Persistent Nerves Revisited

Nicholas J. Cavanna and Donald R. Sheehy

University of Connecticut

1 Introduction

A nerve is a simplicial complex derived from a cover of a topological space. Nerves appear all over computational topology and geometry, e.g. a Delauney triangulation is the nerve of a Voronoi diagram and the Čech complex is the nerve of a collection of metric balls. They are used to solve problems concerning surface reconstruction, homology inference, and homological sensor networks, among other areas.

If one has an open cover of a paracompact space in which all non-empty intersections of the cover elements are contractible, i.e. it is a *good cover*, then its nerve is homotopy equivalent to the covered space. This result is known as the Nerve Theorem. It has a natural extension to the setting of persistent homology called the Persistent Nerve Lemma (PNL), due to Chazal and Oudot [2]. The PNL implies that given a filtration of covers of a filtration of spaces such that at each time the cover is good, then the persistent homology of the space filtration is that of the nerve filtration. Good covers are not always an option, e.g. if a metric space is not convex then the metric balls of a finite point sample may cover the space, but they won't be a good cover without adding other conditions. The requirement of having a good cover in order to invoke the PNL is the motivation for our work – instead we assume the cover elements' homology is trivial when included into a later scale.

Recently, Botnan and Spreemann assumed an interleaving between two cover filtrations to prove a bound on the bottleneck distance between the persistence diagrams of their nerve filtrations [1]. Govc and Skraba [3] considered a simplicial filtration and a cover of the terminal complex, using it to construct a cover filtration by taking the intersection of the cover with the simplicial filtration at each scale. They assumed that the persistence modules of all the non-empty k-wise intersections of the filtration's cover elements were ε-interleaved with 0 and proved a tight bound on the bottleneck distance between the persistence diagrams of the filtered simplicial complex and the nerve filtration, linear with respect to dimension and ε.

We consider a more general cover assumption—that we have an arbitrary cover filtration consisting of simplicial complexes which collectively cover a filtered simplicial complex, rather than a cover filtration induced by intersection like Govc and Skraba. We also assume the cover filtration is what we define as ε-*good*. The major contributions are that there are interleavings between the filtered simplicial complex and the nerve filtration's homology groups, induced by chain maps, which imply a tight dimension-dependent bound on the bottleneck distance between the two persistence diagrams, linear with respect to dimension and ε. A noteworthy corollary of our result is the Persistent Nerve Lemma by considering a 0-good cover.

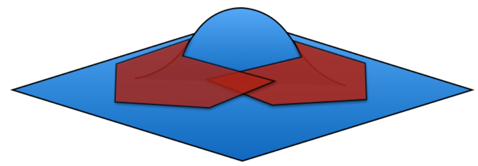

Figure 1: A bump in $\mathbb{R}^3$ and subsets with non-contractible intersection.

2 Background

Let $\mathcal{U} := \{U_1, \ldots, U_n\}$ be an arbitrary collection of *filtrations*, growing sequences of spaces, where $U_i := (U_i^\alpha)_{\alpha \geq 0}$, and each U_i^α is a simplicial complex. We call $\mathcal{U}$ a *cover filtration*. For each $\alpha \geq 0$, define $\mathcal{U}^\alpha := \{U_1^\alpha, \ldots, U_n^\alpha\}$ and $W^\alpha := \bigcup_{i\in[n]} U_i^\alpha$. For each non-empty $v \subseteq [n] = \{1, \ldots, n\}$, let $U_v^\alpha := \bigcap_{i\in v} U_i^\alpha$.

The *nerve* of the cover $\mathcal{U}^\alpha$ is defined as $\mathrm{Nrv}\,\mathcal{U}^\alpha := \{v \subseteq [n] \mid U_v^\alpha \neq \emptyset\}$. One can check this is a simplicial complex. The *nerve filtration* is defined as $\mathrm{Nrv}\,\mathcal{U} := (\mathrm{Nrv}\,\mathcal{U}^\alpha)_{\alpha\geq 0}$. When we consider the collection of spaces that each $\mathcal{U}^\alpha$ covers over all $\alpha \geq 0$, we get the *union filtration*, $\mathcal{W} := (W^\alpha)_{\alpha\geq 0}$. $\mathcal{U}^\alpha$ is a *good cover* of W^α if for all subsets $v \subseteq [n]$, we have U_v^α is empty or contractible. For filtrations, we say $\mathcal{U}$ is a *good cover* of $\mathcal{W}$ if $\mathcal{U}^\alpha$ is a good cover of W^α for all $\alpha \geq 0$. The Persistent Nerve Lemma implies that if $\mathcal{U}$ is a good cover of $\mathcal{W}$, then $\mathrm{Dgm}(\mathrm{Nrv}\,\mathcal{U}) = \mathrm{Dgm}(\mathcal{W})$, where $\mathrm{Dgm}(\cdot)$ is the *persistence diagram* over all dimen-

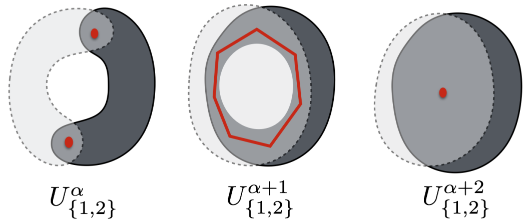

Figure 2: A cover filtration that is not good, but is 1-good.

sions of the input filtration– a multiset representing the "birth" and "deaths" of homological features as $\alpha \to \infty$. When we refer to the persistence diagram of just the k-dimensional homological features we will write $\mathrm{Dgm}_k(\cdot)$.

Given a cover filtration $\mathcal{U}$, we say it is *ε-good* if for all non-empty $v \subseteq [n]$, and for all $\alpha \geq 0$, $\tilde{H}_*(U_v^\alpha \hookrightarrow U_v^{\alpha+\varepsilon}) = 0$, so any nontrivial homology classes of U_v^α are trivial when mapped into $U_v^{\alpha+\varepsilon}$. Note that due to the definition of contractibility, $\mathcal{U}$ being a good cover of $\mathcal{W}$ implies that it is 0-good.

For each $\mathcal{U}^\alpha$ there is a corresponding commutative diagram $D\mathcal{U}^\alpha$, where the spaces are the non-empty sets U_v^α for non-empty $v \subseteq [n]$ and there is an inclusion map $U_v^\alpha \hookrightarrow U_{v'}^\alpha$ whenever $v' \subset v$. Let N^α be the barycentric subdivision of $\mathrm{Nrv}\,\mathcal{U}^\alpha$ which has simplices of the form $\sigma = v_0 \to \ldots \to v_k$, where $v_i \subset v_{i+1}$, and each v_i corresponds to a simplex of $\mathrm{Nrv}\,\mathcal{U}^\alpha$. This is an abstract simplicial complex and we denote the associated geometric filtration as $\mathcal{N} := (|N^\alpha|)_{\alpha \geq 0}$. We define the *homotopy colimit* of $D\mathcal{U}^\alpha$ as

$$\mathrm{hocolim}\, D\mathcal{U}^\alpha := \bigcup_{N^\alpha \ni \sigma = v_0 \to \ldots} U_{v_0}^\alpha \times |\sigma|,$$

where $|\cdot|$ is the geometric realization functor. This homotopy colimit is also known as the Meyer-Vietoris blowup complex [4]. It yields another filtration, $\mathcal{B} = (B^\alpha)_{\alpha \geq 0}$, where $B^\alpha = \mathrm{hocolim}\, D\mathcal{U}^\alpha$.

We note that there is a (pseudo)-metric between two persistence diagrams D and D' called the *bottleneck distance*, denoted $\mathrm{d}_B(D, D')$, which is the standard measure of the similarity of two persistence diagrams, and with that, the persistent homology of two filtrations.

3 Results

Theorem 1 *If $\mathcal{U} = \{U_1, \ldots, U_n\}$ is a set of simplicial filtrations that is an ε-good cover of the simplicial filtration $\mathcal{W} = \bigcup_{i=1}^n U_i$, then*

$$\mathrm{d}_B(Dgm_k(\mathcal{W}), Dgm_k(\mathrm{Nrv}\,\mathcal{U})) \leq \frac{(k+1)\varepsilon}{2}.$$

Furthermore, there is an upper-bound of $\frac{(D+1)\varepsilon}{2}$, where D is the dimension of the nerve filtration.

An overview of the proof is as follows. As for all α, $B^\alpha \subseteq W^\alpha \times |N^\alpha|$, we have natural chain maps induced by projection $b^\alpha : C_*(B^\alpha) \to C_*(W^\alpha)$ and $p^\alpha : C_*(B^\alpha) \to C_*(|N^\alpha|)$, where b^α at the space level is a homotopy equivalence, so $\mathrm{Dgm}(\mathcal{W}) = \mathrm{Dgm}(\mathcal{B})$.

Since $|N^\alpha|$ is homeomorphic to $|\mathrm{Nrv}\,\mathcal{U}^\alpha|$, it follows that $\mathrm{Dgm}(\mathcal{N}) = \mathrm{Dgm}(|\mathrm{Nrv}\,\mathcal{U}|)$, and as simplicial homology is equivalent to singular homology, we have that $\mathrm{Dgm}(\mathcal{N}) = \mathrm{Dgm}(\mathrm{Nrv}\,\mathcal{U})$. Define $t := (k+1)\varepsilon$, where k is the maximal dimension of homology groups being considered. We create a chain map $q^\alpha : C_*(|N^\alpha|) \to C_*(W^{\alpha+t})$ such that $a^\alpha := q^\alpha \circ p^\alpha$ is chain homotopic to $i_{\mathcal{B}}^{\alpha,\alpha+t} \circ b^\alpha$, via chain homotopy $c^\alpha : C_*(|N^\alpha|) \to C_*(W^{\alpha+t})$. These maps can be viewed in diagram 1.

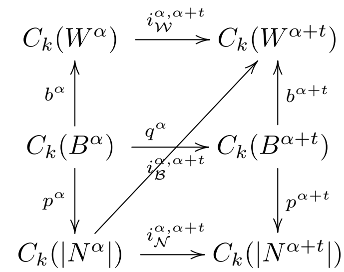

We use q^α to define a chain map $\bar{q}^\alpha : C_*(|N^\alpha|) \to C_*(B^{\alpha+t})$, where $\bar{q}^\alpha(\sigma) := \sum_{i=0}^{k} q(\sigma_i) \otimes \bar{\sigma}_i$, with $\sigma_i := v_0 \to \ldots \to v_i$ and $\bar{\sigma}_i := v_i \to \ldots \to v_k$, such that $p^{\alpha+t} \circ \bar{q}^\alpha$ commutes with $i_{\mathcal{N}}^{\alpha,\alpha+t}$ and $\bar{q}^\alpha \circ p^\alpha$ is chain homotopic to $i_{\mathcal{B}}^{\alpha,\alpha+t}$, via chain homotopy $\bar{c}^\alpha$, defined analogously to $\bar{q}^\alpha$.

By applying the homology functor to the diagram at all α, the chain maps p^α and $\bar{q}^\alpha$ commute with all the inclusion homomorphisms, forming interleaving homomorphisms between $\mathcal{N}$ and $\mathcal{B}$ thus implying our result.

References

[1] Magnus Bakke Botnan and Gard Spreemann. Approximating persistent homology in Euclidean space through collapses. *Applicable Algebra in Engineering, Communication and Computing*, pages 1–29, 2015.

[2] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 232–241, New York, NY, USA, 2008. ACM.

[3] Dejan Govc and Primoz Skraba. An approximate nerve theorem. *arXiv*, https://arxiv.org/pdf/1608.06956v2.pdf, 2016.

[4] Afra Zomorodian and Gunnar Carlsson. Localized homology. *Computational Geometry: Theory and Applications*, 41(3):126–148, 2008.