A Multicover Nerve for Geometric Inference

Donald R. Sheehy

Abstract

We show that filtering the barycentric decomposition of a Čech complex by the cardinality of the vertices captures precisely the topology of k-covered regions among a collection of balls for all values of k. Moreover, we relate this result to the Vietoris-Rips complex to get an approximation in terms of the persistent homology.

1 Introduction

Computational geometers use topology to certify correctness of geometric constructions and inferences. For example, in surface reconstruction one often wants a homeomorphic reconstruction [8] or in medial axis approximation, one might seek a homotopy equivalence between the approximation and the true medial axis[4]. In some sensor network problems, topological guarantees can certify that the network covers a geometric domain [7]. A growing literature deals explicitly with the inference of topological structure in data sets (see Carlsson [1] for a survey).

Many of these examples depend on the Nerve Theorem or variations thereof to extract topological information from geometry. This classic result in algebraic topology relates the topology of a union of sets to that of a simplicial complex called the nerve (under certain conditions on the intersections of the sets).

In this paper, we extend the Nerve Theorem to consider regions covered by at least k different sets. In the language of sensor networks, this new nerve captures the notion of k-coverage. Whereas the Nerve Theorem can be applied directly for any fixed k, there is little correspondence between the nerves computed for different values of k. We show that a natural filtration of the barycentric decomposition of the nerve can capture this information for all values of k.

Noise and outliers are a major problem in topological data analysis. Even a single outlier can appear as a significant topological feature using standard methods. By considering k-covered regions only, our filtration ignores up to k points locally. This is closely related to a common approach to de-noising data for topological data analysis points are treated as noise if the distance to their kth nearest neighbor is at least some threshold α (see [11] and [2] for two notable examples). Our method has the added advantage that it is easy to relate results for different choices of k.

We prove our results in the setting of persistent homology. This allows us to relate the main result also to sets in general metric spaces where it may be difficult to compute k-wise intersections directly.

The specific case we are interested is the (k, α)-offsets of a point set $P \subset \mathbb{R}^d$, defined as the α-sublevel set of the kth nearest neighbor distance function. Equivalently, this is the subset of $\mathbb{R}^d$ covered by at least k balls of radius α centered at points in P (see Figure 1).

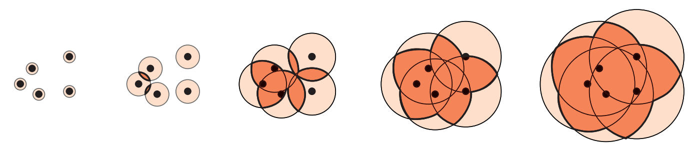

Figure 1: The α-offsets overlaid with the $(2, \alpha)$-offsets from growing values of α.

2 Background

Topology. We will assume a basic knowledge of standard definitions in topology including topological spaces, homotopy equivalence, and homology. The book by Munkres [10] is a good source for all the necessary background.

Simplicial Complexes. A **simplicial complex** S is family of subsets of a vertex set $V(S)$ that is closed under taking subsets. That is, if $\sigma' \subset \sigma \in S$ then $\sigma' \in S$. The elements of a simplicial complex are called **simplices** and the elements of the simplices are called **vertices**. The **dimension** of a simplex σ is defined as $|\sigma| - 1$. In this paper, we deal purely with abstract simplicial complexes and do not make any assumptions about how they are embedded.

Given a subset U of the vertices of S, the **induced subcomplex** of S on the vertex set U is the set of simplices whose vertices are all in U.

Filtrations. A **filtration** is a nested sequence of topological spaces. In this paper, we deal primarily with filtrations parameterized by the nonnegative real numbers. So, a filtration $\mathcal{G} = \{G^\alpha\}_{\alpha \geq 0}$ is a family of spaces such that $G^\alpha \subseteq G^\beta$ whenever $0 \leq \alpha \leq \beta$. For brevity, we omit the parameter set and write $\mathcal{G} = \{G^\alpha\}$ when it is obvious that α ranges over $\mathbb{R}_{\geq 0}$. If the spaces in a filtration are all simplicial complexes then we call it a **filtered simplicial complex**. Throughout the paper, superscripts are always used to index into a filtration.

Persistent Homology and Persistence Diagrams The theory of persistent homology describes the way the topology of the spaces in a filtration change as α ranges over $\mathbb{R}_{\geq 0}$. Given a filtered simplicial complex, there is an efficient algorithm for computing its so-called persistence diagram [13]. This diagram is a multiset of points in the extended plane $(\mathbb{R} \cup \{\infty\})^2$ where every point represents a topological feature. The x and y coordinates of a point in the persistence diagram represent the values of α for which that particular feature appeared and disappeared respectively in the filtration. For example, a cycle may form at $\alpha = x$ and then be filled in (killed) by triangles at $\alpha = y$. These are sometimes called the birth and death times of the feature.

By convention the diagonal $x = y$ is included in every persistence diagram. The distance from this diagonal is a measure of how long a feature persisted before being killed.

It is beyond the scope of this paper to give a full treatment of persistent homology; the book by Edelsbrunner and Harer gives a complete background[9].

From Sets to Filtrations. Persistent homology extends homology theory from spaces to filtrations. Below, we present some basic definitions and known results about persistence with an emphasis on the generalization of results from spaces to filtrations. Often, this means we will overload notation so that the same notations apply to both spaces and filtrations.

First we define the basic set operations on filtrations, defining $\{F^\alpha\} \cup \{G^\alpha\} = \{F^\alpha \cup G^\alpha\}$ and $\{F^\alpha\} \cap \{G^\alpha\} = \{F^\alpha \cap G^\alpha\}$. For any collection T of sets (or filtrations), we use the shorthand notation $\bigcup T = \bigcup_{S \in T} S$ and $\bigcap T = \bigcap_{S \in T} S$.

The first task is to extend a notion of topological equivalence from spaces to filtrations. Since the persistence diagram is a complete invariant of the filtration so filtrations $\mathcal{F}$ and $\mathcal{G}$ have isomorphic persistent homology if $\operatorname{Dgm} \mathcal{F} = \operatorname{Dgm} \mathcal{G}$. Unfortunately, to prove $\operatorname{Dgm} \mathcal{F} = \operatorname{Dgm} \mathcal{G}$, it does not suffice to have $\mathrm{H}_*(F^\alpha) \cong \mathrm{H}_*(G^\alpha)$ or even $F^\alpha \simeq G^\alpha$. The following lemma gives a sufficient condition. It is a special case of the persistence equivalence theorem [9, page 159]

Lemma 1 *Let $\mathcal{F} = \{F^\alpha\}$ and $\mathcal{G} = \{G^\alpha\}$ be filtrations. If for all $0 \leq \alpha \leq \beta$, there are isomorphisms $\mathrm{H}_*(F^\alpha) \to \mathrm{H}_*(G^\alpha)$ and $\mathrm{H}_*(F^\beta) \to \mathrm{H}_*(G^\beta)$ that commute with the homomorphisms $\mathrm{H}_*(F^\alpha) \to \mathrm{H}_*(F^\beta)$ and $\mathrm{H}_*(G^\alpha) \to \mathrm{H}_*(G^\beta)$ induced by inclusion, then $\operatorname{Dgm} \mathcal{F} = \operatorname{Dgm} \mathcal{G}$.*

Simplicial Maps. Let S and T be simplicial complexes. A map $f : S \to T$ is a **simplicial map** if f maps vertices to vertices and for every $\sigma \in S$, $f(\sigma) \in T$. A simplicial map is defined entirely by how it maps vertices to vertices. A simplicial map that is both injective and surjective is an **isomorphism of simplicial complexes**. We say that $\mathcal{F} = \{F^\alpha\}$ and $\mathcal{G} = \{G^\alpha\}$ are **isomorphic filtered simplicial complexes** if there exists a family of isomorphisms $\phi_\alpha : F^\alpha \to G^\alpha$ such that for all $0 \leq \alpha \leq \beta$, ϕ_α is the restriction of ϕ_β to F^α, denoted $\phi_\alpha = \phi_\beta|_{F^\alpha}$. The following Lemma follows directly from the definition of isomorphic filtrations and Lemma 1.

Lemma 2 *If $\mathcal{F}$ and $\mathcal{G}$ are isomorphic filtered simplicial complexes then $\operatorname{Dgm} \mathcal{F} = \operatorname{Dgm} \mathcal{G}$.*

When $S \subset T$, a map $f : S \to T$ is a **retraction** if $f(\sigma) = \sigma$ for all $\sigma \in S$. A pair of simplicial maps $f, g : S \to T$ are **contiguous** if $f(\sigma) \cup g(\sigma) \in T$ for all $\sigma \in S$. The theory of contiguity is a simplicial analogue of homotopy theory. The following lemma gives a homology analogue of a deformation retraction.

Lemma 3 *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 $\pi : Y \to X$ such that $i \circ \pi$ and id_Y are contiguous, then i induces an isomorphism $i_\star : \mathrm{H}_*(X) \to \mathrm{H}_*(Y)$ between the corresponding homology groups.*

Barycentric Decomposition. Let S be a simplicial complex. A **flag** in S is an ordered subset of simplices $\{\sigma_1, \ldots, \sigma_t\} \subseteq S$ such that $\sigma_1 \subset \cdots \subset \sigma_t$. The **barycentric decomposition** of S is the simplicial complex formed by the set of flags of S:

$$\text{Bary}\, S := \{U \subset S : U \text{ is a flag of } S\}.$$

We also define the barycentric decomposition of a filtered simplicial complex $\{S^\alpha\}$ to be the filtered simplicial complex $\text{Bary}\,\{S^\alpha\} := \{\text{Bary}\, S^\alpha\}$.

There is a natural filtration on a barycentric decomposition induced by considering only the flags of some minimum cardinality. We define the complexes in this filtration as

$$k\text{-Bary}\, S := \{\gamma \in \text{Bary}\, S : \min_{\sigma \in \gamma} |\sigma| \geq k\}.$$

As before, this definition is extended to filtered complexes $\{S^\alpha\}$ as $k\text{-Bary}\,\{S^\alpha\} := \{k\text{-Bary}\, S^\alpha\}$.

The operation of taking barycentric decompositions does not change the underlying topology. This fact is expressed in the following lemma, whose proof is trivial and omitted.

Lemma 4 *If $\mathcal{S}$ is a filtered simplicial complex then*

$$\operatorname{Dgm} \mathcal{S} = \operatorname{Dgm}\,(\text{Bary}\, \mathcal{S})$$

Note that this lemma is not true if we replace Bary with k-Bary .

Nerves. Let $\mathcal{F} = \{\{F_1^\alpha\}, \ldots, \{F_n^\alpha\}\}$ be a collection of filtrations. Define $\mathcal{F}_\alpha$ to be the collection of sets $\{F_1^\alpha, \ldots, F_n^\alpha\}$.

We say that $\mathcal{F}_\alpha$ is a **good open cover** of $\bigcup \mathcal{F}_\alpha$ if all F_i^α and their intersections are empty or contractible. This condition is easily satisfied if the F_i^α are open convex sets. We say that $\mathcal{F}$ is a **good filtered cover** if $\mathcal{F}^\alpha$ is a good open cover for all $\alpha \geq 0$.

The **nerve** of a collection of sets $\mathcal{F}_\alpha$ is the abstract simplicial complex $\mathrm{Nerve}\,\mathcal{F}_\alpha := \{U \subseteq \mathcal{F}_\alpha : \bigcap U \neq \emptyset\}$. The nerve of a collection of filtrations $\mathcal{F}$ is the filtered simplicial complex $\mathrm{Nerve}\,\mathcal{F} := \{\mathrm{Nerve}\,\mathcal{F}_\alpha\}_{\alpha\geq 0}$.

The following is a classic result in algebraic topology called the nerve theorem.

Theorem 5 (The Nerve Theorem) *If $\mathcal{F}_\alpha$ is a good open cover then*

$$\bigcup \mathcal{F}_\alpha \simeq \mathrm{Nerve}\,\mathcal{F}_\alpha.$$

The extension of the nerve theorem to the persistence setting follows from the Persistent Nerve Lemma of Chazal and Oudot [5] and Lemma 1:

Theorem 6 (Persistent Nerves) *If $\mathcal{F}$ is a good filtered cover then*

$$\mathrm{Dgm}\,\left\{\bigcup \mathcal{F}_\alpha\right\} = \mathrm{Dgm}\,(\mathrm{Nerve}\,\mathcal{F}).$$

k-Covers. Given a collection $\mathcal{F}$ of sets (or filtrations), the k-Cover of the collection is the set of k-wise intersections:

$$k\text{-Cover}\,\mathcal{F} := \left\{\bigcap U\right\}_{U\in\binom{\mathcal{F}}{k}}.$$

The k-cover of a collection of sets is a new collection of sets. The k-cover of a collection of filtrations is a new collection of filtrations.

3 Barycentric Bifiltration

The Barycentric Čech Filtration. Consider a set of points $P \subset \mathbb{R}^d$. The **Čech complex** at scale α is the nerve of the set of α-balls centered at the points of P. The collection of these complexes at all scales is the **Čech filtration** $\mathcal{C} = \{\mathcal{C}^\alpha\}$. The k-barycentric decomposition of the Čech filtration is denoted $\tilde{\mathcal{C}}_k = k\text{-Bary}\,\mathcal{C}$.

Since $\tilde{\mathcal{C}}_k^\alpha \subseteq \tilde{\mathcal{C}}_{k-1}^\alpha$ for any $\alpha \geq 0$ and $k \in \mathbb{N}$, this gives a filtration in two variables known as a bifiltration, where one dimension is parameterized by (increasing) α and the other is parameterized by (decreasing) k. In fact, the construction of $\tilde{\mathcal{C}}_k$ gives a general recipe for deriving a bifiltration from a filtered simplicial complex.

Our goal is to show that the filtration $\tilde{\mathcal{C}}_k$ has the same persistent homology as the (k, α)-offsets, P_k^α.

Theorem 7 *For any finite set of points $P \subset \mathbb{R}^d$ and any $k \in \mathbb{N}$, the persistence diagrams of the (k, α)-offsets of P and the k-barycentric decomposition of the Čech filtration are identical:*

$$\mathrm{Dgm}\,\{\tilde{\mathcal{C}}_k(P)\} = \mathrm{Dgm}\,\{P_k^\alpha\}.$$

This theorem follows from a more general result about good filtered covers, Theorem 10 below. It is the special case when the good filtered cover is the collection of balls of radius α centered at the points of P.

The Main Result Before getting to the main result, we set up some definitions and prove two necessary lemmas. Let $\mathcal{F}$ be a good filtered cover and let $k \in N$ be a fixed constant. Define the following filtrations:

$$\tilde{\mathcal{J}}_k := k\text{-Bary}\,(\mathrm{Nerve}\,\mathcal{F})$$
$$\mathcal{N}_k := \mathrm{Nerve}\,(k\text{-Cover}\,\mathcal{F})$$
$$\tilde{\mathcal{N}}_k := \mathrm{Bary}\,(\mathcal{N}_k)$$

Formally, the vertices of $\tilde{\mathcal{N}}_k^\alpha$ are the simplices of $\mathcal{N}_k^\alpha$, those collections of k-wise intersections of sets in $\mathcal{F}_\alpha$ that have a nonempty intersection. However, we will instead identify this vertex set with the corresponding collection of k-tuples from $\mathcal{F}_\alpha$. Letting X^α and Y^α be the vertex sets of $\tilde{\mathcal{J}}_k^\alpha$ and $\tilde{\mathcal{N}}_k^\alpha$ respectively, we have

$$X^\alpha = \left\{U \subseteq \mathcal{F}_\alpha : |U| \geq k \text{ and } \bigcap U_\alpha \neq \emptyset\right\}$$
$$Y^\alpha = \left\{V \subseteq \binom{\mathcal{F}_\alpha}{k} : \bigcap_{V'\in V} \bigcap V' \neq \emptyset\right\}$$

The complex $\tilde{\mathcal{N}}_k$ contains redundant information. The map $\pi : Y \to Y$ induces a simplicial map that "projects out" this redundant information. It is defined by

$$\pi(V) = \binom{\bigcup V}{k}.$$

Figure 2 demonstrates the construction of some of the simplicial complexes described above for the special case of the Čech filtration and $k = 2$.

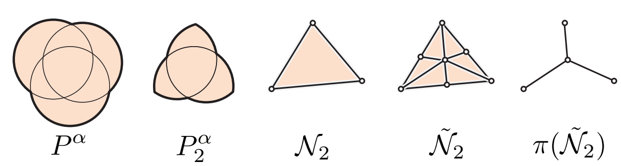

Figure 2: The construction of $\mathcal{N}_2$, its barycentric decomposition, and its image under π.

The following Lemma sows that the persistence diagram of $\tilde{\mathcal{N}}_k$ is unchanged by π.

Lemma 8 $\mathrm{Dgm}\,\tilde{\mathcal{N}}_k = \mathrm{Dgm}\,\pi(\tilde{\mathcal{N}}_k)$.

Proof. By Lemma 1, it suffices to show that for all $\alpha \geq 0$, the inclusion $\psi : \pi(\tilde{\mathcal{N}}_k^\alpha) \hookrightarrow \tilde{\mathcal{N}}_k^\alpha$ induces an isomorphism at the homology level. As above, let Y^α be the set of vertices of $\tilde{\mathcal{N}}_k$ and let $s = \max_{V \in Y^\alpha} |V|$. For $i = 0 \ldots s$, define

$$Y_i = \{V \in Y^\alpha : |V| \leq i\} \cup \{\pi(V) : V \in Y^\alpha, |V| > i\}.$$

Let A_i be the subcomplex of $\tilde{\mathcal{N}}_k^\alpha$ induced on Y_i. So

$$\pi(\tilde{\mathcal{N}}_k^\alpha) = A_0 \subset \cdots \subset A_s = \tilde{\mathcal{N}}_k^\alpha.$$

It will suffice to show that the inclusion $\psi_i : A_{i-1} \hookrightarrow A_i$ induces an isomorphism at the homology level for all $i = 1 \ldots s$. Let $\pi_i : Y_i \to Y_{i-1}$ be defined as

$$\pi_i(V) = \begin{cases} V & \text{if } |V| < i \\ \pi(V) & \text{if } |V| \geq i \end{cases}$$

So $\psi = \psi_s \circ \cdots \circ \psi_1$ and $\pi = \pi_1 \circ \cdots \circ \pi_s$. Lemma 3 will give the desired isomorphism if π_i is a simplicial retraction such that $\psi_i \circ \pi_i$ is contiguous with the identity map, i.e. that (1) π_i restricts to the identity on Y_{i-1}, (2) $\pi_i(\sigma) \in A_{i-1}$ for all $\sigma \in A_i$, and (3) $(\sigma \cup \pi_i(\sigma)) \in A_i$ for all $\sigma \in A_i$.

Item (1) is obvious from the definitions. To prove (2) and (3), fix a simplex $\sigma = \{V_0, \ldots, V_t\} \in A_i$ and let $\sigma' = \sigma \cup \pi_i(\sigma)$. If $\sigma = \sigma'$ then we are done, so we may assume that for some vertex $V_j \in \sigma$, $\pi(V_j) \notin \sigma$. Recall that the simplices of $\tilde{\mathcal{N}}_k$ (and also A_i) are strictly nested sequences of vertices. So, there is at most one vertex V_j such that $\pi(V_j) \notin \sigma$, namely the one with cardinality i. We may therefore express σ' as $\sigma \cup \{\pi(V_j)\}$. Since $\sigma \in A_i \subset \tilde{\mathcal{N}}_k$, $V_0 \subset \cdots \subset V_t$. Observe that $V \subseteq \pi(V)$ for all $V \in Y^\alpha$ and moreover that $U \subset V$ implies $\pi(U) \subseteq \pi(V)$. So, it follows that

$$V_0 \subset \cdots \subset V_j \subset \pi(V_j) \subset \pi(V_{j+1}) = V_{j+1} \subset \cdots \subset V_t.$$

The inclusion of $\pi(V_j) \subset \pi(V_{j+1})$ is strict because of the assumption that $\pi(V_j) \notin \sigma$. This is a strictly nested sequence of the vertices of σ' so $\sigma' \in A_i$, proving (3). Moreover, $\pi_i(\sigma) = \sigma' \setminus \{V_j\}$ so $\pi_i(\sigma) \in \tilde{\mathcal{N}}_k$ as well. Since $\pi_i(\sigma) \subset Y_{i-1}$, we conclude that $\pi_i(\sigma) \in A_{i-1}$, proving (2). □

Next, we prove that $\tilde{\mathcal{J}}_k$ and $\pi(\tilde{\mathcal{N}}_k)$ have identical persistence diagrams.

Lemma 9 $\operatorname{Dgm} \tilde{\mathcal{J}}_k = \operatorname{Dgm} \pi(\tilde{\mathcal{N}}_k)$

Proof. We will show that $\tilde{\mathcal{J}}_k$ and $\pi(\tilde{\mathcal{N}}_k)$ are isomorphic filtered simplicial complexes and so the result will follow from Lemma 2. It will suffice to show that for all $\alpha \geq 0$, $\tilde{\mathcal{J}}_k^\alpha$ and $\tilde{\mathcal{N}}_k^\alpha$ are isomorphic and that the isomorphism does not depend on α.

The desired isomorphism is the map $\phi : X^\alpha \to \pi(Y^\alpha)$ defined as $\phi(U) = \binom{U}{k}$. The inverse of this map is $\phi^{-1}(V) = \bigcup V$. So, ϕ takes subsets $U \subset \mathcal{F}^\alpha$ of size at least k such that $\bigcap U \neq \emptyset$ to the family of k-element subsets of U. It is easy to check that ϕ is a bijection.

To show that ϕ is an isomorphism, we will prove that σ is a simplex of $\tilde{\mathcal{J}}_k^\alpha$ if and only if $\phi(\sigma)$ is a simplex of $\pi(\tilde{\mathcal{N}}_k)$. Let $\sigma = (U_0, \ldots, U_j) \in \tilde{\mathcal{J}}_k^\alpha$ be any simplex. By the definition of $\tilde{\mathcal{J}}_k^\alpha$, $U_0 \subset \cdots \subset U_j$. For any pair of vertices U_a and U_b, $U_a \subset U_b$ if and only if $\phi(U_a) \subset \phi(U_b)$. So, $\sigma \in \tilde{\mathcal{J}}_k^\alpha$ if and only if $\phi(U_0) \subset \cdots \subset \phi(U_j)$, which holds if and only if $\phi(\sigma) \in \pi(\tilde{\mathcal{N}}_k)$. □

We are now ready to prove the main theorem relating the persistence diagrams of k-Bary (Nerve $\mathcal{F}$) and $\bigcup k$-Cover $\mathcal{F}$. The basic strategy is illustrated in Figure 3.

Theorem 10 *If $\mathcal{F}$ is a good filtered cover and $k \in \mathbb{N}$ then* $\operatorname{Dgm}(k\text{-Bary}(\text{Nerve}\,\mathcal{F})) = \operatorname{Dgm}\left(\bigcup k\text{-Cover}\,\mathcal{F}\right)$.

Proof. Recall the notations $\tilde{\mathcal{J}}_k$, $\mathcal{N}_k$, and $\tilde{\mathcal{N}}_k$ defined above.

$$\begin{aligned}
\operatorname{Dgm} \tilde{\mathcal{J}}_k &= \operatorname{Dgm} \pi(\tilde{\mathcal{N}}_k) && \text{[by Lemma 9]}\\
&= \operatorname{Dgm} \tilde{\mathcal{N}}_k && \text{[by Lemma 8]}\\
&= \operatorname{Dgm} \mathcal{N}_k && \text{[by Lemma 4]}\\
&= \operatorname{Dgm}(\text{Nerve}\,(k\text{-Cover}\,\mathcal{F})) && \text{[by definition]}\\
&= \operatorname{Dgm}\left(\bigcup k\text{-Cover}\,\mathcal{F}\right) && \text{[by Theorem 6]}
\end{aligned}$$

□

The Barycentric Vietoris-Rips Filtration One drawback of the Čech filtration is that it requires testing sets of balls for common intersections. An alternative approach is to construct the edges only and include simplices for every clique. This is known as the Vietoris-Rips filtration $\mathcal{R} = \{\mathcal{R}^\alpha\}$, where

$$\mathcal{R}^\alpha := \{Q \subseteq P : \text{diameter}(Q) \leq 2\alpha\}.$$

This can be computed using only the pairwise distances between points and therefore is well-defined for any metric space. We can apply the same barycentric bifiltration approach used above to yield a bifiltration $\{\tilde{\mathcal{R}}_k^\alpha\}$.

Given filtrations $\mathcal{F}$ and $\mathcal{G}$, we say $\operatorname{Dgm} \mathcal{F}$ is c-approximation for $\operatorname{Dgm} \mathcal{G}$ if there is a 1-1 correspondence that maps each $(x, y) \in \operatorname{Dgm} \mathcal{F}$ to $(u, v) \in \operatorname{Dgm} \mathcal{G}$ such that $u/c \leq x \leq cu$ and $v/c \leq y \leq cv$. A sufficient condition for $\operatorname{Dgm} \mathcal{F}$ to be a c-approximation to $\operatorname{Dgm} \mathcal{G}$ is that $\mathcal{F}^{\alpha/c} \subseteq \mathcal{G}^\alpha \subseteq \mathcal{F}^{c\alpha}$ for all $\alpha \geq 0$. This is a simple corollary to the Strong Stability Theorem of Chazal et al. [3].

The Vietoris-Rips filtration gives a good approximation to the Čech filtration. It was shown by de Silva and Ghrist that $\mathcal{C}^\alpha \subseteq \mathcal{R}^\alpha \subseteq \mathcal{C}^{c\alpha}$, where $c = 2$ for general metric spaces and $c = \sqrt{2}$ for Euclidean spaces[6].

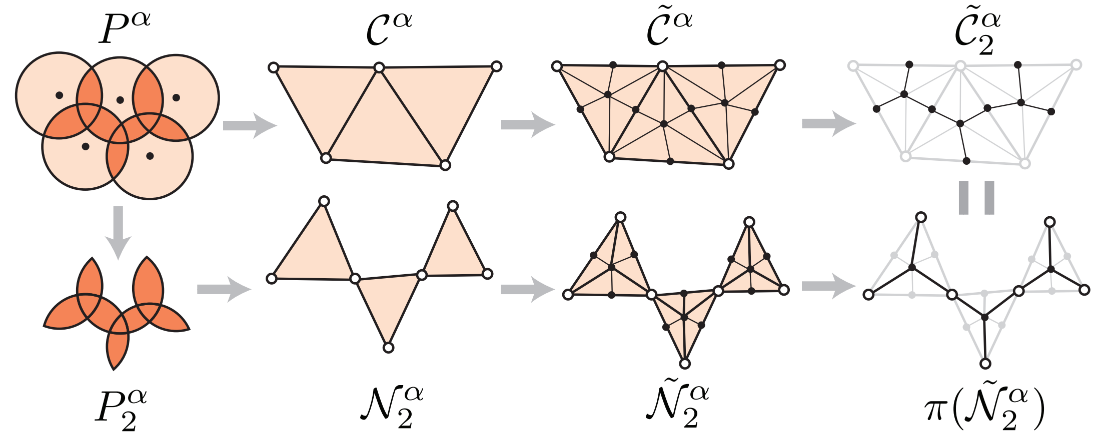

Figure 3: We transform the collection of balls in two different ways to get equivalent complexes, $\tilde{\mathcal{C}}_k^\alpha$ (top) and $\pi(\tilde{\mathcal{N}}_k^\alpha)$ (bottom) for $k = 2$.

So, the Vietoris-Rips filtration gives a c-approximation to the /v Cech filtration for persistent homology. The interleaving also implies the following extension to the Vietoris-Rips bifiltration $\{\tilde{\mathcal{R}}_k^\alpha\}$, where $\mathcal{R}_k = k\text{-Bary}\,\mathcal{R}$.

Theorem 11 *For any fixed k, the persistence diagram of the barycentric Rips filtration, $\{\tilde{\mathcal{R}}_k^\alpha\}$, is a $\sqrt{2}$-approximation to the persistence diagram of the (k,α)-offsets $\{P_k^\alpha\}$ when the underlying space is Euclidean, and is a 2-approximation for general metrics.*

Proof. It suffices to observe that $\mathcal{C}^\alpha \subseteq \mathcal{R}^\alpha \subseteq \mathcal{C}^{c\alpha}$ implies $k\text{-Bary}\,\mathcal{C}^\alpha \subseteq k\text{-Bary}\,\mathcal{R}^\alpha \subseteq k\text{-Bary}\,\mathcal{C}^{c\alpha}$. □

4 Conclusions and Future Work

We have presented a nerve construction to capture the topology of the k-covered regions of a collection of well-behaved sets. Our focus was on guaranteeing the correct persistent homology, when the sets are filtrations, but it is also possible to consider the case of just a single good open cover $\mathcal{S}$. In that case, using a slightly stronger version of Lemma 3, it is possible to prove that the $k\text{-Bary}(\text{Nerve}\mathcal{S})$ is homotopy equivalent to $\bigcup k\text{-Cover}\mathcal{S}$.

In practice, it is common to truncate Čech filtrations at some maximum scale to avoid the huge complexity blowup. The method of barycentric bifiltrations naturally adapts to this setting. In recent work, we proposed an alternative approach to controlling the complexity of distance based filtrations using hierarchical net trees [12]. It may be possible to combine those ideas with those presented in this paper to give sparse approximations of the (k,α)-offsets. This is the subject of future work.

References

[1] G. Carlsson. Topology and data. *Bull. Amer. Math. Soc.*, 46:255–308, 2009.

[2] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, 2008.

[3] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th ACM Symposium on Computational Geometry*, pages 237–246, 2009.

[4] F. Chazal and A. Lieutier. The "λ-medial axis". *Graphical Models*, 67(4):304–331, 2005.

[5] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the 24th ACM Symposium on Computational Geometry*, pages 232–241, 2008.

[6] V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology. *Algorithmic & Geometric Topology*, 7:339–358, 2007.

[7] V. de Silva and R. Ghrist. Homological sensor networks. *Notices Amer. Math. Soc.*, 54(1):10–17, 2007.

[8] T. K. Dey. *Curve and Surface Reconstruction : Algorithms with Mathematical Analysis*. Cambridge University Press, 2007.

[9] H. Edelsbrunner and J. L. Harer. *Computational Topology: An Introduction*. Amer. Math. Soc., 2009.

[10] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984.

[11] P. Niyogi, S. Smale, and S. Weinberger. A topological view of unsupervised learning from noisy data. 2008.

[12] D. R. Sheehy. Linear-size approximations to the vietoris-rips filtration. In *Proceedings of the 28th ACM Symposium on Computational Geometry*, 2012.

[13] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.