

Linear-Size Approximations to the Vietoris–Rips Filtration

Donald R. Sheehy

Received: 13 July 2012 / Revised: 15 February 2013 / Accepted: 18 April 2013 /
Published online: 25 May 2013
© Springer Science+Business Media New York 2013

Abstract The Vietoris–Rips filtration is a versatile tool in topological data analysis. It is a sequence of simplicial complexes built on a metric space to add topological structure to an otherwise disconnected set of points. It is widely used because it encodes useful information about the topology of the underlying metric space. This information is often extracted from its so-called persistence diagram. Unfortunately, this filtration is often too large to construct in full. We show how to construct an $O(n)$ -size filtered simplicial complex on an n -point metric space such that its persistence diagram is a good approximation to that of the Vietoris–Rips filtration. This new filtration can be constructed in $O(n \log n)$ time. The constant factors in both the size and the running time depend only on the doubling dimension of the metric space and the desired tightness of the approximation. For the first time, this makes it computationally tractable to approximate the persistence diagram of the Vietoris–Rips filtration across all scales for large data sets. We describe two different sparse filtrations. The first is a zigzag filtration that removes points as the scale increases. The second is a (non-zigzag) filtration that yields the same persistence diagram. Both methods are based on a hierarchical net-tree and yield the same guarantees.

Keywords Persistent Homology · Vietoris–Rips filtration · Net-trees

1 Introduction

There is an extensive literature on the problem of computing sparse approximations to metric spaces (see the book [29] and references therein). There is also a growing

D. R. Sheehy (✉)
Inria Saclay Île-de-France, Bâtiment Alan Turing, 1 rue Honoré d’Estienne d’Orves,
91120 Palaiseau, France
e-mail: don.r.sheehy@gmail.com

literature on topological data analysis and its efforts to extract topological information from metric data (see the survey [3] and references therein). One might expect that topological data analysis would be a major user of metric approximation algorithms, especially given that topological data analysis often considers simplicial complexes that grow exponentially in the number of input points. Unfortunately, this is not the case. The benefits of a sparser representation are sorely needed, but it is not obvious how an approximation to the metric will affect the underlying topology. The goal of this paper is to bring together these two research areas and to show how to build sparse metric approximations that come with topological guarantees.

The target for approximation is the Vietoris–Rips complex, which has a simplex for every subset of input points with diameter at most some parameter α . The collection of Vietoris–Rips complexes at all scales yields the Vietoris–Rips filtration. The persistence algorithm takes this filtration and produces a persistence diagram representing the changes in topology corresponding to changes in scale [33]. The Vietoris–Rips filtration has become a standard tool in topological data analysis because it encodes relevant and useful information about the topology of the underlying metric space [8]. It also extends easily to high dimensional data, general metric spaces, or even non-metric distance functions.

Unfortunately, the Vietoris–Rips filtration has a major drawback: It’s huge! Even the k -skeleton (the simplices up to dimension k) has size $O(n^{k+1})$ for n points.

This paper proposes an alternative filtration called the sparse Vietoris–Rips filtration, which has size $O(n)$ and can be computed in $O(n \log n)$ time. Moreover, the persistence diagram of this new filtration is provably close to that of the Vietoris–Rips filtration. The constants depend only on the doubling dimension of the metric (defined below) and a user-defined parameter ε governing the tightness of the approximation. For the k -skeleton, the constants are bounded by $(\frac{1}{\varepsilon})^{O(kd)}$.

The main tool we use to construct the sparse filtration is the net-tree of Har-Peled and Mendel [25]. Net-trees are closely related to hierarchical metric spanners [22, 23] and their construction is analogous to data structures used for nearest neighbor search in metric spaces [12, 13, 16].

Outline After reviewing some related work and definitions in Sects. 2 and 3, we explain how to perturb the input metric using weighted distances in Sect. 4. This perturbation is used in the definition of a sparse zigzag filtration in Sect. 5, i.e. one in which simplices are both added and removed as the scale increases. The full definition of the net-trees is given in Sect. 6. Using the properties of the net-tree and the perturbed distances, we prove in Sect. 7 that removing points from the filtration does not change the topology. This implies that the zigzag filtration does not actually zigzag at the homology level (Sect. 8.1). The zigzag filtration can then be converted into an ordinary (i.e. non-zigzag) filtration that also approximates the Vietoris–Rips filtration (Sect. 8.2). The theoretical guarantees are proven in Sect. 9. Section 9.1 proves that the resulting persistence diagrams are good approximations to the persistence diagram of the full Vietoris–Rips filtration. The size complexity of the sparse filtration is shown to be $O(n)$ in Sect. 9.2. Finally, in Sect. 10, we outline the $O(n \log n)$ -time construction, which turns out to be quite easy once you have a net-tree.

2 Related Work

The theory of persistent homology [21,33] gives an algorithm for computing the persistent topological features of a complex that grows over time. It has been applied successfully to many problem domains, including image analysis [6], biology [11,30], and sensor networks [18,19]. See also the survey by Carlsson for background on the topological view of data [3]. It is also possible to consider the complexes that alternate between growing and shrinking in what is known as zigzag persistence [4,5,27,31].

Due to the rapid blowup in the size of the Vietoris–Rips filtration, some attempts have been made to build approximations. Some notable examples include witness complexes [2,17,24] as well as the mesh-based methods of Hudson et al. in Euclidean spaces [26].

The work most similar to the current paper is by Chazal and Oudot [10]. In that paper, they looked at a sequence of persistence diagrams on denser and denser subsamples. However, they were not able to combine these diagrams into a single diagram with a provable guarantee. Moreover, they were not able to prove general guarantees on the size of the filtration except under very strict assumptions on the data.

Recently, Zomorodian [32] and Attali et al. [1] have presented new methods for simplifying Vietoris–Rips complexes. These methods depend only on the combinatorial structure. However, they have not yielded results in simplifying filtrations, only static complexes. In this paper, we exploit the geometry to get topologically equivalent sparsification of an entire filtration.

3 Background

Doubling metrics For a point $p \in P$ and a set $S \subseteq P$, we will write $\mathbf{d}(p, S)$ to denote the minimum distance from p to S , i.e. $\mathbf{d}(p, S) = \min_{q \in S} \mathbf{d}(p, q)$. In a metric space $\mathcal{M} = (P, \mathbf{d})$, a **metric ball** centered at $p \in P$ with radius $r \in \mathbb{R}$ is the set $\mathbf{ball}(p, r) = \{q \in P : \mathbf{d}(p, q) \leq r\}$.

Definition 1 The **doubling constant** λ of a metric space $\mathcal{M} = (P, \mathbf{d})$ is the minimum number of metric balls of radius r required to cover any ball of radius $2r$. The **doubling dimension** is $d = \lceil \lg \lambda \rceil$. A metric space whose doubling dimension is bounded by a constant is called a **doubling metric**.

The **spread** Δ of a metric space $\mathcal{M} = (P, \mathbf{d})$ is the ratio of the largest to smallest interpoint distances. A metric with doubling dimension d and spread Δ has at most $\Delta^{O(d)}$ points. This is easily seen by starting with a ball of radius equal to the largest pairwise distance and covering it with λ balls of half this radius. Covering all of the resulting balls by yet smaller balls and repeating $O(\log \Delta)$ times results in balls that can contain at most one point each because the radii are smaller than the minimum interpoint distance. The number of such balls is $\lambda^{O(\log \Delta)} = \Delta^{O(\log \lambda)} = \Delta^{O(d)}$.

Simplicial Complexes A **simplicial complex** X is a collection of **vertices** denoted $V(X)$ and a collection of subsets of $V(X)$ called **simplices** that is closed under the subset operation, i.e. $\sigma \subset \psi$ and $\psi \in X$ together imply $\sigma \in X$. The **dimension**




Fig. 1 Top: A filtered simplicial complex. Bottom: A zigzag filtration of simplicial complexes

of a simplex σ is $|\sigma| - 1$, where $|\cdot|$ denotes cardinality. Note that this definition is combinatorial rather than geometric. These abstract simplicial complexes are not necessarily embedded in a geometric space.

Homology In this paper we will use simplicial homology over a field (see Munkres [28] for an accessible introduction to algebraic topology). Thus, given a space X , the homology groups $H_i(X)$ are vector spaces for each i . Let $H_*(X)$ denote the collection of these homology groups for all i .

The star subscript denotes the homomorphism of homology groups induced by a map between spaces, i.e. $f : X \rightarrow Y$ induces $f_* : H_*(X) \rightarrow H_*(Y)$. We recall the functorial properties of the Homology operator, $H_*(\cdot)$. In particular, $(f \circ g)_* = f_* \circ g_*$ and $\text{id}_{X_*} = \text{id}_{H_*(X)}$, where id indicates the identity map.

Persistence Modules and Diagrams A **filtration** is a nested sequence of topological spaces: $X_1 \subseteq X_2 \subseteq \dots \subseteq X_n$. If the spaces are simplicial complexes (as with all the filtrations in this paper), then it is called a **filtered simplicial complex** (see Fig. 1, top).

A **persistence module** is a sequence of Homology groups connected by homomorphisms:

$$H_*(X_1) \rightarrow H_*(X_2) \rightarrow \dots \rightarrow H_*(X_n).$$

The homology functor turns a filtration with inclusion maps $X_1 \hookrightarrow X_2 \hookrightarrow \dots$ into a persistence module, but as we will see, this is not the only way to get one.

One can also consider **zigzag filtrations**, which allow the inclusions to go in both directions: $X_1 \subseteq X_2 \supseteq X_3 \subseteq \dots$. The resulting module is called a **zigzag module**.

$$H_*(X_1) \rightarrow H_*(X_2) \leftarrow H_*(X_3) \rightarrow \dots$$

The **persistence diagram** of a persistence module is a multiset of points in $(\mathbb{R} \cup \{\infty\})^2$. Each point of the diagram represents a topological feature. The x and y coordinates of the points are the *birth* and *death* times of the feature and correspond to the indices in the persistence module where that feature appears and disappears. Points far from the diagonal persisted for a long time, while those “non-persistent” points

near the diagonal may be considered *topological noise*. By convention, the persistence diagram also contains every point (x, x) of the diagonal with infinite multiplicity.

Given a filtration \mathcal{F} , we let $D\mathcal{F}$ denote the persistence diagram of the persistence module generated by \mathcal{F} . The persistence algorithm computes a persistence diagram from \mathcal{F} [33]. It is also known how to compute a persistence diagram when \mathcal{F} is a zigzag filtration [4,27].

Approximating Persistence Diagrams Given two filtrations \mathcal{F} and \mathcal{G} , we say that the persistence diagram $D\mathcal{F}$ is a **c -approximation to the diagram $D\mathcal{G}$** if there is a bijection $\phi : D\mathcal{F} \rightarrow D\mathcal{G}$ such that for each $p \in D\mathcal{F}$, the birth times of p and $\phi(p)$ differ by at most a factor of c and the death times also differ by at most a factor of c . The reader familiar with stability results for persistent homology [7,15] will recognize this as bounding the ℓ_∞ -bottleneck distance between the persistence diagrams after reparameterizing the filtrations on a log-scale.

We will make use of two standard results on persistence diagrams. The first gives a sufficient condition for two persistence modules to yield identical persistence diagrams.

Theorem 1 (Persistence Equivalence Theorem [20], p. 159)

Consider two sequences of vector spaces connected by homomorphisms $\phi_i : U_i \rightarrow V_i$:

$$\begin{array}{ccccccc}
 V_0 & \longrightarrow & V_1 & \longrightarrow & \cdots & \longrightarrow & V_{n-1} & \longrightarrow & V_n \\
 \uparrow & & \uparrow & & & & \uparrow & & \uparrow \\
 U_0 & \longrightarrow & U_1 & \longrightarrow & \cdots & \longrightarrow & U_{n-1} & \longrightarrow & U_n
 \end{array}$$

If the vertical maps are isomorphisms and all squares commute then the persistence diagram defined by the U_i is the same as that defined by the V_i .

We prove approximation guarantees for persistence diagrams using the following lemma, which is a direct corollary of the Strong Stability Theorem of Chazal et al. [7] rephrased in the language of approximate persistence diagrams.

Lemma 1 (Persistence Approximation Lemma)

For any two filtrations $\mathcal{A} = \{A_\alpha\}_{\alpha \geq 0}$ and $\mathcal{B} = \{B_\alpha\}_{\alpha \geq 0}$, if $A_{\alpha/c} \subseteq B_\alpha \subseteq A_{c\alpha}$ for all $\alpha \geq 0$ then the persistence diagram $D\mathcal{A}$ is a c -approximation to the persistence diagram of $D\mathcal{B}$.

Contiguous Simplicial Maps Contiguity gives a discrete version of homotopy theory for simplicial complexes.

Definition 2 Let X and Y be simplicial complexes. A **simplicial map** $f : X \rightarrow Y$ is a function that maps vertices of X to vertices of Y and $f(\sigma) := \bigcup_{v \in \sigma} f(v)$ is a simplex of Y for all $\sigma \in X$.

A simplicial map is determined by its behavior on the vertex set. Consequently, we will abuse notation slightly and identify maps between vertex sets and maps between

simplices. When it is relevant and non-obvious, we will always prove that the resulting map between simplicial complexes is simplicial.

Definition 3 Two simplicial maps $f, g : X \rightarrow Y$ are **contiguous** if $f(\sigma) \cup g(\sigma) \in Y$ for all $\sigma \in X$.

Definition 4 For any pair of topological spaces $X \subseteq Y$, a map $f : Y \rightarrow X$ is a **retraction** if $f(x) = x$ for all $x \in X$. Equivalently, $f \circ i = \text{id}_X$ where $i : X \hookrightarrow Y$ is the inclusion map.

The theory of contiguity is a simplicial analogue of homotopy theory. If two simplicial maps are contiguous then they induce identical homomorphisms at the homology level [28, Sect. 12]. The following lemma gives a homology analogue of a deformation retraction.

Lemma 2 *Let X and Y be simplicial complexes such that $X \subseteq Y$ and let $i : X \hookrightarrow Y$ be the canonical inclusion map. If there exists a simplicial retraction $\pi : Y \rightarrow X$ such that $i \circ \pi$ and id_Y are contiguous, then i induces an isomorphism $i_\star : H_\star(X) \rightarrow H_\star(Y)$ between the corresponding homology groups.*

Proof Since $i \circ \pi$ and id_Y are contiguous, the induced homomorphisms $(i \circ \pi)_\star : H_\star(Y) \rightarrow H_\star(Y)$ and $\text{id}_{Y_\star} : H_\star(Y) \rightarrow H_\star(Y)$ are identical [28, Sect. 12]. Since $\text{id}_{Y_\star} = (i \circ \pi)_\star = i_\star \circ \pi_\star$ is an isomorphism, it follows that i_\star is surjective.

Since π is a retraction, $\pi \circ i = \text{id}_X$ and thus $(\pi \circ i)_\star : H_\star(X) \rightarrow H_\star(X)$ and $\text{id}_{X_\star} : H_\star(X) \rightarrow H_\star(X)$ are identical. Since $\text{id}_{X_\star} = (\pi \circ i)_\star = \pi_\star \circ i_\star$ is an isomorphism, it follows that i_\star is injective.

Thus, i_\star is an isomorphism because it is both injective and surjective. □

4 The Relaxed Vietoris–Rips Filtration

In this section, we relax the input metric so that it is no longer a metric, but it will still be provably close to the input. The new distance adds a small weight to each point that grows with α . The intuition behind this process is illustrated in Fig. 2. The weighted




Fig. 2 *Top:* Some points on a line. The *white point* contributes little to the union of α -balls. *Bottom:* Using the relaxed distance, the new α -ball is completely contained in the union of the other balls. Later, we use this property to prove that removing the *white point* will not change the topology




Fig. 3 The weight function for a point p . The weight is 0 until just before its removal time t_p . Then there is a period of steeper increase (slope = $1/2$) followed by slower increase (slope = ϵ)

distance effectively shrinks the metric balls locally so that one ball may be covered by nearby balls.

Throughout, we assume the user-defined parameter $\epsilon \leq \frac{1}{3}$ is fixed. Each point p is assigned a **deletion time** $t_p \in \mathbb{R}_{\geq 0}$. The specific choice of t_p will come from the net-tree construction in Sect. 6. For now, we will assume the deletion times are given, assuming only that they are nonnegative. The weight $w_p(\alpha)$ of point p at scale α is defined as follows (Fig. 3)

$$w_p(\alpha) := \begin{cases} 0 & \text{if } \alpha \leq (1 - 2\epsilon)t_p, \\ \frac{1}{2}(\alpha - (1 - 2\epsilon)t_p) & \text{if } (1 - 2\epsilon)t_p < \alpha < t_p, \\ \epsilon\alpha & \text{if } t_p \leq \alpha. \end{cases}$$

The **relaxed distance** at scale α is defined as

$$\hat{\mathbf{d}}_\alpha(p, q) := \mathbf{d}(p, q) + w_p(\alpha) + w_q(\alpha).$$

For any pair $p, q \in P$, the relaxed distance $\hat{\mathbf{d}}_\alpha(p, q)$ is monotonically non-decreasing in α . In particular, $\hat{\mathbf{d}}_\alpha \geq \hat{\mathbf{d}}_0 = \mathbf{d}$ for all $\alpha \geq 0$. Although distances can grow as α grows, this growth is sufficiently slow to allow the following lemma which will be useful later.

Lemma 3 *If $\hat{\mathbf{d}}_\alpha(p, q) \leq \alpha \leq \beta$ then $\hat{\mathbf{d}}_\beta(p, q) \leq \beta$.*

Proof The weight of a point is $\frac{1}{2}$ -Lipschitz in α , so $w_p(\beta) \leq w_p(\alpha) + \frac{1}{2}|\beta - \alpha|$, and similarly, $w_q(\beta) \leq w_q(\alpha) + \frac{1}{2}|\beta - \alpha|$. So,

$$\begin{aligned} \hat{\mathbf{d}}_\beta(p, q) &= \mathbf{d}(p, q) + w_p(\beta) + w_q(\beta) \\ &\leq \mathbf{d}(p, q) + w_p(\alpha) + w_q(\alpha) + (\beta - \alpha) \\ &= \hat{\mathbf{d}}_\alpha(p, q) + \beta - \alpha \\ &\leq \beta \end{aligned}$$

Given a set P , a distance function $\mathbf{d}' : P \times P \rightarrow \mathbb{R}$, and a scale parameter $\alpha \in \mathbb{R}$, we can construct a **Vietoris–Rips complex**

$$\text{VR}(P, \mathbf{d}', \alpha) := \{\sigma \subset P : \mathbf{d}'(p, q) \leq \alpha \text{ for all } p, q \in \sigma\}.$$

The Vietoris–Rips complex associated with the input metric space (P, \mathbf{d}) is $\mathcal{R}_\alpha := \text{VR}(P, \mathbf{d}, \alpha)$. The **relaxed Vietoris–Rips complex** is $\hat{\mathcal{R}}_\alpha := \text{VR}(P, \hat{\mathbf{d}}_\alpha, \alpha)$.

By considering the family of Vietoris–Rips complexes for all values of $\alpha \geq 0$, we get the **Vietoris–Rips filtration**, $\mathcal{R} := \{\mathcal{R}_\alpha\}_{\alpha \geq 0}$. Similarly, we may define the

relaxed Vietoris–Rips filtration, $\hat{\mathcal{R}} := \{\hat{\mathcal{R}}_\alpha\}_{\alpha \geq 0}$. Lemma 3 implies that $\hat{\mathcal{R}}$ is indeed a filtration. The filtrations \mathcal{R} and $\hat{\mathcal{R}}$ are very similar. The following lemma makes this similarity precise via a multiplicative interleaving.

Lemma 4 For all $\alpha \geq 0$, $\mathcal{R}_{\frac{\alpha}{c}} \subseteq \hat{\mathcal{R}}_\alpha \subseteq \mathcal{R}_\alpha$, where $c = \frac{1}{1-2\varepsilon}$.

Proof To prove inclusions between Vietoris–Rips complexes, it suffices to prove inclusion of the edge sets. For the first inclusion, we must prove that for any pair p, q , if $\mathbf{d}(p, q) \leq \frac{\alpha}{c}$ then $\hat{\mathbf{d}}_\alpha(p, q) \leq \alpha$. Fix any such pair p, q . By definition, $w_p(\alpha) \leq \varepsilon\alpha$ and $w_q(\alpha) \leq \varepsilon\alpha$. So,

$$\hat{\mathbf{d}}_\alpha(p, q) = \mathbf{d}(p, q) + w_p(\alpha) + w_q(\alpha) \leq \frac{\alpha}{c} + 2\varepsilon\alpha = \alpha.$$

For the second inclusion, $\hat{\mathbf{d}}_\alpha \geq \mathbf{d}$. So, if $\hat{\mathbf{d}}_\alpha(p, q) \leq \alpha$ then $\mathbf{d}(p, q) \leq \alpha$ as well. Thus any edge of $\hat{\mathcal{R}}_\alpha$ is also an edge of \mathcal{R}_α .

5 The Sparse Zigzag Vietoris–Rips Filtration

We will construct a sparse subcomplex of the relaxed Vietoris–Rips complex $\hat{\mathcal{R}}_\alpha$ that is guaranteed to have linear size for any α . In fact, we will get a zigzag filtration that only has a linear total number of simplices, yet its persistence diagram is identical to that of the relaxed Vietoris–Rips filtration.

We define the **open net** \mathcal{N}_α at scale α to be the subset of P with deletion time greater than α , i.e.

$$\mathcal{N}_\alpha := \{p \in P : t_p > \alpha\}.$$

Similarly, the **closed net** at scale α is

$$\bar{\mathcal{N}}_\alpha := \{p \in P : t_p \geq \alpha\}.$$

The **sparse zigzag Vietoris–Rips complex** \mathcal{Q}_α at scale α is just the subcomplex of $\hat{\mathcal{R}}_\alpha$ induced on the vertices of \mathcal{N}_α . Formally,

$$\mathcal{Q}_\alpha := \{\sigma \in \hat{\mathcal{R}}_\alpha : \sigma \subseteq \mathcal{N}_\alpha\} = \text{VR}(\mathcal{N}_\alpha, \hat{\mathbf{d}}_\alpha, \alpha).$$

We also define a closed version of the sparse zigzag Vietoris–Rips complex:

$$\bar{\mathcal{Q}}_\alpha := \text{VR}(\bar{\mathcal{N}}_\alpha, \hat{\mathbf{d}}_\alpha, \alpha).$$

Note that if $\alpha \neq t_p$ for all $p \in P$ then $\mathcal{N}_\alpha = \bar{\mathcal{N}}_\alpha$ and $\mathcal{Q}_\alpha = \bar{\mathcal{Q}}_\alpha$.

The complexes \mathcal{R}_α , $\hat{\mathcal{R}}_\alpha$, \mathcal{Q}_α , and $\bar{\mathcal{Q}}_\alpha$ are well-defined for all $\alpha \geq 0$, however, they only change at discrete scales. Let $A = \{a_i\}_{i \in \mathbb{N}}$ be an ordered, discrete set of nonnegative real numbers such that $t_p \in A$ for all $p \in P$ and $\alpha \in A$ for any pair p, q such that $\alpha = \hat{\mathbf{d}}_\alpha(p, q)$. That is, A contains every scale at which a combinatorial

changes happens, either a point deletion or an edge insertion. This implies that $\mathcal{N}_{a_i} = \overline{\mathcal{N}}_{a_{i+1}}$ and thus, using Lemma 3, that $\mathcal{Q}_{a_i} \subseteq \overline{\mathcal{Q}}_{a_{i+1}}$.

The sparse Vietoris–Rips complexes can be arranged into a zigzag filtration \mathcal{Q} as follows.

$$\overline{\mathcal{Q}}_{a_1} \leftrightarrow \mathcal{Q}_{a_1} \leftrightarrow \overline{\mathcal{Q}}_{a_2} \leftrightarrow \mathcal{Q}_{a_2} \leftrightarrow \dots$$

We will return to \mathcal{Q} later as it has some interesting properties. However, at this point, it is underspecified as we have not yet shown how to compute the deletion times for the vertices. The next section will fill this gap.

6 Hierarchical Net-Trees

The following treatment of net-trees is adapted from the paper by Har-Peled and Mendel [25].

Definition 5 A **net-tree** of a metric $\mathcal{M} = (P, \mathbf{d})$ is a rooted tree T with vertex $v \in T$ having a **representative point** $\text{rep}(v) \in P$. There are $n = |P|$ leaves, each represented by a different point of P . Each non-root vertex $v \in T$ has a unique **parent** $\text{par}(v)$. The set of vertices with the same parent v are called the **children** of v , denoted $\text{child}(v)$. If $\text{child}(v)$ is nonempty then for some $u \in \text{child}(v)$, $\text{rep}(u) = \text{rep}(v)$. The set $P_v \subseteq P$ denotes the points represented by the leaves of the subtree rooted at v . Each vertex $v \in T$ has an associated radius $\text{rad}(v)$ satisfying the following two conditions.

1. **Covering Condition:** $P_v \subset \text{ball}(\text{rep}(v), \text{rad}(v))$, and
2. **Packing Condition:** if v is not the root, then

$$P \cap \text{ball}(\text{rep}(v), K_p \text{rad}(\text{par}(v))) \subseteq P_v,$$

where K_p (the “ p ” is for “packing”) is a constant independent of \mathcal{M} and n (Fig. 4).

The radii of the net tree nodes are always some constant times larger than the radius of their children. Simple packing arguments guarantee that no node of the tree has more than $\lambda^{O(1)}$ children, where λ is the doubling constant of the metric. The whole tree can be constructed in $O(n \log n)$ randomized time or in $O(n \log \Delta)$ time deterministically [25]. Moreover, it is important to note the construction does not require that we know the doubling dimension in advance.

Given a net-tree T for $\mathcal{M} = (P, \mathbf{d})$ and a point $p \in P$, let v_p denote the least ancestor among the nodes in T represented by p . For each $p \in P$ the deletion time t_p is defined as

$$t_p := \frac{1}{\varepsilon(1 - 2\varepsilon)} \text{rad}(\text{par}(v_p)).$$

This is just the radius of the parent of v_p with a small scaling factor included for technical reasons. When the scale α reaches t_p , we remove point p from the (zigzag) filtration. The choice of weights as a function of t_p guarantees that any point with




Fig. 4 A net-tree is built over the set points from Fig. 2. Each level of the tree represents a sparse approximation to the original point set at a different scale




Fig. 5 Three levels of a net tree for a collection of points in the plane. In each level, the larger disks cover the rest of the points, whereas the smaller disks are disjoint, i.e. they pack

relaxed distance at most t_p from p will also have relaxed distance at most t_p from $\text{rep}(\text{par}(v_p))$. As we prove later, this guarantees that the topology of the Rips complex does not change when we remove p at scale t_p .

For a fixed scale $\alpha \in \mathbb{R}$, the set \mathcal{N}_α is a subset of points of P induced by the net-tree. The sets \mathcal{N}_α are the **nets** of the net-tree. For any α it satisfies a packing condition and a covering condition as defined in the following lemma.

Lemma 5 *Let $\mathcal{M} = (P, \mathbf{d})$ be a metric space and let T be a net-tree for \mathcal{M} . For all $\alpha \geq 0$, the net \mathcal{N}_α induced by T at scale α satisfies the following two conditions.*

1. **Covering Condition:** For all $p \in P$, $\mathbf{d}(p, \mathcal{N}_\alpha) \leq \varepsilon(1 - 2\varepsilon)\alpha$.
2. **Packing Condition:** For all distinct $p, q \in \mathcal{N}_\alpha$, $\mathbf{d}(p, q) \geq K_p\varepsilon(1 - 2\varepsilon)\alpha$.

Proof First, we prove that the covering condition holds. Fix any $p \in P$. The statement is trivial if $p \in \mathcal{N}_\alpha$ so we may assume that $t_p \leq \alpha$.

Let v be the lowest ancestor of p in T such that $\text{rep}(v) \in \mathcal{N}_\alpha$. Let u be the ancestor of p among the children of v . If $q = \text{rep}(u)$ then $t_q = \frac{\text{rad}(u)}{\varepsilon(1-2\varepsilon)}$. By our choice of v , $q \notin \mathcal{N}_\alpha$ and thus $t_q \leq \alpha$. It follows that $\text{rad}(v) \leq \varepsilon(1 - 2\varepsilon)\alpha$. Thus, $\mathbf{d}(p, \mathcal{N}_\alpha) \leq \mathbf{d}(p, \text{rep}(v)) \leq \text{rad}(v) \leq \varepsilon(1 - 2\varepsilon)\alpha$.

We now prove that the packing condition holds. Let p, q be any two distinct points of \mathcal{N}_α . Without loss of generality, assume $t_p \leq t_q$. Thus, $q \notin P_{v_p}$, where (as before) v_p is the least ancestor among the nodes of T represented by p . Since $p \in \mathcal{N}_\alpha$, $\alpha < t_p = \frac{\text{rad}(\text{par}(v_p))}{\varepsilon(1-2\varepsilon)}$. Therefore, using the packing condition on the net-tree T , $\mathbf{d}(p, q) \geq K_p \text{rad}(\text{par}(v_p)) > K_p\varepsilon(1 - 2\varepsilon)\alpha$.

A subset that satisfies this type of packing and covering conditions is sometimes referred to as a metric space net (not to be confused with a range space net) or, more accurately, as a Delone set [14]. An example is given in Fig. 5.

7 Topology-Preserving Sparsification

In this section, we make the intuition of Fig. 2 concrete by showing that deleting a vertex p (and its incident simplices) from the relaxed Vietoris–Rips complex $\hat{\mathcal{R}}_{t_p}$ does not change the topology.

For any $\alpha \geq 0$, we define the “projection” of P onto \mathcal{N}_α as

$$\pi_\alpha(p) := \begin{cases} p & \text{if } p \in \mathcal{N}_\alpha, \\ \operatorname{argmin}_{q \in \mathcal{N}_\alpha} \hat{\mathbf{d}}_\alpha(p, q) & \text{otherwise.} \end{cases}$$

The following lemma shows that the distance from a point to its projection is bounded by the difference in the weights of the point and its projection.

Lemma 6 *For all $p \in P$, $\mathbf{d}(p, \pi_\alpha(p)) \leq w_p(\alpha) - w_{\pi_\alpha(p)}(\alpha)$.*

Proof Fix any $p \in P$. We first prove that if $\mathbf{d}(p, q) \leq w_p(\alpha) - w_q(\alpha)$ for some $q \in \mathcal{N}_\alpha$, then it holds for $q = \pi_\alpha(p)$.

If we have such a q , then the definitions of $\hat{\mathbf{d}}_\alpha$ and π_α imply the following.

$$\begin{aligned} \mathbf{d}(p, \pi_\alpha(p)) &= \hat{\mathbf{d}}_\alpha(p, \pi_\alpha(p)) - w_p(\alpha) - w_{\pi_\alpha(p)}(\alpha) \\ &\leq \hat{\mathbf{d}}_\alpha(p, q) - w_p(\alpha) - w_{\pi_\alpha(p)}(\alpha) \\ &\leq \mathbf{d}(p, q) + w_q(\alpha) - w_{\pi_\alpha(p)}(\alpha) \\ &\leq w_p(\alpha) - w_{\pi_\alpha(p)}(\alpha). \end{aligned}$$

So, it will suffice to find a $q \in \mathcal{N}_\alpha$ such that $\mathbf{d}(p, q) \leq w_p(\alpha) - w_q(\alpha)$. If $p \in \mathcal{N}_\alpha$ then this is trivial. So we may assume $p \notin \mathcal{N}_\alpha$ and therefore $t_p \leq \alpha$ and $w_p(\alpha) = \varepsilon\alpha$.

Let $u \in T$ be the ancestor of p such that $\operatorname{rad}(u) < \frac{\varepsilon\alpha}{1-\varepsilon}$ and $\operatorname{rad}(\operatorname{par}(u)) \geq \frac{\varepsilon\alpha}{1-\varepsilon}$. Let $q = \operatorname{rep}(u)$. Since

$$t_q \geq \frac{1}{\varepsilon(1-2\varepsilon)} \operatorname{rad}(\operatorname{par}(u)) \geq \frac{\alpha}{1-2\varepsilon},$$

it follows that $w_q(\alpha) = 0$ and that $q \in \mathcal{N}_\alpha$. Finally, since $p \in P_u$, $\mathbf{d}(p, q) \leq \operatorname{rad}(u) \leq \varepsilon\alpha = w_p(\alpha) - w_q(\alpha)$.

By bounding the distance between points and their projections, we can now show that distances in the projection do not grow.

Lemma 7 *For all $p, q \in P$ and all $\alpha \geq 0$, $\hat{\mathbf{d}}_\alpha(\pi_\alpha(p), q) \leq \hat{\mathbf{d}}_\alpha(p, q)$.*

Proof The bound follows from the definition of $\hat{\mathbf{d}}_\alpha$, the triangle inequality, and Lemma 6.

Lemma 8 *Let $\alpha \geq 0$ be a fixed constant. Let X be a set of points such that $\mathcal{N}_\alpha \subseteq X \subseteq P$ and let $K = \operatorname{VR}(X, \hat{\mathbf{d}}_\alpha, \alpha)$. The inclusion map $i : \mathcal{Q}_\alpha \hookrightarrow K$ induces an isomorphism at the homology level.*

Proof The map $K \rightarrow Q_\alpha$ induced by π_α is a retraction because $Q_\alpha \subseteq K$ and π_α is a retraction onto \mathcal{N}_α , the vertex set of Q_α . By Lemma 2, it will suffice to prove that π_α is simplicial and that $i \circ \pi_\alpha$ is contiguous to the identity map on K .

Since Q_α and K are Vietoris–Rips complexes, it will suffice to prove these facts for the edges:

1. π_α is simplicial: for all $p, q \in X$, if $\hat{\mathbf{d}}_\alpha(p, q) \leq \alpha$ then $\hat{\mathbf{d}}_\alpha(\pi_\alpha(p), \pi_\alpha(q)) \leq \alpha$, and
2. $i \circ \pi_\alpha$ and id_K are contiguous: for all $p, q \in X$, if $\hat{\mathbf{d}}_\alpha(p, q) \leq \alpha$ then all six edges of the tetrahedron $\{p, q, \pi_\alpha(p), \pi_\alpha(q)\}$ are in K .

The first statement follows from two successive applications of Lemma 7. The second statement follows from Lemma 6 for the edges $\{p, \pi_\alpha(p)\}$ and $\{q, \pi_\alpha(q)\}$ and from Lemma 7 for the other edges.

Corollary 1 *For all $\alpha \in A$, the inclusions $f : Q_\alpha \hookrightarrow \overline{Q}_\alpha$, $g : Q_\alpha \hookrightarrow \hat{\mathcal{R}}_\alpha$, and $h : \overline{Q}_\alpha \hookrightarrow \hat{\mathcal{R}}_\alpha$ induce isomorphisms at the homology level.*

Proof The inclusions f and g induce isomorphisms by applying Lemma 8 with $X = \overline{\mathcal{N}}_\alpha$ and $X = P$ respectively. Composing the inclusions, we get that $g = h \circ f$. Thus, at the homology level, we get $h_\star = g_\star \circ f_\star^{-1}$ is also an isomorphism.

8 Straightening Out the Zigzags

In this section, we show two different ways in which zigzag persistence may be avoided. First, in Sect. 8.1, we show that the sparse zigzag filtration \mathcal{Q} does not zigzag at the homology level. Then, in Sect. 8.2, we show how to modify the zigzag filtration so it does not zigzag as a filtration either.

The advantage of the non-zigzagging filtration is that it allows one to use the standard persistence algorithm, but it has larger size in the intermediate complexes. As we will see in Sect. 9.2, the total size is still linear.

8.1 Reversing Homology Isomorphisms

The backwards arrows in the zigzag filtration \mathcal{Q} all induce isomorphisms. At the homology level, these isomorphisms can be replaced by their inverses to give a persistence module that does not zigzag. That is, the zigzag module

$$\dots \rightarrow H_*(\overline{Q}_\alpha) \xleftarrow{\cong} H_*(Q_\alpha) \rightarrow H_*(\overline{Q}_\beta) \xleftarrow{\cong} H_*(Q_\beta) \rightarrow \dots$$

can be transformed into

$$\dots \rightarrow H_*(\overline{Q}_\alpha) \xrightarrow{\cong} H_*(Q_\alpha) \rightarrow H_*(\overline{Q}_\beta) \xrightarrow{\cong} H_*(Q_\beta) \rightarrow \dots$$

The latter module implies the existence of another that only uses the closed sparse Vietoris–Rips complexes:

$$\dots \rightarrow H_*(\overline{Q}_\alpha) \rightarrow H_*(\overline{Q}_\beta) \rightarrow \dots$$

Note that this module does not duplicate the indices in the zigzag. In these various transformations, we have only reversed or concatenated isomorphisms, thus we have not changed the rank of any induced map $H_*(\overline{Q}_\alpha) \rightarrow H_*(\overline{Q}_\beta)$. As a result the persistence diagram DQ is unchanged.

This is novel in that we construct a zigzag filtration and we apply the zigzag persistence algorithm, but we are really computing the diagram of a persistence module that does not zigzag. The zigzagging can then be interpreted as sparsifying the complex without changing the topology.

8.2 A Sparse Filtration Without the Zigzag

The preceding section showed that the sparse zigzag Vietoris–Rips filtration does not zigzag as a persistence module. This hints that it is possible to construct a filtration that does not zigzag with the same persistence diagram. Indeed, this is possible using the filtration

$$S := \{S_{a_k}\}_{a_k \in A}, \text{ where } S_{a_k} := \bigcup_{i=1}^k \overline{Q}_{a_i}.$$

We first prove that $H_*(\overline{Q}_{a_k})$ and $H_*(S_{a_k})$ are isomorphic.

Lemma 9 *For all $a_k \in A$, the inclusion $h : \overline{Q}_{a_k} \hookrightarrow S_{a_k}$ induces a homology isomorphism.*

Proof We define some intermediate complexes that interpolate between \overline{Q}_{a_k} and S_{a_k} .

$$T_{i,k} := \bigcup_{j=i}^k \overline{Q}_{a_j}.$$

In particular, we have that $\overline{Q}_{a_k} = T_{k,k}$ and $S_{a_k} = T_{1,k}$. The map h can be expressed as $h = h_1 \circ \dots \circ h_{k-1}$, where $h_i : T_{i+1,k} \hookrightarrow T_{i,k}$ is an inclusion. It will suffice to prove that each h_i induces an isomorphism at the homology level for each $i = 1 \dots k - 1$. By Lemma 2, it will suffice to show that the projection $\pi_{a_i} : T_{i,k} \rightarrow T_{i+1,k}$ is a simplicial retraction and $h_i \circ \pi_{a_i}$ and $\text{id}_{T_{i,k}}$ are contiguous.

Let $\sigma \in T_{i,k}$ be any simplex. So, $\sigma \in \overline{Q}_{a_j}$ for some integer j such that $i \leq j \leq k$.

First, we prove that π_{a_i} is a retraction. If $\sigma \in T_{i+1,k}$ then $j \geq i + 1$. So, $\sigma \subseteq \overline{N}_{a_j} \subseteq \mathcal{N}_{a_i}$ and thus $\pi_{a_i}(\sigma) = \sigma$ because π_{a_i} is a retraction onto \mathcal{N}_{a_i} by definition when viewed as a function on the vertex sets.

Second, we show that π_{a_i} is a simplicial map from $T_{i,k}$ to $T_{i+1,k}$. Since it is a retraction, it only remains to show that $\pi_{a_i}(\sigma) \in T_{i,k}$ when $\sigma \in T_{i,k} \setminus T_{i+1,k}$, i.e. when $j = i$. In this case, $\pi_{a_i}(\sigma) \in \overline{Q}_{a_i}$ because $\pi_{a_i} : \overline{Q}_{a_i} \rightarrow \overline{Q}_{a_i}$ is simplicial (as shown in the proof of Lemma 8). Since $\overline{Q}_{a_i} \subseteq \overline{Q}_{a_{i+1}} \subseteq T_{i,k}$, it follows that $\pi_{a_i}(\sigma) \in T_{i,k}$ as desired.

Last, we prove contiguity. We need to prove that $\sigma \cup \pi_{a_i}(\sigma) \in T_{i,k}$. If $j > i$, then $\sigma \cup \pi_{a_i}(\sigma) = \sigma \in T_{i,k}$ as desired. If $i = j$, then $\sigma \cup \pi_{a_i}(\sigma) \in \overline{Q}_{a_i}$ as shown in the proof of Lemma 8. Since $\overline{Q}_{a_i} \subseteq T_{i,k}$, it follows that $\sigma \cup \pi_{a_i}(\sigma) \in T_{i,k}$ as desired.

Theorem 2 *The persistence diagrams of \mathcal{Q} and \mathcal{S} are identical.*

Proof For any $a_i, a_{i+1} \in A$, we get the following commutative diagram where all maps are induced by inclusions.

$$\begin{array}{ccccc}
 H_*(\overline{Q}_{a_i}) & \xleftarrow{\cong} & H_*(Q_{a_i}) & \xrightarrow{\quad} & H_*(\overline{Q}_{a_{i+1}}) \\
 \cong \searrow & & \swarrow & & \cong \swarrow \\
 & & H_*(S_{a_i}) & \xrightarrow{\quad} & H_*(S_{a_{i+1}})
 \end{array}$$

Lemma 9 and Corollary 1 show that the indicated maps are isomorphisms. As in Sect. 8.1, we reverse the isomorphism $H_*(Q_{a_i}) \rightarrow H_*(\overline{Q}_{a_i})$ to get the following diagrams, which also commutes.

$$\begin{array}{ccc}
 H_*(\overline{Q}_{a_i}) & \xrightarrow{\quad} & H_*(\overline{Q}_{a_{i+1}}) \\
 \cong \downarrow & & \cong \downarrow \\
 H_*(S_{a_i}) & \xrightarrow{\quad} & H_*(S_{a_{i+1}})
 \end{array}$$

Therefore, by the Persistence Equivalence Theorem, $D\mathcal{Q} = D\mathcal{S}$.

8.3 The Connection with Extended Persistence

The sparse Rips zigzag has the property that every other space is the intersection of its neighbors on either side. That is, by a simple exercise, one can show that $Q_{a_i} = \overline{Q}_{a_i} \cap \overline{Q}_{a_{i+1}}$. Carlsson et al. give a general method for comparing such zigzags to filtrations that do not zigzag in their work on levelset zigzags induced by real-valued functions [5]. They proved that the Mayer-Vietoris diamond principle from a paper by Carlsson and De Silva [4] allows one to relate the persistence diagram of such a zigzag with the so-called extended persistence diagram of the union filtration. In our case, this result implies that the persistence diagram of the sparse zigzag Vietoris–Rips filtration can be derived from the persistence diagram of the extended filtration

$$\begin{array}{ccccccc}
 H_*(T_{1,1}) & \xrightarrow{\quad} & \cdots & \xrightarrow{\quad} & H_*(T_{1,N}) & & \\
 & & & & \downarrow & & \\
 H_*(T_{1,N}, T_{N,N}) & \xleftarrow{\quad} & \cdots & \xleftarrow{\quad} & H_*(T_{N,N}, T_{1,N}) & &
 \end{array}$$

where $N = |A|$, $T_{i,k} := \bigcup_{j=i}^k \overline{Q}_{a_j}$ as in the proof of Lemma 9, and $H_*(T_{N,N}, T_{i,N})$ is the homology of $T_{N,N}$ relative to $T_{i,N}$, or equivalently, the homology of the quotient

$T_{N,N}/T_{i,N}$. The first half of this filtration is precisely the sparse Vietoris–Rips filtration \mathcal{S} .

In light of this result, we see that Lemma 9 implies that there is nothing interesting happening in the second half of the extended filtration. In the language of extended persistence, this means that there are no extended or relative pairs. Thus, as shown in Theorem 2, there is no need to compute the extended persistence to compute the persistence of the sparse Vietoris–Rips filtration.

9 Theoretical Guarantees

There are two main theoretical guarantees regarding the sparse Vietoris–Rips filtrations. First, in Sect. 9.1, we show that the resulting persistence diagrams are good approximations to the true Vietoris–Rips filtration. Second, in Sect. 9.2, we show that the filtrations have linear size.

9.1 The Approximation Guarantee

In this section we prove that the persistence diagram of the sparse Vietoris–Rips filtration is a multiplicative c -approximation to the persistence diagram of the standard Vietoris–Rips filtration, where $c = \frac{1}{1-2\varepsilon}$. The approach has two parts. First, we show that the relaxed filtration is a multiplicative c -approximation to the classical Vietoris–Rips filtration. Second, we show that the sparse and relaxed Vietoris–Rips filtrations have the same persistence diagrams, i.e. that $D\mathcal{Q} = D\hat{\mathcal{R}}$. By passing through the filtration $\hat{\mathcal{R}}$, we obviate the need to develop new stability results for zigzag persistence.

Theorem 3 *For any metric space $\mathcal{M} = (P, \mathbf{d})$, the persistence diagrams of the corresponding sparse Vietoris–Rips filtrations $\mathcal{Q} = \mathcal{Q}(\mathcal{M})$ and $\mathcal{S} = \mathcal{S}(\mathcal{M})$ both yield c -approximations to the persistence diagram of the Vietoris–Rips filtration $\mathcal{R} = \mathcal{R}(\mathcal{M})$, where $c = \frac{1}{1-2\varepsilon}$ and $\varepsilon \leq \frac{1}{3}$ is a user-defined constant.*

Proof By Lemma 4, we have a multiplicative c -interleaving between \mathcal{R} and $\hat{\mathcal{R}}$. Thus, the Persistence Approximation Lemma implies that $D\hat{\mathcal{R}}$ is a c -approximation to $D\mathcal{R}$.

We have shown in Theorem 2 that $D\mathcal{Q} = D\mathcal{S}$, so it will suffice to prove that $D\mathcal{Q} = D\hat{\mathcal{R}}$. The rest of the proof follows the same pattern as in Theorem 2. For any $a_i, a_{i+1} \in A$, we get the following commutative diagram induced by inclusion maps.

$$\begin{array}{ccccc}
 H_*(\bar{\mathcal{Q}}_{a_i}) & \xleftarrow{\cong} & H_*(\mathcal{Q}_{a_i}) & \xrightarrow{\quad} & H_*(\bar{\mathcal{Q}}_{a_{i+1}}) \\
 & \searrow \cong & \swarrow \cong & & \swarrow \cong \\
 & & H_*(\hat{\mathcal{R}}_{a_i}) & \xrightarrow{\quad} & H_*(\hat{\mathcal{R}}_{a_{i+1}})
 \end{array}$$

Corollary 1 implies that many of these inclusions induce isomorphisms at the homology level (as indicated in the diagram). As a consequence, the following diagram also

commutes and the vertical maps are isomorphisms.

$$\begin{array}{ccc}
 H_*(\overline{Q}_{a_i}) & \longrightarrow & H_*(\overline{Q}_{a_{i+1}}) \\
 \cong \downarrow & & \cong \downarrow \\
 H_*(\hat{R}_{a_i}) & \longrightarrow & H_*(\hat{R}_{a_{i+1}})
 \end{array}$$

So, the Persistence Equivalence Theorem implies that $DQ = D\hat{R}$ as desired.

9.2 The Linear Complexity of the Sparse Filtration

In this section, we prove that the total number of simplices in the sparse Vietoris–Rips filtration is only linear in the number of input points. We start by showing that the graph of all edges appearing in the filtration has only a linear number of edges.

For a point $p \in P$, let $E(p)$ be the set of neighbors of p whose removal time is at least as large as that of p :

$$E(p) := \{q \in P : t_p \leq t_q \text{ and } (p, q) \in \overline{Q}_{t_p}\}.$$

To compute the filtrations Q and S , it suffices to compute $E(p)$ for each $p \in P$. In fact S_∞ is just the clique complex on the graph of all edges (p, q) such that $q \in E(p)$.

Lemma 10 *Given a set of n points in a metric space $\mathcal{M} = (P, \mathbf{d})$ with doubling dimension at most d and a net-tree with parameter $\varepsilon \leq \frac{1}{3}$, the cardinality $|E(p)|$ is at most $\frac{1}{\varepsilon} O(d)$ for each $p \in P$.*

Proof Let $\Delta(E(p))$ denote the spread of $E(p)$. Since $E(p)$ is a finite metric with doubling dimension at most d , the number of points is at most $\Delta(E(p))^{O(d)}$. So, it will suffice to prove that for all $p \in P$, $\Delta(E(p)) = O(\frac{1}{\varepsilon})$.

The definition of $E(p)$ implies that $E(p) \subseteq \mathcal{N}_{t_p}$ and so by Lemma 5, the nearest pair in $E(p)$ are at least $K_p\varepsilon(1 - 2\varepsilon)t_p$ apart. For $q \in E(p)$, since $(p, q) \in \overline{Q}_{t_p}$, $\mathbf{d}(p, q) \leq \hat{\mathbf{d}}_{t_p}(p, q) \leq t_p$. It follows that the farthest pair in $E(p)$ are at most $2t_p$ apart. So, we get that $\Delta(E(p)) \leq \frac{2t_p}{K_p\varepsilon(1-2\varepsilon)t_p} = O(\frac{1}{\varepsilon})$ as desired.

We see that the size of the graph in the filtration is governed by three variables: the doubling dimension, d ; the packing constant of the net-tree, K_p ; and the desired tightness of the approximation, ε . The preceding Lemma easily implies the following bound on the higher order simplices.

Theorem 4 *Given a set of n points in a metric space $\mathcal{M} = (P, \mathbf{d})$ with doubling dimension d , the total number of k -simplices in the sparse Vietoris–Rips filtrations Q and S is at most $(\frac{1}{\varepsilon})^{O(kd)} n$.*

10 An Algorithm to Construct the Sparse Filtration

The net-tree defines the deletion times of the input points and thus determines the perturbed metric. It also gives the necessary data structure to efficiently find the neighbors of a point in the perturbed metric in order to compute the filtration. In fact, this is exactly the kind of search that the net-tree makes easy. Then we find all cliques, which takes linear time because each is subset of $E(p)$ for some $p \in P$ and each $E(p)$ has constant size.

As explained in the Har-Peled and Mendel paper [25], it is often useful to augment the net-tree with “cross” edges connecting nodes at the same level in the tree that are represented by geometrically close points. The set of *relatives* of a node $u \in T$ is defined as

$$\text{Rel}(u) := \{v \in T : \text{rad}(v) \leq \text{rad}(u) < \text{rad}(\text{par}(v)) \text{ and} \\ \mathbf{d}(\text{rep}(u), \text{rep}(v)) \leq C \text{rad}(u)\},$$

where C is a constant bigger than 3.¹ The size of $\text{Rel}(u)$ is a constant using the same packing arguments as in Lemma 10.

This makes it easy to do a range search to find the points of $E(p)$. In fact, we will find the slightly larger set $E'(p) = \overline{\mathcal{N}}_{t_p} \cap \mathbf{ball}(p, t_p)$. The search starts by finding u the highest ancestor of v_p whose radius is at most some fixed constant times t_p . Since the radius increases by a constant factor on each level, this is only a constant number of levels. Then the subtrees rooted at each $v \in \text{Rel}(u)$ are searched down to the level of v_p . Thus, we search a constant number of trees of constant degree down a constant number of levels. The resulting search finds all of the points of $E(p)$ in constant time (see Fig. 6).

Since the work is only constant time per point, the only superlinear work is in the computation of the net-tree. As noted before, this requires only $O(n \log n)$ time.

11 Conclusions and Directions for Future Work

We have presented an efficient method for approximating the persistent homology of the Vietoris–Rips filtration. Computing these approximate persistence diagrams at all scales has the potential to make persistence-based methods on metric spaces tractable for much larger inputs.

Adapting the proofs given in this paper to the Čech filtration is a simple exercise. Moreover, it may be possible to apply a similar sparsification to complexes filtered by alternative distance-like functions like the distance to a measure introduced by Chazal et al. [9].

¹ The precise value of C depends on some constants chosen in the construction of the net-tree and can be extracted from the Har-Peled and Mendel paper. For our purposes, we only need the fact that it is bigger than 3.




Fig. 6 To find the nearby points at or above the level of v_p in the net-tree only requires traveling up a constant number of levels and then searching the relative trees

Another direction for future work is to identify a more general class of hierarchical structures that may be used in such a construction. The net-tree used in this paper is just one example chosen primarily because it can be computed efficiently.

The analytic technique used in this paper may find more uses in the future. We effectively bounded the difference between the persistence diagrams of a filtration and a zigzag filtration by embedding the zigzag filtration in a topologically equivalent filtration that does not zigzag at the homology level. This is very similar to the relationship between the levelset zigzag and extended persistence demonstrated by Carlsson et al. [5]. In that paper, such a technique gave some stability results for levelset zigzags of real-valued Morse functions on manifolds. It may be that other zigzag filtrations can be analyzed in this way.

Acknowledgments This work was partially supported by GIGA grant ANR-09-BLAN-0331-01 and the European project CG-Learning No. 255827. I would like to thank Steve Oudot and Frédéric Chazal for their helpful comments. I would also like to acknowledge the keen insight of David Cohen-Steiner who suggested that I look for a non-zigzagging filtration. This insight became the filtration S of Sect. 8.2.

References

1. Attali, D., Lieutier, A., Salinas, D.: Efficient data structure for representing and simplifying simplicial complexes in high dimension. In: Proceedings of the 27th ACM Symposium on Computational Geometry, pp. 501–509. Paris (2011)
2. Boissonat, J.-D., Guibas, L.J., Oudot, S.Y.: Manifold reconstruction in arbitrary dimensions using witness complexes. In: Proceedings of the 23rd ACM Symposium on Computational Geometry, pp. 194–203. Gyeongju (2007)
3. Carlsson, G.: Topology and data. *Bull. Am. Math. Soc.* **46**, 255–308 (2009)
4. Carlsson, G., de Silva, V.: Zigzag persistence. *Found. Comput. Math.* **10**(4), 367–405 (2010)
5. Carlsson, G., de Silva, V., Morozov, D.: Zigzag persistent homology and real-valued functions. In: Proceedings of the 25th ACM Symposium on Computational Geometry, pp. 247–256 (2009)
6. Carlsson, G., Ishkhanov, T., de Silva, V., Zomorodian, A.: On the local behavior of spaces of natural images. *Int. J. Comput. Vis.* **76**(1), 1–12 (2008)

7. Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J., Oudot, S.Y.: Proximity of persistence modules and their diagrams. In: Proceedings of the 25th ACM Symposium on Computational Geometry, pp. 237–246 (2009)
8. Chazal, F., Cohen-Steiner, D., Guibas, L.J., Mémoli, F., Oudot, S.Y.: Gromov-Hausdorff stable signatures for shapes using persistence. *Comput. Graph. Forum* **28**(5), 1393–1403 (2009)
9. Chazal, F., Cohen-Steiner, D., Mériqot, Q.: Geometric inference for probability measures. *Found. Comput. Math.* **11**, 733–751 (2011)
10. Chazal, F., Oudot, S.Y.: Towards persistence-based reconstruction in Euclidean spaces. In: Proceedings of the 24th ACM Symposium on Computational Geometry, pp. 232–241 (2008)
11. Chung, M.K., Bubenik, P., Kim, P.T.: Persistence diagrams of cortical surface data. *Inf. Process. Med. Imaging* **5636**, 386–397 (2009)
12. Clarkson, K.L.: Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.* **22**(1), 63–93 (1999)
13. Clarkson, K.L.: Nearest neighbor searching in metric spaces: experimental results for $sb(S)$. Preliminary version presented at ALENEX99, pp. 63–93 (2003)
14. Clarkson, K.L.: Building triangulations using epsilon-nets. In: STOC: ACM Symposium on Theory of Computing, pp. 326–335, New York (2006)
15. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. *Discrete Comput. Geom.* **37**(1), 103–120 (2007)
16. Cole, R., Gottlieb, L.-A.: Searching dynamic point sets in spaces with bounded doubling dimension. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, pp. 574–583. Seattle (2006)
17. de Silva, V., Carlsson, G.: Topological estimation using witness complexes. In: Proceedings of Symposium on Point-Based Graphics, pp. 157–166 (2004)
18. de Silva, V., Ghrist, R.: Coverage in sensor networks via persistent homology. *Algorithmic Geom. Topol.* **7**, 339–358 (2007)
19. de Silva, V., Ghrist, R.: Homological sensor networks. *Notices Am. Math. Soc.* **54**(1), 10–17 (2007)
20. Edelsbrunner, H., Harer, J.L.: *Computational Topology: An Introduction*. American Mathematical Society, Providence, RI (2009)
21. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. *Discret. Comput. Geom.* **4**(28), 511–533 (2002)
22. Gao, J., Guibas, L.J., An T.N.: Deformable spanners and applications. *Comput. Geom.* **35**(1–2), 2–19 (2006)
23. Gottlieb, L.-A., Roditty, L.: An optimal dynamic spanner for doubling metric spaces. In: Proceedings of the 16th Annual European Symposium on Algorithms, pp. 468–489 (2008)
24. Guibas, L.J., Oudot, S.Y.: Reconstruction using witness complexes. In: Proceedings 18th ACM-SIAM Symposium: Discrete Algorithms, pp. 1076–1085 (2007)
25. Har-Peled, S., Mendel, M.: Fast construction of nets in low dimensional metrics, and their applications. *SIAM J. Comput.* **35**(5), 1148–1184 (2006)
26. Hudson, Benoît, Miller, G.L., Oudot, S.Y., Sheehy, D.R.: Topological inference via meshing. In: Proceedings of the 26th ACM Symposium on Computational Geometry, pp. 277–286. Chapel Hill (2010)
27. Milosavljevic, N., Morozov, D., Skraba, P.: Zigzag persistent homology in matrix multiplication time. In: Proceedings of the 27th ACM Symposium on Computational Geometry, Paris (2011)
28. Munkres, J.R.: *Elements of Algebraic Topology*. Munkres. Addison-Wesley, Reading (1984)
29. Narasimhan, G., Smid, M.H.M.: *Geometric Spanner Networks*. Cambridge University Press, Cambridge (2007)
30. Singh, G., Mémoli, F., Ishkhanov, T., Sapiro, G., Carlsson, G., Ringach, D.L.: Topological analysis of population activity in visual cortex. *J. Vis.* **8**(8), 1–18 (2008)
31. Tausz, A., Carlsson, G.: Applications of zigzag persistence to topological data, analysis. <http://arxiv/1108.3545> (2011)
32. Zomorodian, A.: The tidy set. In: Proceedings of the 26th ACM Symposium on Computational Geometry, pp. 257–266, Aarhus (2010)
33. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 249–274 (2005)