Computational Topology for Data Analysis: Notes from Book by

Tamal Krishna Dey
Department of Computer Science
Purdue University
West Lafayette, Indiana, USA 46907

Yusu Wang
Halıcıoğlu Data Science Institute
University of California, San Diego
La Jolla, California, USA 92093

Topic 4: Topological persistence

Suppose we have a noisy point set (data) P sampled from a space, say a curve in $\mathbb{R}^2$ as in Figure 4.2. Our goal is to get the information that the sampled space had two loops, one bigger and more prominent than the other. The notion of persistence captures this information. Consider the distance function $r : \mathbb{R}^2 \to \mathbb{R}$ defined over $\mathbb{R}^2$ where $r(x)$ equals $d(x, P)$, that is, the minimum distance of x to the points in P. Now let us look at the sublevel sets of r, that is, $r^{-1}[-\infty, a]$ for some $a \in \mathbb{R}^+ \cup \{0\}$. These sublevel sets are union of closed balls of radius a centering the points.

We can observe from Figure 4.2 that if we increase a starting from zero, we come across different holes surrounded by the union of these balls which ultimately get filled up at different times. However, the two holes corresponding to the original two loops persist longer than the others. We can abstract out this observation by looking at how long a feature (homological class) survives when we scan over the increasing sublevel sets. This weeds out the 'false' features (noise) from the true ones. The notion of persistent homology formalizes this idea – *It takes a function defined on a topological space and quantizes the changes in homology classes as the sublevel sets grow with increasing value of the function.*

There are two predominant scenarios where persistence appears though in slightly different contexts. One is when the function is defined on a topological space which requires considering singular homology groups of the sublevel sets. The other is when the function is defined on a simplicial complex and the sequence of sublevel sets are implicitly given by a nested sequence of subcomplexes called a *filtration*. This involves simplicial homology. Section 4.1 introduces persistence in both of these contexts though we focus mainly on the simplicial setting which is availed most commonly for computational purposes.

The birth and death of homological classes give rise to intervals during which a class remains alive. These intervals together called a *barcode* summarize the topological persistence of a filtration. An equivalent notion called *persistence diagrams* plots the intervals as points in the extended plane $\bar{\mathbb{R}}^2 := (\mathbb{R} \cup \{\pm\infty\})^2$; specifically, the birth and death constitutes the x- and y-coordinates of a point. The stability of the persistence diagrams over the perturbation of the functions giving rise to filtrations is an important result that makes topological persistence robust against noise. When filtrations are given without any explicit mention of a function, we can still talk about the stability of the persistence diagrams with respect to the so-called *interleaving distance* between the induced *persistence modules*. Sections 4.2 and 4.3 are devoted to these concepts.

4.1 Filtrations and persistence

At the core of topological persistence is the notion of filtrations which can arise in the context of topological spaces or simplicial complexes.

4.1.1 Space filtration

Consider a real-valued function $f : \mathbb{T} \to \mathbb{R}$ defined on a topological space $\mathbb{T}$. Let $\mathbb{T}_a = f^{-1}(-\infty, a]$ denote the sublevel set for the function value a. Certainly, we have inclusions:

$$\mathbb{T}_a \subseteq \mathbb{T}_b \text{ for } a \leq b.$$

Now consider a sequence of distinct values $a_1 < a_2 < \ldots, < a_n$ which are often chosen to be critical values where the homology group of the sublevel sets change as illustrated in Figure 4.1. Considering the sublevel sets at these values and a dummy value $a_0 = -\infty$ with $\mathbb{T}_{a_0} = \emptyset$, we obtain a nested sequence of subspaces of $\mathbb{T}$ connected by inclusions which gives a *filtration* $\mathcal{F}_f$:

$$\mathcal{F}_f : \emptyset = \mathbb{T}_{a_0} \hookrightarrow \mathbb{T}_{a_1} \hookrightarrow \mathbb{T}_{a_2} \hookrightarrow \cdots \hookrightarrow \mathbb{T}_{a_n}. \tag{4.1}$$

Figure 4.1 shows an example of the inclusions of the sublevel sets. The inclusions in a filtration induce linear maps in the singular homology groups of the subspaces involved. So, if $\iota : \mathbb{T}_{a_i} \to \mathbb{T}_{a_j}$, $i \le j$, denotes the inclusion map $x \mapsto x$, we have an induced homomorphism

$$h_p^{i,j} = \iota_* : \mathsf{H}_p(\mathbb{T}_{a_i}) \to \mathsf{H}_p(\mathbb{T}_{a_j}) \tag{4.2}$$

for all p and $0 \le i \le j \le n$. Therefore, we have a sequence of homomorphisms induced by inclusions forming what we call a *homology module*:

$$0 = \mathsf{H}_p(\mathbb{T}_{a_0}) \to \mathsf{H}_p(\mathbb{T}_{a_1}) \to \mathsf{H}_p(\mathbb{T}_{a_2}) \to \cdots \to \mathsf{H}_p(\mathbb{T}_{a_n}).$$

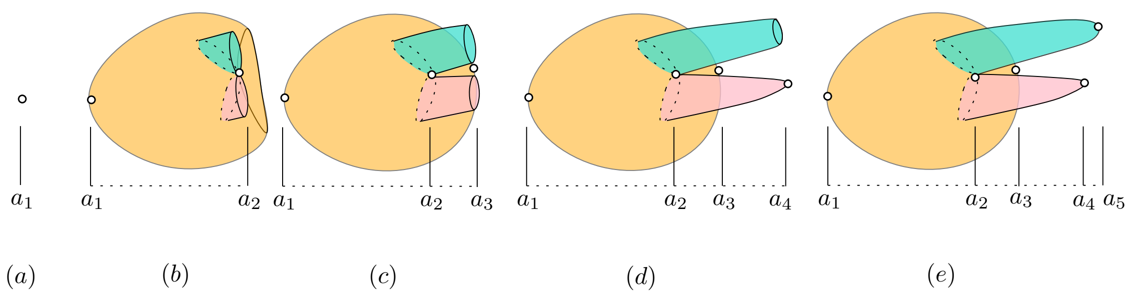

Figure 4.1: Persistence of a function on a topological space that has five critical values: (a) $\mathbb{T}_{a_1}$: only a new class in H_0 is created, (b) $\mathbb{T}_{a_2}$: two new independent classes in H_1 are created, (c) $\mathbb{T}_{a_3}$: one of the two classes in H_1 dies, (d) $\mathbb{T}_{a_4}$: the single remaining class in H_1 dies, (e) $\mathbb{T}_{a_5}$: a new class in H_2 is created.

The homomorphism $h_p^{i,j}$ takes the homology classes of the sublevel set $\mathbb{T}_{a_i}$ to those of the sublevel sets of $\mathbb{T}_{a_j}$. Some of these classes may die or get merged with other classes while the others survive. The image Im $h_p^{i,j}$ contains this information.

The inclusions of sublevel sets give rise to persistence also in the context of point clouds, a common input form in data analysis.

Point cloud. For a point set P in a metric space (M, d), we define the distance function $f : M \to \mathbb{R}$, $x \mapsto \mathsf{d}(x, p)$ where $p \in \operatorname{argmin}_{q \in P} \mathsf{d}(x, q)$. Observe that the sublevel sets $f^{-1}(-\infty, a]$ are the union of closed metric balls of radius a centering points in P. Now we have exactly the same setting as we described for general topological spaces above where $\mathbb{T}$ is replaced by the union of metric balls that grows with increasing value of a. Figure 4.2 illustrates an example where M is the Euclidean space $\mathbb{R}^2$.

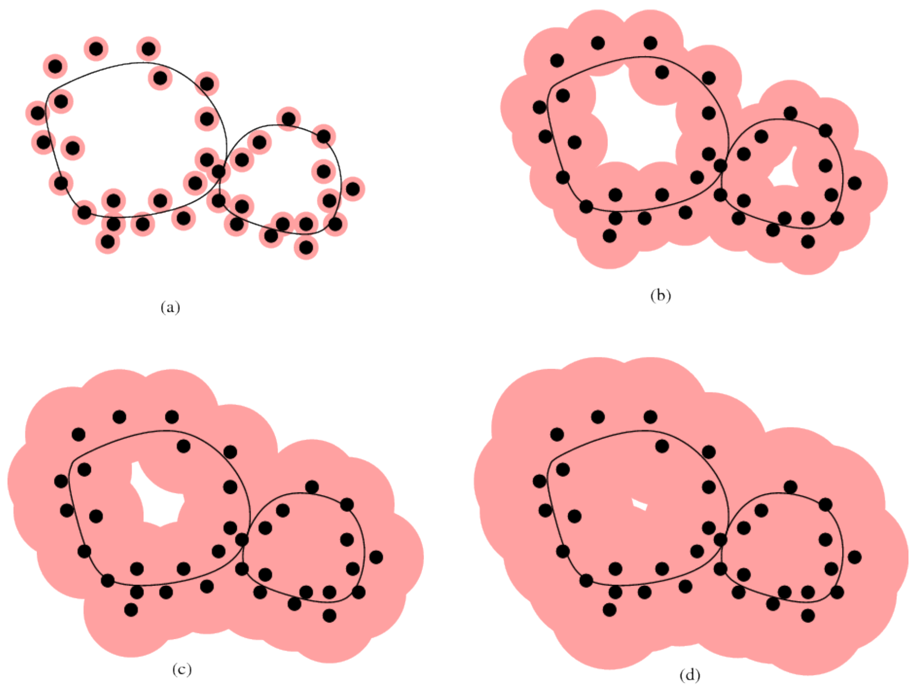

Figure 4.2: Noisy sample of a curve with two loops and the growing sublevel sets of the distance function to the sample points: The larger loop appearing as the bigger hole in the complement of the union of balls persists longer than the same for the smaller loop while other spurious holes persist even shorter.

4.1.2 Simplicial filtrations and persistence

Persistence on topological spaces involve computing singular homology groups for sublevel sets. Computationally, this is cumbersome. So, we take refuge in the discrete analogue of the topological persistence. This involves two important adaptations: first, the topological space is replaced with a simplicial complex; second, singular homology groups are replaced with simplicial homology groups. This means that the topological space $\mathbb{T}$ considered before is replaced with one of its triangulations as Figure 4.3 illustrates. For point cloud data, the union of balls can be replaced by their nerve, the Čech complex or its cousin Vietoris-Rips complex introduced earlier. Figure 4.4 illustrates this conversion for example in Figure 4.2. Of course, these replacements need to preserve the original persistence in some sense, which is addressed by the notion of stability of persistence in Section 4.3.

The nested sequence of topological spaces that arise with growing sublevel sets translates into a nested sequence of simplicial complexes in the discrete analogue. This brings in the concept of filtration of simplicial complexes that allows defining the persistence using simplicial homology groups.

Definition 1 (Simplicial filtration)**.** A *filtration* $\mathcal{F} = \mathcal{F}(K)$ of a simplicial complex K is a nested

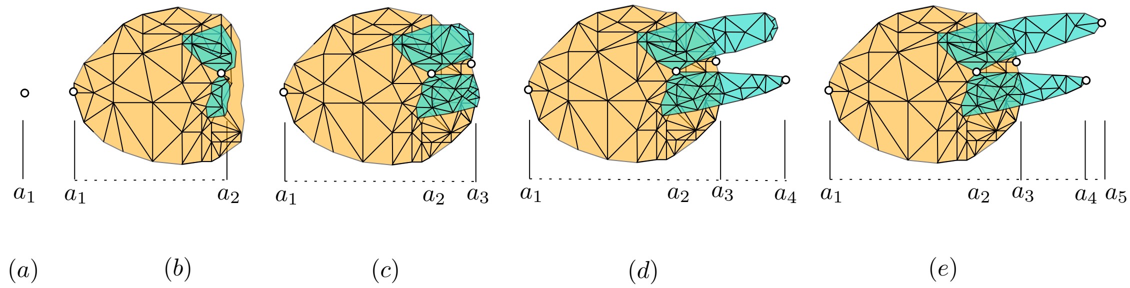

Figure 4.3: Persistence of the piecewise linear version of the function on a triangulation of the topological space considered in Figure 4.1.

sequence of its subcomplexes

$$\mathcal{F} : \emptyset = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_n = K.$$

$\mathcal{F}$ is called *simplex-wise* if $K_i \setminus K_{i-1}$ is either empty or a single simplex for every $i \in [1, n]$. Notice that the possibility of difference being empty allows two consecutive complexes to be the same.

Simplicial filtrations can appear in various contexts.

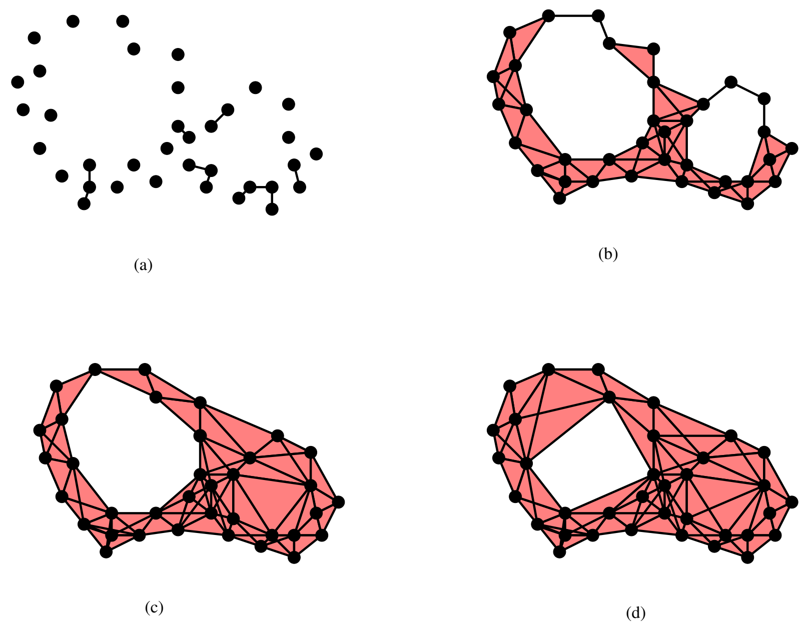

Figure 4.4: Čech complex of the union of balls considered in Figure 4.2. Homology classes in H_1 are being born and die as the union grows. The two most prominent holes appear as two most persistent homology classes in H_1. Other classes appear and disappear quickly with relatively much shorter persistence.

Simplex-wise monotone function. Consider a simplicial complex K and a (simplex-wise) function $f : K \to \mathbb{R}$ on it. We call the function f *simplex-wise monotone* if for every $\sigma' \subseteq \sigma$, we have $f(\sigma') \leq f(\sigma)$. This property ensures that the sublevel sets $f^{-1}(-\infty, a]$ are subcomplexes of K for every $a \in \mathbb{R}$. Denoting $K_i = f^{-1}(-\infty, a_i]$ and a dummy value $a_0 = -\infty$, we get a filtration:

$$\emptyset = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_n = K.$$

Vertex function. In this case, a *vertex function* $f : V(K) \to \mathbb{R}$ is defined on the vertex set $V(K)$ of the complex K. We can construct a filtration $\mathcal{F}$ from such a function.

Lower/upper stars. Recall that we have already defined the star and link of a vertex $v \in K$ which intuitively captures the concept of local neighborhood of v in K. We infuse the information about a vertex function f into these structures. First, we fix a total order on vertices $V = \{v_1, \ldots, v_n\}$ of K so that their f-values are in non-decreasing order, that is, $f(v_1) \leq f(v_2) \leq \cdots \leq f(v_n)$. The *lower-star* of a vertex $v \in V$, denoted by $\mathrm{Lst}(v)$, is the set of simplices in $\mathrm{St}(v)$ whose vertices except v appear before v in this order. The *closed lower-star* $\overline{\mathrm{Lst}}(v)$ is the closure of $\mathrm{Lst}(v)$ (i.e, consisting of simplices in $\mathrm{Lst}(v)$ and their faces). The *lower-link* $\mathrm{Llk}(v)$ is the set of simplices in $\overline{\mathrm{Lst}}(v)$ disjoint from v. Symmetrically, we can define the *upper star* $\mathrm{Ust}(v)$, *closed upper star* $\overline{\mathrm{Ust}}(v)$, and *upper link* $\mathrm{Ulk}(v)$, spanned by vertices in the star of v which appear after v in the chosen order.

One gets a filtration using the lower stars of the vertices: $K_{f(v_i)}$ in the following filtration denotes all simplices in K spanned by vertices in $\{v_1, \ldots, v_i\}$. Let v_0 denote a dummy vertex with $f(v_0) = -\infty$.

$$\emptyset = K_{f(v_0)} \subseteq K_{f(v_1)} \subseteq K_{f(v_2)} \subseteq \cdots \subseteq K_{f(v_n)} = K$$

Observe that the $K_{f(v_i)} \setminus K_{f(v_{i-1})} = \mathrm{Lst}(v_i)$ in the above filtration, that is, each time we add the lower star of the next vertex in the filtration. This filtration called the *lower star filtration* for f is studied in the book in more details. A lower stat filtration can be made simplex-wise by adding the simplices in a lower star in any order that puts a simplex after all of its faces. Figure 4.5 shows a simplex-wise lower star filtration.

Alternatively, we may consider the vertices in non-increasing order of f values and obtain an upper star filtration. For this we take $K_{f(v_i)}$ to be all simplices that have vertices in $\{v_i, v_{i+1}, \ldots, v_n\}$. Assuming a dummy vertex v_{n+1} with $f(v_{n+1}) = \infty$, one gets a filtration

$$\emptyset = K_{f(v_{n+1})} \subseteq K_{f(v_n)} \subseteq K_{f(v_{n-1})} \subseteq \cdots \subseteq K_{f(v_1)} = K$$

Observe that the $K_{f(v_i)} \setminus K_{f(v_{i+1})} = \mathrm{Ust}(v_i)$ in the above filtration, that is, each time we add the upper star of the next vertex in the filtration. This filtration called the *upper star filtration* for f is in some sense a symmetric version of the lower star filtration though they may provide different persistence pairs. An upper stat filtration can also be made simplex-wise by adding the simplices in an upper star in any order that puts a simplex after all of its faces. In this book, by default, we will assume that the function values along a filtration in non-decreasing. This means that we consider only lower filtrations by default.

Vertex functions are closely related to the so called piecewise linear functions (PL-functions). A vertex function $f : K \to \mathbb{R}$ defines a piecewise linear function (PL-function) on the underlying

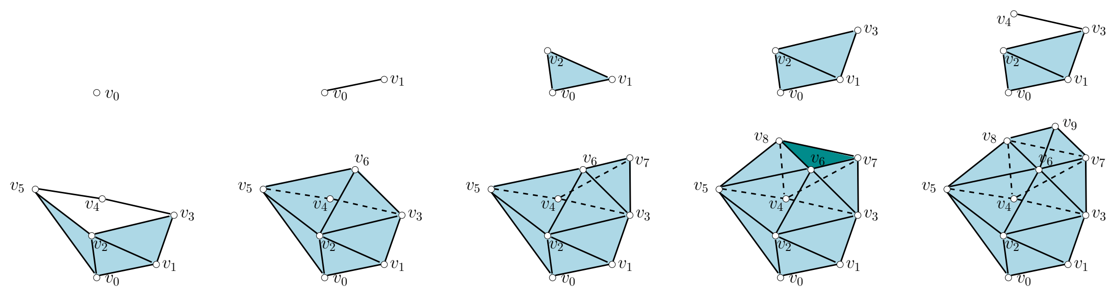

Figure 4.5: The sequence shows a *lower-star* filtration of K induced by a vertex function which is a 'height function' that records the vertical height of a vertex increasing from bottom to top here.

space $|K|$ of K which is obtained by linearly interpolating f over all simplices. On the other hand, the restriction of a PL-function to vertices trivially provides a vertex function.

Definition 2 (PL-functions)**.** Given a simplicial function K, a *piecewise-linear(PL) function* $\bar{f} : |K| \to \mathbb{R}$ is defined to be the linear extension of a vertex function $f : V(K) \to \mathbb{R}$ defined on vertices $V(K)$ of K so that for every point $x \in |K|$, $\bar{f}(x) = \sum_{i=1}^{k+1} \alpha_i f(v_i)$ where $\sigma = \{v_1, \ldots, v_{k+1}\}$ is the unique lowest dimensional simplex of dimension $k \geq 0$ containing x and $\alpha_1, \ldots, \alpha_{k+1}$ are the barycentric coordinates of x in σ.

Fact 1. *A PL-function $f : |K| \to \mathbb{R}$ naturally provides a vertex function $f : V(K) \to \mathbb{R}$. A simplex-wise lower star filtration for f is also a filtration for the simplex-wise monotonic function $\bar{f} : K \to \mathbb{R}$ where $\bar{f}(\sigma) = \max_{v \in \sigma} f(v)$. Similarly, a simplex-wise upper star filtration for f is also a filtration for the simplex-wise monotonic function $\bar{f}(\sigma) = \max_{v \in \sigma}(-f(v))$.*

Observe that a given vertex function $f : K \to \mathbb{R}$ induces a PL-function $\bar{f} : |K| \to \mathbb{R}$ whose persistence on the topological space $|K|$ can be defined by taking sublevel sets at critical values and then applying Definition 4. The relation of this persistence to the persistence of the lower star filtration of K induced by f is studied in the book. Indeed, the persistence of $\bar{f}$ can be read from the persistence of f.

Finally, we note that any simplicial filtration $\mathcal{F}$ can naturally be induced by a function. We introduce this association for unifying the definition of persistence pairing later in Definition 7.

Definition 3 (Filtration function)**.** If a simplicial filtration $\mathcal{F}$ is obtained from a simplex-wise monotone function or a vertex function f, then $\mathcal{F}$ is induced by f. Conversely, if $\mathcal{F}$ is given without any explicit input function, we say $\mathcal{F}$ is induced by the simplex-wise monotone function f where every simplex $\sigma \in (K_i \setminus K_{i-1})$ is given the value $f(\sigma) = i$.

Naturally, every simplicial filtration gives rise to a sequence of homomorphisms $h_p^{i,j}$ as in Equation 4.2 induced by inclusions again forming a homology module

$$0 = \mathsf{H}_p(K_0) \to \mathsf{H}_p(K_1) \to \cdots \to \mathsf{H}_p(K_i) \xrightarrow{h_p^{i,j}} \cdots \to \mathsf{H}_p(K_j) \cdots \to \mathsf{H}_p(K_n) = \mathsf{H}_p(K).$$

4.2 Persistence

In both cases of space and simplicial filtration $\mathcal{F}$, we arrive at a homology module:

$$\mathsf{H}_p\mathcal{F} : 0 = \mathsf{H}_p(X_0) \to \mathsf{H}_p(X_1) \to \cdots \to \mathsf{H}_p(X_i) \xrightarrow{h_p^{i,j}} \mathsf{H}_p(X_j) \cdots \to \mathsf{H}_p(X_n) = \mathsf{H}_p(X) \quad (4.3)$$

where $X_i = \mathbb{T}_{a_i}$ if $\mathcal{F}$ is a space filtration of a topological space $X = \mathbb{T}$ or $X_i = K_i$ if $\mathcal{F}$ is a simplicial filtration of a simplicial complex $X = K$. Persistent homology groups for a homology module are algebraic structures capturing the survival of the homology classes through this sequence.In general, we will call homology modules as *persistence modules* in Section 4.3 recognizing that we can replace homology groups with vector spaces.

Definition 4 (Persistent Betti number)**.** The *p-th persistent homology groups* are the images of the homomorphisms; $\mathsf{H}_p^{i,j} = \text{im } h_p^{i,j}$, for $0 \le i \le j \le n$. The *p-th persistent Betti numbers* are the dimensions $\beta_p^{i,j} = \dim \mathsf{H}_p^{i,j}$ of the vector spaces $\mathsf{H}_p^{i,j}$.

The p-th persistent homology groups contain the important information of when a homology class is born or when it dies. The issue of birth and death of a class becomes more subtle because when a new class is born, many other classes that are sum of this new class and any other existing class also are born. Similarly, when a class ceases to exist, many other classes also do so along with it. Therefore, we need a mechanism to pair births and deaths canonically.

Observe that the non trivial elements of p-th persistent homology groups $\mathsf{H}_p^{i,j}$ consist of classes that survive from X_i to X_j, that is, the classes which do not get 'quotient out' by the boundaries in X_j. So, one can observe

Fact 2. $\mathsf{H}_p^{i,j} = \mathsf{Z}_p(X_i)/(\mathsf{B}_p(X_j) \cap \mathsf{Z}_p(X_i))$ *and* $\beta_p^{i,j} = \dim \mathsf{H}_p^{ij}$.

We now formally state when a class is born or dies.

Definition 5 (Birth and death)**.** A p-th homology class $\xi \in \mathsf{H}_p(X_a)$ is born at X_i, $i \le a$, if $\xi \in \mathsf{H}_p^{i,a}$ but $\xi \notin \mathsf{H}_p^{i-1,a}$. Similarly, a p-th homology class $\xi \in \mathsf{H}_p(X_a)$ dies entering X_j, $a < j$, if $h_p^{a,j-1}(\xi)$ is not zero (non-trivial) but $h_p^{a,j}(\xi) = 0$.

Observe that not all classes that are born at X_i necessarily die entering X_j though more than one such may do so.

Fact 3. *Let* $[c] \in \mathsf{H}_p(X_{j-1})$ *be a p-th homology class that dies entering* X_j*. Then, it is born at* X_i *if and only if there exists a sequence* $i_1 \le i_2 \le \cdots \le i_k = i$ *for some* $k \ge 1$ *so that (i)* $0 \ne [c_{i_\ell}] \in \mathsf{H}_p(X_{j-1})$ *is born at* X_{i_ℓ} *for every* $\ell \in \{1, \ldots, k\}$ *and (ii)* $[c] = [c_{i_1}] + \cdots + [c_{i_k}]$.

One may interpret the above fact as follows. When a class dies, it may be thought of as a merge of several classes including the trivial one among which the youngest one ($[c_{i_k}]$) determines the birth point. This viewpoint is particularly helpful while pairing simplices in the persistence algorithm (PairPersistence) presented later.

Notice that each X_i, $i = 0, \ldots, n$, is associated with a value of the function f that induces $\mathcal{F}$. For a space filtration, we say $f(X_i) = a_i$ where $X_i = \mathbb{T}_{a_i}$. For a simplicial filtration, we say $f(X_i) = a_i$ where $a_i = f(\sigma)$ for any $\sigma \in X_i$ when the filtration function (Definition 3) is simplex-wise monotone. When it is a vertex function f, then we extend f to a simplex-wise monotone function as defined in Fact 1.

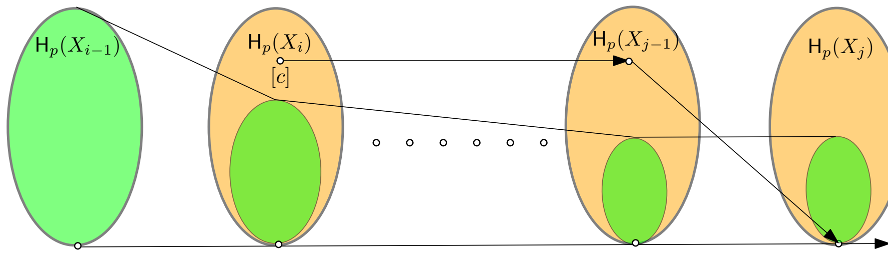

Figure 4.6: A simplistic view of birth and death of classes: The class $[c]$ is born at X_i since it is not in the image of $\mathsf{H}_p(X_{i-1})$. It dies entering X_j since this is the first time its image becomes trivial.

4.2.1 Persistence diagram

A visual representation of the persistent homology can be created by drawing a collection of points in the plane; see Figure 4.7(left). Consider the extended plane $\bar{\mathbb{R}}^2 := (\mathbb{R} \cup \{\pm\infty\})^2$ on which we represent a birth at a_i paired with the death at a_j as a point (a_i, a_j). This pairing uses a *persistence pairing function* μ_p^{ij} defined below. Strictly positive values of this function correspond to points in the persistence diagram defined later. In what follows, to account for classes that never die, we extend the induced module in Eqn.(4.3) on the right end by assuming that $\mathsf{H}_p(X_{n+1}) = 0$.

Definition 6. For $0 < i < j \leq n+1$, define

$$\mu_p^{i,j} = (\beta_p^{i,j-1} - \beta_p^{i,j}) - (\beta_p^{i-1,j-1} - \beta_p^{i-1,j}). \tag{4.4}$$

The first difference on the RHS counts the number of independent classes that are born at or before X_i and die entering X_j. The second difference counts the number of independent classes that are born at or before X_{i-1} and die entering X_j. The difference between the two differences thus counts the number of independent classes that are born at X_i and die entering X_j. When $j = n+1$, $\mu_p^{i,n+1}$ counts the number of independent classes that are born at X_i and die entering X_{n+1}. They remain alive till the end in the original filtration without extension, or we say that they never die. To emphasize that classes which exist in X_n actually never die, we equate $n+1$ with ∞ and take $a_{n+1} = a_\infty = \infty$. Observe that, with this assumption, we have $\beta^{i,n+1} = \beta^{i,\infty} = 0$ for every $i \leq n$.

Remark 1. *The p-th homology classes in $\mathsf{H}_p(X_{j-1})$ that get born at X_i and die entering X_j may not form a vector space. Hence, we cannot talk about its dimension. In fact, definition of μ_p^{ij}, in some sense, compensates for this limitation. This definition involves alternating sums of dimensions (β_{ij}'s) of vector spaces. The dimensions appearing with the negative signs lead to this anomaly. However, one can express μ_p^{ij} as the dimension of a vector space which is a quotient of a subspace.*

Definition 7 (Class persistence)**.** For $\mu_p^{i,j} \neq 0$, the persistence $\mathrm{Pers}([c])$ of a class $[c]$ that is born at X_i and dies at X_j is defined as $\mathrm{Pers}([c]) = a_j - a_i$. When $j = n+1 = \infty$, $\mathrm{Pers}([c])$ equals $a_{n+1} - a_i = \infty$.

Notice that, values a_is can be the index i when no explicit function is given (Definition 3). In that case, persistence of a class sometimes referred as *index persistence* which is $j - i$.

Definition 8 (Persistence diagram)**.** The persistence diagram $\mathrm{Dgm}_p(\mathcal{F})$ (also written $\mathrm{Dgm}_p f$) of a filtration induced by a function f is obtained by drawing a point (a_i, a_j) with non-zero multiplicity μ_p^{ij}, $i < j$, on the extended plane where the diagonal $\Delta : \{(x, x)\}$ is added with infinite multiplicity.

The addition of the diagonal is a technical necessity for results that we will see afterward.

A class born at a_i and never dying is represented as a point (a_i, ∞) (point v in Figure 4.7) – we call such point in the persistence diagram as *essential persistent point*, and their corresponding homology classes as *essential homology classes*. Classes may have the same coordinates because they may be born and die at the same time. This happens only when we allow mutiple homology classes being created or destroyed at the same function value or filtration point. In general, this also opens up the possibility of creating infinitely many birth-death pairs even if the filtration is finite. To avoid such pathological cases, we always assume that the linear maps in the homology modules have finite rank, a condition known as q-tameness in the literature [4].

There is also an alternative representation of persistence called *barcode* where each birth-death pair (a_i, a_j) is represented by a line segment $[a_i, a_j)$ called a *bar* which is open on the right. The open end signifies that the class dying entering X_j does not exist in X_j. Points at infinity such as (a_i, ∞) are represented with a ray $[a_i, \infty)$ giving an *infinite bar*. See Figure 4.7(right).

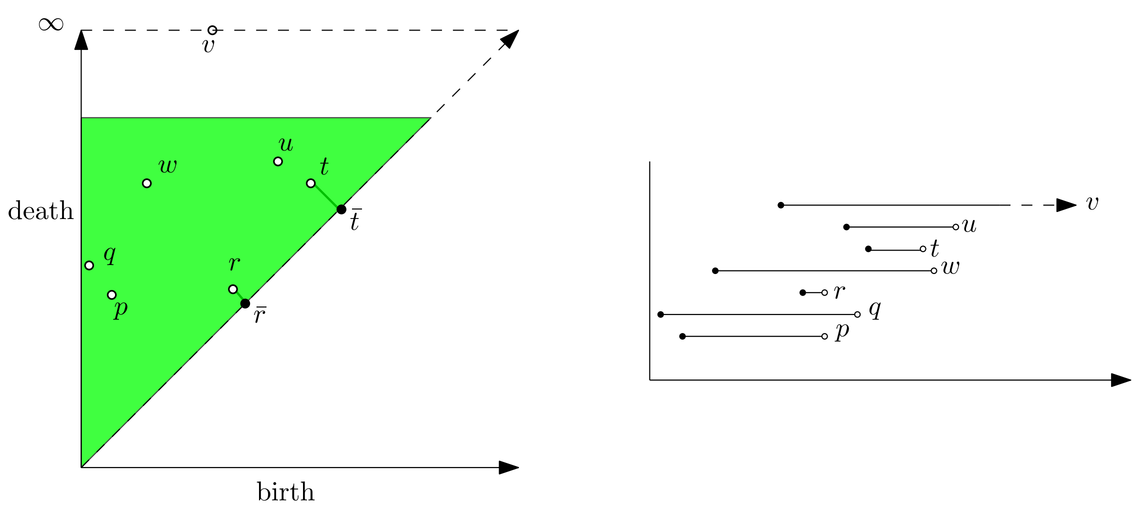

Figure 4.7: Persistence diagram with non-diagonal points only in the positive quadrant and the corresponding barcode.

Fact 4.

1. *If a class has persistence s, then the point representing it will be at a Euclidean distance $s/\sqrt{2}$ from the diagonal Δ*

2. *For sublevel set filtrations, all points (a_i, a_j) representing a class have $a_i \leq a_j$, so they lie on or above the diagonal.*

3. *If m_i denote the multiplicity of an essential point (a_i, ∞) in $\mathrm{Dgm}_p(\mathcal{F})$, where $\mathcal{F}$ is a filtration of X, one has $\Sigma_i m_i = \dim \mathsf{H}_p(X)$, the p-th Betti number of $X = X_n$.*

Here is one important fact relating persistent Betti numbers and persistence diagrams.

Theorem 1. *For every pair of indices $0 \leq k \leq \ell \leq n$ and every p, the p-th persistent Betti number satisfies $\beta_p^{k,\ell} = \sum_{i \leq k} \sum_{j > \ell} \mu_p^{i,j}$.*

Observe that $\beta_p^{k,\ell}$ is the number of points in the upper left quadrant of the corner (a_k, a_ℓ). A class that is born at X_i and dies enetering X_j is counted for $\beta_p^{k,\ell}$ iff $i \leq k$ and $j > \ell$. The quadrant is therefore closed on the right and open on the bottom.

Stability of persistence diagrams. A persistence diagram $\mathrm{Dgm}_p(\mathcal{F}_f)$, as a set of points in the extended plane $\bar{\mathbb{R}}^2 := (\mathbb{R} \cup \{\pm\infty\})^2$, summarizes certain topological information of a simplicial complex (space) in relation to the function f that induces the filtration $\mathcal{F}_f$. However, this is not useful in practice unless we can be certain that a slight change in f does not change this diagram dramatically. In practice f is seldom measured accurately, and if its persistence diagram can be approximated from a slightly perturbed version, it becomes useful. Fortunately, persistence diagrams are stable. To formulate this stability, we need a notion of distances between persistence diagrams.

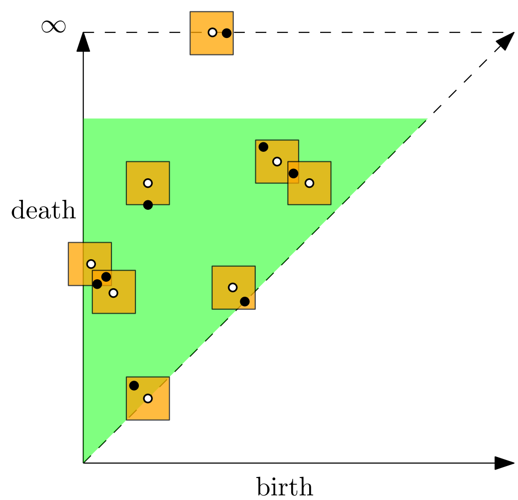

Figure 4.8: Two persistence diagrams and their bottleneck distance which is half of the side lengths of the squares representing bijections.

Let $\mathrm{Dgm}_p(\mathcal{F}_f)$ and $\mathrm{Dgm}_p(\mathcal{F}_g)$ be two persistence diagrams for two functions f and g. We want to consider bijections between points from $\mathrm{Dgm}_p(\mathcal{F}_f)$ and $\mathrm{Dgm}_p(\mathcal{F}_g)$. However, they may have different caridinality of off-diagonal points. Recall that persistence diagrams include the points on the diagonal Δ each with infinite multiplicity. This addition allows us to borrow points from the diagonal when necessary to define the bijections. We note that we are here considering only filtrations of finite complexes which also makes each homology group finite.

Definition 9 (Bottleneck distance)**.** Let $\Pi = \{\pi\}$ denote the set of all bijections $\pi : \mathrm{Dgm}_p(\mathcal{F}_f) \to \mathrm{Dgm}_p(\mathcal{F}_g)$. Consider the distance between two points $x = (x_1, x_2)$ and $y = (y_1, y_2)$ in L_∞-norm $\|x - y\|_\infty = max\{|x_1 - x_2|, |y_1 - y_2|\}$. The bottleneck distance between the two diagrams is:

$$\mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) = \inf_{\pi\in\Pi} \sup_{x\in\mathrm{Dgm}_p(\mathcal{F}_f)} \|x - \pi(x)\|_\infty.$$

Fact 5. d_b *is a metric on the space of persistence diagrams. Clearly,* $\mathsf{d}_b(X, Y) = 0$ *iff* $X = Y$. *Moreover,* $\mathsf{d}_b(X, Y) = \mathsf{d}_b(Y, X)$ *and* $\mathsf{d}_b(X, Y) \le \mathsf{d}_b(X, Z) + \mathsf{d}_b(Z, Y)$.

There is a caveat for the above fact. If d_b is taken as a distance on the space of homology modules $\mathsf{H}_p\mathcal{F}$ instead of the persistence diagrams $\mathrm{Dgm}_p(\mathcal{F})$ they generate, that is, if we define $\mathsf{d}_b(\mathsf{H}_p\mathcal{F}_f, \mathsf{H}_p\mathcal{F}_g) := \mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g))$, then it may not be a metric. The first axiom for metric becomes false if the homology modules are allowed to have classes created and destroyed at the same function values. These classes of zero persistence generate points on the diagonal Δ in the diagram. Since points on the diagonal have infinite multiplicity, two modules differing in the number of such classes of zero persistence may have diagrams with zero bottleneck distance. If we allow such cases, d_b becomes a pseudometric on the space of homology modules meaning that it satisfies all axioms of a metric except the first one.

The following theorems originally proved in [5] and further detailed in [8] quantify the notion of the stabilty of the persistence diagram. There are two versions, one involves simplicial filtrations and another involves space filtrations. For two functions, $f, g : X \to \mathbb{R}$, the infinity norm is defined as $\|f - g\|_\infty := \sup_{x\in X} |f(x) - g(x)|$.

Theorem 2 (Stability for filtrations)**.** *Let* $f, g : K \to \mathbb{R}$ *be two simplex-wise monotone functions giving rise to two simplicial filtrations* $\mathcal{F}_f$ *and* $\mathcal{F}_g$. *Then, for every* $p \ge 0$,

$$\mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) \le \|f - g\|_\infty.$$

For the second version of the stability theorem, we require that the functions referred in the theorem are ‘nice’ in the sense that they are tame. A function $f : X \to \mathbb{R}$ is *tame* if the homology groups of its sublevel sets have finite rank and these ranks change only at finitely many values called *critical*.

Theorem 3 (Stability for spaces)**.** *Let* X *be a triangulable space and* $f, g : X \to \mathbb{R}$ *be two tame functions giving rise to two space filtrations* $\mathcal{F}_f$ *and* $\mathcal{F}_g$ *where the sublevel sets are taken for critical values. Then, for every* $p \ge 0$,

$$\mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) \le \|f - g\|_\infty.$$

There is another distance called q-Wasserstein distance with which also persistence diagrams are often compared.

Definition 10 (Wasserstein distance)**.** Let Π be the set of bijections as defined in Definition 9. For any $p \geq 0$, $q \geq 1$, the q-Wasserstein distance is define as

$$\mathsf{d}_{W,q}(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) = \inf_{\pi \in \Pi} \left[\Sigma_{x \in \mathrm{Dgm}_p(\mathcal{F}_f)} \left(\|x - \pi(x)\|_q \right)^q \right]^{1/q}.$$

The distance $\mathsf{d}_{W,q}$ also is a metric on the space of persistence diagrams just like the bottleneck distance. It also enjoys a stability property though it is not as strong as in Theorem 3.

Fact 6. *Let $f, g : X \to \mathbb{R}$ be two Lipschitz functions defined on a triangulable compact metric space X. Then, there exist constants C and k depending on X and the Lipschitz constants of f and g so that for every $p \geq 0$ and $q \geq k$,*

$$\mathsf{d}_{W,q}(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) \leq C \cdot \|f - g\|_\infty^{1-\frac{k}{q}}.$$

Recently, a better stability result has appeared in the literature [17] for Wasserstein distance which considers the L^q-distance between functions defined on a common domain X:

$$\|f - g\|_q = \left(\Sigma_{x \in X} |f(x) - g(x)|^q\right)^{1/q}.$$

Theorem 4 (Stability for Wasserstein distance)**.** *Let $f, g : K \to \mathbb{R}$ be two simplex-wise monotone functions on a simplicial complex K. Then, one has*

$$\mathsf{d}_{W,q}(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) \leq \|f - g\|_q.$$

Bottleneck distances can be computed using perfect matchings in bipartite graphs. Computing Wasserstein distances become more difficult. It can be computed using an algorithm for minimum weight perfect matching in weighted bipartite graphs. We leave it as an Exercise question (Exercise 2).

Computing bottleneck distances.

Let A and B be the non-diagonal points in two persistence diagrams $\mathrm{Dgm}_p(\mathcal{F}_f)$ and $\mathrm{Dgm}_p(\mathcal{F}_g)$ respectively. For a point $a \in A$, let $\bar{a}$ denote the nearest point of a on the diagonal. Define $\bar{b}$ for every point $b \in B$ similarly. Let $\bar{A} = \{\bar{a}\}$ and $\bar{B} = \{\bar{b}\}$. Let $\tilde{A} = A \cup \bar{B}$ and $\tilde{B} = B \cup \bar{A}$. We want to bijectively match points in $\tilde{A}$ and $\tilde{B}$. Let $\Pi = \{\pi\}$ denote such a matching. It follows from the definition that

$$\mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g)) = \min_{\pi \in \Pi} \sup_{a \in \tilde{A}, \pi(a) \in \tilde{B}} \|a - \pi(a)\|_\infty.$$

Then, the bottleneck distance we want to compute must be L_∞ distance $\max\{|x_a - x_b|, |y_a - y_b|\}$ for two points $a \in \tilde{A}$ and $b \in \tilde{B}$. We do a binary search on all such possible $O(n^2)$ distances where $|\tilde{A}| = |\tilde{B}| = n$. Let $\delta_0, \delta_1, \cdots, \delta_{n'}$ be the sorted sequence of these distances in a non-decreasing order.

Given a $\delta = \delta_i \geq 0$ where i is the median of the index in the binary search interval $[\ell, u]$, we construct a bipartite graph $G = (\tilde{A} \cup \tilde{B}, E)$ where an edge $e = (a, b)_{\{a \in \tilde{A}, b \in \tilde{B}\}}$ is in E if and only if either both $a \in \bar{A}$ and $b \in \bar{B}$ (weight$(e) = 0$) or $\|a - b\|_\infty \leq \delta$ (weight$(e) = \|a - b\|_\infty$). A complete matching in G is a set of n edges so that every vertex in $\tilde{A}$ and $\tilde{B}$ is incident to exactly one edge in the set. To determine if G has a complete matching, one can use an $O(n^{2.5})$ algorithm of Hopcroft and Karp [13] for complete matching in a bipartite graph. However, exploiting the geometric embedding of the points in the persistence diagrams, we can apply an $O(n^{1.5})$ time algorithm of Efrat et al. [10] for the purpose. If such an algorithm affirms that a complete matching exists, we do the following: if $\ell = u$ we output δ, otherwise we set $u = i$ and repeat. If no matching exists, we set $\ell = i$ and repeat. Observe that matching has to exist for some value of δ, in particular for $\delta_{n'}$ and thus the binary search always succeeds. Algorithm 1: Bottleneck lays out the pseudocode for this matching. The algorithm runs in $O(n^{1.5} \log n)$ time accounting for the $O(\log n)$ probes for binary search each applying $O(n^{1.5})$ time matching algorithm. However, to achieve this complexity, we have to avoid sorting $n' = O(n^2)$ values taking $O(n^2 \log n)$ time. Again, using the geometric embedding of the points, one can perform the binary probes without incurring the cost for sorting. For details and an efficient implementation of this algorithm see [14].

Algorithm 1 Bottleneck($\mathrm{Dgm}_p(\mathcal{F}_f)$,$\mathrm{Dgm}_p(\mathcal{F}_g)$)

Input:
Two persistent diagrams $\mathrm{Dgm}_p(\mathcal{F}_f)$, $\mathrm{Dgm}_p(\mathcal{F}_g)$
Output:
Bottleneck distance $\mathsf{d}_b(\mathrm{Dgm}_p(\mathcal{F}_f), \mathrm{Dgm}_p(\mathcal{F}_g))$

Compute sorted distances $\delta_0 \leq \delta_1 \leq \cdots \leq \delta_{n'}$ from $\mathrm{Dgm}_p(\mathcal{F}_f)$ and $\mathrm{Dgm}_p(\mathcal{F}_g)$;
$\ell := 0$; $u = n'$;
while $\ell < u$ **do**
  $i := \lfloor \frac{(u+\ell)}{2} \rfloor$; $\delta := \delta_i$
  Compute graph $G = (\tilde{A} \cup \tilde{B}, E)$ where $\forall e \in E$, $weight(e) \leq \delta$
  if $\exists$ complete matching in G **then**
    $u := i$
  else
    $\ell := i$
  end if
end while
Output δ

4.3 Persistence modules

We have seen in Section 4.2.1 that persistence diagrams are stable with respect to the perturbation of the function that defines the filtration on a given simplicial complex or a space. This requires the domain of the function to be fixed. The result depends on the observation that perturbations in the filtrations are bounded by the perturbations in the function which in turn also results into bounded perturbations at the homology level. A natural follow up is to derive a bound of the perturbations of the persistence diagrams directly in terms of the perturbations at the homology level. Toward this goal, we now define a generalized notion of homology modules called *persistence modules* and a distance among them called the *interleaving distance*.

Recall that a filtration gives rise to a homology module which is a sequence of homology groups connected by homomorphisms that are induced by inclusions defining the filtration. These homology groups when defined over a field (e.g. $\mathbb{Z}_2$) can be thought of as vector spaces connected by linear maps. Persistence modules extend homology modules by taking vector spaces in place of homology groups and linear maps in place of inclusion induced homomorphisms.

We make one more extension. So far, the sequences in a filtration and homology modules have been indexed over a finite subset of natural numbers. It turns out that we can enlarge the index set to be any poset, say a subset of $\mathbb{R}$.

Definition 11 (persistence module)**.** A *persistence module* over a poset $A \subseteq \mathbb{R}$ is any collection $\mathbb{V} = \{\mathrm{V}_a\}_{a\in A}$ of vector spaces V_a together with linear maps $v_{a,a'} : \mathrm{V}_a \to \mathrm{V}_{a'}$ so that $v_{a,a} = \mathrm{id}$ and $v_{a',a''} \circ v_{a,a'} = v_{a,a''}$ for all $a, a', a'' \in A$ where $a \leq a' \leq a''$. Sometimes we write $\mathbb{V} = \{\mathrm{V}_a \xrightarrow{v_{a,a'}} \mathrm{V}_{a'}\}_{a\leq a'}$ to denote this collection with the maps.

Remark 2. *A persistence module defined over a subposet A of* $\mathbb{R}$ *can be 'extended' into a module over* $\mathbb{R}$*. For this, for any* $a < a' \in A$ *where the open interval* (a, a') *is not in A and for any* $a \leq b < b' < a'$*, assume that* $v_{b,b'}$ *is an isomorphism and* $\lim_{a\to-\infty} V_a = 0$ *if it is not given.*

Our goal is to define a distance between two persistence modules with respect to which we would bound the distance between their persistence diagrams. Given two persistence modules defined over the index set $\mathbb{R}$, we define a distance between them by identifying maps between constituent vector spaces of the modules.

Definition 12 (ε-interleaving)**.** Let $\mathbb{U}$ and $\mathbb{V}$ be two persistence modules over the index set $\mathbb{R}$. We say $\mathbb{U}$ and $\mathbb{V}$ are ε-interleaved if there exist two families of maps $\varphi_a : \mathrm{U_a} \to \mathrm{V}_{a+\varepsilon}$ and $\psi_a : \mathrm{V_a} \to \mathrm{U}_{a+\varepsilon}$ satisfying the following two conditions:

1. $v_{a+\varepsilon,a'+\varepsilon} \circ \varphi_a = \varphi_{a'} \circ u_{a,a'}$ and $u_{a+\varepsilon,a'+\varepsilon} \circ \psi_a = \psi_{a'} \circ v_{a,a'}$ [rectangular commutativity]

2. $\psi_{a+\varepsilon} \circ \varphi_a = u_{a,a+2\varepsilon}$ and $\varphi_{a+\varepsilon} \circ \psi_a = v_{a,a+2\varepsilon}$ [triangular commutativity]

The two parallelograms and the two triangles below depict the rectangular and the triangular commutativities respectively.

$$\begin{tikzcd}
\mathrm{U}_a \arrow[rr, "u_{a,a'}"] \arrow[dr, "\varphi_a"] & & \mathrm{U}_{a'} \arrow[dr, "\varphi_{a'}"] & \\
& \mathrm{V}_{a+\varepsilon} \arrow[rr, "v_{a+\varepsilon,a'+\varepsilon}"'] & & \mathrm{V}_{a'+\varepsilon}
\end{tikzcd}
\qquad
\begin{tikzcd}
& \mathrm{U}_{a+\varepsilon} \arrow[rr, "u_{a+\varepsilon,a'+\varepsilon}"] & & \mathrm{V}_{a'+\varepsilon} \\
\mathrm{V}_a \arrow[ur, "\psi_a"] \arrow[rr, "v_{a,a'}"'] & & \mathrm{V}_{a'} \arrow[ur, "\psi_{a'}"] &
\end{tikzcd}$$

$$\begin{tikzcd}
\mathrm{U}_a \arrow[rr, "u_{a,a+2\varepsilon}"] \arrow[dr, "\varphi_a"] & & \mathrm{U}_{a+2\varepsilon} \\
& \mathrm{V}_{a+\varepsilon} \arrow[ur, "\psi_{a+\varepsilon}"'] &
\end{tikzcd}
\qquad
\begin{tikzcd}
& \mathrm{U}_{a+\varepsilon} \arrow[dr, "\varphi_{a+\varepsilon}"] & \\
\mathrm{V}_a \arrow[ur, "\psi_a"] \arrow[rr, "v_{a,a+2\varepsilon}"] & & \mathrm{V}_{a+2\varepsilon}
\end{tikzcd}$$

Definition 13 (interleaving distance)**.** Given two persistence modules $\mathbb{U}$ and $\mathbb{V}$, their interleaving distance is defined as

$$d_I(\mathbb{U}, \mathbb{V}) = \inf\{\varepsilon \mid \mathbb{U} \text{ and } \mathbb{V} \text{ are } \varepsilon\text{-interleaved}\}$$

Observe that, when $\varepsilon = 0$, Definition 12 implies that the maps $\varphi_a : \mathrm{U}_a \to \mathrm{V}_a$ and $\psi_a : \mathrm{V}_a \to \mathrm{U}_a$ are isomorphisms. In that case, we get the following diagrams where each vertical map is an isomorphism and each square commutes. We get two isomorphic persistence modules.

$$\begin{tikzcd}
\mathbb{U}: & \cdots\cdots \arrow[r] & \mathrm{U}_a \arrow[r] \arrow[d, equal] & \cdots & \mathrm{U}_{a'} \arrow[r] \arrow[d, equal] & \cdots\cdots \\
\mathbb{V}: & \cdots\cdots \arrow[r] & \mathrm{V}_a \arrow[r] & \cdots & \mathrm{V}_{a'} \arrow[r] & \cdots\cdots
\end{tikzcd}$$

Definition 14 (isomorphic persistence modules)**.** We say two persistence modules $\mathbb{U}$ and $\mathbb{V}$ indexed over an index set A are isomorphic if the following two conditions hold (illustrated by the diagram above).

1. $\mathrm{U}_a \cong \mathrm{V}_a$ for every $a \in \mathbb{R}$, and
2. for every $x \in \mathrm{U}_a$, if x is mapped to $y \in \mathrm{V}_a$ by the isomorphism, then $u_{a,a'}(x) \in \mathrm{U}_{a'}$ is mapped to $v_{a,a'}(y) \in \mathrm{V}_{a'}$ also by the isomorphism.

Fact 7. *If two persistence modules arising from two filtrations $\mathcal{F}_f$ and $\mathcal{F}_g$ are isomophic, the persistence digrams* $\mathrm{Dgm}_p(\mathcal{F}_f)$ *and* $\mathrm{Dgm}_p(\mathcal{F}_g)$ *are identical.*

Now we relate the interleaving distance between two persistence modules and the persistence diagrams they define. For this, we consider a special type of a persistence module called *interval module*. Below, we use the standard convention that an open end of an interval is denoted with the first brcakets '(' or ')' and a closed end of an interval with the third brackets '[' or ']'.

Definition 15 (Interval module)**.** Given an index set $A \subseteq \mathbb{R}$ and a pair of indices $b, d \in A$, $b \le d$, four types of *interval modules* denoted $\mathbb{I}[b,d)$, $\mathbb{I}(b,d]$, $\mathbb{I}[b,d]$, $\mathbb{I}(b,d)$ respectively are special persistence modules defined as:

- (closed-open): $\mathbb{I}[b,d) : \{\mathrm{V_a} \xrightarrow{v_{a,a'}} \mathrm{V_{a'}}\}_{a,a' \in A}$ where (i) $\mathrm{V}_a = \mathbb{Z}_2$ for all $a \in [b,d)$ and $V_a = 0$ otherwise, (ii) $v_{a,a'}$ is identity map for $b \le a \le a' < d$ and zero map otherwise.
- (open-closed): $\mathbb{I}(b,d] : \{\mathrm{V_a} \xrightarrow{v_{a,a'}} \mathrm{V_{a'}}\}_{a,a' \in A}$ where (i) $\mathrm{V}_a = \mathbb{Z}_2$ for all $a \in (b,d]$ and $V_a = 0$ otherwise, (ii) $v_{a,a'}$ is identity map for $b < a \le a' \le d$ and zero map otherwise.
- (closed-closed): $\mathbb{I}[b,d] : \{\mathrm{V_a} \xrightarrow{v_{a,a'}} \mathrm{V_{a'}}\}_{a,a' \in A}$ where (i) $\mathrm{V}_a = \mathbb{Z}_2$ for all $a \in [b,d]$ and $V_a = 0$ otherwise, (ii) $v_{a,a'}$ is identity map for $b \le a \le a' \le d$ and zero map otherwise.
- (open-open): $\mathbb{I}(b,d) : \{\mathrm{V_a} \xrightarrow{v_{a,a'}} \mathrm{V_{a'}}\}_{a,a' \in A}$ where (i) $\mathrm{V}_a = \mathbb{Z}_2$ for all $a \in (b,d)$ and $V_a = 0$ otherwise, (ii) $v_{a,a'}$ is identity map for $b < a \le a' < d$ and zero map otherwise.

In general, we denote the four types of interval modules as $\mathbb{I}\langle b, d\rangle$ being oblivious to the particular type. The two end points b, d signify the birth and the death points of the interval in analogy to the bars we have seen for persistence diagrams. This is why sometimes we also write $\mathbb{I}\langle b, d\rangle = \langle b, d\rangle$. Gabriel [12] showed that a persistence module decomposes uniquely into interval modules when the index set is finite. This condition can be relaxed further as stated in proposition below. A persistence module $\mathbb{U}$ for which each of the vectors spaces U_a, $a \in \mathbb{R}$ has finite dimension is called a *pointwise finite dimensional* (p.f.d. in short) persistence module. A persistence module for which the connecting linear maps have finite rank is called *q-tame*. The results below are part of a more general concept called *quiver theory*.

Proposition 5.

- *Any persistence module over a finite index set decomposes uniquely into closed-closed interval modules, that is,* $\mathbb{U} \cong \bigoplus_{j \in J} \mathbb{I}[b_j, d_j]$ *[12].*
- *Any p.f.d. persistence module decomposes uniquely into interval modules, that is,* $\mathbb{U} \cong \bigoplus_{j \in J} \mathbb{I}\langle b_j, d_j\rangle$ *[7, 18].*
- *Any q-tame persistence module decomposes uniquely into interval modules [4].*

The birth and death points of the interval modules that a given persistence module $\mathbb{U}$ decomposes into (Proposition 5) can be plotted as points in $\mathbb{R}^2$. This can define a persistence diagram $\mathrm{Dgm}\mathbb{U}$ for a persistence module $\mathbb{U}$. We aim to relate the interleaving distance between persistence modules and the bottleneck distance between their persistence diagram thus defined.

Definition 16 (PD for persistence module)**.** Let $\mathbb{U} \cong \bigoplus_j \langle b_j, d_j\rangle$ be the interval decomposition of a given persistence module $\mathbb{U}$ (Proposition 5). The collection of points $\{(b_j, d_j)\}$ with proper multiplicity and the points on the diagonal $\Delta : \{(x, x)\}$ with infinite multiplicity constitute the persistence diagram $\mathrm{Dgm}\mathbb{U}$ of the persistence module $\mathbb{U}$.

For the index set $A = \mathbb{R}$, Chazal et al. [3] showed that the bottleneck distance between two persistence diagrams of p.f.d. modules is bounded from above by their interleaving distance. The result also holds for q-tame modules. Lesnick [15] and later Lesnick and Bauer [1] proved that the two distances are indeed equal.

Theorem 6. *Given two q-tame persistence modules defined over the totally ordered index set* $\mathbb{R}$, $\mathsf{d}_I(\mathbb{U}, \mathbb{V}) = \mathsf{d}_b(\mathrm{Dgm}\,\mathbb{U}, \mathrm{Dgm}\,\mathbb{V})$.

Remark 3. *The isometry theorem stated for the index set* $\mathbb{R}$ *does not apply directly to the persistence modules that are not defined over the index set* $\mathbb{R}$*. In this case, to define the interleaving distance, we can extend the module to be indexed over* $\mathbb{R}$ *as described in Remark 2. For example, consider a persistence module* $\mathsf{H}_p\mathcal{F}$ *obtained from a filtration* $\mathcal{F}$ *defined on a finite index set A or when* $A = \mathbb{Z}$*. Observe that, all interval modules for* $\mathsf{H}_p\mathcal{F}$ *(without extension) are of closed-closed type* $[b, d]$ *for some* $b, d \in A$*. This brings out a subtlety. The intervals of the form* $[b, d]$ *where* $b = d$ *are mapped to the diagonal* Δ *in the persistence diagram (Definition 8). These points get ignored while computing the bottleneck distance as both diagrams have the diagonal points with infinite multiplicity. In fact, the isometry theorem (Theorem 6) does not hold if this is not* *taken care of. To address the issue, for persistence modules* $\mathsf{H}_p\mathcal{F}$ *generated by a finite filtration* $\mathcal{F}$*, we can map each interval* $[b, d]$ *in the decomposition of* $\mathsf{H}_p\mathcal{F}$ *to a point* $(b, d+1)$ *in* $\mathrm{Dgm}_p(\mathcal{F})$*. This aligns with the observation that, after the extension, the interval* $[b, d]$ *indeed stretches to* $[b, d+1)$.

The concept of topological persistence came to the fore in early 2000 with the paper by Edelsbrunner, Letscher, and Zomorodian [9] though the concept was proposed in a rudimentary form (for 0-dimensional homology) in other papers by Frosini [11] and Robins [16].

The concept of bottleneck distance for persistence diagrams was first proposed by Cohen-Steiner et al. [5] who also showed the stability of such diagrams in terms of bottleneck distances with respect to the infinity norm of the difference between functions generating them. This result was extended to Wassrstein distance though in a weaker form in [6] which got improved recently [17]. The more general concept of interleaving distance between persistence modules and the stability of persistence diagrams with respect them was presented by Chazal et al. [3]. The fact that bottleneck distance between persistence diagrams is not only bounded from above by interleaving distance but is indeed equal to it was shown by Lesnick [15]. Also see [2] for more generalization at algebraic level.

Exercises

1. Prove Theorem 1.

2. Give a polynomial time algorithm for computing $\mathsf{d}_{W,q}$.

3. Let $\mathcal{F}$ be a simplex-wise filtration $\mathcal{F}$ of complex K induced by the sequence of simplicies: $\sigma_1, \ldots, \sigma_N$. Let $\mathcal{F}'$ be a modification of $\mathcal{F}$ where only two consecutive simplices σ_k and σ_{k+1} swap their order; that is, $\mathcal{F}'$ is induced by the sequence:

$$\sigma_1, \ldots, \sigma_{k-1}, \sigma_{k+1}, \sigma_k, \sigma_{k+2}, \ldots, \sigma_N.$$

   Describe the relation between their corresponding persistence diagrams $\mathrm{Dgm}(\mathcal{F})$ and $\mathrm{Dgm}(\mathcal{F}')$.

4. Consider the two persistence modules $\mathbb{U}$ and $\mathbb{V}$ as shown below and a sequence of linear maps $f_i : U_i \to V_i$ so that all squares commute.

$$\begin{tikzcd} \mathbb{U}: & U_1 \arrow[r] \arrow[d, "f_1"] & U_2 \arrow[r] \arrow[d, "f_2"] & U_3 \arrow[r] \arrow[d, "f_3"] & \cdots\cdots \arrow[r] & U_m \arrow[d, "f_m"] \\ \mathbb{V}: & V_1 \arrow[r] & V_2 \arrow[r] & V_3 \arrow[r] & \cdots\cdots \arrow[r] & V_m \end{tikzcd}$$

   Consider the sequences

$$\ker f : \{\ker f_i \subseteq U_i \to \ker f_{i+1} \subseteq U_{i+1}\}$$

   where the maps are induced from the module $\mathbb{U}$. Prove that $\ker f$ is a persistence module. Show the same for the sequences

$$\mathrm{im}\, f : \{\mathrm{im}\, f_i \subseteq V_i \to \mathrm{im}\, f_{i+1} \subseteq \mathrm{im}\, V_{i+1}\} \text{ and}$$

$$\mathrm{coker} f : \{\mathrm{coker} f_i = V_i/\mathrm{im}\, f_i \to \mathrm{coker} f_{i+1} = V_{i+1}/\mathrm{im}\, f_{i+1}\}.$$

Bibliography

[1] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, SOCG'14, pages 355:355–355:364, 2014.

[2] Peter Bubenik, Jonathan A. Scott, and Donald Stanley. Wasserstein distance for generalized persistence modules and abelian categories. *arXiv: Rings and Algebras*, arxiv:1809.09654, 2018.

[3] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas Guibas, and Steve Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, SCG '09, pages 237–246, 2009.

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

[5] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, Jan 2007.

[6] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have L_p-stable persistence. *Found. Comput. Math.*, 10(2):127–139, 2010.

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

[8] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. Applied Mathematics. American Mathematical Society, 2010.

[9] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.

[10] Alon Efrat, Alon Itai, and Matthew J. Katz. Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31(1):1–28, 2001.

[11] Patrizio Frosini. A distance for similarity classes of submanifolds of a euclidean space. *Bulletin of the Australian Mathematical Society*, 42(3):407–415, 1990.

[12] Peter Gabriel. Unzerlegbare darstellungen i. *manuscripta mathematica*, 6(1):71–103, Mar 1972.

[13] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.

[14] Michael Kerber, Dmitriy Morozov, and Arnur Nigmetov. Geometry helps to compare persistence diagrams. *Journal of Experimental Algorithmics (JEA)*, 22(1):1–4, 2017.

[15] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015.

[16] Vanessa Robins. Towards computing homology from finite approximations. *Topology Proceedings*, 24(1):503–532, 1999.

[17] Primoz Skraba and Katharine Turner. Wasserstein stability for persistence diagrams. *CoRR*, arXiv:2006.16824, 2021.

[18] Carry Webb. Decomposition of graded modules. *Proc. American Math. Soc.*, 94(4):565–571, 1985.