set in \mathbb{R}^d , then $|\mathcal{R}^{16\varepsilon}(L)| \leq 2^{33d}|L|$. If in addition L lies on a m -submanifold X of \mathbb{R}^d , with $\text{rch}(X) > 256\varepsilon$, then $|\mathcal{R}^{16\varepsilon}(L)| \leq 2^{35m}|L|$. These upper bounds are tight in order of magnitude.*

Whenever the input point cloud W lies on a smooth m -submanifold X of \mathbb{R}^d , the lemma suggests⁷ that the algorithm goes through two consecutive phases. First, a transition phase where the landmark set L is too coarse for the dimensionality of X to have an influence on the shapes and sizes of the stars of the vertices of $\mathcal{R}^{16\varepsilon}(L)$. For instance, if X is an embedded curve that roughly fills in the unit ball in \mathbb{R}^d , then, for large values of ε , L is nothing but a sampling of the d -ball. Then comes a second, stable phase, where L is dense enough for the dimensionality of X to play a role. Denoting by i_0 the last iteration of the transition phase, we deduce from Lemmas 5.2 and 5.3 that the running time of the algorithm is $O(|W||L(i_0)|^2 + 8^{33d}|L(i_0)|^5 + 8^{35m}|W|^5)$. There remains to get rid of the term depending on d , which we do using a backtracking strategy. Specifically, we first run the algorithm without maintaining $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$, which simply sorts the points of W according to their order of insertion in L . Then, we run the algorithm backwards, starting with $L = L(|W|) = W$ and considering at each iteration j the landmark set $L(|W| - j)$. During this second phase, we do maintain $\mathcal{R}^{4\varepsilon}(L)$ and $\mathcal{R}^{16\varepsilon}(L)$ and compute their persistent Betti numbers. If W samples X densely enough, then Theorem 5.1 ensures that the relevant plateaus will be computed before the transition phase starts, and thus before the size of the data structure becomes independent of the dimension of X . It is then up to the user to stop the process when the space complexity becomes too large. This variant of the algorithm has the following complexity:

THEOREM 5.4. *If W is a point cloud in Euclidean space \mathbb{R}^d , then the running time of the algorithm is $O(8^{33d}|W|^5)$,*

⁷Note that, at every iteration i of the process, $L(i)$ is an $\varepsilon(i)$ -sparse point set, since the algorithm always inserts in L the point of W lying furthest away from L [26, Lemma 4.1].

where $|W|$ denotes the size of W . If in addition W is a δ -sample of some smooth m -submanifold of \mathbb{R}^d with δ small enough, then the running time becomes $O(8^{35m} |W|^5)$.

By Lemma 5.3, the bounds in Theorem 5.4 are tight in order of magnitude. Thus, the algorithm can have a doubly exponential complexity in m when the input point set is densely sampled from a m -dimensional smooth manifold. However, it can be shown that the m -skeleton of the Rips complex is in fact simply exponential in m . Hence, when a reasonable upper bound b on m is known, one can reduce the running time of the algorithm to $2^{O(m^2)} |W|^5$ by considering only the b -skeleton of the Rips complex. Similarly, the running time reduces to $2^{O(d^2)} |W|^5$ if only the d -skeleton is considered.

6. CONCLUSION

This paper makes effective the approach developed in [12, 15] by providing an efficient, provably-good and easy-to-implement algorithm for the topological and geometric analysis of point cloud data in arbitrary dimensions. Addressing a weaker version of the classical reconstruction problem, the algorithm ultimately outputs a nested pair of complexes at a user-defined scale, from which the homology of the underlying shape X can be inferred. When X is a smooth submanifold of \mathbb{R}^d , the complexity of the algorithm scales up with the intrinsic dimension of X and not with the ambient dimension d , assuming that the pairwise distances between the data points have been pre-computed. Thus, a new step is made towards reconstructing (low-dimensional) manifolds in high-dimensional spaces in reasonable time with guarantees. However, there still remains the challenging problem of constructing an embedded complex that is topologically equivalent and geometrically close to the sampled shape.

The theoretical framework developed in the paper can be used for the analysis of various persistence-based methods in Euclidean spaces. It can also virtually be applied in any metric space (provided that the result of [12, 15] on unions of balls can be extended), thanks to the genericity of Lemma 3.4 and of the arguments of Section 3.2.1. A class of spaces of particular interest to us is the class of compact Riemannian manifolds, possibly with boundaries, with applications in machine learning and sensor networks.

Acknowledgements.

The authors wish to thank Gunnar Carlsson, Vin de Silva, Leonidas Guibas, Tamal Dey, and the anonymous referees for their insightful comments and suggestions. Our experiments were carried out using Matlab and the Plex library. This work was partially supported by ANR grant GeoTopAl.

7. REFERENCES

- [1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Comput. Geom. Appl.*, 22(4):481–504, 1999.
- [2] N. Amenta, M. Bern, and D. Eppstein. The crust and the β -skeleton: Combinatorial curve reconstruction. *Comput. Geom. Appl.*, 60:125–135, 1998.
- [3] D. Attali, H. Edelsbrunner, and Y. Mileyko. Weak witnesses for Delaunay triangulations of submanifolds. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2007. To appear.
- [4] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *Comput. Geom. Appl.*, 23(3):345–405, September 1991.
- [5] P. Bendich, D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and D. Morozov. Inferring local homology from sampled stratified spaces. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 536–546, 2007.
- [6] J.-D. Boissonnat, L. J. Guibas, and S. Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 194–203, 2007.
- [7] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, June 2007.
- [8] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J.D. Boissonnat and M. Teillaud, editors, *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 231–273. Springer, 2006.
- [9] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 319–326, 2006.
- [10] F. Chazal and A. Lieutier. The λ -medial axis. *Comput. Geom. Appl.*, 67(4):304–331, July 2005.
- [11] F. Chazal and A. Lieutier. Topology guaranteeing manifold reconstruction using distance function to noisy data. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 112–118, 2006.
- [12] F. Chazal and A. Lieutier. Stability and computation of topological invariants of solids in \mathbb{R}^n . *Comput. Geom. Appl.*, 37(4):601–617, 2007.
- [13] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S.-H. Teng. Sliver exudation. *Comput. Geom. Appl.*, 47(5):883–904, 2000.
- [14] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 1018–1027, 2005.
- [15] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 263–271, 2005.
- [16] V. de Silva. A weak definition of Delaunay triangulation. Technical report, Stanford University, October 2003. To appear in *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*.
- [17] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 157–166, 2004.
- [18] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Comput. Geom. Appl.*, 7:339–358, 2007.
- [19] J. Dugundji. An extension of Tietze’s theorem. *Ann. of Math.*, 1:353–367, 1951.
- [20] H. Edelsbrunner. The union of balls and its dual shape. *Comput. Geom. Appl.*, 13:415–440, 1995.
- [21] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Comput. Geom. Appl.*, 28:511–533, 2002.
- [22] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Comput. Geom. Appl.*, 7:365–378, 1997.
- [23] H. Federer. Curvature measures. *Comm. Pure Appl. Math.*, 93:418–491, 1959.
- [24] R. Ghrist. Barcodes: The persistent topology of data. *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, October 2007.
- [25] K. Grove. Critical point theory for distance functions. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, volume 54, 1993.
- [26] L. G. Guibas and S. Y. Oudot. Reconstruction using witness complexes. In *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 1076–1085, 2007.
- [27] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2001.
- [28] T. Kaczynski, K. Mischaikow, and M. Mrozek. *On the Topology of Manifolds*. Number 157 in Applied Mathematical Sciences. Springer-Verlag, 2004.
- [29] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Proceedings of the 12th ACM SIGPLAN Conference on Programming Language Design and Implementation*, to appear.
- [30] S. Y. Oudot. On the topology of the restricted Delaunay triangulation and witness complex in higher dimensions. Technical report, Stanford University, November 2006. LANL arXiv:0803.1296v1 [cs.CG], <http://arxiv.org/abs/0803.1296>.
- [31] V. Robins. Towards computing homology from approximations. *Comput. Geom. Appl.*, 24:503–532, 1999.
- [32] S. Tsukiyama, M. Ide, H. Ariyoshi, and I. Shirakawa. A new algorithm for generating all the maximal independent sets. *Comput. Geom. Appl.*, 6:505–517, 1977.
- [33] A. Zomorodian and G. Carlsson. Computing persistent homology. *Comput. Geom. Appl.*, 33(2):249–274, 2005.