Persistent Homology

Francis Lazarus and Arnaud de Mesmay

January 9, 2017

Contents

The theory of persistent homology developed from 2000, motivated by practical problems related to approximation and reverse engineering [Rob99, ELZ00]. The main objective is to infer the topology of an object given by a finite cloud of points that approximates the object. In a typical application one is given a sampling P of the surface S of a manufactured object captured by some probing device. The question is to recover topological invariants (e.g., the number of connected components) of S with the sole knowledge of P. Although S and P may look very different, their ε-neighborhood have a similar topology for an appropriate range of ε. Recall that ε-neighborhood of an object is the union of balls of radius ε centered at every point of the object. See Figure 1. This crucial observation is the basis of persistent homology. Since

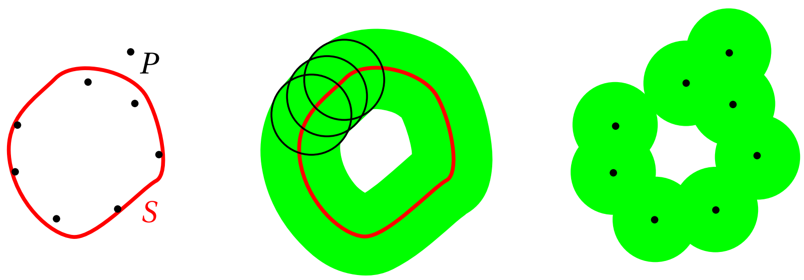

Figure 1: Left, an approximate sampling of a curve S. Middle, the ε-neighborhood of S. Right, the ε-neighborhood of P.

the correct range of ε is unknown and depends on the density of the sampling with

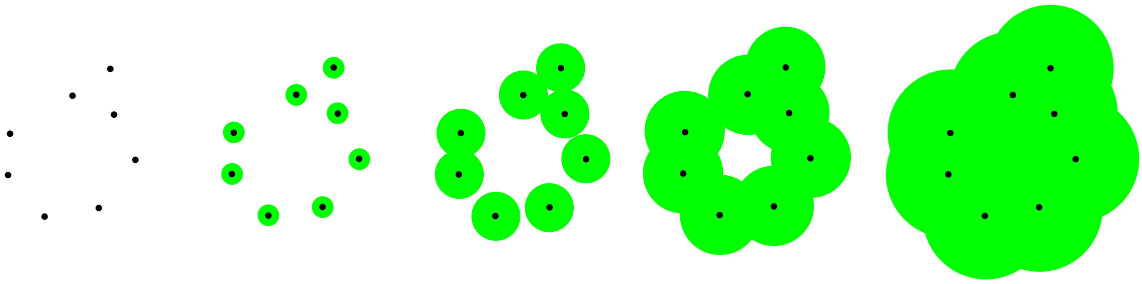

Figure 2: As ε increases, the topology of the ε-neighborhoods of P changes.

respect to S, one is led to study the topology of the whole sequence of ε-neighborhoods of P for ε ranging from zero to infinity. See Figure 2. Note that $\varepsilon < \eta$ implies the inclusion of the ε-neighborhood in the η-neighborhood. Such a nested sequence of spaces is called a **filtration**. By applying the homology functor, each inclusion $X \subset Y$ in the filtration induces a linear map $H_*(X) \to H_*(Y)$. The idea of persistent homology is to apply the homology functor to the filtration and study the resulting sequence of maps as a whole rather than the homology of each space individually. This sequence of maps not only provides topological information on each space in the filtration but also indicates how the spaces are nested. As a simple example consider the inclusions of a circle in a 2-dimensional torus as on Figure 3.

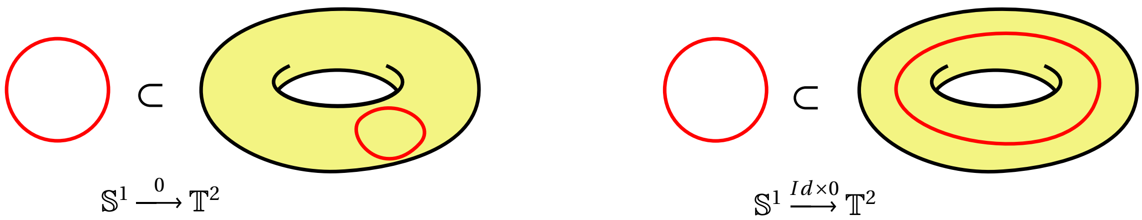

Figure 3: The circle may be included as a zero homologous cycle (left) or a non-trivial cycle (right) in the torus. While the induced maps in homology have the same domain and codomain, the maps themselves are distinct.

1 Persistence Modules

For computational reasons we shall only consider homology with coefficients in a field $\mathbb{F}$. Hence, a filtration $X_1 \subset X_2 \subset \cdots \subset X_n$ gives rise to a sequence $H_*(X_1) \to H_*(X_2) \to \cdots \to H_*(X_n)$ of linear maps between the vector spaces $H_*(X_i)$. In general, a sequence of linear maps between spaces indexed by an ordered set (typically $[1,n]$, or a subset of $\mathbb{R}$) is called a **persistence module**. The persistence modules over a fixed set of indices form a category. Here, taking $[1,n]$ as indices, a morphism between persistence modules $(f_i : E_i \to E_{i+1})_{1\le i<n}$ and $(g_i : F_i \to F_{i+1})_{1\le i<n}$ is a sequence of linear maps $\phi_i : E_i \to F_i$

that makes the diagram

$$\begin{tikzcd}[ampersand replacement=\&] (f_i): \& E_1 \arrow[r, "f_1"] \arrow[d, "\phi_1"] \& E_2 \arrow[r, "f_2"] \arrow[d, "\phi_2"] \& \dots \arrow[r, "f_{n-1}"] \& E_n \arrow[d, "\phi_n"] \\ (g_i): \& F_1 \arrow[r, "g_1"] \& F_2 \arrow[r, "g_2"] \& \dots \arrow[r, "g_{n-1}"] \& F_n \end{tikzcd}$$

commute (*i.e.*, $\phi_{i+1} \circ f_i = g_i \circ \phi_i$). The modules (f_i) and (g_i) are isomorphic if we can choose the ϕ_i to be isomorphisms. The **direct sum** of persistence modules (f_i) and (g_i) is the persistence module

$$(f_i) \oplus (g_i) : E_1 \oplus F_1 \xrightarrow{f_1 \oplus g_1} E_2 \oplus F_2 \xrightarrow{f_2 \oplus g_2} \cdots \xrightarrow{f_{n-1} \oplus g_{n-1}} E_n \oplus F_n,$$

where, as usual, $f_i \oplus g_i$ maps $(x, y) \in E_i \oplus F_i$ to $(f_i(x), g_i(y))$.

1.1 Classification of Persistence Modules

A persistence module is **decomposable** if it is isomorphic to the direct sum of two non-trivial persistence modules. It is otherwise **indecomposable**. In this section we only consider finite persistence modules indexed over $[1, n]$, and for $1 \leq a \leq b \leq n+1$ we denote by $\mathbb{I}[a, b[$ the persistence module

$$\overset{1}{0} \longrightarrow \cdots \longrightarrow 0 \longrightarrow \overset{a}{\mathbb{F}} \xrightarrow{Id} \cdots \xrightarrow{Id} \overset{b-1}{\mathbb{F}} \longrightarrow \overset{b}{0} \cdots \longrightarrow \overset{n}{0}$$

whose ith space is the 1-dimensional vector space over $\mathbb{F}$ if $i \in [a, b[$ and 0 otherwise.

Exercise 1.1. Show that $\mathbb{I}[a, b[$ is indecomposable.

The main result about the classification of persistence modules is the uniqueness of the decomposition into indecomposables.

Theorem 1.2. *Let $(f_i)_{1 \leq i < n}$ be a persistence module. There exists a unique multiset I of subintervals of $[1, n+1[$ such that*

$$(f_i)_{1 \leq i < n} \cong \bigoplus_{[a,b[\in I} \mathbb{I}[a, b[,$$

where each interval in this sum occurs with its multiplicity in I.

The multiset I is the **barcode** of $(f_i)_{1 \leq i < n}$. It is composed of **persistence intervals**.

Corollary 1.3. *The barcode is a complete invariant for the isomorphism classes of persistence modules.*

Given a persistence module $E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n$, we let $f_{i,j} : E_i \to E_j$ be the composition of the f_k's between e_i and e_j. Precisely, we set

- $\forall i \in [1, n] : f_{i,i} = Id_{E_i}$

- $\forall 1 \le i < j \le n : f_{i,j} = f_{i+1,j} \circ f_i$ and $f_{j,i} = 0$.

We also denote the rank of $f_{i,j}$ by $\beta_{i,j}$. The multiplicity of interval $[i,j[$ in the barcode I is denoted by $m_{i,j}$.

Lemma 1.4. $m_{i,j} = (\beta_{i,j-1} - \beta_{i-1,j-1}) - (\beta_{i,j} - \beta_{i-1,j})$

PROOF. Suppose that $E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n \cong \bigoplus_{[a,b[\in I} \mathbb{I}[a,b[$. We easily compute

$$\beta_{i,j}(\mathbb{I}[a,b[) = \begin{cases} 1 & \text{if } [i,j] \subset [a,b[\\ 0 & \text{else.} \end{cases}$$

Note that for any persistence modules (g_k) and (h_k) we have $\beta_{i,j}((g_k)\oplus(h_k)) = \beta_{i,j}((g_k)) + \beta_{i,j}((h_k))$. It follows that $\beta_{i,j}((f_k))$ counts the number of persistence intervals of (f_k) that contain $[i,j]$. Hence, $\delta_{i,j} := \beta_{i,j} - \beta_{i-1,j}$ counts the number of persistence intervals of the form $[i,\ell[, \ell > j$. We infer that $m_{i,j} = \delta_{i,j-1} - \delta_{i,j} = (\beta_{i,j-1} - \beta_{i-1,j-1}) - (\beta_{i,j} - \beta_{i-1,j})$. □

Consider a vector $x \in E_i$ in the persistence module $(f_i)_{1\le i<n} = E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n$. We put $x(j) := f_{i,j}(x)$. A **compatible basis** is a family of vectors $X \subset \bigcup_i E_i$, the union being considered as disjoint, so that

$$X(i) := \{x(i) \mid (x \in X) \wedge (x(i) \neq 0)\}$$

is a basis of E_i for $1 \le i \le n$. In particular, $x, y \in X$ and $x(i) \neq 0$ implies $y(i) \neq x(i)$. The **persistence interval** of $x \in X$ is defined as $I_x = \{i \mid x(i) \neq 0\}$. For convenience, we introduce the *activation function* $a : \bigcup_i E_i \to [1,n]$ such that $x \in E_{a(x)}$ for all $x \in \bigcup_i E_i$. Hence, the lower bound of I_x is $a(x)$.

Lemma 1.5. *If $(f_i)_{1\le i<n}$ admits a compatible basis X, then $(f_i)_{1\le i<n}$ has a decomposition whose barcode is the multiset of persistence intervals $\{I_x \mid x \in X\}$.*

PROOF. $\bigoplus_{x\in X} \mathbb{I}(I_x)$ has an obvious compatible basis Y obtained by choosing for every $x \in X$ a generator of $\mathbb{F}$ at index $a(x)$. It remains to check that the persistence modules $(f_i)_i$ and $\bigoplus_{x\in X} \mathbb{I}(I_x)$ are isomorphic by constructing an isomorphism sending the bases $X(i)$ to $Y(i)$. □

Proposition 1.6. *Every persistence module admits a compatible basis.*

PROOF. For a persistence module $E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n$ we build a compatible basis by induction on n. If $n = 1$, a compatible basis is provided by any basis of E_1. We next assume to have constructed a compatible basis X for

$$E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-2}} E_{n-1}. \tag{1}$$

Let $k = |X|$ be the number of basis vectors in X. We recursively define compatible bases $X_1 = X, X_2, \ldots, X_k$ for (1). The goal is to get a compatible basis X_k such that

$\{x(n) \mid x \in X_k \wedge x(n) \neq 0\}$ is an independent family in E_n. To this end we first order the elements $x_1, x_2, \ldots, x_k$ of X in a non-decreasing fashion with respect to activation, *i.e.* such that $1 \leq j < k$ implies $a(x_j) \leq a(x_{j+1})$. Suppose we have constructed $X_{i-1} = \{y_1, y_2, \ldots, y_k\}$ for some $k \geq i > 1$, such that the y_j are indexed in non-decreasing order for activation, and such that the nonzero vectors in $\{y_1(n), y_2(n), \ldots, y_{i-1}(n)\}$ form an independent family in E_n.

- If $y_i(n) = 0$ or if $\{y_1(n), y_2(n), \ldots, y_i(n)\}$ is independent, we set $X_i = X_{i-1}$,
- otherwise, we may write $y_i(n) = \sum_{j<i} \lambda_j y_j(n)$. We then put $y'_i = y_i - \sum_{j<i} \lambda_j y_j(i)$, so that $y'_i(n) = 0$, and set $X_i = X_{i-1} \setminus \{y_i\} \cup \{y'_i\}$.

In both cases it is easily seen that X_i is a compatible basis for (1). By construction the nonzero images in E_n of the i first vectors in X_i form an independent family. By induction, X_k satisfies our goal. It remains to complete X_k with a basis of a complementary space of $f_{n-1}(E_{n-1})$ in E_n to obtain a compatible basis for $E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n$. □

PROOF OF THEOREM 1.2. By Proposition 1.6, the persistence module $(f_i)_i$ has a compatible basis, hence a decomposition into indecomposable modules of the form $\mathbb{I}[a, b[$ by Lemma 1.5. This decomposition is determined by its barcode which is uniquely defined according to Lemma 1.4. □

1.2 Restrictions of Persistence Modules

The barcode of a persistence module and of its sub-sequences can be easily related. This relationship will be used in the proof of the stability theorem in Section 4.1. In order to formalize the relation, consider a strictly increasing map $\kappa : [1, m] \to [1, n]$. The restriction to κ of the persistence module $(f_i) : E_1 \xrightarrow{f_1} E_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} E_n$ is the persistence module

$$(f_i)|_\kappa : E_{\kappa(1)} \xrightarrow{f_{\kappa(1),\kappa(2)}} E_{\kappa(2)} \cdots \xrightarrow{f_{\kappa(m-1),\kappa(m)}} E_{\kappa(m)}$$

where $f_{i,j}$ was defined below Corollary 1.3. Consider the map

$$\begin{array}{rcl} \mu : [1, n+1] & \to & [1, m+1] \\ i & \mapsto & \min\{j \in [1, m+1] \mid \kappa(j) \geq i\} \end{array}$$

where by convention $\kappa(m+1) = n+1$. It is not difficult to see that

$$\mathbb{I}[a, b[|_\kappa = \begin{cases} \mathbb{I}[\mu(a), \mu(b)[& \text{if } \mu(a) < \mu(b) \\ 0 & \text{otherwise.} \end{cases}$$

As an immediate consequence:

Lemma 1.7. *Let I be the barcode of a persistence module. The barcode of its restriction to κ is the multiset $\{[\mu(a), \mu(b)[\}_{[a,b[\in I \text{ and } \mu(a)<\mu(b)}$.*

Exercise 1.8. Prove the above evaluation for $\mathbb{I}[a, b[|_\kappa$.

2 Application to Topological Inference

As explained in the introduction one of the main motivation for the persistence homology theory is the ability to recover the topology of a shape from a sampled set of points, say P. We further remarked that it is appropriate to study the filtration $(P^{\varepsilon})_{\varepsilon\in\mathbb{R}_+}$ of the ε-neighborhoods for ϵ ranging from 0 to infinity. We are thus faced with the computation of the barcode of the corresponding induced persistence module. In general, it is more convenient to use simplicial complexes to represent topological spaces in a computer. In particular, the computation of homology groups becomes relatively easy, as seen in a previous lecture. In order to reduce the filtration (P^{ε}) to a filtration of a simplicial complexes, we can rely on the following nerve theorem. The **nerve** of a cover $(U_i)_{i\in I}$ of a space X is the abstract simplicial complex whose set of vertices is I and whose simplices are subsets $J \subset I$ such that $\cap_{j\in J} U_j \neq \emptyset$. See Figure 4 for an illustration. A cover $(U_i)_{i\in I}$ is **good** if its parts U_i are open sets and if any nonempty intersection of U_i's is contractible[1].

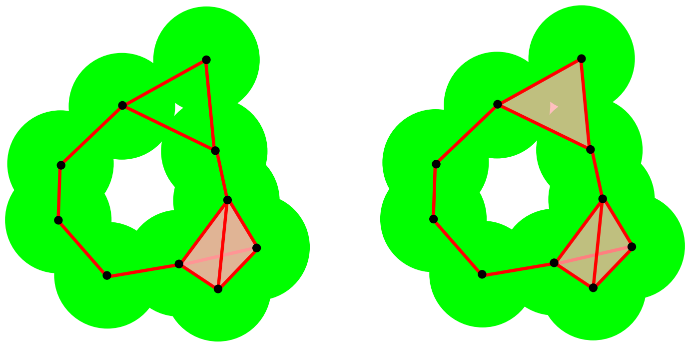

Figure 4: Left, the nerve, or Čech complex of a union of balls. Right, the Rips complexe with parameter the diameter of the balls.

Theorem 2.1 (Nerve –, Leray'1945, Borsuk'1948)**.** *Let* $(U_i)_{i\in I}$ *be a good cover of* X*, then the nerve of* $(U_i)_{i\in I}$ *has the same homotopy type as* X*.*

Considering the open balls of radius ε as a cover of P^{ε}, we observe thanks to the convexity of the balls that they constitute a good cover. Their nerve $C^{\varepsilon}(P)$ thus has the same homotopy type as P^{ε}. This nerve is sometimes called the **Čech complex** of P of parameter ε. What is more, if $\varepsilon < \eta$ the inclusions $P^{\varepsilon} \hookrightarrow P^{\eta}$ and $C^{\varepsilon}(P) \hookrightarrow C^{\eta}(P)$ (as a subcomplex) form a commutative diagram with the homotopy equivalences provided by the nerve theorem $C^{\varepsilon}(P) \longleftrightarrow P^{\varepsilon}$ and $C^{\eta}(P) \longleftrightarrow P^{\eta}$:

$$\begin{tikzcd} C^{\varepsilon}(P) \arrow[r, hook] \arrow[d, leftrightarrow] & C^{\eta}(P) \arrow[d, leftrightarrow] \\ P^{\varepsilon} \arrow[r, hook] & P^{\eta} \end{tikzcd}$$

[1] *i.e.*, has the homotopy type of a point, where two spaces have the same **homotopy** type if there exist maps $f : X \to Y$ and $g : Y \to X$, called **homotopy equivalences**, such that $g \circ f$ is homotopic to the identity on X and $f \circ g$ is homotopic to the identity on Y.

See [CO08] for a proof. It follows that the persistence modules induced by the filtrations $(P^{\varepsilon})_{\varepsilon\in\mathbb{R}_+}$ and $(C^{\varepsilon}(P))_{\varepsilon\in\mathbb{R}_+}$ are isomorphic, hence have the same barcode.

In practice, the construction of $C^{\varepsilon}(P)$ from P and ε is not very efficient. One should check for every subset of P if the corresponding ε-balls have a common intersection. This is why the Rips complex is sometimes preferred. The **Rips complex** $R^{\varepsilon}(P)$ of parameter ε associated with P is the **clique complex** of the graph over P where two points are linked by an edge if they are at distance less than ϵ. Recall that the clique complex of a graph is a simplicial complex over the vertices of the graph and has a simplex for every clique of the graph. Hence, event though the Rips complexe can be bigger than the Čech complex it is much easier to compute and can be concisely encoded by its graph. Furthermore, it is easily seen that

$$C^{\varepsilon/2}(P) \subset R^{\varepsilon}(P) \subset C^{\varepsilon}(P)$$

In fact, for $P \subset \mathbb{R}^d$ it was shown [DSG07] that

$$R^{\varepsilon}(P) \subset C^{\eta/2}(P) \subset R^{\eta}(P)$$

with $\eta = \varepsilon\sqrt{\frac{2d}{d+1}}$. Such relations can be used to replace the Čech complex by the Rips complex in the computations of the barcode. See [DSG07] for more details. The next section explains how to compute the barcode of a filtration of simplicial complexes.

3 Computing the Barcode

Consider a filtration $\mathscr{K} : K_1 \subset K_2 \subset \ldots \subset K_n$ of a simplicial complex $K = K_n$. We want to compute the barcode $I(\mathscr{K})$ of the induced persistence module. In practice we restrict to **simple filtrations** for which each $K_i = K_{i-1} \cup \sigma_i$ is obtained by adding a single simplex σ_i to K_{i-1} (by convention $K_0 = \emptyset$). Thanks to Lemma 1.7, this actually allows to compute the barcode of non-simple filtrations.

We fix a coefficient field $\mathbb{F}$ and denote by $C(K_i)$, $Z(K_i)$ and $B(K_i)$ the $\mathbb{F}$-vector spaces of chains, cycles and boundaries of K_i, respectively. Hence, the homology group of K_i (actually an $\mathbb{F}$-vector space) is given by $H(K_i) = \ker\partial/\mathrm{Im}\,\partial = Z(K_i)/B(K_i)$, where $\partial : C(K_i) \to C(K_i)$ is the boundary operator. We omit the dimension of the relevant simplices in $C(K_i), Z(K_i), B(K_i)$ and $H(K_i)$, considering that $C(K_i)$ (resp. $Z(K_i)\ldots$) is the direct sum of the chain spaces $C_k(K_i)$ for each dimension k. By the rank-nullity theorem applied to the boundary operator:

$$\dim C(K_i) = \dim Z(K_i) + \dim B(K_i)$$

Noting that $\dim C(K_i)$ is the number of simplices in K_i, we get

$$\big(\dim Z(K_i) - \dim Z(K_{i-1})\big) + \big(\dim B(K_i) - \dim B(K_{i-1})\big) = 1$$

Since $\dim Z(K_i) \geq Z(K_{i-1})$ and $\dim B(K_i) \geq B(K_{i-1})$, we have

1. either $\dim Z(K_i) = \dim Z(K_{i-1}) + 1$ and $B(K_i) = B(K_{i-1})$,
2. or $Z(K_i) = Z(K_{i-1})$ and $\dim B(K_i) = \dim B(K_{i-1}) + 1$.

We say that index i (or simplex σ_i) is **positive** in the first case and **negative** in the other case. We denote by $\mathscr{P}(\mathscr{K})$ and $\mathscr{N}(\mathscr{K})$ the set of positive, respectively negative, indices.

Lemma 3.1. *The following are equivalent:*

- σ_i *is positive,*
- σ_i *is in the support of a cycle* $z \in Z(K_i)$*. Moreover,* $Z(K_i) = Z(K_{i-1}) \oplus \mathbb{F}z$,
- $\partial \sigma_i \in B(K_{i-1})$,

The proof is left as an exercise. See Figure 5. Note that in any case,

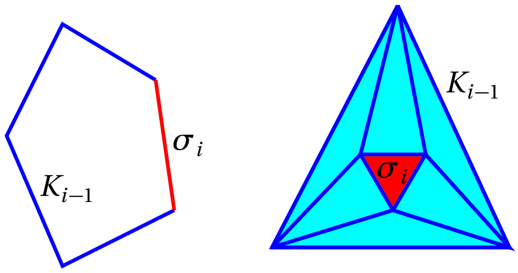

Figure 5: Left, σ_i belongs to a cycle of K_i and is thus positive. Right, σ_i is negative.

$$B(K_i) = B(K_{i-1}) + \mathbb{F}\partial\sigma_i. \tag{2}$$

The above sum is direct if and only if σ_i est negative. The endpoints a and b of the persistence interval $[a, b[$ are respectively called its lower and upper bound.

Lemma 3.2. *Every persistence interval* $[i, j[\in I(\mathscr{K})$ *satisfies*

$$(i, j) \in \mathscr{P}(\mathscr{K}) \times (\mathscr{N}(\mathscr{K}) \cup \{n+1\}).$$

Moreover,

- *Each positive index is the lower bound of a unique interval in* $I(\mathscr{K})$.
- *Each negative index is the upper bound of a unique interval in* $I(\mathscr{K})$.

Note that $n+1$ is not an index of the filtration and that it may be the upper bound of several persistence intervals.

PROOF. The morphism $\varphi_{i-1} : H(K_{i-1}) \to H(K_i)$ is a quotient of the inclusion $Z(K_{i-1}) \subset Z(K_i)$ by $B(K_{i-1})$ at the domain and by $B(K_i)$ at the codomain. From the definition of a positive simplex it follows that φ_{i-1} is one-to-one and that $\dim H(K_i) = \dim H(K_{i-1}) + 1$ when σ_i is positive. In this case i cannot be the upper bound of a persistence interval of the form $[a, i[$. Indeed, the corresponding indecomposable module $\mathbb{I}[a, i[$ would appear in the decomposition of $(H(K_i))_i$. However, the segment of $\mathbb{I}[a, i[$ between index $i-1$ and i is the map $\mathbb{F} \to 0$, which is obviously not injective. Similarly, if σ_i

is negative then φ_{i-1} is onto and $\dim H(K_i) = \dim H(K_{i-1}) - 1$. As a consequence, i cannot be the lower bound of any persistence interval. On the other hand, $\dim H(K_i)$ is the number of persistence intervals that contain i. It easily follows that exactly one interval starts when σ_i is positive and one interval ends when σ_i is negative. □

We can thus define the **birth function** as the map $b : \mathcal{N}(\mathcal{K}) \to \mathcal{P}(\mathcal{K})$ such that for all $j \in \mathcal{N}(\mathcal{K})$, $[b(j), j[\in I(\mathcal{K})$. In particular,

$$I(\mathcal{K}) = \{[b(j), j[\}_{j\in\mathcal{N}(\mathcal{K})} \cup \{[i, n+1[\}_{i\in\mathcal{P}(\mathcal{K})\setminus \mathrm{Im}\, b} \tag{3}$$

Hence, we may recover the barcode $I(\mathcal{K})$ from the knowledge of the signs of the simplices and of the birth function.

3.1 Compatible Boundary Basis

A **Compatible boundary basis** is a family of cycles $\mathcal{B}(\mathcal{K}) = \{x_j\}_{j\in J} \subset Z(K)$, with $J \subset [1, n]$, such that:

1. $\forall i \in [1, n], \{x_j\}_{j\in J\cap[1,i]}$ is a basis of $B(K_i)$,
2. the map $\beta : J \to [1, n]$, $j \mapsto$ (maximum index of the simplices in x_j) is injective.

Lemma 3.3. *Suppose that $\mathcal{K}$ has a compatible boundary basis, then β coincides with the birth function b.*

Proof. The above Condition 1 and the remark after Equation (2) show that $J = \mathcal{N}(\mathcal{K})$. Lemma 3.1 also implies that $\beta(j) \in \mathcal{P}(\mathcal{K})$ for all $j \in J$. For every $i \in \mathcal{P}(\mathcal{K})$, define $z_i \in Z(K_i)$ as follows.

- If $i = \beta(j)$ for some $j \in J$, then $z_i = x_j$.
- Else, choose z_i such that $Z(K_i) = Z(K_{i-1}) \oplus \mathbb{F} z_i$ (cf. Lemma 3.1).

Remark that the simplex with maximum index in the support of z_i is σ_i. Hence, (cf. Lemma 3.1) $\{z_j\}_{j\in\mathcal{P}(\mathcal{K}), j\le i}$ is a basis of $Z(K_i)$. Let $[z]_i$ denote the homology class of cycle z in $H(K_i)$. We need to check that $([z_j]_j)_{j\in\mathcal{P}(\mathcal{K})}$ is a compatible basis for the homology sequence of $\mathcal{K}$ and that the persistence interval of each $[z_{\beta(j)}]_{\beta(j)}$ is $[\beta(j), j[$, while the persistence interval of $[z_j]_j$, $j \in \mathcal{P}(\mathcal{K}) \setminus \beta(J)$, is $[j, n+1[$. We claim that

$$Z(i) := \{[z_{\beta(j)}]_i\}_{(j\in J)\wedge(\beta(j)\le i)\wedge(j>i)} \cup \{[z_j]_i\}_{(j\le i)\wedge(j\in\mathcal{P}(\mathcal{K})\setminus\beta(J))}$$

is a basis of $H(K_i)$. Since $[z_{\beta(j)}]_j = [x_j]_j = 0$, we also have $[z_{\beta(j)}]_i = 0$ for $i \ge j$ and it follows from the above remark that $Z(i)$ spans $H(K_i)$. To see that $Z(i)$ is an independent set, consider a linear combination $\sum_{(j\in J)\wedge(\beta(j)\le i)\wedge(j>i)} \alpha_j [z_{\beta(j)}]_i + \sum_{(j\le i)\wedge(j\in\mathcal{P}(\mathcal{K})\setminus\beta(J))} \alpha_j [z_j]_i$ of elements in $Z(i)$. If it is zero, then the combination $c := \sum_{(j\in J)\wedge(\beta(j)\le i)\wedge(j>i)} \alpha_j z_{\beta(j)} + \sum_{(j\le i)\wedge(j\in\mathcal{P}(\mathcal{K})\setminus\beta(J))} \alpha_j z_j$ of the corresponding cycles must lie in $B(K_i)$. By the first condition in the definition of a compatible boundary basis, cycle c must be equal to a linear combination of $\{x_j \mid (j \in J) \wedge (j \le i)\}$. Because the maximum index of the simplices in the support of each $z_{\beta(j)}, z_j$ and x_j are pairwise distinct, it must be that all the coefficients α_j in c are null, thus concluding the proof of the claim. We finally observe that the persistence interval of $[z_j]_j$ is the set of i's for which $[z_j]_i \in Z(i)$. Whence, for $j \in J$ the persistence interval of $[z_{\beta(j)}]_{\beta(j)}$ is $[\beta(j), j[$, while for $j \in \mathcal{P}(\mathcal{K}) \setminus \beta(J)$ the persistence interval of $[z_j]_j$ is $[j, n+1[$. □

3.2 Algorithm

Lemma 3.3 and Equation (3) show that it is enough to construct a compatible boundary basis for $\mathscr{K}$ to derive the sign of each simplex and the barcode of $\mathscr{K}$. We can construct a compatible boundary basis by induction on the size n of the filtration. The base case $n = 1$ is trivial because the unique simplex in the filtration must be a (positive) vertex. We thus assume that we have computed a compatible boundary basis $\mathscr{B}(\mathscr{K}') = \{x_j\}_{j\in J}$ for the sub-filtration $\mathscr{K}'$:

$$K_1 \subset K_2 \subset \ldots \subset K_{n-1}$$

We denote by $b : J \to \mathscr{P}(\mathscr{K}')$ the corresponding birth function. Suppose that we can write

$$\partial \sigma_n = \sum_{j\in J} \alpha_j x_j + y, \tag{4}$$

where

1. either $y = 0$,

2. or the maximum index of the simplices in y is not in $b(J)$.

In case 1, we have $B(K_n) = B(K_{n-1})$ and $\mathscr{B}(\mathscr{K}')$ remains a compatible boundary basis for $\mathscr{K}$. In case 2, n is negative and $\mathscr{B}(\mathscr{K}') \cup \{y\}$ is a compatible boundary basis for $\mathscr{K}$.

By the second condition in its definition, every compatible boundary basis is in echelon form when the cycles are written as combination of simplices. We can thus apply Gaussian elimination as in the following pseudocode to obtain a decomposition as in (4).

$y := \partial \sigma_n$
$i :=$ maximum index of the simplices in y
While ($(y \neq 0) \wedge (i \in b(J))$)
  $j := b^{-1}(i)$
  $\alpha :=$ coefficient of σ_i in y
  $\beta :=$ coefficient of σ_i in x_j
  $y := y - (\alpha/\beta)x_j$
  $i :=$ maximum index of the simplices in y * *undefined if* $y = 0$ *\
End while
* $y = 0$ *or* $y = x_n$ *when leaving the while loop* *\

We can store each x_j as a table of coefficients indexed by the n simplices of the filtration. We represent the birth function as a table of length n; the jth entry contains $b(j)$ if j is negative and 0 otherwise. We also store the inverse map b^{-1} in a table of length n. The computation of x_n by the above loop takes $O(n^2)$ time. Hence,

Proposition 3.4. *We can compute a compatible boundary basis and the birth function of a filtration of length n in $O(n^3)$ time on an $\mathbb{F}$-RAM machine. We can moreover compute the barcode of the filtration in the same amount of time.*

4 Persistence Diagrams

A function $f : K \to \mathbb{R}$ over a simplicial complex K is **non-decreasing** if

$$\forall \sigma, \tau \in K : \sigma \prec \tau \implies f(\sigma) \leq f(\tau)$$

where $\sigma \prec \tau$ means "*σ is a face of τ*". A filtration of K can be equivalently described by a non-decreasing function f over K. Indeed, if $f_1 < f_2 < \ldots < f_n$ is the sequence of *values* of f, then the sequence

$$f^{-1}([-\infty, f_1]) \subset f^{-1}([-\infty, f_2]) \subset \ldots \subset f^{-1}([-\infty, f_n]) \tag{5}$$

is a filtration of K, which we denote by $\mathscr{K}_f$. Conversely, any filtration $K_1 \subset K_2 \subset \ldots \subset K_n = K$ has the form $\mathscr{K}_f$ for f defined over K by $f(\sigma) = i \Leftrightarrow \sigma \in K_i \setminus K_{i-1}$.

We set $f_{n+1} = +\infty$. The **persistence diagram** $D(f)$ of f is the multiset of points in the extended plane $(\mathbb{R} \cup \{-\infty, +\infty\})^2$ given by

$$D(f) = \{(f_i, f_j)\}_{[i,j[\in I(\mathscr{K}_f)} \cup \Delta^\infty,$$

where Δ^∞ is the multiset of points on the diagonal $\{x = y\}$, each counted with countably infinite multiplicity. We say that the filtration $\mathscr{K} : K_1 \subset K_2 \subset \ldots \subset K_m = K$ is **compatible** with $f : K \to \mathbb{R}$ if $\mathscr{K}_f$ is a sub-filtration of $\mathscr{K}$. In other words, $\mathscr{K}$ is compatible with f if f is constant over each $K_i \setminus K_{i-1}$ and if $f_{\mathscr{K}} : [1, m] \to \mathbb{R}$, $i \mapsto f(K_i \setminus K_{i-1})$ is non-decreasing. In this case we define the **persistence diagram of f relatively to $\mathscr{K}$** as the multiset:

$$D(f, \mathscr{K}) = \{\big(f_{\mathscr{K}}(i), f_{\mathscr{K}}(j)\big)\}_{[i,j[\in I(\mathscr{K})} \cup \Delta^\infty$$

where we have put $f_{\mathscr{K}}(m+1) = +\infty$. In particular, $D(f) = D(f, \mathscr{K}_f)$.

Lemma 4.1. *$D(f) = D(f, \mathscr{K})$ for any filtration $\mathscr{K}$ compatible with f.*

PROOF. Let $f_1 < f_2 < \ldots < f_n$ be the sequence of values of f over K. We set

$$\kappa : [1, n] \to [1, m], i \mapsto \max\{j \mid f_{\mathscr{K}}(j) = f_i\}.$$

Hence, $f^{-1}([-\infty, f_i]) = K_{\kappa(i)}$ and the persistence module induced by the homology of $\mathscr{K}_f$ is the restriction to κ of the persistence module induced by $\mathscr{K}$ (see Section 1.2). By Lemma 1.7 we have $I(\mathscr{K}_f) = \big\{[\mu(i), \mu(j)[\big\}_{[i,j[\in I(\mathscr{K}) \text{ and } \mu(i)<\mu(j)}$, with μ as in Lemma 1.7. It follows that

$$D(f) = D(f, \mathscr{K}_f) = \{\big(f_{\mu(i)}, f_{\mu(j)}\big) \mid [i, j[\in I(\mathscr{K}) \text{ and } \mu(i) < \mu(j)\} \cup \Delta^\infty$$

We easily check from the definitions of κ and μ that $f_{\mu(i)} = f_{\mathscr{K}}(i)$. Hence,

$$D(f) = \{\big(f_{\mathscr{K}}(i), f_{\mathscr{K}}(j)\big) \mid [i, j[\in I(\mathscr{K}) \text{ and } \mu(i) < \mu(j)\} \cup \Delta^\infty$$

Now, if $\mu(i) \geq \mu(j)$ for some interval $[i, j[\in I(\mathscr{K})$ then $f_{\mathscr{K}}(i) = f_{\mathscr{K}}(j)$ and the corresponding point $(f_{\mathscr{K}}(i), f_{\mathscr{K}}(i))$ is "absorbed" by the diagonal Δ^∞. We finally conclude $D(f) = \Delta^\infty \cup \{\big(f_{\mathscr{K}}(i), f_{\mathscr{K}}(j)\big)\}_{[i,j[\in I(\mathscr{K})} = D(f, \mathscr{K})$. $\square$

4.1 Stability of Persistence Diagrams

The stability of the persistence diagram $D(f)$ with respect to f is the main result of Persistence theory. We first introduce the **bottleneck distance** d_B between persistence diagrams. Note that thanks to the diagonal Δ^∞ any two diagrams D, D' are in bijection. We set

$$d_B(D, D') = \inf_{\phi} \sup_{p \in D} \|p - \phi(p)\|_\infty$$

where $\phi : D \to D'$ runs over the bijections between D and D' and $\|p-q\|_\infty = \max\{|x_p - x_q|, |y_p - y_q|\}$ (by convention, $|+\infty - x| = 0$ if $x = +\infty$ and $|+\infty - x| = +\infty$ otherwise). See Figure 6. Note that d_B is not a distance properly speaking: it can take infinite

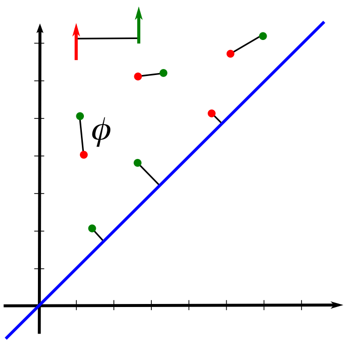

Figure 6: The bottleneck distance is computed by minimizing over all bijections ϕ the largest distance in each pairing.

values but otherwise satisfies the triangular inequality.

As usual, for any functions $f, g : K \to \mathbb{R}$, we denote their L_∞ distance by

$$\|f - g\|_\infty = \sup_{\sigma \in K} |f(\sigma) - g(\sigma)|$$

Theorem 4.2 (Stability –, [CSEH07, CSEM06])**.** $d_B(D(f), D(g)) \leq \|f - g\|_\infty$

Proof. Put $f_t = f + t(g - f)$. note that if f, g are non-decreasing over K, so is f_t. For every two simplices $\sigma, \tau \in K$, there exists $u \in [0,1]$ such that the sign of $f_t(\sigma) - f_t(\tau)$ is constant for $t \in [0, u]$ and the same is true for $t \in [u, 1]$. There is thus a finite partition $0 = t_0 < t_1 < \ldots < t_r = 1$ of[2] $[0,1]$ so that the relative order of the f_t-values of the simplices is independent of t over each interval $[t_i, t_{i+1}]$. It follows that for each $i \in [0, r-1]$ we can exhibit a simple filtration $\mathcal{K}_i$ compatible with *every* function f_t for $t \in [t_i, t_{i+1}]$. By Lemma 4.1, we have

$$D(f_t) = D(f_t, \mathcal{K}_i) = \Delta^\infty \cup \{\big(f_t(\sigma_a), f_t(\sigma_b)\big)\}_{[a,b[\in I(\mathcal{K}_i)}$$

where σ_a is the ath simplex of $\mathcal{K}_i$. Considering the obvious correspondence between $D(f_{t_i})$ and $D(f_{t_{i+1}})$ that restricts to the identity over Δ^∞ and sends $\big(f_{t_i}(\sigma_a), f_{t_i}(\sigma_b)\big)$ to $\big(f_{t_{i+1}}(\sigma_a), f_{t_{i+1}}(\sigma_b)\big)$, we infer $d_B(D(f_{t_i}), D(f_{t_{i+1}})) \leq (t_{i+1} - t_i)\|f - g\|_\infty$. Applying the triangular inequality we finally conclude

$$d_B(D(f), D(g)) \leq \sum_i (t_{i+1} - t_i)\|f - g\|_\infty = \|f - g\|_\infty$$

[2] $r \leq \binom{m}{2} + 1$ where m is the number of simplices of K

□

The Stability theorem was refined in a more general context by Chazal et al. and Bubenik and Scott [CCSG+09, CDSGO12, BS14]. A first generalization is to consider "continuous" persistence module indexed over $\mathbb{R}$. This is a family of vector spaces $(V_x)_{x\in\mathbb{R}}$ and a family of linear maps $(v_{x,y} : V_x \to V_y)_{x\le y}$ satisfying $v_{x,x} = Id_{V_x}$ and $v_{x,z} = v_{x,y} \circ v_{y,z}$ for all $x \le y \le z$. We denote it by $\mathbb{V}$. Given two persistence modules $\mathbb{V}$ and $\mathbb{W}$ over $\mathbb{R}$ and a real number d, a **degree d morphism** $\varphi : \mathbb{V} \to \mathbb{W}$ is a family of linear maps $(\varphi_x : V_x \to W_{x+d})_{x\in\mathbb{R}}$ such that the following diagram:

$$\begin{tikzcd} V_x \arrow[r] \arrow[dr, "\varphi_x"] & V_y \arrow[dr, "\varphi_y"] & \\ & W_{x+\varepsilon} \arrow[r] & W_{y+\varepsilon} \end{tikzcd}$$

commutes for all $x \in \mathbb{R}$. An **ε-interleaving** is a pair of morphisms $\varphi : \mathbb{V} \to \mathbb{W}$ and $\psi : \mathbb{W} \to \mathbb{V}$, each of degree ε, such that the following diagrams:

$$\begin{tikzcd} V_x \arrow[rr] \arrow[dr, "\varphi_x"'] & & V_{x+2\varepsilon} \\ & W_{x+\varepsilon} \arrow[ur, "\psi_{x+\varepsilon}"'] & \end{tikzcd} \qquad \begin{tikzcd} & V_{x+\varepsilon} \arrow[dr, "\varphi_{x+\varepsilon}"] & \\ W_x \arrow[ur, "\psi_x"] \arrow[rr] & & W_{x+2\varepsilon} \end{tikzcd}$$

commute. The **interleaving distance** between $\mathbb{V}$, $\mathbb{W}$ is

$$d_i(\mathbb{V}, \mathbb{W}) = \inf\{\varepsilon \mid \exists \varepsilon\text{-interleaving between } \mathbb{V}, \mathbb{W}\}$$

When is $\mathbb{V}$ is pointwise finite dimensional, *i.e.*, when each space V_x if finite dimensional, it can be shown that the decomposition Theorem 1.2 remains valid [CB15]. This time each indecomposable module $\mathbb{I}(\iota)$ may apply to any type of interval ι (half-open, closed, semi-infinite,...) and satisfies $\mathbb{I}(\iota)_x = \mathbb{F}$ for $x \in \iota$ and $\mathbb{I}(\iota)_x = 0$ otherwise, with identity or zero maps wherever it applies. The persistence diagram is then defined as the multiset of points (u, v) where u, v runs over the endpoints of the persistence intervals.

Theorem 4.3 (Isometry –, [CCSG+09, CDSGO12])**.** *Let $\mathbb{V}$ and $\mathbb{W}$ be pointwise finite dimensional persistence modules over $\mathbb{R}$ such that rank $v_{x,y}$ and rank $w_{x,y}$ is finite for every $x < y$. Then,*

$$d_B(D(\mathbb{V}), D(\mathbb{W})) = d_i(\mathbb{V}, \mathbb{W})$$

The stability theorem can be deduced from the isometry theorem as follows. Let X be a topological space and $f : X \to \mathbb{R}$. Put $X_t^f := f^{-1}(-\infty, t]$. The filtration $\mathbb{X}^f := (X_t^f)_{t\in\mathbb{R}}$ induces a persistence module $H(\mathbb{X}^f)$ by applying the homology functor. Now, $\|f - g\|_\infty \le \varepsilon$ implies $X_t^f \subset X_{t+\varepsilon}^g \subset X_{t+2\varepsilon}^f$. It follows that $H(\mathbb{X}^f), H(\mathbb{X}^g)$ are ε-interleaved, whence by the isometry Theorem $d_i(H(\mathbb{X}^f), H(\mathbb{X}^g)) \le \varepsilon$. In turn, this implies $d_B(D(f), D(g)) = d_i(H_*(X_f), H_*(X_g)) \le \|f - g\|_\infty$.

Exercise 4.4. Check that a persistence module over an ordered set $(X, \le)$, e.g. $X = \mathbb{R}$, is the same as a functor from the category $(X, \le)$ to the category of $\mathbb{F}$-vector spaces. Here, the objects of $(X, \le)$ are the elements of X and each pair (x, y) of objects has exactly one morphism if $x \le y$ (and none otherwise).

References

[BS14] Peter Bubenik and Jonathan A Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51(3):600–627, 2014.

[CB15] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and Its Applications*, 14(05):1550066, 2015.

[CCSG+09] F. Chazal, D. Cohen-Steiner, M. Glisse, L. Guibas, and S. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 237–246. ACM, 2009.

[CDSGO12] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. The structure and stability of persistence modules. arXiv preprint arXiv:1207.3674, 2012.

[CO08] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 232–241. ACM, 2008.

[CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of Persistence Diagrams. *Discrete & Computational Ceometry*, 37(1):103–120, 2007.

[CSEM06] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and Vineyards by Updating Persistence in Linear Time. In *22nd Annual ACM Symposium on Computational Geometry*, pages 119–126, 2006.

[DSG07] Vin De Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7(1):339–358, 2007.

[ELZ00] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological Persistence and Simplification. In *IEEE Symposium on Foundations of Computer Science*, pages 454–463, 2000.

[Rob99] Vanessa Robins. Towards computing homology from finite approximations. *Topology proceedings*, 24, 1999.