Parameterized Vietoris-Rips Filtrations via Covers

Bradley J. Nelson

University of Chicago

Vietoris-Rips Seminar
February 18, 2022

Section 1

Introduction

Cover Complexes

We'll consider *Vietoris-Rips Cover Complexes*. Let $\mathcal{U}$ be a cover of a data set $\mathbf{X}$, and $\mathcal{R}(\mathbf{X}; r)$ denote the Vietoris-Rips complex of $\mathbf{X}$ at parameter r. We define

$$\mathcal{R}(\mathbf{X}, \mathcal{U}; r) = \bigcup_{U \in \mathcal{U}} \mathcal{R}(U; r)$$

Why consider this?

- Smaller than standard VR complex, more fine-grained structure than Nerve
- Parallelization of persistent homology computation [Yoon, 2018]
- Parameterization using $\mathcal{U}$ - like mapper [Singh et al., 2007]

Note: GUDHI has something called a `CoverComplex` which is more related to the Nerve. This is different.

Vietoris-Rips Filtrations

Let $(\mathbf{X}, d)$ be a dissimilarity space. We extend the dissimilarity to tuples of points $x_0, \dots, x_k \subseteq \mathbf{X}$ as

$$d(x_0, \dots, x_k) = \max_{0 \le i < j \le k} d(x_i, x_j)$$

The Vietoris-Rips complex of $\mathbf{X}$ at parameter r is the simplicial complex

$$\mathcal{R}(\mathbf{X}; r) = \{(x_0, \dots, x_k) \mid x_0, \dots, x_k \in \mathbf{X}, d(x_0, \dots, x_k) \le r\}.$$

The Vietoris-Rips filtration of $\mathbf{X}$ is the nested sequence of complexes

$$\mathcal{R}(\mathbf{X}; 0) \subseteq \cdots \subseteq \mathcal{R}(\mathbf{X}; r) \subseteq \dots$$

Using Vietoris-Rips Filtrations

Pros:

- Easy to define, construct algorithmically
- Persistent homological features capture structure in the sample $\mathbf{X}$.
- Persistent homology is stable to perturbations of $(\mathbf{X}, d)$.

Cons:

- Size. Number of k-simplices on n points is $\binom{n}{k} \approx O(n^k)$.

Nerves and Covers

Another idea to get a topological signature of $\mathbf{X}$: Produce a cover $\mathcal{U} = \{U_i\}_{i \in \mathcal{I}}$ of $\mathbf{X}$. Construct the Nerve

$$\mathcal{N}(\mathcal{U}) = \{(U_0, \ldots, U_k) \mid U_0 \cap U_k \neq \emptyset\}$$

Nerve Theorem

If $\mathbf{X}$ is paracompact and $\mathcal{U}$ is a good cover, then $\mathcal{N}(\mathcal{U}) \simeq \mathbf{X}$.

A good cover means that $\bigcap_{i \in \mathcal{J}} U_i$ is either empty or contractible for all $\mathcal{J} \subseteq \mathcal{I}$.

Using Nerves

Pros:

- Easy to define, construct algorithmically
- Typically much smaller than $\mathcal{R}$

Cons:

- When $\mathbf{X}$ is a discrete sample, the only good cover has the same discrete topology as $\mathbf{X}$ - no large scale geometric structure.
- How to choose cover? Lose some geometric structure inside sets.
- Not an immediately obvious way to filter - this would depend on parameterizing construction of $\mathcal{U}$

Persistent Homology

Persistent homology can be used to compute topological features of a filtration.

$$H_k(\mathcal{X}(\cdot)) = H_k(\mathcal{X}(0)) \to \cdots \to H_k(\mathcal{X}(r)) \to \dots$$

Is a persistence vector space, and is classified up to isomrorphism by its *barcode* $\{(b_i, d_i)\}$ where each b_i indicates the birth parameter of a new homology class which maps through inclusions until entering a kernel at parameter d_i.

Interleavings

An interleaving relates two persistence vector spaces. Let $\alpha, \beta : \mathbb{R} \to \mathbb{R}$ be non-decreasing maps.
An (α, β)-interleaving between $H_k(\mathcal{X}(\cdot))$ and $H_k(\mathcal{Y}(\cdot))$ consists of two shift maps

$$F^{\alpha} : H_k(\mathcal{X}(r)) \to H_k(\mathcal{Y}(\alpha(r))$$
$$G^{\beta} : H_k(\mathcal{Y}(r)) \to H_k(\mathcal{X}(\beta(r))$$

so that the following diagram commutes

$$\begin{array}{ccccc} H_k(\mathcal{X}(r)) & \to & H_k(\mathcal{X}(\beta(r))) & \to & H_k(\mathcal{X}(\beta \circ \alpha(r))) \\ & \searrow \nearrow & & \searrow \nearrow & \\ H_k(\mathcal{Y}(r)) & \to & H_k(\mathcal{Y}(\alpha(r))) & \to & H_k(\mathcal{Y}(\alpha \circ \beta(r))) \end{array}$$

Nerve Theorems and Persistence

If a filtration of Nerves always satisfies the good cover property, then there is a natural extension of the Nerve theorem [Chazal and Oudot, 2008]

Results on extending the Nerve theorem to interleavings: [Govc and Skraba, 2018, Cavanna and Sheehy, 2018] Relaxation of good cover to having ϵ-acyclic intersections.

Parameterization

In classical topology, can study spaces *parameterized* by base space $\mathcal{B}$.

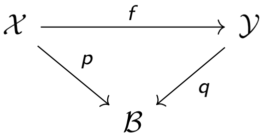

In setting of discrete samples, makes sense to work with covers of $\mathcal{B}$ and pullbacks.

The mapper construction [Singh et al., 2007] is an example of this.

[Yoon, 2018] studies $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$ where $\mathcal{U}$ is the pullback of a cover of $\mathbb{R}$.

Contributions

This talk:

- Interleaving relationship between $\mathcal{R}(\mathbf{X}; r)$ and $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$.
- To do this, we develop an extension of the method of acyclic carriers to interleavings.

Other applications of these techniques:

- Stability of $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$
- Approximate nerve theorem for cover complexes
- Alternative approach geometric stability results e.g. [Chazal et al., 2014]

References: [Nelson, 2020], [Nelson, 2022] in preparation.

Section 2

Examples

Pullback Covers

In a variety of situations is can be interesting to use a map $f : \mathbf{X} \to Y$ to investigate data. This helps with dimensionality reduction, and can also help incorporate known structure.

This is very similar to the approach of the mapper construction [Singh et al., 2007], which is the nerve of a refinement of a pullback cover.

Torus 1

Flat torus in 4 dimensions, 500 samples on a coil.

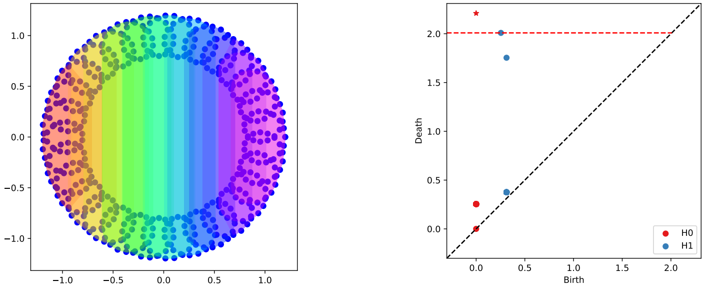

Figure: Cover pulled back from projection onto first coordinate

Torus 2

Flat torus in 4 dimensions, 500 samples on a coil.

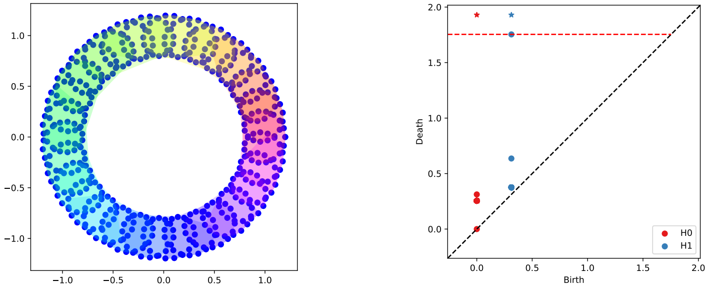

Figure: Cover pulled back from projection onto first two coordinates

Torus 3

Flat torus in 4 dimensions, 500 samples on a coil.

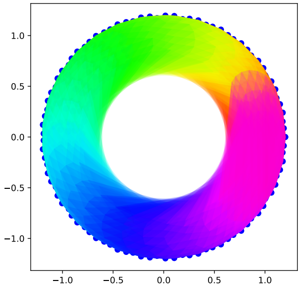

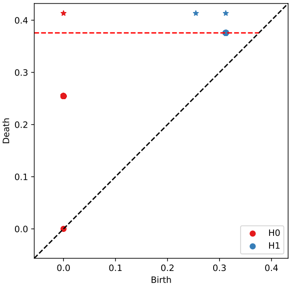

Figure: Cover based on 20-nearest neighbors of each point

Torus 4

Flat torus in 4 dimensions, 500 samples on a coil.

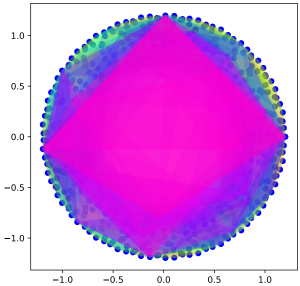

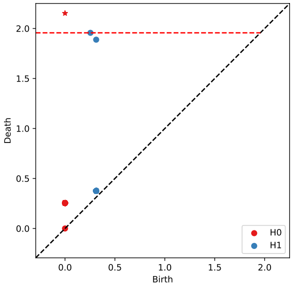

Figure: Cover based on nested landmarking procedure

Section 3

A Filtered Acyclic Carrier Theorem

Acyclic Carrier Theorem

If a carrier $\mathcal{C} : C_* \to D_*$ is acyclic, and $L_* \subset C_*$ is a sub-chain complex of C_*, then any chain map $\hat{F}_* : L_* \to D_*$ can be extended to a chain map $F_* : C_* \to D_*$. Furthermore, this extension is unique up to chain homotopy.

We'll talk about what a carrier is, but the take-away is that there is a way to construct maps between chain complexes from some initial data.

Classical introductions/proofs: [Eilenberg and Steenrod, 1952, Mosher and Tangora, 1968, Munkres, 1984]

Filtered Objects

Filtered Objects

A filtered object in a category over a poset T is a collection of objects $\mathcal{X}^T = \{\mathcal{X}^t\}_{t \in T}$ where $\mathcal{X}^{t_1} \subseteq \mathcal{X}^{t_2}$ if $t_1 \leq t_2$.

Shift Maps

Let $\mathcal{X}^S, \mathcal{Y}^T$ be filtered objects in a category over posets S, T respectively. Let $\alpha : S \to T$ be a non-decreasing map. An α-shift map $f^\alpha : \mathcal{X}^S \to \mathcal{Y}^T$ is a collection of maps $f^s : \mathcal{X}^s \to \mathcal{Y}^{\alpha(s)}$ for each $s \in S$ so that the following diagram commutes.

$$\begin{array}{ccc} \mathcal{X}^s & \longrightarrow & \mathcal{X}^{s'} \\ \downarrow{\scriptstyle f^s} & & \downarrow{\scriptstyle f^{s'}} \\ \mathcal{Y}^{\alpha(s)} & \longrightarrow & \mathcal{Y}^{\alpha(s')} \end{array} \tag{1}$$

Filtered Carriers

Filtered Carriers

A filtered carrier of chain complexes over a poset T, denoted $\mathcal{C}^T : C_*^S \to D_*^T$ is an assignment of basis vectors of C_*^S to filtered sub-complexes of D_*^T. In situations where T is understood, we will drop the superscript, and simply write $\mathcal{C} : C_*^S \to D_*^T$.

We can also define a filtered carrier of cell complexes $\mathcal{C}^T : \mathcal{X}^S \to \mathcal{Y}^T$ by assigning cells of $\mathcal{X}^S$ to sub-cell complexes of $\mathcal{Y}^T$.

To get the classical version of carriers, just take the trivial posets $S = T = \{0\}$.

Carrying a Map

Let $\mathcal{C}^T : C_*^S \to D_*^T$ be a filtered carrier, and F_*^α be an α-shift chain map. We say that $F_*^\alpha : C_*^S \to D_*^T$ is carried by $\mathcal{C}^T$ if $F^\alpha(x) \in \mathcal{C}^T(x)$ at parameter $\alpha(s)$ for all basis elements $x \in C_*^s$.

Acyclic Carriers

Acyclic

A chain complex C_* is acyclic if $\tilde{H}_*(C_*) = 0$
($\ker \partial_k = \operatorname{img} \partial_{k+1}$, or all cycles are boundaries).

β-Acyclic

Let $\beta : T \to T$ be a nondecreasing map. C_*^T is β-acyclic if every cycle in C_*^t has a boundary in $C_*^{\alpha(t)}$.

If C_*^T is β-acyclic, then its reduced persistent homology has no bars that survive a shift by β.

Acyclic Carriers

(α, β)-Acyclic Carrier

Let C_*^S, D_*^T be filtered chain complexes, $\mathcal{C}^T : C_*^S \to D_*^T$ be a filtered carrier, and $\alpha : S \to T$, $\beta : T \to T$ be non-decreasing maps. We say $\mathcal{C}^T$ is (α, β)-acyclic if $\mathcal{C}^T(x)$ is β-acyclic after $t = \alpha(s)$ for all $x \in C_*^s$ and for all $s \in S$. In the case where $\beta = \text{id}$, then we just say $\mathcal{C}^T$ is α-acyclic.

(After initial shift by α, the carrier is β-acyclic).

Filtered Acyclic Carrier Theorem

Theorem

Let $\mathcal{C}^T : C_*^S \to D_*^T$ be an (α, β)-acyclic carrier of filtered chain complexes, with S a strict total order with an initial object $0 \in S$. Let $L_*^S \subseteq C_*^S$ be a filtered sub-complex generated by a filtered sub-basis of C_*^S, and $\tilde{F}^\alpha : L_*^S \to D_*^T$ be an α-filtered chain map carried by $\mathcal{C}^T$. Then $\tilde{F}^\alpha$ extends to a filtered chain map $F^{\beta^k \circ \alpha} : C_*^S \to D_*^T$, where k is the maximal dimension of the chain map, and the extension is unique up to β-chain homotopy.

Proof Sketch I

We use induction on the dimension k of the map, and on the total order on S.

Base case: $s = 0, k = 0$

First, we start with $\tilde{F}_0^{0,\alpha(0)} : L_0^0 \to D_0^{\alpha(0)}$. From the acyclic carrier theorem, we can extend to a chain map $F_0^{0,\alpha(0)} \to C_0^0 \to D_0^{\alpha(0)}$.

Proof Sketch II

$s > 0, k = 0$

Now, let $s > 0$. Assume that we have extended F_0^α for all $r < s$ so that if $r' < r$,

$$F_0^{r,\alpha(r)} |_{C_*^{r'}} = F_0^{r',\alpha(r')} \quad (2)$$

Let $L_0'^S = L_0^S \cup \bigcup_{r<s} C_0^r$, and $\tilde{F}_0^\alpha$ denote the extended map up to all $r < s$. We can now apply the acyclic carrier theorem again to extend to $F^{s,\alpha(s)}$ to C_0^s.

Because S is a strict total order, eq. (2) continues to be satisfied because the function is extended on each basis element exactly once. By induction, we can extend to a map of 0-chains $F^\alpha : C_0^S \to D_0^T$.

Proof Sketch III

Start a β-chain homotopy.

Because the extension is not necessarily unique, suppose that F_0^α and G_0^α are both extensions of $\tilde{F}_0^\alpha$ carried by $\mathcal{C}$.

$\partial_0(F_0^\alpha - G_0^\alpha) = 0$, so can be expressed as the boundary of $K_0^{\beta\circ\alpha} : C_0^S \to D_1^T$ after shifting by an additional factor of β (since the image of the carrier is β-acyclic). This gives a β homotopy of 0-chain maps.

Proof Sketch IV

Extension to higher dimensions, $s = 0$

Suppose we have extended the map to dimension k:
$F_k : C_k^0 \to D_k^{\beta^k \circ \alpha(0)}$.

Let $x \in C_{k+1}$ be a basis element that we must extend at filtration parameter $s = 0$.

We need $\partial_{k+1} F_{k+1} x = F_k \partial_{k+1} x$. The image of the boundary $F_k \partial_{k+1} x$ lies in $D_k^{\beta^k \circ \alpha(0)}$, but since $\mathcal{C}$ is (α, β)-acyclic, the cycle need not have a boundary until we increase the filtration parameter T by another factor of β.

We can then choose some boundary y to be $F_{k+1}(x)$.

Proof Sketch V

Extension to higher dimensions, $s > 0$

Assume that so far we have satisfied for $r' < r < s$

$$F_{k+1}^{r,\beta^{k+1}\circ\alpha(r)} \mid_{C_k^{r'}} = F_{k+1}^{r',\beta^{k+1}\circ\alpha(r')} \tag{3}$$

and we have shifted the chain maps in lower dimensions via $F^{\beta^{k+1}\circ\alpha} = \iota^{\beta} F^{\beta^k\circ\alpha}$.

Let $x \in C_{k+1}$ via a basis element that we must extend at filtration parameter s. The image of the boundary $F_k\partial_{k+1}x$ lies in $D_k^{\beta^k\circ\alpha(s)}$, and we have already shifted the grade to $\beta^{k+1}\circ\alpha(s)$ at which point the cycle is a boundary of some $y \in D_{k+1}^{\beta^{k+1}\circ\alpha(s)}$ in $\mathcal{C}(x)$.

We can choose this y to be $F_{k+1}(x)$.

Proof Sketch VI

Extension of β-chain homotopy

Following a similar inductive argument, we can extend a β homotopy of extended chain maps $F_k^{\beta^k \circ \alpha}$, $G_k^{\beta^k \circ \alpha}$ to a β homotopy of $F_{k+1}^{\beta^{k+1} \circ \alpha}$ and $G_k^{\beta^{k+1} \circ \alpha}$, incurring an additional shift of β in each dimension.

Comments

To compute induced maps on homology in dimension k, it is only necessary to extend to dimension $k+1$. Do not need to incur additional shifts by β for higher dimensions.

In a variety of cases, $\beta = \text{id}$, or $\beta^k = \text{id}$ for $k \geq k_0$. This can happen if the carrier of chain complexes is obtained from filtered cell complex that becomes contractible.

If S is not a strict total ordering, then additional restrictions on the extension are needed.

Augmentation-Preserving Maps

We say a carrier $\mathcal{C} : C_* \to D_*$ is propoer with respect to a basis of D_* if $\mathcal{C}(x)$ is generated by a sub-basis of D_* for each x in the basis of C_*.

Proposition

Let $\mathcal{C} : C_*^S \to D_*^T$ be an (α, β)-acyclic carrier that is proper with respect to a T-filtered basis B_*^D of D_*. Then there exists a chain map $F_0^\alpha : C_0^S \to D_0^T$ carried by $\mathcal{C}$ which preserves the canonical augmentation $\epsilon : x \mapsto 1$ for basis elements $x \in C_0^S$.

Proposition

Suppose $F_*^\alpha, G_*^\alpha : C_*^S \to D_*^T$ are augmentation-preserving chain maps carried by an (α, β)-acyclic carrier $\mathcal{C}$. Then F_* and G_* are β-chain-homotopic.

Filtered Acylic Carriers to Interleavings

Proposition

Let $\mathcal{X}^S$ and $\mathcal{Y}^T$ be filtered cell complexes, and suppose that $\mathcal{C} : \mathcal{X}^S \to \mathcal{Y}^T$ is an α-acyclic carrier, $\mathcal{D} : \mathcal{Y}^T \to \mathcal{X}^S$ is a β-acyclic carrier, $\mathcal{A} \supseteq \mathcal{D} \circ \mathcal{C}$ is a $(\beta \circ \alpha)$-acyclic carrier that carries the inclusion map on $\mathcal{Y}^T$, and $\mathcal{B} \supseteq \mathcal{C} \circ \mathcal{D}$ is $(\alpha \circ \beta)$-acyclic and carries the inclusion map on $\mathcal{X}^S$. Then $H_q(\mathcal{X}^S)$ and $H_q(\mathcal{Y}^T)$ are (α, β)-interleaved for any $q = 0, 1, \ldots$.

Sketch: Construct augementation-preserving chain maps, and note that they are homotopic to inclusions, which are also augmentation-preserving.

Section 4

Vietoris-Rips Cover Complexes

Interleaving

We wish to compare to compare $\mathcal{R}(\mathbf{X}; r)$ and $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$. Because $\mathcal{R}(\mathbf{X}, \mathcal{U}; r) \subseteq \mathcal{R}(\mathbf{X}; r)$, we just need to worry about when we can construct a map f^{α} below

$$\begin{array}{ccc} \mathcal{R}(\mathbf{X}, \mathcal{U}; r) & \hookrightarrow & \mathcal{R}(\mathbf{X}, \mathcal{U}; \alpha(r)) \\ \downarrow & \nearrow^{f^{\alpha}} & \downarrow \\ \mathcal{R}(\mathbf{X}; r) & \hookrightarrow & \mathcal{R}(\mathbf{X}; \alpha(r)) \end{array}$$

Passing to homology, this will give an (id, α)-interleaving.

Carrier

We will focus on a carrier $\mathcal{C} : \mathcal{R}(\mathbf{X}; r) \to \mathcal{R}(\mathbf{X}, \mathcal{U}; r)$ generated from witness sets

$$\mathbf{X}(x_0, \ldots, x_k) = \{y \in \mathbf{X} \mid d(y, x_i) \leq d(x_0, \ldots, x_k)\ \forall i = 0, \ldots, k\}$$

and their union, denoted

$$\bar{\mathbf{X}}(x_0, \ldots, x_k) = \bigcup_{S \in \mathcal{P}(\{x_0, \ldots, x_k\})} \mathbf{X}(S)$$

We define the carrier $\mathcal{C} : \mathcal{R}(\mathbf{X}; r) \to \mathcal{R}(\mathbf{X}, \mathcal{U}; r)$ via

$$\mathcal{C} : (x_0, \ldots, x_k) \mapsto \langle \bar{\mathbf{X}}(x_0, \ldots, x_k) \rangle$$

Restriction of Cover

We also want to consider the restriction of the cover to $\bar{\mathbf{X}}(x_0, \ldots, x_k)$

$$\bar{\mathcal{U}}(x_0, \ldots, x_k) = \{V \cap \bar{\mathbf{X}}(x_0, \ldots, x_k) \mid V \in \bar{\mathcal{U}}, \bar{\mathbf{X}}(x_0, \ldots, x_k) \cap V \neq \emptyset\}$$

Regimes

We can show three regimes for α, determined by parameters $0 \leq R_1 \leq R_2 \leq R_3$.

1. For $r \leq R_1$, $\alpha = \mathrm{id}$.
2. For $r \leq R_2$, $\alpha \leq r \mapsto 2r$.
3. For $r \leq R_3$, $\alpha \leq r \mapsto 3r$.

These regimes are determined by properties of the cover $\mathcal{U}$.

Regime 1

Proposition

Let R_1 be the largest value so that if $d(x_0, \ldots, x_k) \leq R_1$ then there exists some $U \in \mathcal{U}$ so that $x_0, \ldots, x_k \in U$. Then for $r \leq R_1$,

$\mathcal{R}(\mathbf{X}, \mathcal{U}; r) = \mathcal{R}(\mathbf{X}; r)$ (they are $(\mathrm{id}, \mathrm{id})$-interleaved).

Proof: This follows from the definition of $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$.

Regime 2

Proposition

Let R_2 be the largest value so that if $d(x_0, \ldots, x_k) \leq R_2$ then $\mathbf{X}(x_0, \ldots, x_k) \cap V)$ is non-empty for each $V \in \bar{\mathcal{U}}(x_0, \ldots, x_k)$, and $\bar{\mathcal{U}}(x_0, \ldots, x_k)$ is acyclic.

Then for $r \leq R_2$, $H_k(\mathcal{R}(\mathbf{X}, \mathcal{U}; r))$ and $H_k(\mathcal{R}(\mathbf{X}; r))$ are (id, α)-interleaved where $\alpha : r \to 3r$.

Proof: In each $V \in \bar{\mathcal{U}}(x_0 \ldots, xk)$, this condition means that there is some $y \in V$ where $d(y, x_i) \leq r$ for all $i = 0, \ldots, k$. Any other $y' \in V$ has $d(y', x_i) \leq r$ for some $i \in \{0, \ldots, k\}$, so $\mathcal{R}(V; 2r)$ forms a cone with y by triangle inequality, and is constrctible. The carrier is then acyclic by the Nerve theorem.

Regime 3

Proposition

Let R_3 be the largest value so that if $d(x_0, \ldots, x_k) \leq R_3$ then $\mathcal{N}(\bar{\mathcal{U}}(x_0, \ldots, x_k))$ is acyclic. Then for $r \leq R_3$, $H_k(\mathcal{R}(\mathbf{X}, \mathcal{U}; r))$ and $H_k(\mathcal{R}(\mathbf{X}; r))$ are (id, α)-interleaved where $\alpha : r \to 3r$.

Proof: Now, we may not be able to cone with some $y \in V$ at parameter $2r$, but $\mathcal{R}(V; 3r)$ forms a clique (thus contractible) through the fact that for any $y, y' \in V$, $d(y, x_i), d(y', x_j) \leq r$ for some $i, j \in \{0, \ldots, k\}$ and triangle inequality:

$$d(y, y') \leq d(y, x_i) + d(x_i, x_j) + d(x_j, y') = r + r + r = 3r$$

The carrier is thus acyclic by the Nerve theorem.

Outside Regime 3

At some parameter $R_\star$, $\mathcal{R}(U; r)$ becomes acyclic for every set $U \in \mathcal{U}$ and all $r > R_\star$. This means $H_*(\mathcal{R}(\mathbf{X}, \mathcal{U}; r)) = H_*(\mathcal{N}(\mathcal{U}))$ for any $r > R_\star$.

Unless $\tilde{H}_*(\mathcal{N}(\mathcal{U})) = 0 = \tilde{H}_*(\mathcal{R}(\mathbf{X}; r))$, there is no interleaving beyond this point.

How to Choose a Cover?

In general, we may wish to choose a cover that increases R_1, R_2, and R_3 as much as possible, while not adding too many points to each set in $\mathcal{U}$.

1. To maximize R_1, want to include all r-nearest neighbors in some set
2. To maximize R_2, want to ensure that there are witnesses to simplices. May require sets covering large distances in sparse regions.
3. To maximize R_3, want to make $\mathcal{N}(\bar{\mathcal{U}}(x_0, \ldots, x_k))$ acyclic. Want sufficient overlap of sets in cover.

Landmarking Procedure

A heuristic way to produce a cover with the desired properties:

1. Obtain a nested sequence of landmarks with $\mathbf{X} = \mathbf{X}_n \supset \cdots \supset \mathbf{X}_1 \supset \mathbf{X}_0$. Take $n = i_0 > i_1 > \ldots$
2. Create covers of $\mathbf{X}_{i_j}$, $\mathcal{U}_{i_n}$, where each $U_\ell \in \mathcal{U}_{i_j}$ consists of points in $\mathbf{X}_{i_j}$ which have $x_\ell \in \mathbf{X}_{i_{n+1}}$ in their k-closest landmarks.
3. Take $\mathcal{U} = \bigcup_{i_0, i_1, \ldots} \mathcal{U}_{i_j}$

$\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$ obtained in this way is similar to a sparse filtration [Sheehy, 2013]. The main difference is that the longer edges are not re-weighted to tighten interleaving.

Landmark Cover

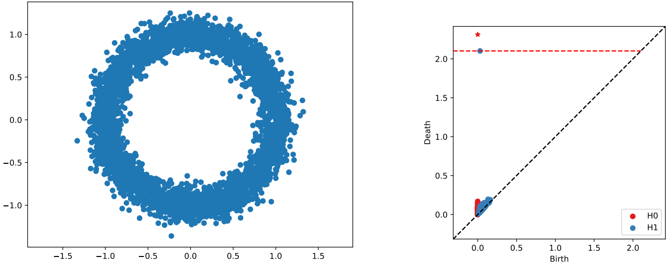

Figure: $n = 4000, i_j = n/(2^j), j = 0, 1, \ldots$. 191152 simplices in 2-skeleton. ≈ 0.5 seconds to compute in BATS.

Conclusion

Extension of acyclic carrier theorem to interleavings:

- Procedural way to obtain shift maps from correspondences
- Can be applied to a variety of situations (including non-simplicial)

Interleaving of $\mathcal{R}(\mathbf{X}; r)$ and $\mathcal{R}(\mathbf{X}, \mathcal{U}; r)$

- Quality of interleaving depends on cover/data
- Motivates r-NN, k-NN based landmarking procedures

What next?

- Use of covers can be quite general, application specific
- Algorithmic use of carriers
- Parallelization of PH for cover complexes

Acknowledgements

Funding:

- UnBox AI
- DARPA HR00112190040

People:

- Gunnar Carlsson
- Jonathan Taylor

References I

Cavanna, N. J. and Sheehy, D. R. (2018).
The generalized persistent nerve theorem.
Preprint: http://arxiv.org/abs/1807.07920.

Chazal, F., de Silva, V., and Oudot, S. (2014).
Persistence stability for geometric complexes.
Geometriae Dedicata, 173:193–214.

Chazal, F. and Oudot, S. Y. (2008).
Towards persistence-based reconstruction in euclidean spaces.
In *Proceedings of the twenty-fourth annual symposium on Computational geometry - SCG '08*, page 232. ACM Press.

Eilenberg, S. and Steenrod, N. E. (1952).
Foundations of Algebraic Topology.
Princeton Mathematical Series. Princeton University Press.

Govc, D. and Skraba, P. (2018).
An approximate nerve theorem.
Foundations of Computational Mathematics, 18:1245–1297.

Mosher, R. and Tangora, M. (1968).
Cohomology Operations and Applications in Homotopy Theory.
Harper's Series in Modern Mathematics. Harper & Row.

Munkres, J. R. (1984).
Elements of Algebraic Topology.
Addison-Wesley Publishing Company.

References II

Nelson, B. J. (2020).
Parameterized Topological Data Analysis.
Ph.D. dissertation, Stanford University.

Nelson, B. J. (2022).
Parameterized Vietoris-Rips Filtrations via Covers.
In preparation.

Sheehy, D. R. (2013).
Linear-size approximations to the vietoris–rips filtration.
Discrete & Computational Geometry, 49(4):778–796.

Singh, G., Memoli, F., and Carlsson, G. (2007).
Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition.
In Botsch, M., Pajarola, R., Chen, B., and Zwicker, M., editors, *Eurographics Symposium on Point-Based Graphics*. The Eurographics Association.

Yoon, H. R. (2018).
Cellular Sheaves And Cosheaves For Distributed Topological Data Analysis.
Ph.D. dissertation, University of Pennsylvania.