

HELMHOLTZ
MUNICH

AIH Institute of AI for Health

Topological Machine Learning: The (W)Hole Truth

Lecture 2

Bastian Rieck (@Pseudomanifold)

Preliminaries

Do you have feedback or any questions? Write to bastian.riek@helmholtz-muenchen.de or reach out to [@Pseudomanifold](https://twitter.com/Pseudomanifold) on Twitter. You can find the slides and additional information with links to more literature here:

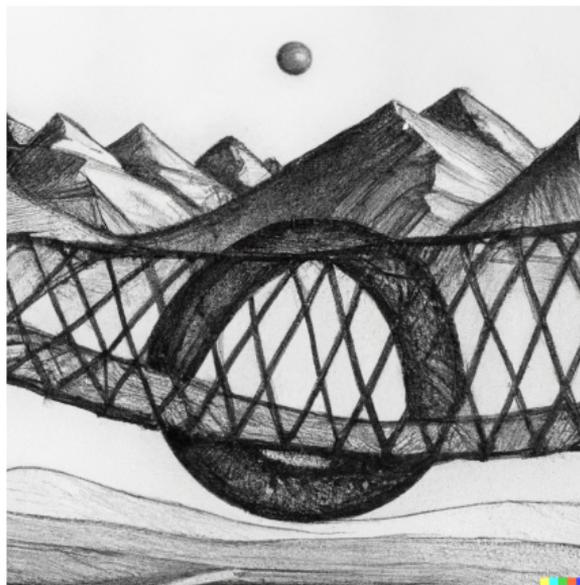
<https://heidelberg.topology.rocks>

Recap

- ☆ To distinguish between topological objects, we can use *Betti numbers*.
- ☆ Betti numbers count high-dimensional holes.
- ☆ Their calculation requires a simplicial complex and some linear algebra.

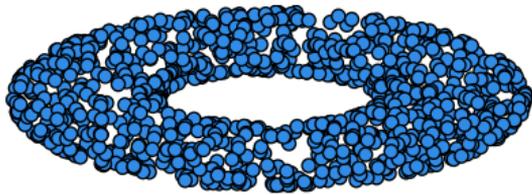
In this lecture

Going from theory to practice



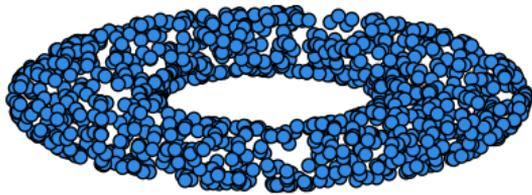
Real-world objects are typically not described in terms of simplicial complexes. How to bridge this gap?

Example

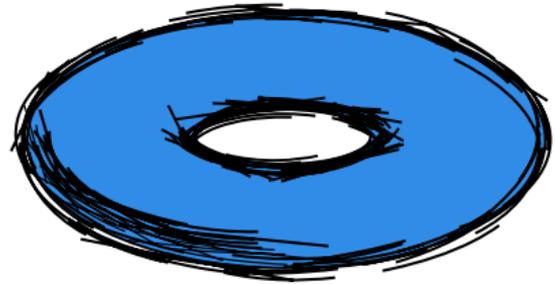


What we get

Example



What we get



What we see

From point clouds to simplicial complexes

Vietoris–Rips complex

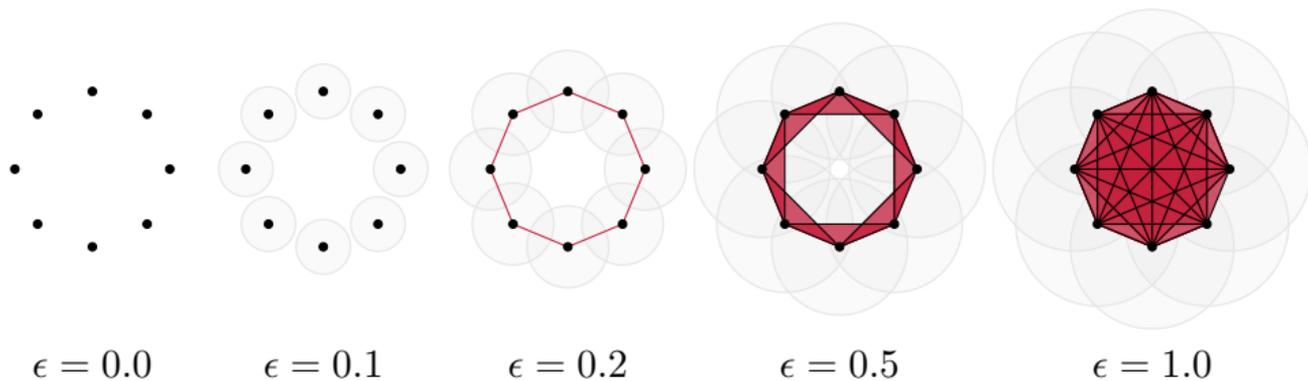
Given a set of points $\mathcal{X} = \{x_1, \dots, x_n\}$ and a metric dist such as the Euclidean distance, pick a threshold ϵ and build the Vietoris–Rips complex \mathcal{V}_ϵ defined as:

$$\mathcal{V}_\epsilon(\mathcal{X}) := \{\sigma \subseteq \mathcal{X} \mid \forall u, v \in \sigma : \text{dist}(u, v) \leq \epsilon\}$$

Equivalently, \mathcal{V}_ϵ contains all simplices whose *diameter* is less than or equal to ϵ .

Example

Vietoris–Rips construction



Draw Euclidean balls (circles) of diameter ϵ and create a k -simplex σ for each subset of $k + 1$ points that intersect pairwise.

Some details about this construction



- ☆ This construction dates back to a 1927 article by Leopold Vietoris,¹ who is shown on the left.

¹L. Vietoris, 'Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen', *Mathematische Annalen* 97.1, 1927, pp. 454–472

Some details about this construction



- ☆ This construction dates back to a 1927 article by Leopold Vietoris,¹ who is shown on the left.
- ☆ A 2010 paper by Afra Zomorodian² describes several construction algorithms.

¹L. Vietoris, 'Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen', *Mathematische Annalen* 97.1, 1927, pp. 454–472

²A. J. Zomorodian, 'Fast construction of the Vietoris–Rips complex', *Computers & Graphics* 34.3, 2010, pp. 263–271

Some details about this construction



- ☆ This construction dates back to a 1927 article by Leopold Vietoris,¹ who is shown on the left.
- ☆ A 2010 paper by Afra Zomorodian² describes several construction algorithms.
- ☆ The basic idea is to build higher-dimensional simplices *inductively* from lower-dimensional ones.

¹L. Vietoris, 'Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen', *Mathematische Annalen* 97.1, 1927, pp. 454–472

²A. J. Zomorodian, 'Fast construction of the Vietoris–Rips complex', *Computers & Graphics* 34.3, 2010, pp. 263–271

Some details about this construction



- ☆ This construction dates back to a 1927 article by Leopold Vietoris,¹ who is shown on the left.
- ☆ A 2010 paper by Afra Zomorodian² describes several construction algorithms.
- ☆ The basic idea is to build higher-dimensional simplices *inductively* from lower-dimensional ones.
- ☆ In the worst case, the Vietoris–Rips complex will contain all 2^n subsets of its underlying point cloud \mathcal{X} !

¹L. Vietoris, 'Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen', *Mathematische Annalen* 97.1, 1927, pp. 454–472

²A. J. Zomorodian, 'Fast construction of the Vietoris–Rips complex', *Computers & Graphics* 34.3, 2010, pp. 263–271

Interlude

Other simplicial complexes

☆ Alpha shapes and alpha complexes³

³H. Edelsbrunner, D. Kirkpatrick and R. Seidel, 'On the shape of a set of points in the plane', *IEEE Transactions on Information Theory* 29.4, 1983, pp. 551–559

Interlude

Other simplicial complexes

- ☆ Alpha shapes and alpha complexes³
- ☆ Dowker complexes⁴

³H. Edelsbrunner, D. Kirkpatrick and R. Seidel, 'On the shape of a set of points in the plane', *IEEE Transactions on Information Theory* 29.4, 1983, pp. 551–559

⁴C. H. Dowker, 'Homology Groups of Relations', *Annals of Mathematics* 56.1, 1952, pp. 84–95

Interlude

Other simplicial complexes

- ☆ Alpha shapes and alpha complexes³
- ☆ Dowker complexes⁴
- ☆ k -NN complexes⁵

³H. Edelsbrunner, D. Kirkpatrick and R. Seidel, 'On the shape of a set of points in the plane', *IEEE Transactions on Information Theory* 29.4, 1983, pp. 551–559

⁴C. H. Dowker, 'Homology Groups of Relations', *Annals of Mathematics* 56.1, 1952, pp. 84–95

⁵M. Q. Le and D. Taylor, 'Persistent Homology with k -nearest-neighbor Filtrations reveals Topological Convergence of PageRank', 2022, arXiv: 2206.04725 [math . AT]

Interlude

Other simplicial complexes

- ☆ Alpha shapes and alpha complexes³
- ☆ Dowker complexes⁴
- ☆ k -NN complexes⁵

³H. Edelsbrunner, D. Kirkpatrick and R. Seidel, 'On the shape of a set of points in the plane', *IEEE Transactions on Information Theory* 29.4, 1983, pp. 551–559

⁴C. H. Dowker, 'Homology Groups of Relations', *Annals of Mathematics* 56.1, 1952, pp. 84–95

⁵M. Q. Le and D. Taylor, 'Persistent Homology with k -nearest-neighbor Filtrations reveals Topological Convergence of PageRank', 2022, arXiv: 2206.04725 [math . AT]

Interlude

Other simplicial complexes

- ☆ Alpha shapes and alpha complexes³
- ☆ Dowker complexes⁴
- ☆ k -NN complexes⁵

Still a very active research area: how can we have sparse complexes whose worst-case complexity is improved?

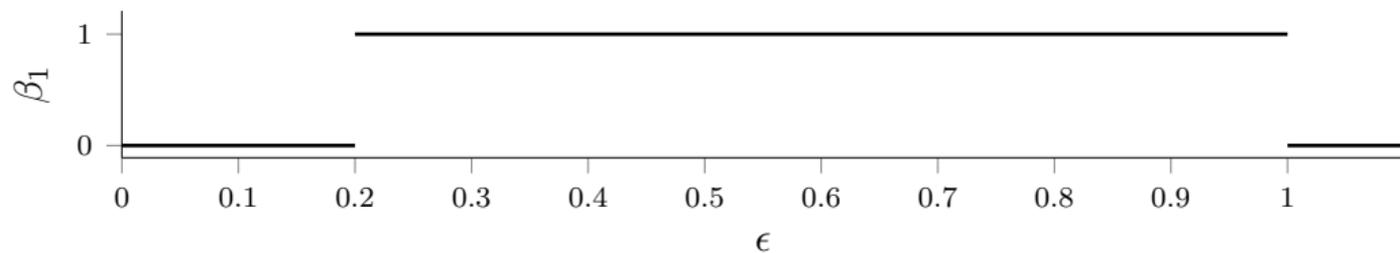
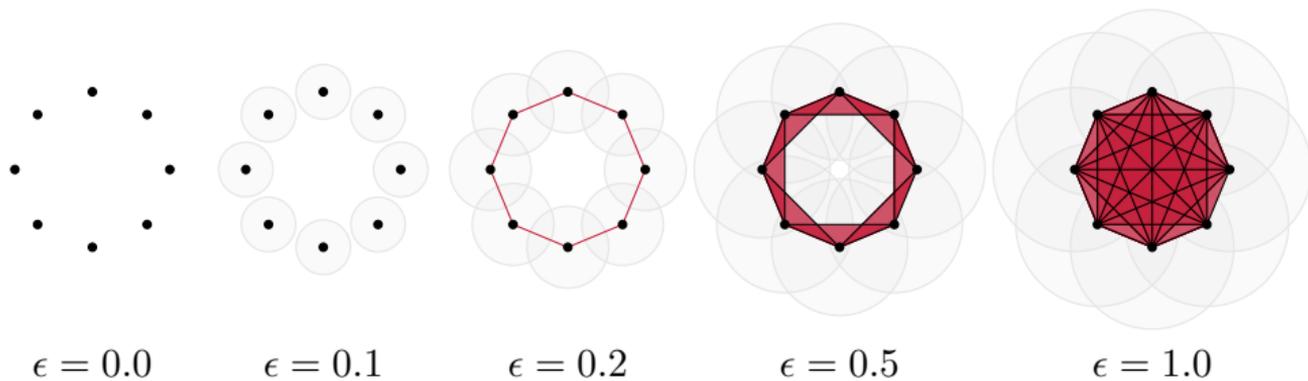
³H. Edelsbrunner, D. Kirkpatrick and R. Seidel, 'On the shape of a set of points in the plane', *IEEE Transactions on Information Theory* 29.4, 1983, pp. 551–559

⁴C. H. Dowker, 'Homology Groups of Relations', *Annals of Mathematics* 56.1, 1952, pp. 84–95

⁵M. Q. Le and D. Taylor, 'Persistent Homology with k -nearest-neighbor Filtrations reveals Topological Convergence of PageRank', 2022, arXiv: 2206.04725 [math . AT]

Example

The Betti numbers of a Vietoris–Rips complex



Some issues with this approach

☆ How to pick ϵ ?

Some issues with this approach

- ☆ How to pick ϵ ?
- ☆ There might not be one 'correct' value for ϵ .

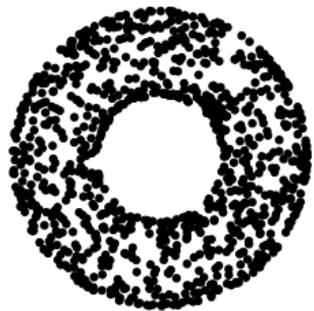
Some issues with this approach

- ☆ How to pick ϵ ?
- ☆ There might not be one 'correct' value for ϵ .
- ☆ Computationally inefficient; matrix reduction has to be performed for *every* simplicial complex.

Just calculate topological features for *all* possible scales and track their changes!

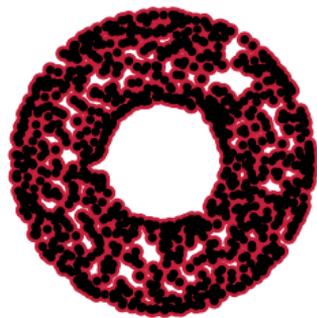
Intuition and preview

Go through all scales and *track* topological features



Intuition and preview

Go through all scales and *track* topological features



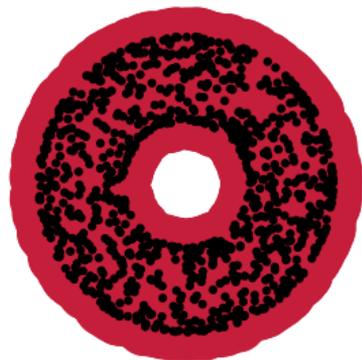
Intuition and preview

Go through all scales and *track* topological features



Intuition and preview

Go through all scales and *track* topological features



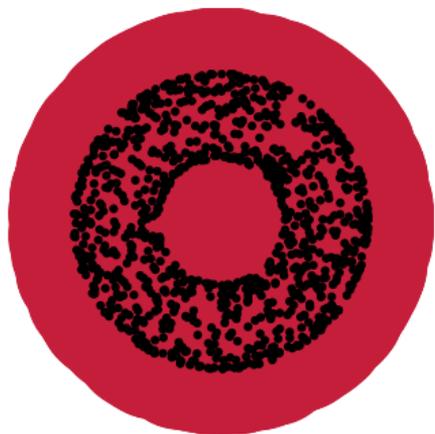
Intuition and preview

Go through all scales and *track* topological features



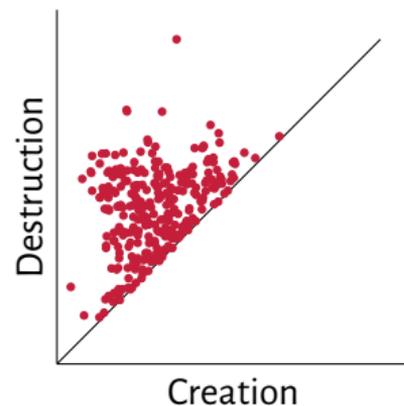
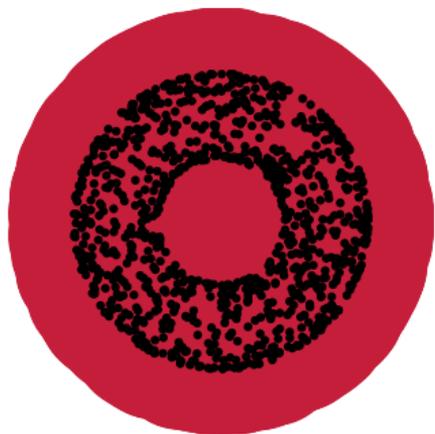
Intuition and preview

Go through all scales and *track* topological features



Intuition and preview

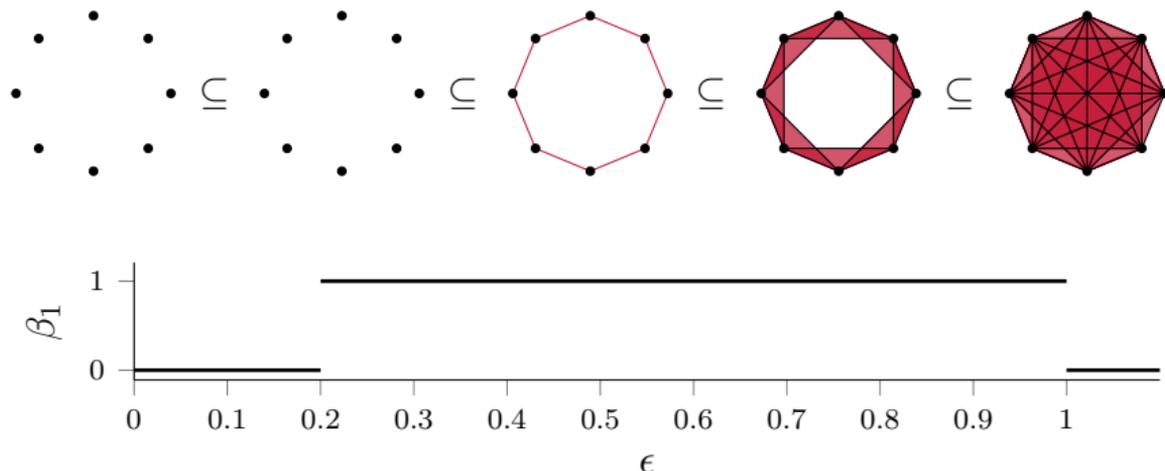
Go through all scales and *track* topological features



Nesting property of Vietoris–Rips complexes

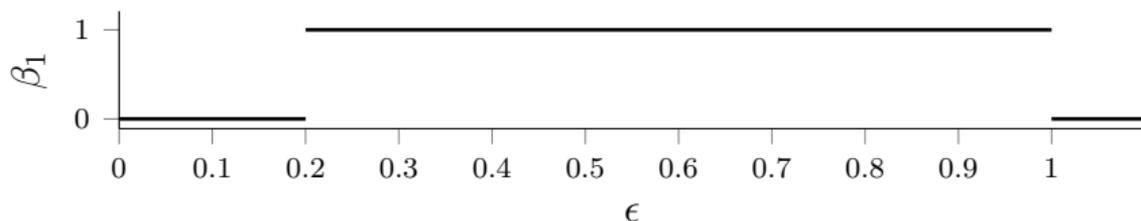
Given $\epsilon_1 \leq \epsilon_2$, we have $\mathcal{V}_{\epsilon_1} \subseteq \mathcal{V}_{\epsilon_2}$. This *nesting property* is the key concept for improving all calculations!

Example



Filtrations

A general concept in (computational) topology



The Betti number of the data *persists* over a range of the threshold parameter ϵ . To formalise this, assume that simplices in the Vietoris–Rips complex are added one after the other. This gives rise to a *filtration*, i.e.

$$\emptyset = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{n-1} \subseteq K_n = \mathcal{V}_\epsilon,$$

where each K_i is a valid simplicial subcomplex of \mathcal{V}_ϵ .

Chain complexes and filtrations

Connection to simplicial homology

Since $K_i \subseteq K_j$ for $i \leq j$, we obtain a sequence of homomorphisms connecting the homology groups of each simplicial complex, i.e.

$$f_p^{i,j} : H_p(K_i) \rightarrow H_p(K_j),$$

which in turn gives rise to a sequence of homology groups, i.e.

$$0 = H_p(K_0) \xrightarrow{f_p^{0,1}} H_p(K_1) \xrightarrow{f_p^{1,2}} \dots \xrightarrow{f_p^{n-2,n-1}} H_p(K_{n-1}) \xrightarrow{f_p^{n-1,n}} H_p(K_n) = H_p(\mathcal{V}_\epsilon),$$

with p denoting the dimension of the corresponding homology group.

Interlude

This is not ‘abstract nonsense’

The fact that we can *reformulate* the previously-seen concepts in the context of a filtration illustrates how *generic* this formulation is!

Sneak preview

- ☆ L. O’Bray*, **B. Rieck*** and K. Borgwardt, ‘Filtration Curves for Graph Representation’, *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, New York, NY, USA: Association for Computing Machinery, 2021, pp. 1267–1275
- ☆ K. Turner, S. Mukherjee and D. M. Boyer, ‘Persistent homology transform for modeling shapes and surfaces’, *Information and Inference: A Journal of the IMA* 3.4, 2014, pp. 310–344

Persistent homology group

Given two indices $i \leq j$, the p^{th} persistent homology group $H_p^{i,j}$ is defined as

$$H_p^{i,j} := Z_p(K_i) / (B_p(K_j) \cap Z_p(K_i)),$$

which contains all the homology classes of K_i that are still present in K_j .

Implication

We can calculate a new set of homology groups alongside the filtration and assign a ‘duration’ to each topological feature.

Persistent homology

Tracking of topological features

- ☆ *Creation* in K_i : $c \in H_p(K_i)$, but $c \notin H_p^{i-1,i}(K_i)$
- ☆ *Destruction* in K_j : c is created in K_i , with $f_p^{i,j-1}(c) \notin H_p^{i-1,j-1}$ and $f_p^{i,j}(c) \in H_p^{i-1,j}$

Persistence

The *persistence* of a class c that is created in K_i and destroyed in K_j is defined as

$$\text{pers}(c) := |w(j) - w(i)|,$$

where $w: \mathbb{Z} \rightarrow \mathbb{R}$ assigns each simplicial complex of the filtration a weight, such as an associated distance, or an index. Persistence thus measures the ‘scale’ at which a certain topological feature occurs.

Standard filtrations

The distance filtration

Given a distance metric dist , such as the Euclidean metric, the *distance filtration* assigns weights based on pairwise distances between points:

$$w(\sigma) := \begin{cases} 0 & \text{if } \sigma \text{ is a vertex} \\ \text{dist}(u, v) & \text{if } \sigma = \{u, v\} \\ \max_{\tau \subseteq \sigma} w(\tau) & \text{else} \end{cases}$$

Simplices need to be sorted in *ascending* order of their weights; in case of a tie, faces precede co-faces.

Persistent homology is capable of *preserving* distances under random projections.⁶

⁶D. R. Sheehy, 'The Persistent Homology of Distance Functions under Random Projection', *Proceedings of the 30th Annual Symposium on Computational Geometry*, 2014, pp. 328–334

Standard filtrations

The sublevel set filtration

Given a scalar function $f: \text{vert}(K) \rightarrow \mathbb{R}$ defined on the vertices of a simplicial complex, such as a temperature measurement, the *sublevel set filtration* propagates those weights through a simplicial complex:

$$w(\sigma) := \begin{cases} f(v) & \text{if } \sigma = \{v\} \\ \max_{\tau \subseteq \sigma} w(\tau) & \text{else} \end{cases}$$

Simplices need to be sorted in *ascending* order of their weights; in case of a tie, faces precede co-faces.

Conversely, one can calculate the *superlevel set filtration* by using min instead of max and sorting in simplices in *descending* order of their weights.

How to calculate this?

Example

Boundary matrix calculation alongside a filtration

a •

$$M = \begin{matrix} & \begin{matrix} a & b & c & ab & bc & ac & abc \end{matrix} \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} & \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix} \end{matrix}$$

Example

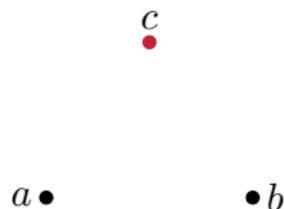
Boundary matrix calculation alongside a filtration

$a \bullet$ $\bullet b$

$$M = \begin{matrix} & a & b & c & ab & bc & ac & abc \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} & \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix} \end{matrix}$$

Example

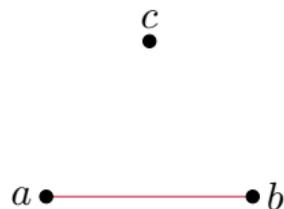
Boundary matrix calculation alongside a filtration



$$M = \begin{matrix} & a & b & c & ab & bc & ac & abc \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} & \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix} \end{matrix}$$

Example

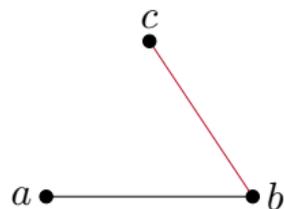
Boundary matrix calculation alongside a filtration



$$M = \begin{pmatrix} a & b & c & ab & bc & ac & abc \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix}$$

Example

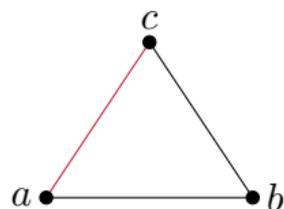
Boundary matrix calculation alongside a filtration



$$M = \begin{pmatrix} a & b & c & ab & bc & ac & abc \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix}$$

Example

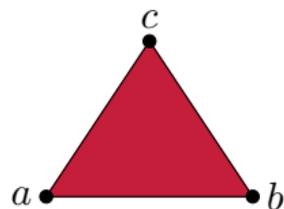
Boundary matrix calculation alongside a filtration



$$M = \begin{pmatrix} a & b & c & ab & bc & ac & abc \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix}$$

Example

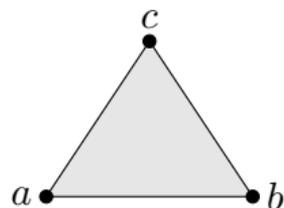
Boundary matrix calculation alongside a filtration



$$M = \begin{pmatrix} a & b & c & ab & bc & ac & abc \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix}$$

Example

Boundary matrix calculation alongside a filtration



$$M = \begin{pmatrix} a & b & c & ab & bc & ac & abc \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{matrix}$$

Boundary matrix reduction by column operations

Let M be a *boundary matrix*

for $i = 1$ **do**

while $\exists i' < i : \text{low}(i') = \text{low}(i) \neq 0$ **do**

$M(i) = M(i) \oplus M(i')$

end while

end for

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Boundary matrix reduction by column operations

Let M be a *boundary matrix*

for $i = 1$ **do**

while $\exists i' < i : \text{low}(i') = \text{low}(i) \neq 0$ **do**

$M(i) = M(i) \oplus M(i')$

end while

end for

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Boundary matrix reduction by column operations

Let M be a *boundary matrix*

for $i = 1$ **do**

while $\exists i' < i : \text{low}(i') = \text{low}(i) \neq 0$ **do**

$M(i) = M(i) \oplus M(i')$

end while

end for

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Boundary matrix reduction by column operations

Let M be a *boundary matrix*

for $i = 1$ **do**

while $\exists i' < i : \text{low}(i') = \text{low}(i) \neq 0$ **do**

$M(i) = M(i) \oplus M(i')$

end while

end for

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Boundary matrix reduction by column operations

Let M be a *boundary matrix*

for $i = 1$ **do**

while $\exists i' < i : \text{low}(i') = \text{low}(i) \neq 0$ **do**

$M(i) = M(i) \oplus M(i')$

end while

end for

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Using the reduced boundary matrix

$$\begin{array}{cccccc} a & b & c & ab & bc & ac & abc \\ \left(\begin{array}{cccccc} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) \begin{array}{l} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{array} \end{array}$$

- ☆ If column i is empty, then σ_i is a *positive* simplex that *creates* a topological feature.

Using the reduced boundary matrix

$$\begin{array}{cccccc} a & b & c & ab & bc & ac & abc \\ \left(\begin{array}{cccccc} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) \begin{array}{l} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{array} \end{array}$$

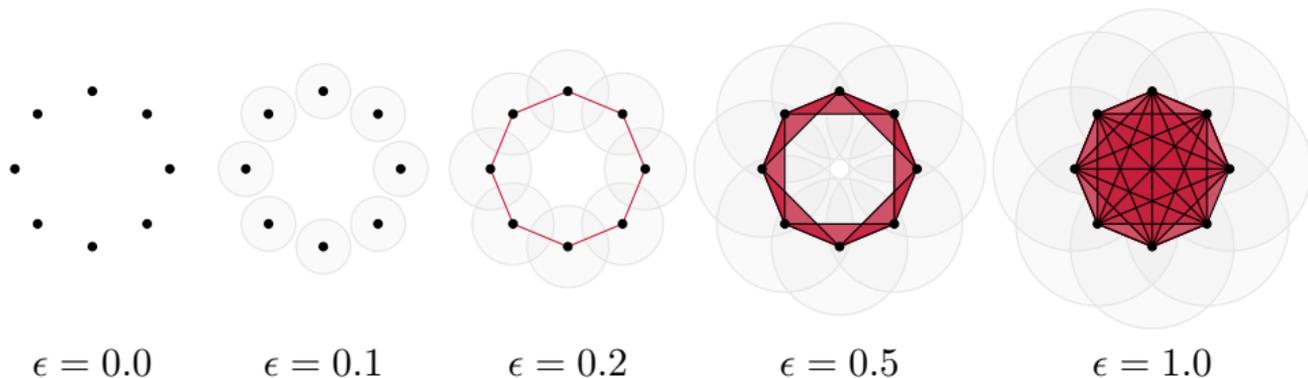
- ☆ If column i is empty, then σ_i is a *positive* simplex that *creates* a topological feature.
- ☆ If column j is non-empty with $\text{low}(j) = k$, then σ_j is a *negative* simplex that *destroys* the topological feature created by σ_k .

Using the reduced boundary matrix

$$\begin{array}{ccccccc} a & b & c & ab & bc & ac & abc \\ \left(\begin{array}{ccccccc} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) \begin{array}{l} a \\ b \\ c \\ ab \\ bc \\ ac \\ abc \end{array} \end{array}$$

- ☆ If column i is empty, then σ_i is a *positive* simplex that *creates* a topological feature.
- ☆ If column j is non-empty with $\text{low}(j) = k$, then σ_j is a *negative* simplex that *destroys* the topological feature created by σ_k .
- ☆ For example, simplex abc destroys the cycle created by ac .

Illustrative example I



Here, the topological feature is the circle that underlies that data. Since it persists from $\epsilon = 0.20$ to $\epsilon = 1.0$, its persistence is $\text{pers} = 1.0 - 0.20 = 0.80$.

Topological features and how to track them

Types of topological features

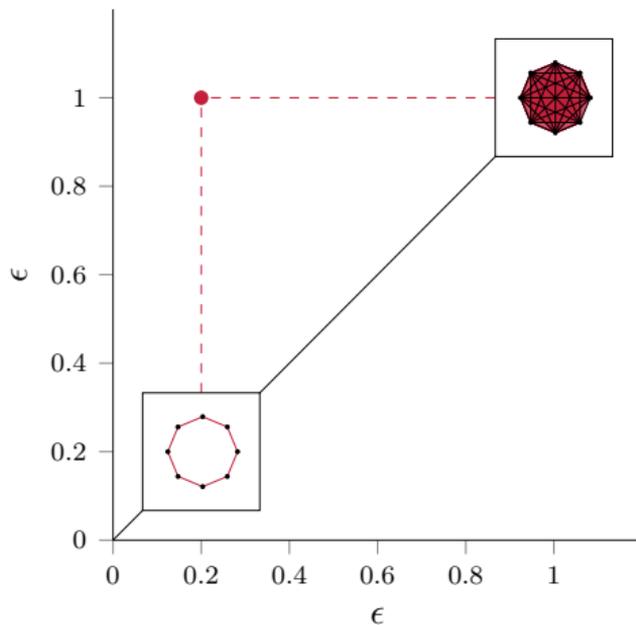
- ☆ Dimension 0: *connected components*
- ☆ Dimension 1: *cycles*
- ☆ Dimension 2: *voids*

Given a topological feature with associated simplicial complexes K_i and K_j , store $(w(i), w(j))$ in a *persistence diagram*.

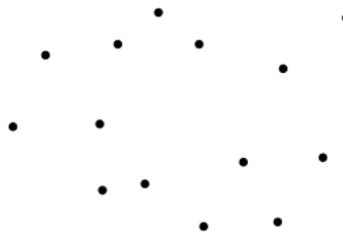


If a feature is *never* destroyed, we assign it a weight of ∞ .

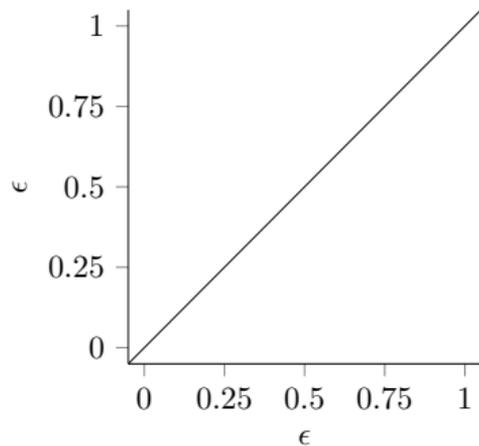
Example persistence diagram



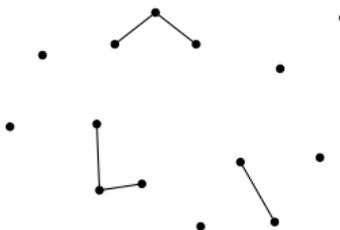
Illustrative example II



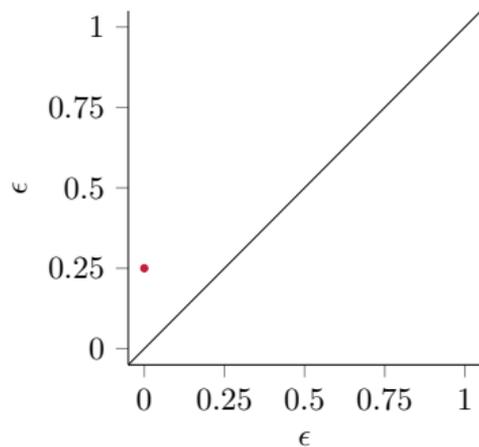
$\epsilon = 0.00$: 16 connected components



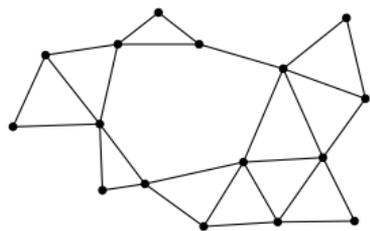
Illustrative example II



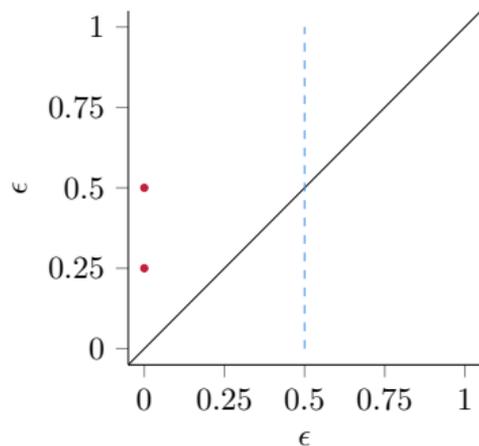
$\epsilon = 0.25$: 11 connected components



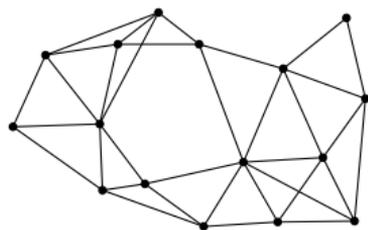
Illustrative example II



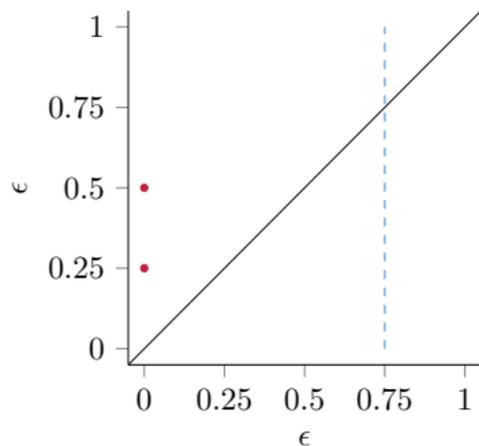
$\epsilon = 0.50$: 1 connected component, 12 cycles



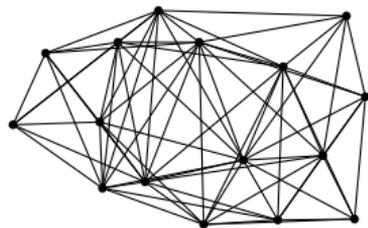
Illustrative example II



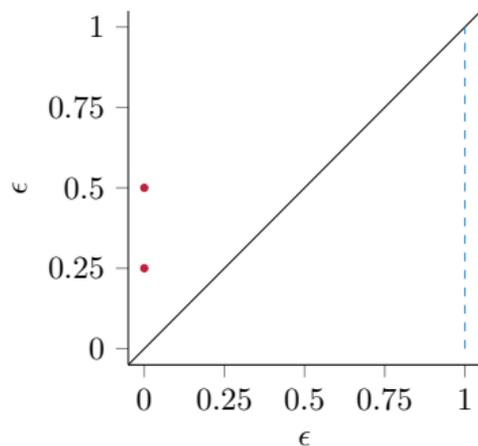
$\epsilon = 0.75$: 1 connected component, 19 cycles



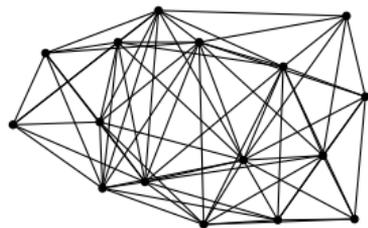
Illustrative example II



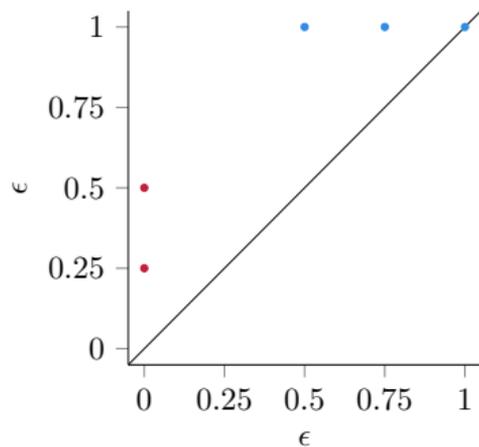
$\epsilon = 1.00$: 1 connected component, 57 cycles



Illustrative example II



$\epsilon = 1.00$: 1 connected component, 57 cycles



Calculation in practice

Ripsper: efficient computation of
Vietoris–Rips persistence barcodes

Ulrich Bauer*

August 8, 2019

We present an algorithm for the computation of Vietoris–Rips persistence barcodes and describe its implementation in the software Ripsper. The method relies on implicit representations of the coboundary operator and of the filtration order of the simplices, avoiding explicit construction and storage of the filtration coboundary matrix. Our implementation shows substantial improvements over previous software both in time and memory usage.

1 Introduction

Persistence homology is a central tool of computational topology and topological data analysis. It captures topological features of a dataset, a combinatorial family of growing topological spaces, and tracks their life span throughout the parameter range in the form of a collection of intervals called the persistence barcode. One of the most common geometric constructions for such a filtration is the Vietoris–Rips complex, which is constructed from a finite metric space by connecting a subset of the points to a simplex whenever the diameter of the subset does not exceed a specified threshold.

The computation of persistence homology has attracted strong interest in recent years [1, 2], with at least 10 different implementations publicly available to date [3–14, 15, 17, 21, 24, 25, 26, 30, 31, 32–40]. Over the years, dramatic improvements in performance have been achieved, as demonstrated in recent benchmarks [20].

The predominant approach to persistence computation consists of two steps: the generation of a filtration boundary matrix, and the computation of persistence barcodes using matrix reduction. Among the latter codes for the matrix reduction step is PHAT [3], which has been created with the goal of assessing and understanding the relative efficiency of the various optimizations proposed in the literature for the matrix reduction problem. In the course of that project, it became evident that often the construction of the filtration boundary matrix becomes the bottleneck for the computation of Vietoris–Rips barcodes.

The approach followed in Ripsper [1] is to avoid the construction and storage of the filtration boundary matrix as a whole, discarding and recomputing parts of it when necessary. In particular, instead of representing the coboundary map explicitly as a matrix data structure, it is represented as an algorithmic operator, accumulating the coboundary of a simplex whenever needed. The filtration is specified using another algorithmic operator for computing simplices with respect to their appearance in the filtration order. The overall reduction for persisting this strategy was found to reduce the memory usage. Perhaps surprisingly, the method turns out to be substantially faster than accessing the coboundary from memory. This effect can be explained by the fact that on current computer architectures, memory access is much more expensive than elementary arithmetic operations.

The computation of persistence homology implemented in Ripsper is based on matrix reduction and uses four key optimizations in order to achieve an efficient implementation, two of which have been proposed in the literature before. While one implementation is specific to Vietoris–Rips filtrations, these ideas are also applicable to persistence computations for other filtrations as well.

Clearing zero columns. The standard matrix reduction algorithm does not make use of the special structure of boundary matrix ∂_k , which satisfies $\partial_k^2 = 0$, i.e., boundaries are always cycles. Spurring this observation leads to a large number of unnecessary matrix operations in the matrix reduction, computing a large number of cycles that get re-used subsequently. The clearing optimization takes called the zero optimization, suggested by Chen and Kerber [1], avoids the computation of these cycles.

*Former teaching assistant of the course <https://www.mpl.mpg.de/~bau1>



Calculations are bounded by matrix reduction complexity.

Calculation in practice

Ripser: efficient computation of
Vietoris–Rips persistence barcodes

Ulrich Bauer*

August 8, 2019

We present an algorithm for the computation of Vietoris–Rips persistence barcodes and describe its implementation in the software Ripser. The method relies on implicit representations of the coboundary operator and of the filtration order of the simplices, avoiding explicit construction and storage of the filtration coboundary matrix. Our implementation shows substantial improvements over previous software both in time and memory usage.

1 Introduction

Persistence homology is a central tool of computational topology and topological data analysis. It captures topological features of a dataset, a continuous family of growing topological spaces, and tracks their life span throughout the parameter range in the form of a collection of intervals called the persistence barcode. One of the most common geometric constructions for such a filtration is the Vietoris–Rips complex, which is constructed from a finite metric space by connecting a subset of the points by a simplex whenever the diameter of the subset does not exceed a specified threshold.

The computation of persistence homology has attracted strong interest in recent years [1, 2], with at least 10 different implementations publicly available as of June 15, 17, 21, 24, 28, 30, 31, 1–10. Over the years, dramatic improvements in performance have been achieved, as demonstrated in recent benchmarks [3].

The predominant approach to persistence computation consists of two steps: the generation of a filtration boundary matrix, and the computation of persistence barcodes using matrix reduction. Among the latter codes for the matrix reduction step is PHAT [3], which has been created with the goal of assessing and understanding the relative efficiency of the various optimizations proposed in the literature for the matrix reduction problem. In the course of that project, it became evident that often the construction of the filtration boundary matrix becomes the bottleneck for the computation of Vietoris–Rips barcodes.

The approach followed in Ripser [1] is to avoid the construction and storage of the filtration boundary matrix as a whole, discarding and recomputing parts of it when necessary. In particular, instead of representing the coboundary map explicitly as a matrix data structure, it is represented as an algorithmic operator, accumulating the coboundary of a simplex whenever needed. The filtration is specified using an order algorithm, operator for computing simplices with respect to their appearance in the filtration order. The central motivation for pursuing this strategy was partly to reduce the memory usage. Perhaps surprisingly, the method turned out to be substantially faster than accessing the coboundary from memory. This effect can be explained by the fact that an extreme complex architecture, namely access to much more expensive than elementary arithmetic operations.

The computation of persistence barcodes implemented in Ripser is based on matrix reduction and uses four key optimizations in order to achieve an efficient implementation, two of which have been proposed in the literature before. While our implementation is specific to Vietoris–Rips filtrations, these ideas are also applicable to persistence computations for other filtrations as well.

Clearing Mark cycles. The standard matrix reduction algorithm does not make use of the special structure of boundary matrix ∂ , which satisfies $\partial^2 = 0$, i.e., boundaries are always cycles. Spurring this observation leads to a large number of unnecessary matrix operations in the matrix reduction, computing a large number of cycles that get re-used subsequently. The clearing optimization takes called the *clear optimization*, suggested by Chen and Kerber [4], avoids the computation of these cycles.

*Formerly University of Mainz (TUO), Germany. ulrich.bauer@uni-mainz.de



Calculations are bounded by matrix reduction complexity.



Smart implementations make a difference.

Calculation in practice

Ripser: efficient computation of
Vietoris–Rips persistence barcodes

Ulrich Bauer*

August 8, 2019

We present an algorithm for the computation of Vietoris–Rips persistence barcodes and describe its implementation in the software Ripser. The method relies on implicit representations of the cohomology operator and of the filtration order of the simplices, avoiding explicit construction and storage of the filtration cohomology matrix. Our implementation shows substantial improvements over previous software both in time and memory usage.

1 Introduction

Persistence homology is a central tool of computational topology and topological data analysis. It captures topological features of a filtration, a countable family of growing topological spaces, and tracks their life span throughout the parameter range in the form of a collection of intervals called the persistence barcode. One of the most common geometric constructions for such a filtration is the Vietoris–Rips complex, which is constructed from a finite metric space by connecting a subset of the points to a simplex whenever the diameter of the subset does not exceed a specified threshold.

The computation of persistence homology has attracted strong interest in recent years [1, 2], with at least 10 different implementations publicly available in date [3–14, 15, 17, 21, 24, 25, 26, 30, 31, 32–40]. Over the years, dramatic improvements in performance have been achieved, as demonstrated in recent benchmarks [21].

The predominant approach to persistence computation consists of two steps: the generation of a filtration boundary matrix, and the computation of persistence barcodes using matrix reduction. Among the latter codes for the matrix reduction step PHAT [3], which has been created with the goal of assessing and understanding the stability and accuracy of the various optimization proposals in the literature for the matrix reduction algorithm. In the course of that project, it became evident that often the construction of the filtration boundary matrix becomes the bottleneck for the computation of Vietoris–Rips barcodes.

The approach followed in Ripser [1] is to avoid the construction and storage of the filtration boundary matrix as a whole, discarding and recomputing parts of it when necessary. In particular, instead of representing the cohomology map explicitly as a matrix data structure, it is represented as an algebraic operator, accumulating the cohomology of a simplex whenever needed. The filtration is specified using another algebraic operator for computing simplices with respect to their appearance in the filtration order. The central motivation for pursuing this strategy was partly to reduce the memory usage. Perhaps surprisingly, the modified turns out to be substantially faster than accessing the cohomology from memory. This effect can be explained by the fact that on certain complex architectures, memory access is much more expensive than elementary arithmetic operations.

The computation of persistence homology implemented in Ripser is based on matrix reduction and uses four key optimizations in order to achieve an efficient implementation, two of which have been proposed in the literature before. While one implementation is specific to Vietoris–Rips filtrations, these ideas are also applicable to persistence computations for other filtrations as well.

Clearing Mark cycles. The standard matrix reduction algorithm does not make use of the special structure of boundary matrix B , which satisfies $B^2 = 0$, i.e., boundaries are always cycles. Spurring this observation leads to a large number of unnecessary matrix operations in the matrix reduction, computing a large number of cycles that get re-used subsequently. The clearing operation takes called the *clear operation*, suggested by Chen and Kerber [1], avoids the computation of these cycles.

*Former teaching assistant of the course <https://www.mpi-inf.mpg.de/~ulbau/>.

- ✧ Calculations are bounded by matrix reduction complexity.
- ✧ Smart implementations make a difference.
- ✧ U. Bauer’s Ripser⁷ is extremely fast.

⁷U. Bauer, ‘Ripser: efficient computation of Vietoris–Rips persistence barcodes’, *Journal of Applied and Computational Topology*, 2021

Calculation in practice

Ripser: efficient computation of
Vietoris–Rips persistence barcodes

Ulrich Bauer*

August 8, 2019

We present an algorithm for the computation of Vietoris–Rips persistence barcodes and describe its implementation in the software Ripser. The method relies on implicit representations of the cohomology operator and of the filtration order of the simplices, avoiding explicit construction and storage of the filtration cohomology matrix. Our implementation shows substantial improvements over previous software both in time and memory usage.

1 Introduction

Persistence homology is a central tool of computational topology and topological data analysis. It captures topological features of a filtration, a countable family of growing topological spaces, and tracks their life span throughout the persistence range in the form of a collection of intervals called the persistence barcode. One of the most common geometric constructions for such a filtration is the Vietoris–Rips complex, which is constructed from a finite metric space by connecting a subset of the points to a simplex whenever the diameter of the subset does not exceed a specified threshold.

The computation of persistence homology has attracted strong interest in recent years [1, 2], with at least 10 different implementations publicly available in date [3–14, 15, 17, 21, 24, 25, 26, 30, 31, 32–40]. Over the years, dramatic improvements in performance have been achieved, as demonstrated in recent benchmarks [21].

The predominant approach to persistence computation consists of two steps: the generation of a filtration boundary matrix, and the computation of persistence barcodes using matrix reduction. Among the latter codes for the matrix reduction step is PHAT [3], which has been created with the goal of assessing and understanding the relative efficiency of the various optimizations proposed in the literature for the matrix reduction algorithm. In the course of that project, it became evident that often the construction of the filtration boundary matrix becomes the bottleneck for the computation of Vietoris–Rips barcodes.

The approach followed in Ripser [1] is to avoid the construction and storage of the filtration boundary matrix as a whole, discarding and recomputing parts of it when necessary. In particular, instead of representing the cohomology map explicitly as a matrix data structure, it is represented as an algebraic operator, accumulating the cohomology of a simplex whenever needed. The filtration is specified using another algebraic operator for computing simplices with respect to their appearance in the filtration order. The central motivation for pursuing this strategy was partly to reduce the memory usage. Perhaps surprisingly, the modified turns out to be substantially faster than accessing the cohomology from memory. This effect can be explained by the fact that on certain complex architectures, memory access is much more expensive than elementary arithmetic operations.

The computation of persistence homology implemented in Ripser is based on matrix reduction and uses four key optimizations in order to achieve an efficient implementation, two of which have been proposed in the literature before. While one implementation is specific to Vietoris–Rips filtrations, these ideas are also applicable to persistence computations for other filtrations as well.

Clearing Mark cycles. The standard matrix reduction algorithm does not make use of the special structure of boundary matrix B , which satisfies $B^2 = 0$, i.e., boundaries are always cycles. Spurring this observation leads to a large number of unnecessary matrix operations in the matrix reduction, computing a large number of cycles that get re-used subsequently. The clearing operation takes called the *clear operation*, suggested by Chen and Kerber [1], avoids the computation of these cycles.

*Former teaching assistant of the course <https://www.mpi-inf.mpg.de/~ulrich>.



Calculations are bounded by matrix reduction complexity.



Smart implementations make a difference.



U. Bauer's Ripser⁷ is extremely fast.

⁷U. Bauer, 'Ripser: efficient computation of Vietoris–Rips persistence barcodes', *Journal of Applied and Computational Topology*, 2021

Calculation in practice

Ripser: efficient computation of
Vietoris–Rips persistence barcodes

Ulrich Bauer*

August 8, 2019

We present an algorithm for the computation of Vietoris–Rips persistence barcodes and describe its implementation in the software Ripser. The method relies on implicit representations of the cohomology operator and on the filtration order of the simplices, avoiding explicit construction and storage of the filtration cohomology matrix. Our implementation shows substantial improvements over previous software both in time and memory usage.

1 Introduction

Persistence homology is a central tool of computational topology and topological data analysis. It captures topological features of a dataset, a continuous family of growing topological spaces, and tracks their life span throughout the parameter range in the form of a collection of intervals called the persistence barcode. One of the most common geometric constructions for such a filtration is the Vietoris–Rips complex, which is constructed from a finite metric space by connecting a subset of the points by a simplex whenever the diameter of the subset does not exceed a specified threshold.

The computation of persistence homology has attracted strong interest in recent years [1, 2], with at least 10 different implementations publicly available in date [3–12, 15, 17, 21, 24, 25, 26, 30, 31, 32–40]. Over the years, dramatic improvements in performance have been achieved, as demonstrated in recent benchmarks [20].

The predominant approach to persistence computation consists of two steps: the generation of a filtration boundary matrix, and the computation of persistence barcodes using matrix reduction. Among the latter codes for the matrix reduction step is PHAT [3], which has been created with the goal of assessing and understanding the stability and complexity of the various optimizations proposed in the literature for the matrix reduction algorithm. In the course of that project, it became evident that often the construction of the filtration boundary matrix becomes the bottleneck for the computation of Vietoris–Rips barcodes.

The approach followed in Ripser [1] is to avoid the construction and storage of the filtration boundary matrix as a whole, discarding and recomputing parts of it when necessary. In particular, instead of representing the cohomology map explicitly as a matrix data structure, it is represented as an algorithmic operator, accumulating the cohomology of a simplex whenever needed. The filtration is specified using another algorithmic operator for computing simplices with respect to their appearance in the filtration order. The central motivation for pursuing this strategy was partly to reduce the memory usage. Perhaps surprisingly, the modifications did not substantially increase the complexity of the cohomology computation. This effect can be explained by the fact that on various computer architectures, memory access is much more expensive than elementary arithmetic operations.

The computation of persistence homology implemented in Ripser is based on matrix reduction and uses two key optimizations in order to achieve an efficient implementation, two of which have been proposed in the literature before. While one implementation is specific to Vietoris–Rips filtrations, these ideas are also applicable to persistence computations for other filtrations as well.

Clearing non-zero entries. The standard matrix reduction algorithm does not make use of the special structure of boundary matrix B , which satisfies $B^2 = 0$, i.e., boundaries are always cycles. Sparsing this matrix leads to a large number of unnecessary matrix operations in the matrix reduction, computing a large number of cycles that get re-used subsequently. The clearing operations take called the zero operation, suggested by Chen and Kerber [1], avoids the computation of these cycles.

*Former teaching assistant of the course <https://www.mpi-inf.mpg.de/~ulbau/>

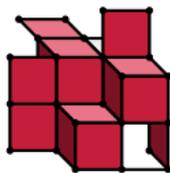
- ✧ Calculations are bounded by matrix reduction complexity.
- ✧ Smart implementations make a difference.
- ✧ U. Bauer’s Ripser⁷ is extremely fast.

Speed-ups involve different orderings for the column reduction steps, using implicit representations (which are more efficient than explicit ones), and much more.

⁷U. Bauer, ‘Ripser: efficient computation of Vietoris–Rips persistence barcodes’, *Journal of Applied and Computational Topology*, 2021

Topological features of non-simplicial domains

Persistent homology is *not* restricted to the ‘triangular’ setting. It is also possible to define a filtration over cubical domains.⁸



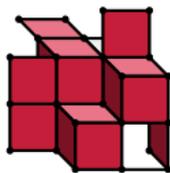
Potential data sources

☆ Images

⁸The image of the ‘cubical complex’ is modelled after R. Christ’s excellent book ‘Elementary Applied Topology’

Topological features of non-simplicial domains

Persistent homology is *not* restricted to the ‘triangular’ setting. It is also possible to define a filtration over cubical domains.⁸



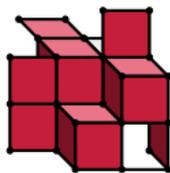
Potential data sources

- ☆ Images
- ☆ Finite element simulations

⁸The image of the ‘cubical complex’ is modelled after R. Ghrist’s excellent book ‘Elementary Applied Topology’

Topological features of non-simplicial domains

Persistent homology is *not* restricted to the ‘triangular’ setting. It is also possible to define a filtration over cubical domains.⁸



Potential data sources

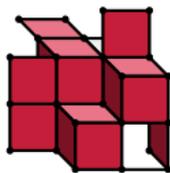
- ☆ Images
- ☆ Finite element simulations
- ☆ Voxel spaces⁹

⁸The image of the ‘cubical complex’ is modelled after R. Christ’s excellent book ‘Elementary Applied Topology’

⁹**B. Rieck*** et al., ‘Uncovering the Topology of Time-Varying fMRI Data using Cubical Persistence’, *Advances in Neural Information Processing Systems*, ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan and H. Lin, vol. 33, Curran Associates, Inc., 2020, pp. 6900–6912, arXiv: 2006.07882 [q-bio.NC]

Topological features of non-simplicial domains

Persistent homology is *not* restricted to the ‘triangular’ setting. It is also possible to define a filtration over cubical domains.⁸



Potential data sources

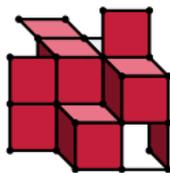
- ☆ Images
- ☆ Finite element simulations
- ☆ Voxel spaces⁹

⁸The image of the ‘cubical complex’ is modelled after R. Christ’s excellent book ‘Elementary Applied Topology’

⁹**B. Rieck*** et al., ‘Uncovering the Topology of Time-Varying fMRI Data using Cubical Persistence’, *Advances in Neural Information Processing Systems*, ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan and H. Lin, vol. 33, Curran Associates, Inc., 2020, pp. 6900–6912, arXiv: 2006.07882 [q-bio.NC]

Topological features of non-simplicial domains

Persistent homology is *not* restricted to the ‘triangular’ setting. It is also possible to define a filtration over cubical domains.⁸



Potential data sources

- ☆ Images
- ☆ Finite element simulations
- ☆ Voxel spaces⁹

The nice thing is that all considerations and concepts apply virtually unchanged!

⁸The image of the ‘cubical complex’ is modelled after R. Ghrist’s excellent book ‘Elementary Applied Topology’

⁹**B. Rieck*** et al., ‘Uncovering the Topology of Time-Varying fMRI Data using Cubical Persistence’, *Advances in Neural Information Processing Systems*, ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan and H. Lin, vol. 33, Curran Associates, Inc., 2020, pp. 6900–6912, arXiv: 2006.07882 [q-bio.NC]

Current research directions

Properties of filtrations

Can we find filtrations that are more *robust* to noise, easier to calculate, and more expressive?¹⁰

Sparse filtrations

Simplices are not equally important. Can we find *sparse* filtrations consisting of fewer simplices but with essentially the same topological features?¹¹

¹⁰H. Anai et al., 'DTM-Based Filtrations', *Topological Data Analysis*, ed. by N. A. Baas, G. E. Carlsson, G. Quick, M. Szymik and M. Thauale, Cham, Switzerland: Springer, 2020, pp. 33–66

¹¹D. R. Sheehy, 'Linear-size approximations to the Vietoris–Rips filtration', *Discrete & Computational Geometry* 49.4, 2013, pp. 778–796

Take-away messages

- ☆ Point clouds can be converted into simplicial complexes.
- ☆ Persistent homology is the multi-scale equivalent of simplicial homology.
- ☆ The calculation of persistent homology *also* boils down to linear algebra.
- ☆ Filtrations are the key for tracking topological features.