

Edit Distance and Persistence Diagrams Over Lattices ^{*}

Alexander McCleary¹ and Amit Patel¹

¹Department of Mathematics, Colorado State University

Abstract

We build a functorial pipeline for persistent homology. The input to this pipeline is a filtered simplicial complex indexed by any finite metric lattice and the output is a persistence diagram defined as the Möbius inversion of its birth-death function. We adapt the Reeb graph edit distance to each of our categories and prove that both functors in our pipeline are 1-Lipschitz making our pipeline stable. Our constructions generalize the classical persistence diagram and, in this setting, the bottleneck distance is strongly equivalent to the edit distance.

1 Introduction

In the most basic setting, persistent homology takes as input a finite 1-parameter filtration $K_0 \subseteq K_1 \subseteq \cdots \subseteq K_n = K$ of a finite simplicial complex K and outputs a persistence diagram. The persistence diagram, as originally defined in [10] and equivalently in [14, 19], is roughly defined as follows. Fix a field k . For every pair of indices $\mathbf{a} \leq \mathbf{b}$, let $f_*[\mathbf{a}, \mathbf{b}]$ be the rank of the k -linear map on homology $H_*(K_{\mathbf{a}}; k) \rightarrow H_*(K_{\mathbf{b}}; k)$ induced by the inclusion $K_{\mathbf{a}} \subseteq K_{\mathbf{b}}$. The persistence diagram g_* is the assignment to every pair $\mathbf{a} \leq \mathbf{b}$ the following signed sum:

$$g_*[\mathbf{a}, \mathbf{b}] := f_*[\mathbf{a}, \mathbf{b} - 1] - f_*[\mathbf{a}, \mathbf{b}] + f_*[\mathbf{a} - 1, \mathbf{b}] - f_*[\mathbf{a} - 1, \mathbf{b} - 1]. \quad (1)$$

We interpret the integer $g_*[\mathbf{a}, \mathbf{b}]$ as the number of independent cycles that appear at $\mathbf{a}$ and become boundaries at $\mathbf{b}$. The most important property of the persistence diagram is that it is stable to perturbations of the filtration on K [10]. See §2 for a discussion of stability and applications.

In [21], the second author identifies Equation (1) as a special case of the Möbius inversion formula. This observation allows for generalizations of the persistence diagram in at least two directions. First, we may consider any coefficient ring and still get a well defined persistence diagram. In fact, any constructible 1-parameter persistence module valued in any essentially

^{*}This material is based upon work supported by the National Science Foundation under Grant No. 1717159.

small abelian category admits a persistence diagram. Furthermore, the bottleneck stability theorem of [10] generalizes to this setting [20]. Second, we may consider as input a filtration indexed over any locally finite poset and still get a well defined persistence diagram as demonstrated by [18]. However, it is only now, in this paper, that we are able to establish a statement of stability for these more general persistence diagrams thus opening a path to applications; see §2.

The persistence diagram is just one invariant of a filtered space. For example, the persistence landscape [5] is a second invariant that is better suited for machine learning and statistical methods [6, 22]. The work of [4] recasts persistence landscapes as a Möbius inversion and, in the process, uncovers previously unknown structure.

Contributions We establish and study the following pipeline of functors for persistent homology with coefficients in a fixed field k :

$$\mathrm{Fil}(K) \xrightarrow{\mathrm{ZB}_*} \mathrm{Mon} \xrightarrow{\mathrm{MI}} \mathrm{Fnc}.$$

The input to the pipeline is the category of filtrations $\mathrm{Fil}(K)$. Its objects are filtrations of finite simplicial complexes indexed over finite metric lattices. The second category Mon consists of monotone integral functions over finite metric lattices. To every object in $\mathrm{Fil}(K)$, the functor ZB_* assigns a monotone integral function. Instead of using the rank function of a filtration, as mentioned earlier, we use the birth-death function ZB_* of a filtration. The birth-death function is inspired by the work of Henselman-Petrusek and Ghrist [15, 16]. In a way, the two functions are equivalent; see §9. The third category Fnc consists of integral, but not necessarily monotone, functions over finite metric lattices. To every object in Mon , the Möbius inversion functor MI assigns its Möbius inversion, which is an integral function but not necessarily monotone. Morphisms in the first two categories are inspired by the definition of the *Reeb graph elementary edit* of Di Fabio and Landi [13]. We think of the morphisms in $\mathrm{Fil}(K)$ as a deformation of one filtration to another. This deformation is formally a Kan extension. Morphisms in Mon are defined in the same way, but morphisms in Fnc are inspired by morphisms between signed measure spaces. Our main theorem, Theorem 8.4, says that both functors in the pipeline are 1-Lipschitz. The three metrics, one for each category, are all inspired by the categorification of the *Reeb graph edit distance* by Bauer, Landi, and Mémoli [3]. Hence, we call all three metrics the *edit distance*. Finally in §9, we prove that our definition of the persistence diagram generalizes the classical definition [10, 14, 19] and that, in this setting, the edit distance between persistence diagrams is strongly equivalent to the bottleneck distance; see Theorem 9.1.

Outline We start with a discussion of stability and its importance in applications of persistent homology. In §3, we establish basic definitions and properties of metric lattices. The next three sections §4, §5, and §6 establish the three categories $\mathrm{Fil}(K)$, Mon , and Fnc , respectively, along with the functors ZB_* and MI . In §7, we define the edit distance in all three categories and prove that both functors ZB_* and MI are 1-Lipschitz. We put the pieces

together in §8 and state our main theorem. Finally in §9, we verify that our framework generalizes the classical persistence diagram, and we prove that the bottleneck distance is strongly equivalent to the edit distance.

2 Applications of Stability

Scientists study natural phenomena by collecting and analyzing data. Data is a finite set usually with additional structure such as, for example, a metric or an embedding into a vector space. The goal is to extract information and then use this information to build models or theories. However, measurements are noisy and so any information that is extracted must be stable to noise. Persistent homology is an attractive tool for studying the shape of data precisely because it extracts stable information in the form of a persistence diagram.

A major drawback of classical persistent homology is its sensitivity to outliers. An outlier is, roughly speaking, a data point that is too far away from where it should be. The persistence diagram changes drastically by the introduction of just one outlier. This is not desirable as outliers are common in many data sets. The solution requires a theory of persistent homology that can handle a two-parameter filtration of a space. Such a theory not only requires a generalization of the classical persistence diagram but also a statement of stability for the reasons described above. In this paper, we present a generalization of the persistence diagram for simplicial complexes filtered over any finite lattice, but, as mentioned in the last section, this is not entirely new. One of the achievements of this paper is a first ever statement of stability for this more general setting; see Theorem 8.4. To our surprise, our stability theorem closely resembles the bottleneck stability theorem; see Theorem 9.1.

Functoriality of persistent homology is the second main achievement of this paper. Considering that category theory has its roots in algebraic topology, it may be surprising that the assignment of a persistence diagram to a filtration was, until now, not known to be functorial. Since its inception in the 1940's, category theory has permeated every field of mathematics, logic, and parts of computer science. It is only reasonable to expect that functoriality will play a major role in the future development of applications for persistent homology. For example, the 1-parameter family of persistence diagrams associated to a time-varying data set can now be described as a constructible cosheaf [12] of persistence diagrams.

We now review a few applications of classical persistent homology that are well known within the applied topology community. All rely on stability.

Homological Inference One of the first applications of the bottleneck stability theorem is homological inference, which appears in the same paper as the theorem [10]. Data often lives along a lower dimensional subspace of a higher dimensional vector space. In this case, it is useful to assume that the data is sampled from some sufficiently nice subspace, say $X \hookrightarrow \mathbb{R}^n$. For example, X could be a smooth manifold, a compact Whitney stratified space, or a piecewise-linear embedding of a finite simplicial complex. A natural question to ask is the following: How can one infer the homology of X from a finite sample $Y \subseteq X$? The answer to this question lies in a quantity called the *homological feature size* of X . The statement is

roughly as follows. Suppose Y is chosen so that the Hausdorff distance between X and Y in $\mathbb{R}^n$ is at most $r/4$, where $r > 0$ is the homological feature size of X . Then the homology of X can be read from the classical persistence diagrams associated to Y .

The above sampling condition relies on the distance between X and the sample Y , which is highly sensitive to outliers. One way of reducing the impact of outliers is by thinking of X and Y as measures and then using the Wasserstein distance between the two measures [9]. One can then imagine setting up a two-filtration: the first parameter filters by distance, as in the case above, and the second parameter filters by mass. We suspect Theorem 8.4 implies a homological inference theorem for this two-parameter setting.

Machine Learning The growing field of data science offers a wide variety of tools for extracting information from data. However, most of these tools require as input a vector, but a persistence diagram is far from a vector. In order to make use of these tools, one must vectorize the persistence diagram. Persistence images is one popular way of vectorizing the classical persistence diagram [1]. Again, because data is inherently noise, it is crucial that any vectorization is stable to noise. Persistence images are stable.

In [17], the authors develop an input layer for deep neural networks that takes a classical persistence diagram and computes a parametrized projection that can be learned during network training. This layer is designed in a way that is stable to perturbations of the input persistence diagrams.

Our main theorem, Theorem 8.4, lays the foundation for using our Möbius inversion based, multi-parameter persistence diagrams for applications in machine learning.

3 Preliminaries

We start with an introduction to bounded lattices and bounded lattice functions. From here, we equip our lattices with a metric and discuss the distortion of a lattice function between two metric lattices.

3.1 Lattices

A *poset* is a set P with a reflexive, antisymmetric, and transitive relation $\leq$. For two elements $a, b \in P$ in a poset, we write $a < b$ to mean $a \leq b$ and $a \neq b$. For any $a \leq c$, the *interval* $[a, c] \subseteq P$ is the subposet consisting of all $b \in P$ such that $a \leq b \leq c$. The poset P has a *bottom* if there is an element $\perp \in P$ such that $\perp \leq a$ for all $a \in P$. The poset P has a *top* if there is an element $\top \in P$ such that $a \leq \top$ for all $a \in P$. A function $\alpha : P \rightarrow Q$ between two posets is *monotone* if for all $a \leq b$, $\alpha(a) \leq \alpha(b)$.

The *meet* of two elements $a, b \in P$ in a poset, written $a \wedge b$, is the greatest lower bound of a and b . The *join* of two elements $a, b \in P$ in a poset, written $a \vee b$, is the least upper bound of a and b . The poset P is a *lattice* if both joins and meets exist for all pairs of elements in P . A lattice is *bounded* if it contains both a top and a bottom. If P is a finite lattice, then the existence of meets and joins implies that P has a bottom and a top,

respectively. Therefore all finite lattices are bounded. A function $\alpha : P \rightarrow Q$ between two bounded lattices is a *bounded lattice function* if $\alpha(\top) = \top$, $\alpha(\perp) = \perp$, and for all $a, b \in P$, $\alpha(a \vee b) = \alpha(a) \vee \alpha(b)$ and $\alpha(a \wedge b) = \alpha(a) \wedge \alpha(b)$. Note that bounded lattice functions are monotone. This is because $a \leq b$ if and only if $a \wedge b = a$, and therefore

$$\alpha(a) = \alpha(a \wedge b) = \alpha(a) \wedge \alpha(b) \implies \alpha(a) \leq \alpha(b).$$

Proposition 3.1: Let P and Q be finite lattices and $\alpha : P \rightarrow Q$ a bounded lattice function. Then for all $a \in Q$, the pre-image $\alpha^{-1}[\perp, a]$ has a maximal element.

Proof. The pre-image is non-empty because $\alpha(\perp) = \perp$. The pre-image is finite because P is finite. For any two elements b and c in the pre-image, both $b \vee c$ and $b \wedge c$ are also in the pre-image because

$$\alpha(b \vee c) = \alpha(b) \vee \alpha(c) \leq a \vee a = a \quad \alpha(b \wedge c) = \alpha(b) \wedge \alpha(c) \leq a \wedge a = a.$$

Thus $\alpha^{-1}[\perp, a]$ is a finite lattice and all finite lattices have a unique maximal element. $\square$

For a finite lattice P , let $\mathbb{I}P := \{[a, b] \subseteq P : a \leq b\}$ be its set of intervals. The product order on $P \times P$ restricts to a partial order $\preceq$ on $\mathbb{I}P$ as follows: $[a, b] \preceq [c, d]$ if $a \leq c$ and $b \leq d$. The join of two intervals is $[a, b] \vee [c, d] = [a \vee c, b \vee d]$, and the meet of two intervals is $[a, b] \wedge [c, d] = [a \wedge c, b \wedge d]$. All this makes $\mathbb{I}P$ a finite lattice. Its bottom element is $[\perp, \perp]$ and its top element is $[\top, \top]$.

A bounded lattice function $\alpha : P \rightarrow Q$ between two finite lattices induces a bounded lattice function $\mathbb{I}\alpha : \mathbb{I}P \rightarrow \mathbb{I}Q$ as follows. For an interval $[a, b] \in \mathbb{I}P$, let $\mathbb{I}\alpha([a, b]) := [\alpha(a), \alpha(b)]$. We have

$$\begin{aligned} \mathbb{I}\alpha([a, b] \wedge [c, d]) &= \mathbb{I}\alpha([a \wedge c, b \wedge d]) = [\alpha(a \wedge c), \alpha(b \wedge d)] \\ &= [\alpha(a) \wedge \alpha(c), \alpha(b) \wedge \alpha(d)] = \mathbb{I}\alpha([a, b]) \wedge \mathbb{I}\alpha([c, d]). \\ \mathbb{I}\alpha([a, b] \vee [c, d]) &= \mathbb{I}\alpha([a \vee c, b \vee d]) = [\alpha(a \vee c), \alpha(b \vee d)] \\ &= [\alpha(a) \vee \alpha(c), \alpha(b) \vee \alpha(d)] = \mathbb{I}\alpha([a, b]) \vee \mathbb{I}\alpha([c, d]) \\ \mathbb{I}\alpha([\perp, \perp]) &= [\alpha(\perp), \alpha(\perp)] = [\perp, \perp] \\ \mathbb{I}\alpha([\top, \top]) &= [\alpha(\top), \alpha(\top)] = [\top, \top]. \end{aligned}$$

Thus $\mathbb{I}\alpha$ is a bounded lattice function. Further, for any pair of bounded lattice functions $\alpha : P \rightarrow Q$ and $\beta : Q \rightarrow R$, $\mathbb{I}(\beta \circ \alpha) = \mathbb{I}\beta \circ \mathbb{I}\alpha$. All this makes $\mathbb{I}$ an endofunctor on the category of finite lattices and bounded lattice functions. To minimize notation, we will write $\bar{P}$ for $\mathbb{I}P$ and $\bar{\alpha}$ for $\mathbb{I}\alpha$.

3.2 Metric Lattices

A *finite (extended) metric lattice* is a tuple (P, d_P) where P is a finite lattice and $d_P : P \times P \rightarrow \mathbb{R}^{\geq 0} \cup \{\infty\}$ an extended metric. A *morphism of finite metric lattices* $\alpha : (P, d_P) \rightarrow (Q, d_Q)$




Figure 1: Above are the Hasse diagrams of two finite metric lattices P and Q . The metrics d_P and d_Q assigns to every pair of elements the length (i.e. the number of edges) of the shortest path between them. For example, $d_P(a, d) = 2$ and $d_Q(p, q) = 1$. The function $\alpha : P \rightarrow Q$ defined as $\alpha(a) = \alpha(b) = p$ and $\alpha(c) = \alpha(d) = r$ is a bounded lattice function. The distortion of α is $\|\alpha\| = 1$.

is a bounded lattice function $\alpha : P \rightarrow Q$. The *distortion* (see [7, Definition 7.1.4]) of a morphism $\alpha : (P, d_P) \rightarrow (Q, d_Q)$, denoted $\|\alpha\|$, is

$$\|\alpha\| := \max_{a, b \in P} |d_P(a, b) - d_Q(\alpha(a), \alpha(b))|.$$

To minimize notation, we will write finite metric lattices (P, d_P) simply as P with the implied metric d_P .

Example 3.2: See Figure 1 for two examples of finite metric lattices P and Q and a morphism of finite metric lattices $\alpha : P \rightarrow Q$. The distortion of α is $\|\alpha\| = 1$. Forthcoming examples will build on this one example.

For every finite metric lattice P , we have the finite metric lattice of intervals $\bar{P}$ where $d_{\bar{P}}([a, b], [c, d]) := \max\{d_P(a, c), d_P(b, d)\}$. A morphism $\alpha : P \rightarrow Q$ of finite metric lattices induces a morphism of finite metric lattices $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$. Define the *distortion* of $\bar{\alpha}$ as

$$\|\bar{\alpha}\| := \max_{[a, b], [c, d] \in \bar{P}} |\max\{d_P(a, c), d_P(b, d)\} - \max\{d_Q(\alpha(a), \alpha(c)), d_Q(\alpha(b), \alpha(d))\}|.$$

Proposition 3.4 says that the two distortions $\|\alpha\|$ and $\|\bar{\alpha}\|$ are equal. Its proof requires the following lemma.

Lemma 3.3: [8, Lemma 3 page 31] For all non-negative real numbers $w, x, y, z \in \mathbb{R}^{\geq 0}$,

$$|\max(w, x) - \max(y, z)| \leq \max(|w - y|, |x - z|).$$

Proposition 3.4: Let $\alpha : P \rightarrow Q$ be a bounded lattice function between two finite metric lattices and let $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ be the induced bounded lattice function on intervals. Then $\|\bar{\alpha}\| = \|\alpha\|$.

Proof. First we show $\|\bar{\alpha}\| \geq \|\alpha\|$. If $\|\alpha\| = \varepsilon$, then there are elements $a, b \in P$ such that $\varepsilon = |d_P(a, b) - d_Q(\alpha(a), \alpha(b))|$. For the intervals $[a, a]$ and $[b, b]$, we have

$$|\max\{d_P(a, b), d_P(a, b)\} - \max\{d_Q(\alpha(a), \alpha(b)), d_Q(\alpha(a), \alpha(b))\}| = \varepsilon$$




Figure 2: Above are the Hasse diagrams of the lattices $\bar{P}$ and $\bar{Q}$ where P and Q are from Example 3.2. The morphism $\alpha : P \rightarrow Q$ from the same example extends to the morphism $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ as follows. The function $\bar{\alpha}$ sends $\{[a, a], [a, b], [b, b]\}$ to $\{[p, p]\}$, $\{[a, c], [a, d], [b, d]\}$ to $\{[p, r]\}$, and $\{[c, c], [c, d], [d, d]\}$ to $\{[r, r]\}$. The distortion of $\bar{\alpha}$ is $\|\bar{\alpha}\| = \|\alpha\| = 1$.

proving the claim. Now we show $\|\bar{\alpha}\| \leq \|\alpha\|$ using Lemma 3.3:

$$\begin{aligned}
\|\bar{\alpha}\| &:= \max_{[a,b],[c,d] \in \bar{P}} \left| \max \{d_P(a, c), d_P(b, d)\} - \max \{d_Q(\alpha(a), \alpha(c)), d_Q(\alpha(b), \alpha(d))\} \right| \\
&\leq \max_{[a,b],[c,d] \in \bar{P}} \left\{ |d_P(a, c) - d_Q(\alpha(a), \alpha(c))|, |d_P(b, d) - d_Q(\alpha(b), \alpha(d))| \right\} \\
&= \|\alpha\|.
\end{aligned}$$

□

Example 3.5: The morphism of finite metric lattices $\alpha : P \rightarrow Q$ in Example 3.2 induces the morphism of finite metric lattices $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ in Figure 2. The distortion of $\bar{\alpha}$ is $\|\bar{\alpha}\| = 1$.

4 Filtrations

We now consider filtrations of a fixed finite simplicial complex indexed by finite metric lattices. Fix a finite simplicial complex K and denote by ΔK the category consisting of all subcomplexes $A \subseteq K$ as its objects and inclusions $A \hookrightarrow B$ as morphisms.

Definition 4.1: Let P be a finite metric lattice and K a finite simplicial complex. A filtration of K indexed by P , or simply a **P -filtration of K** , is a functor $F : P \rightarrow \Delta K$. That is, for all $\mathbf{a} \in P$, $F(\mathbf{a})$ is a subcomplex of K and for all $\mathbf{a} \leq \mathbf{b}$, $F(\mathbf{a} \leq \mathbf{b})$ is the inclusion of $F(\mathbf{a})$ into $F(\mathbf{b})$. Further, we require that $F(\top) = K$.

Definition 4.2: A **filtration-preserving morphism** is a triple (F, G, α) where $F : P \rightarrow \Delta K$ and $G : Q \rightarrow \Delta K$ are P and Q -filtrations of K , respectively, and $\alpha : P \rightarrow Q$ is a bounded lattice function satisfying the following axiom. For all $\mathbf{a} \in Q$, $G(\mathbf{a}) = F(\mathbf{a}^*)$ where $\mathbf{a}^* := \max \alpha^{-1}[\perp, \mathbf{a}]$:

$$\begin{array}{ccc} P & \xrightarrow{\alpha} & Q \\ & \searrow F & \swarrow G \\ & \Delta K & \end{array}$$

Remark 4.3: A more sophisticated but an equivalent definition of a filtration-preserving morphism is the following. A filtration-preserving morphism is a triple (F, G, α) where $F : P \rightarrow \Delta K$ and $G : Q \rightarrow \Delta K$ are P and Q -filtrations of K , respectively, and $\alpha : P \rightarrow Q$ is a bounded lattice function such that G is the left Kan extension of F along α , written $G = \text{Lan}_\alpha F$:

$$\begin{array}{ccc} P & \xrightarrow{F} & \Delta \\ & \searrow \alpha & \downarrow \mu \\ & & Q \end{array} \quad \begin{array}{c} \nearrow G = \text{Lan}_\alpha F \end{array}$$

By construction of the left Kan extension,

$$\text{Lan}_\alpha F(\mathbf{a}) := \text{colim}_{\Delta K} F|_{\alpha^{-1}[\perp, \mathbf{a}]}$$

for all $\mathbf{a} \in Q$. By Proposition 3.1, $\alpha^{-1}[\perp, \mathbf{a}]$ has a maximal element $\mathbf{a}^*$ and therefore $\text{Lan}_\alpha F(\mathbf{a})$ is equal to $F(\mathbf{a}^*)$. For all $\mathbf{a} \leq \mathbf{b}$ in Q , $\mathbf{a}^* \leq \mathbf{b}^*$ inducing the inclusion $\text{Lan}_\alpha F(\mathbf{a} \leq \mathbf{b})$. The natural transformation $\mu : F \Rightarrow G \circ \alpha$ is gotten as follows. For $\mathbf{c} \in P$, let $\mathbf{a} := \alpha(\mathbf{c})$. Since $\mathbf{c} \leq \mathbf{a}^*$ and $\text{Lan}_\alpha F(\mathbf{a})$ is equal to $F(\mathbf{a}^*)$, we get the inclusion $\mu(\mathbf{c}) : F(\mathbf{c}) \hookrightarrow G \circ \alpha(\mathbf{c}) = G(\mathbf{a})$.

Remark 4.4: A zigzag of filtration-preserving morphisms categorifies the notion of a *transposition* introduced in [11]. Consider two filtration-preserving morphisms (F, G, α) and (H, G, β) :

$$\begin{array}{ccccc} P & \xrightarrow{\alpha} & Q & \xleftarrow{\beta} & R \\ & \searrow F & \downarrow G & \swarrow H & \\ & & \Delta K & & \end{array}$$

Suppose for $\mathbf{q} \in Q$, both $\alpha^{-1}(\mathbf{q})$ and $\beta^{-1}(\mathbf{q})$ are nonempty. Then the simplices in K that appear in the filtration F restricted to $\alpha^{-1}(\mathbf{q})$ appear at once in G at $\mathbf{q}$. Further, the same simplices that appear in F restricted to $\alpha^{-1}(\mathbf{q})$ appear in H restricted to $\beta^{-1}(\mathbf{q})$ albeit in a possibly different order. The two morphisms (F, G, α) and (H, G, β) are together a generalization of the notion of a *transposition*.

Example 4.5: Let $\alpha : P \rightarrow Q$ be the bounded lattice function described in Example 3.2. Consider the two filtrations $F : P \rightarrow \Delta K$ and $G : Q \rightarrow \Delta K$ of the 2-simplex K in Figure 3. The triple (F, G, α) is a filtration-preserving morphism $\alpha : F \rightarrow G$.




Figure 3: Filtrations F and G of the 2-simplex along with a filtration-preserving morphism α as described in Example 3.2.

Proposition 4.6: If (F, G, α) and (G, H, β) are filtration-preserving morphisms, then $(F, H, \beta \circ \alpha)$ is a filtration-preserving morphism.

Proof. Suppose $F : P \rightarrow \Delta K$, $G : Q \rightarrow \Delta K$, and $H : R \rightarrow \Delta K$. For all $\mathbf{a} \in R$, $H(\mathbf{a}) = G(\mathbf{a}^*)$ where $\mathbf{a}^* := \max \beta^{-1}[\perp, \mathbf{a}]$. Furthermore, $G(\mathbf{a}^*) = F(\mathbf{a}^{**})$ where $\mathbf{a}^{**} := \max \alpha^{-1}[\perp, \mathbf{a}^*]$. Since $\mathbf{a}^{**} = \max(\beta \circ \alpha)^{-1}[\perp, \mathbf{a}]$, we have that $H(\mathbf{a}) = F(\mathbf{a}^{**})$. Thus $(F, H, \beta \circ \alpha)$ is a filtration-preserving morphism. $\square$

Definition 4.7: Fix a finite simplicial complex K . Let $\text{Fil}(K)$ be the category whose objects are P -filtrations of K , over all finite metric lattices P , and whose morphisms are filtration-preserving morphisms. We call $\text{Fil}(K)$ the **category of filtrations of K** .

5 Monotone Integral Functions

We now define the category of monotone integral functions over finite metric lattices Mon and construct the birth-death functor $\text{ZB}_* : \text{Fil}(K) \rightarrow \text{Mon}$. Let $\mathbb{Z}$ be the poset of integers with the usual total ordering $\leq$.

Definition 5.1: Let P and Q be two finite metric lattices and let $f : P \rightarrow \mathbb{Z}$ and $g : Q \rightarrow \mathbb{Z}$ be two monotone integral functions. A **monotone-preserving morphism** from f to g is a triple $(f, g, \bar{\alpha})$ where $f : \bar{P} \rightarrow \mathbb{Z}$ and $g : \bar{Q} \rightarrow \mathbb{Z}$ are monotone functions and $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is a bounded lattice function satisfying the following axiom. For all $I \in \bar{Q}$ and $I^* := \max \bar{\alpha}^{-1}[\perp, I]$, $g(I) = f(I^*)$:

$$\begin{array}{ccc}
 \bar{P} & \xrightarrow{\bar{\alpha}} & \bar{Q} \\
 \searrow f & & \swarrow g \\
 & \mathbb{Z} &
 \end{array}$$

Note that if $(f, g, \bar{\alpha})$ is a monotone-preserving morphism, then $f[\top, \top] = g[\top, \top]$.

Remark 5.2: A more sophisticated but an equivalent definition of a monotone-preserving morphism is the following. A **monotone-preserving morphism** is a triple $(f, g, \bar{\alpha})$ where




Figure 4: Shown are two monotone integral functions f and g on the metric lattices $\bar{P}$ and $\bar{Q}$ from Example 3.5. The triple $(f, g, \bar{\alpha})$, where $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is from the same example, is a monotone-preserving morphism from f to g .

$f : \bar{P} \rightarrow \mathbb{Z}$ and $g : \bar{Q} \rightarrow \mathbb{Z}$ are monotone functions and $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is a bounded lattice function such that g is the left Kan extension of f along $\bar{\alpha}$, written $g = \text{Lan}_{\bar{\alpha}} f$:

$$\begin{array}{ccc}
 \bar{P} & \xrightarrow{f} & \mathbb{Z} \\
 \searrow \bar{\alpha} & \Downarrow \mu & \nearrow g = \text{Lan}_{\bar{\alpha}} f \\
 & \bar{Q} &
 \end{array}$$

Example 5.3: See Figure 4 for an example of monotone integral functions f and g on the lattices $\bar{P}$ and $\bar{Q}$, respectively, from Example 3.5. The triple $(f, g, \bar{\alpha})$, where $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is from the same example, is a monotone-preserving morphism.

Proposition 5.4: If $(f, g, \bar{\alpha})$ and $(g, h, \bar{\beta})$ are monotone-preserving morphisms, then $(f, h, \bar{\beta} \circ \bar{\alpha})$ is a monotone-preserving morphism.

Proof. Suppose $f : \bar{P} \rightarrow \mathbb{Z}$, $g : \bar{Q} \rightarrow \mathbb{Z}$, and $h : \bar{R} \rightarrow \mathbb{Z}$. For all $I \in \bar{R}$, $h(I) = g(I^*)$ where $I^* := \max \bar{\beta}^{-1}[\perp, I]$. Furthermore, $g(I^*) = f(I^{**})$ where $I^{**} := \max \bar{\alpha}^{-1}[\perp, I^*]$. Since $I^{**} = \max(\bar{\beta} \circ \bar{\alpha})^{-1}[\perp, I]$, we have that $h(I) = f(I^{**})$. Thus the composition $(f, h, \bar{\beta} \circ \bar{\alpha})$ is a monotone-preserving morphism. $\square$

Definition 5.5: Let **Mon** be the category consisting of monotone integral functions $f : \bar{P} \rightarrow \mathbb{Z}$, over all finite metric lattices $\bar{P}$, and monotone-preserving morphisms. We call **Mon** the **category of monotone functions**.

5.1 Birth-Death Functor

Fix a field k . Let $\mathbf{Vec}$ be the category of finite-dimensional k -vector spaces and $\mathbf{Ch}(\mathbf{Vec})$ the category of chain complexes over $\mathbf{Vec}$. Let $\mathbf{C}_\bullet : \Delta K \rightarrow \mathbf{Ch}(\mathbf{Vec})$ be the functor that assigns to every subcomplex its simplicial chain complex and to every inclusion of subcomplexes the induced inclusion of chain complexes. For every object $F : P \rightarrow \Delta K$ in $\mathbf{Fil}(K)$, we get a P -filtered chain complex $\mathbf{C}_\bullet F : P \rightarrow \mathbf{Ch}(\mathbf{Vec})$ whose total chain complex is $\mathbf{C}_\bullet F(\top)$. For all dimensions i , denote by $Z_i F : P \rightarrow \mathbf{Vec}$ the functor that assigns to every $\mathbf{a} \in P$ the subspace of i -cycles in $\mathbf{C}_\bullet F(\mathbf{a})$ and assigns to every $\mathbf{a} \leq \mathbf{b}$ the canonical inclusion of $Z_i F(\mathbf{a}) \hookrightarrow Z_i F(\mathbf{b})$. For all dimensions i , denote by $B_i F : P \rightarrow \mathbf{Vec}$ the functor that assigns to every $\mathbf{a} \in P$ the subspace of i -boundaries in $\mathbf{C}_\bullet F(\mathbf{a})$ and to all $\mathbf{a} \leq \mathbf{b}$ the canonical inclusion of $B_i F(\mathbf{a}) \hookrightarrow B_i F(\mathbf{b})$. In summary, for all $\mathbf{a} \leq \mathbf{b}$ in P , we have following commutative diagram of inclusions between cycles and boundaries:

$$\begin{array}{ccccc} B_i F(\mathbf{a}) & \hookrightarrow & B_i F(\mathbf{b}) & \hookrightarrow & B_i F(\top) \\ \downarrow & & \downarrow & & \downarrow \\ Z_i F(\mathbf{a}) & \hookrightarrow & Z_i F(\mathbf{b}) & \hookrightarrow & Z_i F(\top). \end{array}$$

Definition 5.6: Let $F : P \rightarrow \Delta K$ be an object of $\mathbf{Fil}(K)$. For every interval $[\mathbf{a}, \mathbf{b}] \in \bar{P}$, where $\mathbf{b} \neq \top$, let

$$ZB_i F[\mathbf{a}, \mathbf{b}] := \dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b}))$$

where the intersection is taken inside $Z_i F(\top)$. For all other intervals $[\mathbf{a}, \top]$, let

$$ZB_i F[\mathbf{a}, \top] := \dim Z_i F(\mathbf{a}).$$

The **i -th birth-death function** of F is the function $f_i : \bar{P} \rightarrow \mathbb{Z}$ that assigns to every interval $[\mathbf{a}, \mathbf{b}]$ the integer $ZB_i F[\mathbf{a}, \mathbf{b}]$.

The reason we force $ZB_i F[\mathbf{a}, \top]$ to $\dim Z_i F(\mathbf{a})$ instead of $\dim Z_i F(\mathbf{a}) \cap B_i F(\top)$ is because we want all cycles to be dead by $\top$. Otherwise, the persistence diagram for F (see Definition 8.1) would not see cycles that are born and never die.

Proposition 5.7: Let $F : P \rightarrow \Delta K$ be an object in $\mathbf{Fil}(K)$ and $f_i : \bar{P} \rightarrow \mathbb{Z}$ its i -th birth-death function. Then f_i is monotone.

Proof. For any two intervals $I \preceq J$ in $\bar{P}$, we must show that $f_i(I) \leq f_i(J)$. Suppose $I = [\mathbf{a}, \mathbf{b}]$ and $J = [\mathbf{c}, \mathbf{d}]$ and $\mathbf{d} \neq \top$. Then $Z_i F(\mathbf{a}) \subseteq Z_i F(\mathbf{c})$ and $B_i F(\mathbf{b}) \subseteq B_i F(\mathbf{d})$. Thus $Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b})$ is a subspace of $Z_i F(\mathbf{c}) \cap B_i F(\mathbf{d})$, and therefore $ZB_i F[\mathbf{a}, \mathbf{b}] \leq ZB_i F[\mathbf{c}, \mathbf{d}]$. For $J = [\mathbf{c}, \top]$, $ZB_i F[\mathbf{a}, \mathbf{b}] \subseteq Z_i F(\mathbf{c})$, and therefore $ZB_i F[\mathbf{a}, \mathbf{b}] \leq ZB_i F[\mathbf{c}, \top]$. $\square$

Proposition 5.8: Let (F, G, α) be a morphism in $\mathbf{Fil}(K)$ and f_i and g_i the i -th birth-death functions of F and G , respectively. Then $(f_i, g_i, \bar{\alpha})$ is a morphism in $\mathbf{Mon}$.

Proof. Suppose $F : P \rightarrow \Delta K$ and $G : Q \rightarrow \Delta K$. By definition of morphism in $\text{Fil}(K)$, $G(a) = F(a^*)$, for all $a \in Q$, where $a^* = \max \bar{\alpha}^{-1}[\perp, a]$. For all intervals $I \in \bar{Q}$, let $I^* := \max \bar{\alpha}^{-1}[\perp, I]$. If $I = [a, b]$, then $I^* = [a^*, b^*]$ where $b^* = \max \bar{\alpha}^{-1}[\perp, b]$. The definition of a filtration-preserving morphism implies the following canonical isomorphisms of chain complexes:

$$C_\bullet G(b) \cong C_\bullet F(b^*) \quad C_\bullet G(\top) \cong C_\bullet F(\top) \quad C_\bullet G(a) \cong C_\bullet F(a^*),$$

which, in turn, implies canonical isomorphisms $Z_\bullet G(a) \cong Z_\bullet F(a^*)$ and $B_\bullet G(b) \cong B_\bullet F(b^*)$. We have $ZB_i G(I) = ZB_i F(I^*)$ and therefore $g_i(I) = f_i(I^*)$. $\square$

By Propositions 5.4, 5.7 and 5.8, the assignment to each object in $\text{Fil}(K)$ its birth-death monotone function is functorial.

Definition 5.9: Let $ZB_i : \text{Fil}(K) \rightarrow \text{Mon}$ be the functor that assigns to every filtration its i -th birth-death monotone function and to every filtration-preserving morphism the induced monotone-preserving morphism. We call ZB_i the **i -th birth-death functor**.

Example 5.10: The functor ZB_1 applied to the filtration-preserving morphism (F, G, α) in Example 3.5 is the monotone-preserving morphism $(f, g, \bar{\alpha})$ in Example 5.3.

6 Integral Functions

We now define the category of integral functions over finite metric lattices Fnc and construct the Möbius inversion functor $MI : \text{Mon} \rightarrow \text{Fnc}$.

Definition 6.1: Let P and Q be finite metric lattices and let $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ be two integral functions not necessarily monotone. A **charge-preserving morphism** is a triple $(\sigma, \tau, \bar{\alpha})$ where $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ are integral functions and $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is a bounded lattice function satisfying the following axiom. For all $I \in \bar{Q}$ with $I \neq [q, q]$,

$$\tau(I) = \sum_{J \in \bar{\alpha}^{-1}(I)} \sigma(J). \quad (2)$$

If $\bar{\alpha}^{-1}(I)$ is empty, then we interpret the sum as 0.

Remark 6.2: Our definition of a charge-preserving morphism is related to the definition of a morphism between signed measures. Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces, $\phi : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ a measurable map, and $\mu : \Sigma_X \rightarrow \mathbb{R}$ a signed measure. Then the pushforward of μ along ϕ is the signed measure $\phi_\# \mu : \Sigma_Y \rightarrow \mathbb{R}$ defined as $\phi_\# \mu(U) := \mu(\phi^{-1}(U))$. In the category of signed measures, a morphism from (X, Σ_X, μ) to (Y, Σ_Y, ν) is a measurable map $\mu : \Sigma_X \rightarrow \Sigma_Y$ such that $\phi_\# \mu = \nu$.

Example 6.3: See Figure 5 for integral functions σ and τ on the lattices of intervals $\bar{P}$ and $\bar{Q}$, respectively, from Example 3.5. The triple $(\sigma, \tau, \bar{\alpha})$, where $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is from the same example, is a charge-preserving morphism.




Figure 5: Shown are two integral functions σ and τ on $\bar{P}$ and $\bar{Q}$, respectively, from Example 3.5. The bounded lattice function $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ from the same figure is a charge-preserving morphism from σ to τ .

Proposition 6.4: If $(\sigma, \tau, \bar{\alpha})$ and $(\tau, \nu, \bar{\beta})$ are charge-preserving morphisms, then $(\sigma, \nu, \bar{\beta} \circ \bar{\alpha})$ is a charge-preserving morphism.

Proof. Suppose $\sigma : \bar{P} \rightarrow \mathbb{Z}$, $\tau : \bar{Q} \rightarrow \mathbb{Z}$, and $\nu : \bar{R} \rightarrow \mathbb{Z}$. For all $I \in \bar{R}$,

$$\nu(I) = \sum_{J \in \bar{\beta}^{-1}(I)} \tau(J) = \sum_{J \in \bar{\beta}^{-1}(I)} \sum_{K \in \bar{\beta}^{-1}(J)} \sigma(K) = \sum_{K \in (\bar{\beta} \circ \bar{\alpha})^{-1}(I)} \sigma(K).$$

□

Definition 6.5: Let **Fnc** be the category of whose objects are integral functions $\sigma : \bar{P} \rightarrow \mathbb{Z}$, over all finite metric lattices $\bar{P}$, and charge-preserving morphisms. We call **Fnc** the **category of integral functions**.

6.1 Möbius Inversion Functor

Given any monotone integral function $f : \bar{P} \rightarrow \mathbb{Z}$ of **Mon**, there is a unique integral function $\sigma : \bar{P} \rightarrow \mathbb{Z}$ such that

$$f(J) = \sum_{I \in \bar{P} : I \leq J} \sigma(I) \tag{3}$$

for all $J \in \bar{P}$ [2, 23]. The function σ is called the *Möbius inversion* of f .

Proposition 6.6: Let $(f, g, \bar{\alpha})$ be a morphism in **Mon**, and let σ and τ be the Möbius inversions of f and g , respectively. Then $(\sigma, \tau, \bar{\alpha})$ is a morphism in **Fnc**.

Proof. Suppose $f : \bar{P} \rightarrow \mathbb{Z}$ and $g : \bar{Q} \rightarrow \mathbb{Z}$. We show that

$$\tau(J) = \sum_{K \in \bar{\alpha}^{-1}(J)} \sigma(K)$$

for *all* $J \in \bar{Q}$, and thus $(\sigma, \tau, \bar{\alpha})$ is a charge-preserving morphism. The proof is by induction on the finite metric lattice $\bar{Q}$. By Proposition 3.1, the pre-image $\bar{\alpha}^{-1}[\perp, J]$ has a unique maximal element J^* , and $f(J^*) = g(J)$ by definition of a morphism in **Mon**.

For the base case, suppose $J = \perp$. Then by Equation (3), $g(J) = \tau(J)$. By definition of a morphism in **Mon**, $g(J) = f(J^*)$. By Equation (3),

$$f(J^*) = \sum_{K \leq J^*} \sigma(K) = \sum_{K \in \bar{\alpha}^{-1}(J)} \sigma(K)$$

thus proving the base case.

For the inductive step, suppose $\tau(I) = \sum_{K \in \bar{\alpha}^{-1}(I)} \sigma(K)$, for all $I \prec J$. Then

$$\begin{aligned} \tau(J) &= \sum_{I \in \bar{Q}: I \leq J} \tau(I) - \sum_{I \in \bar{Q}: I \prec J} \tau(I) \\ &= g(J) - \sum_{I \in \bar{Q}: I \prec J} \tau(I) && \text{by Equation (3)} \\ &= g(J) - \sum_{I \in \bar{Q}: I \prec J} \sum_{K \in \bar{\alpha}^{-1}(I)} \sigma(K) && \text{by Inductive Hypothesis} \\ &= f(J^*) - \sum_{K \in \bar{P}: \bar{\alpha}(K) \prec J} \sigma(K) \\ &= \sum_{K \in \bar{P}: K \leq J^*} \sigma(K) - \sum_{K \in \bar{P}: \bar{\alpha}(K) \prec J} \sigma(K) && \text{by Equation (3)} \\ &= \sum_{K \in \bar{P}: \bar{\alpha}(K) \leq J} \sigma(K) - \sum_{K \in \bar{P}: \bar{\alpha}(K) \prec J} \sigma(K) \\ &= \sum_{K \in \bar{P}: \bar{\alpha}(K) = J} \sigma(K) = \sum_{K \in \bar{\alpha}^{-1}(J)} \sigma(K). \end{aligned}$$

□

By Propositions 6.4 and 6.6, the assignment to every object in **Mon** its Möbius inversion is functorial.

Definition 6.7: Let $MI : \mathbf{Mon} \rightarrow \mathbf{Fnc}$ be the functor that assigns to every monotone function its Möbius inversion and to every monotone-preserving morphism the induced charge-preserving morphism. We call MI the **Möbius inversion functor**.

Example 6.8: The functor MI applied to the monotone-preserving morphism $(f, g, \bar{\alpha})$ in Example 5.3 is the charge-preserving morphism $(\sigma, \tau, \bar{\alpha})$ in Example 6.3.

7 Edit Distance

We now define the edit distance in each of the three categories $\mathbf{Fil}(K)$, **Mon**, and **Fnc** and show that the two functors ZB_* and MI are 1-Lipschitz. Denote by $\star$ the metric lattice consisting

of just one element.

7.1 Distance Between Filtrations

A *path* between two filtrations F and H in $\text{Fil}(\mathbf{K})$ is a finite sequence

$$F \xleftrightarrow{\alpha_1} G_1 \xleftrightarrow{\alpha_2} \cdots \xleftrightarrow{\alpha_3} G_{n-1} \xleftrightarrow{\alpha_n} H$$

where $\leftrightarrow$ denotes a filtration-preserving morphism in either direction. The *length* of a path is the sum $\sum_{i=1}^n \|\alpha_i\|$ of the distortions of all the bounded lattice functions. Note that the filtration $\Omega : \star \rightarrow \Delta \mathbf{K}$ is terminal in $\text{Fil}(\mathbf{K})$. This implies that any two filtrations in $\text{Fil}(\mathbf{K})$ are connected by a path.

Definition 7.1: The **edit distance** $d_{\text{Fil}(\mathbf{K})}(F, H)$ between any two filtrations in $\text{Fil}(\mathbf{K})$ is the length of the shortest path between F and H .

7.2 Distance Between Monotone Integral Functions

A *path* between two monotone functions f and h in Mon is a finite sequence

$$f \xleftrightarrow{\bar{\alpha}_1} g_1 \xleftrightarrow{\bar{\alpha}_2} \cdots \xleftrightarrow{\bar{\alpha}_3} g_{n-1} \xleftrightarrow{\bar{\alpha}_n} h$$

where $\leftrightarrow$ denotes a filtration-preserving morphism in either direction. The *length* of a path is the sum $\sum_{i=1}^n \|\bar{\alpha}_i\|$ of the distortions of all the bounded lattice functions. Suppose $f[\top, \top] = \mathbf{n}$, and let $e : \star \rightarrow \mathbb{Z}$ be the monotone integral function where $e[\star, \star] = \mathbf{n}$. Then there is a unique monotone-preserving morphism from f to e . Thus there is a path between any two monotone-integral functions f and h such that $f[\top, \top] = h[\top, \top]$.

Definition 7.2: The **edit distance** $d_{\text{Mon}}(f, h)$ between any two monotone functions in Mon is the length of the shortest path between f and h . If there are no paths, then we let $d_{\text{Mon}}(f, h) = \infty$.

Lemma 7.3: Let F and G be two objects of $\text{Fil}(\mathbf{K})$. Then for every dimension i ,

$$d_{\text{Mon}}(\mathbf{ZB}_i F, \mathbf{ZB}_i G) \leq d_{\text{Fil}(\mathbf{K})}(F, G).$$

Proof. Suppose $d_{\text{Fil}(\mathbf{K})}(F, G) = \varepsilon$. Then there is a path in $\text{Fil}(\mathbf{K})$ between F and G with length ε . Apply the functor $\mathbf{ZB}_i$ to this path and the result is a path in Mon between $\mathbf{ZB}_i F$ and $\mathbf{ZB}_i G$ and its length, by Proposition 3.4, is also ε . Since the distance between the two monotone functions is defined as the length of the shortest path between them, we have the desired inequality. $\square$

7.3 Distance Between Integral Functions

A *path* between two integral functions σ and τ in Fnc is a finite sequence

$$\sigma \xleftrightarrow{\bar{\alpha}_1} \theta_1 \xleftrightarrow{\bar{\alpha}_2} \cdots \xleftrightarrow{\bar{\alpha}_3} \theta_{n-1} \xleftrightarrow{\bar{\alpha}_n} \tau$$

where $\leftrightarrow$ denotes a charge-preserving morphism in either direction. The *length* of path is the sum $\sum_{i=1}^n \|\bar{\alpha}_i\|$ of the distortions of all the bounded lattice functions. Note that any integral function $\omega : \bar{\star} \rightarrow \mathbb{Z}$ is terminal in **Fnc**; see Definition 6.1. This means that any two integral function in **Fnc** are connected by a path.

Definition 7.4: Define the distance $d_{\text{Fnc}}(\sigma, \tau)$ between any two integral functions in **Fnc** as the length of the shortest path between σ and τ .

Lemma 7.5: Let f and g be two objects of **Mon**. Then $d_{\text{Fnc}}(\text{MI}(f), \text{MI}(g)) \leq d_{\text{Mon}}(f, g)$.

Proof. Suppose $d_{\text{Mon}}(f, g) = \varepsilon$. Then there is a path in **Mon** between f and g with length ε . Apply the functor MI to this path and the result is a path in **Fnc** between $\text{MI}(f)$ and $\text{MI}(g)$ and its length is also ε . Since the distance between the two functions is defined as the length of the shortest path between them, we have the desired inequality. $\square$

8 Persistence Diagrams

The pieces established in the last four sections fit together into the following pipeline of 1-Lipschitz functors:

$$\text{Fil}(\mathbf{K}) \xrightarrow{\text{ZB}_*} \text{Mon} \xrightarrow{\text{MI}} \text{Fnc}.$$

The birth-death functor ZB_* assigns to an object $F : P \rightarrow \Delta K$ of $\text{Fil}(\mathbf{K})$ a monotone integral function $f_i := \text{ZB}_i(F) : \bar{P} \rightarrow \mathbb{Z}$ for every dimension i . The value of f_i on an interval $[a, b] \subseteq P$ is the dimension of the k -vector space of i -cycles that appear by a and become boundaries by b . The Möbius inversion functor MI assigns to f_i its Möbius inversion, which is an integral function $\sigma := \text{MI}(f_i) : \bar{P} \rightarrow \mathbb{Z}$.

Definition 8.1: Let P be a finite metric lattice and $F : P \rightarrow \Delta K$ a filtered simplicial complex indexed over P . The **i -th persistence diagram** of F is the integral function $\sigma := \text{MI} \circ \text{ZB}_i(F) : \bar{P} \rightarrow \Delta K$.

Example 8.2: Consider the filtrations F and G in Example 4.5. Their 1-dimensional persistence diagrams are the integral functions σ and τ , respectively, in Example 6.3. The integer $\sigma[b, d] = 1$ represents the 1-cycle that is born at b , and the integer $\sigma[c, d] = 1$ represents the 1-cycle that is born at c . The integer $\sigma[d, d] = -1$, represents the 1-cycle that was born twice but contributes to just one dimension of the total cycle space.

Example 8.3: Consider the example of a filtration $F : P \rightarrow \Delta K$ in Figure 6 where P is the lattice from Example 3.2 and K is the 1-simplex. Recall $\bar{P}$ in Example 3.5. Drawn are its zeroth birth-death function $f := \text{ZB}_0 \circ F : \bar{P} \rightarrow \mathbb{Z}$ and its zeroth persistence diagram $\sigma := \text{MI} \circ \text{ZB}_0 : \bar{P} \rightarrow \mathbb{Z}$. The integer $\sigma[a, d] = 1$ represents the 0-cycle that is born at a , and the integer $\sigma[b, d] = 1$ represents the 0-cycle that is born at b . The integer $\sigma[c, c] = 1$ represents the 0-cycle that is born at c and dies immediately. The integer $\sigma[d, d] = -1$ represents the 0-cycle that was born twice but contributes to just one dimension of the total cycle space.




Figure 6: Above is a filtration of the 1-simplex, its zeroth birth-death function, and its zeroth persistence diagram.

Our main theorem follows immediately from Lemmas 7.3 and 7.5.

Theorem 8.4 (Stability): Let $F : P \rightarrow \Delta K$ and $G : Q \rightarrow \Delta K$ be two filtrations of a finite simplicial complex K indexed by finite metric lattices, and $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ their i -th persistence diagrams. Then $d_{\text{Fnc}}(\sigma, \tau) \leq d_{\text{Fil}(K)}(F, G)$.

9 Classical Persistent Homology

We now relate our definitions to that of classical persistent homology. First, we show that our definition of the persistence diagram is the same as the original definitions of [10] and [14, 19]. Second, we show that the bottleneck distance between classical persistence diagrams is strongly equivalent to the edit distance.

9.1 Classical Persistence Diagrams

Fix a finite 1-parameter filtration $K_{r_1} \subseteq K_{r_1} \subseteq \dots \subseteq K_{r_n} = K$ of a finite simplicial complex K indexed by real the numbers $r_1 < \dots < r_n$. Let $P := 1 < \dots < n < \infty$ be the totally ordered lattice with $n + 1$ elements with $d_P(a, b) = |r_a - r_b|$, for $b \neq \infty$, and $d_P(a, \infty) = \infty$. Let $F : P \rightarrow \Delta K$ be the filtration that assigns to every $a \in P \setminus \{\infty\}$ the subcomplex K_{r_a} and to ∞ the total complex K . Cohen-Steiner, Edelsbrunner, and Harer define the i -th persistence diagram of this filtration as the integral function $\sigma_i : \bar{P} \rightarrow \mathbb{Z}$ defined as follows.

For $\mathbf{a} < \mathbf{b} \neq \infty$, $\sigma_i[\mathbf{a}, \mathbf{b}]$ is the following signed sum of ranks:

$$\begin{aligned}
\sigma_i[\mathbf{a}, \mathbf{b}] &:= \text{rk } H_i F(\mathbf{a} \leq \mathbf{b} - 1) - \text{rk } H_i F(\mathbf{a} - 1 \leq \mathbf{b} - 1) - \text{rk } H_i F(\mathbf{a} \leq \mathbf{b}) + \text{rk } H_i F(\mathbf{a} - 1 \leq \mathbf{b}) \\
&= \dim \frac{Z_i F(\mathbf{a})}{Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b} - 1)} - \dim \frac{Z_i F(\mathbf{a} - 1)}{Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b} - 1)} \\
&\quad - \dim \frac{Z_i F(\mathbf{a})}{Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b})} + \dim \frac{Z_i F(\mathbf{a} - 1)}{Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b})} \\
&= \dim Z_i F(\mathbf{a}) - \dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b} - 1)) - \dim Z_i F(\mathbf{a} - 1) + \dim (Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b} - 1)) \\
&\quad - \dim Z_i F(\mathbf{a}) + \dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b})) + \dim Z_i F(\mathbf{a} - 1) \\
&\quad - \dim (Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b})) \\
&= -\dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b} - 1)) + \dim (Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b} - 1)) \\
&\quad + \dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{b})) - \dim (Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{b})) \\
&= -ZB_i F[\mathbf{a}, \mathbf{b} - 1] + ZB_i F[\mathbf{a} - 1, \mathbf{b} - 1] + ZB_i F[\mathbf{a}, \mathbf{b}] - ZB_i F[\mathbf{a} - 1, \mathbf{b}].
\end{aligned}$$

For $\mathbf{a} < \mathbf{b} = \infty$, $\sigma[\mathbf{a}, \infty]$ is the following signed sum of ranks:

$$\begin{aligned}
\sigma_i[\mathbf{a}, \infty] &:= \text{rk } H_i F(\mathbf{a} \leq \infty) - \text{rk } H_i F(\mathbf{a} - 1 \leq \infty) \\
&= \dim \frac{Z_i F(\mathbf{a})}{Z_i F(\mathbf{a}) \cap B_i F(\mathbf{n})} - \dim \frac{Z_i F(\mathbf{a} - 1)}{Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{n})} \\
&= \dim Z_i F(\mathbf{a}) - \dim (Z_i F(\mathbf{a}) \cap B_i F(\mathbf{n})) - \dim Z_i F(\mathbf{a} - 1) + \dim (Z_i F(\mathbf{a} - 1) \cap B_i F(\mathbf{n})) \\
&= ZB_i F[\mathbf{a}, \infty] - ZB_i F[\mathbf{a}, \mathbf{n}] - ZB_i F[\mathbf{a} - 1, \infty] + ZB_i F[\mathbf{a} - 1, \mathbf{n}].
\end{aligned}$$

However, in this paper we define the persistence diagram of F as $\tau := \text{MI} \circ ZB_*(F) : \bar{P} \rightarrow \mathbb{Z}$; see Definition 8.1. It turns out that the two are the same. Since P is totally ordered, the Möbius inversion of $ZB_* F$ has the following simple formula for any $\mathbf{a} \leq \mathbf{b}$:

$$\tau[\mathbf{a}, \mathbf{b}] := ZB_* F[\mathbf{a}, \mathbf{b}] - ZB_* F[\mathbf{a} - 1, \mathbf{b}] - ZB_* F[\mathbf{a}, \mathbf{b} - 1] + ZB_* F[\mathbf{a} - 1, \mathbf{b} - 1].$$

We see that for $\mathbf{a} < \mathbf{b}$, $\sigma[\mathbf{a}, \mathbf{b}] = \tau[\mathbf{a}, \mathbf{b}]$.

9.2 Bottleneck Distance

We prove that the bottleneck distance defined between classical persistence diagrams is strongly equivalent to the edit distance. Let P and Q be finite, totally ordered metric lattices. In order to define the bottleneck distance between two integral functions $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$, we need embeddings of P and Q into $\mathbb{R}$. This is problematic since the edit distance does not depend on the embedding while the bottleneck distance does. We fix this issue by requiring that $\perp_P$ and $\perp_Q$ map to $0 \in \mathbb{R}$ under the embeddings. We identify elements of P and Q with their images in $\mathbb{R}$ under the assumed embeddings. This section culminates in a proof of the following theorem.

Theorem 9.1: Let $\bar{P}$ and $\bar{Q}$ be finite, totally ordered metric lattices embedded in $\mathbb{R}$ with $\perp_{\bar{P}} = \perp_{\bar{Q}} = 0$ and $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ two integral functions. Then $d_B(\sigma, \tau) \leq d_{\text{Fnc}}(\sigma, \tau) \leq 2d_B(\sigma, \tau)$.

Definition 9.2: For any interval $[a, b] \subseteq \mathbb{R}$, let $\|[a, b]\|_\infty := \max\{|a|, |b|\}$. Addition and scalar multiplication of intervals is defined componentwise by $[a, b] + [c, d] = [a + c, b + d]$ and $x[a, b] = [xa, xb]$ for any $x \in \mathbb{R}$.

Definition 9.3: A **matching** between two non-negative integral functions $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ is a non-negative map $\gamma : \bar{P} \times \bar{Q} \rightarrow \mathbb{Z}$ satisfying

$$\begin{aligned}\sigma(I) &= \sum_{J \in \bar{Q}} \gamma(I, J) \text{ for all } I \neq [p, p] \in \bar{P} \\ \tau(J) &= \sum_{I \in \bar{P}} \gamma(I, J) \text{ for all } J \neq [q, q] \in \bar{Q}.\end{aligned}$$

The **norm** of a matching γ is

$$\|\gamma\| := \max_{\{I \in \bar{P}, J \in \bar{Q} \mid \gamma(I, J) > 0\}} \|I - J\|_\infty.$$

The **bottleneck distance** between σ and τ is

$$d_B(\sigma, \tau) := \min_{\gamma} \|\gamma\|$$

over all matchings γ between σ and τ .

Proposition 9.4: Let $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ be non-negative integral functions and γ be an ε -matching between σ and τ . Then γ induces a 1-parameter family of integral functions $\{v_t\}_{t \in [0, 1]}$ with $v_0 = \sigma$ and $v_1 = \tau$.

Proof. Let $\bar{S}_t := \{(1-t)I + tJ \mid I \in \bar{P}, J \in \bar{Q}, \text{ and } \gamma(I, J) > 0\}$. Define $v_t : \bar{S}_t \rightarrow \mathbb{Z}$ to be

$$v_t(K) := \sum_{\substack{I \in \bar{P}, J \in \bar{Q} \\ (1-t)I + tJ = K}} \gamma(I, J).$$

At $t = 0$ this reduces to

$$v_0(K) = \sum_{\substack{I \in \bar{P}, J \in \bar{Q} \\ I = K}} \gamma(I, J) = \sum_{J \in \bar{Q}} \gamma(K, J) = \sigma(K),$$

for all $K \in \bar{S}_0$, and similarly $v_1(I) = \tau(I)$. □

As t varies from 0 to 1, there are only finitely many places where the combinatorial structure of v_t changes. We call these places critical points; see the following definition. These combinatorial changes occur where endpoints of intervals in $\bar{S}_t$ cross or, equivalently, where the cardinality of the set of endpoints changes.

Definition 9.5: Let $S_t = \{w \in \mathbb{R} \mid [w, x] \text{ or } [x, w] \in \bar{S}_t\}$ be the set of endpoints of intervals in $\bar{S}_t$. A point $t \in [0, 1]$ is **critical** if for all sufficiently small $\delta > 0$, there exists $s \in (t-\delta, t+\delta)$ with $|S_t| \neq |S_s|$.

Lemma 9.6: If $t \in [0, 1]$ is not a critical point, then for any $K \in \bar{S}_t$ there is a unique pair of intervals $I \in \bar{P}$ and $J \in \bar{Q}$ with $\gamma(I, J) > 0$ and $(1-t)I + tJ = K$.

Proof. Suppose $t \in [0, 1]$ is not critical and there exists $I, I' \in \bar{P}$ and $J, J' \in \bar{Q}$ with $\gamma(I, J) > 0$, $\gamma(I', J') > 0$ and $(1-t)I + tJ = (1-t)I' + tJ'$. Then for any t' sufficiently close to t , $(1-t')I + t'J = (1-t')I' + t'J'$. Since the interpolation is linear and two lines that intersect in more than one point must be the same line, it follows that $I = I'$ and $J = J'$. $\square$

Lemma 9.7: If $\alpha : P \rightarrow Q$ is a metric lattice map and $\bar{\alpha} : \bar{P} \rightarrow \bar{Q}$ is its induced map on intervals then

$$\max_{I \in \bar{P}} \|I - \bar{\alpha}(I)\|_\infty \leq \|\bar{\alpha}\| \leq 2 \max_{I \in \bar{P}} \|I - \bar{\alpha}(I)\|_\infty.$$

Proof. First note that by Proposition 3.4, $\|\alpha\| = \|\bar{\alpha}\|$ and since $\bar{\alpha}$ is induced by α , $\max_{I \in \bar{P}} \|I - \bar{\alpha}(I)\|_\infty = \max_{a \in P} |a - \alpha(a)|$ so the inequality reduces to

$$\max_{a \in P} |a - \alpha(a)| \leq \max_{a, b \in P} ||a - b| - |\alpha(a) - \alpha(b)|| \leq 2 \max_{a \in P} |a - \alpha(a)|.$$

Note that since $\perp_P = \perp_Q = 0$, each element of P and Q are non-negative. Assume, without loss of generality, that $a \geq b$. Then the middle quantity above reduces to

$$\max_{a \geq b \in P} |a - b - (\alpha(a) - \alpha(b))| = \max_{a \geq b \in P} |a - \alpha(a) - (b - \alpha(b))|$$

Letting $b = 0$ yields the first inequality and the triangle inequality yields the second. $\square$

Lemma 9.8: If $t \in [0, 1]$ is not a critical point and $s \in [0, 1]$ is any point with no critical points strictly between t and s , then there is a charge-preserving morphism $(v_t, v_s, \bar{\alpha}_{t,s})$ with distortion at most $2\varepsilon|s - t|$.

Proof. We start by defining a map $\alpha_{t,s} : S_t \rightarrow S_s$. For any $b \in S_t$ note that since t is not critical, there are unique intervals $I \in \bar{P}$ and $J \in \bar{Q}$ with $\gamma(I, J) > 0$ and either $(1-t)I + tJ = [a, b]$ or $[b, c]$. If $(1-t)I + tJ = [a, b]$ then define $\alpha_{t,s}(b)$ to be the right endpoint of the interval $(1-s)I + sJ$. Similarly, if b is a left endpoint, then we define $\alpha_{t,s}(b)$ to be the left endpoint of $(1-s)I + sJ$. This map is order preserving since as t varies, endpoints of intervals only cross at critical points and there are no critical points strictly between t and s .

To prove that $\bar{\alpha}_{t,s}$ is charge-preserving, observe that

$$\begin{aligned}
\sum_{K \in \bar{\alpha}_{t,s}^{-1}(L)} v_t(K) &= \sum_{\substack{I \in \bar{P}, J \in \bar{Q} \\ (1-s)I + sJ = L}} v_t((1-t)I + tJ) \\
&= \sum_{\substack{I \in \bar{P}, J \in \bar{Q} \\ (1-s)I + sJ = L}} \left(\sum_{\substack{I' \in \bar{P}, J' \in \bar{Q} \\ (1-t)I' + tJ' = (1-t)I + tJ}} \gamma(I', J') \right) \\
&= \sum_{\substack{I \in \bar{P}, J \in \bar{Q} \\ (1-s)I + sJ = L}} \gamma(I, J) = v_s(L)
\end{aligned}$$

where the third equality follows from Lemma 9.6 and the assumption that t is not critical. The distortion of $\bar{\alpha}_{t,s}$ is

$$\begin{aligned}
\|\bar{\alpha}_{t,s}\| &\leq 2 \max_{K \in \bar{S}_t} \|K - \bar{\alpha}_{t,s}(K)\|_\infty \\
&= 2 \max_{\substack{I \in \bar{P}, J \in \bar{Q} \\ \gamma(I, J) > 0}} \|(1-t)I + tJ - (1-s)I - sJ\|_\infty \\
&= 2|s - t| \max_{\substack{I \in \bar{P}, J \in \bar{Q} \\ \gamma(I, J) > 0}} \|I - J\|_\infty \leq 2\varepsilon|s - t|.
\end{aligned}$$

□

Lemma 9.9: For any non-negative integral functions $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ over finite sublattices $P, Q \subseteq \mathbb{R}$, $d_{\text{Fnc}}(\sigma, \tau) \leq 2d_B(\sigma, \tau)$.

Proof. We show that $d_{\text{Fnc}}(\sigma, \tau) \leq 2d_B(\sigma, \tau)$ by showing that an ε -matching between σ and τ induces a path between σ and τ of length at most 2ε . For any ε -matching γ between σ and τ , let $\{v_t\}_{t \in [0,1]}$ be the interpolation induced by γ from Proposition 9.4. Let $\{s_0 = 0 < s_1 \cdots < s_n = 1\} \subseteq [0, 1]$ be the set of critical points of the interpolation and choose $\{t_0 < \cdots < t_{n-1}\} \subseteq [0, 1]$ with $0 < t_0 < s_1 < t_1 \cdots < t_{n-1} < 1$. Then the charge-preserving morphisms $\alpha_{t_i, s_{i \pm 1}}$ from Lemma 9.8 form a path between σ and τ with length at most $\sum_{i=0}^{n-1} 2\varepsilon|s_i - t_i| = 2\varepsilon$. □

Lemma 9.10: For any non-negative integral functions $\sigma : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ over finite sublattices $P, Q \subseteq \mathbb{R}$, $d_{\text{Fnc}}(\sigma, \tau) \geq d_B(\sigma, \tau)$.

Proof. To show that $d_B(\sigma, \tau) \leq d_{\text{Fnc}}(\sigma, \tau)$, we show that a charge-preserving morphism induces a matching. Let (σ, τ, α) be a charge-preserving morphism with distortion ε . Define a matching γ between σ and τ by

$$\gamma(I, J) := \begin{cases} \sigma(I) & \text{if } \alpha(I) = J \\ 0 & \text{otherwise} \end{cases}.$$

Then we have that for any $J \in \bar{Q}$

$$\sum_{I \in \bar{P}} \gamma(I, J) = \sum_{I \in \alpha^{-1}(J)} \sigma(I) = \tau(J)$$

and for any $I \in \bar{P}$

$$\sum_{J \in \bar{Q}} \gamma(I, J) = \alpha(I).$$

Therefore γ is a matching. The norm of γ is

$$\|\gamma\| = \max_{I \in \bar{P}, J \in \bar{Q} : \gamma(I, J) > 0} \|I - J\|_\infty = \max_{I \in \bar{P}} \|I - \alpha(I)\|_\infty \leq \|\alpha\|.$$

□

Theorem 9.1 follows immediately from Lemma 9.9 and Lemma 9.10. The following two examples show that the bounds in Theorem 9.1 are tight.

Example 9.11: Let $P = \{0 < 1 < 2 < 3\}$ be a totally ordered metric lattice where the distance between two elements is the absolute value of their difference. Let $\sigma, \nu : \bar{P} \rightarrow \mathbb{Z}$ be two integral functions defined as

$$\sigma[a, b] := \begin{cases} 1 & \text{if } [a, b] = [0, 1], [2, 3] \\ 0 & \text{otherwise} \end{cases} \quad \nu[a, b] := \begin{cases} 1 & \text{if } [a, b] = [0, 2], [1, 3] \\ 0 & \text{otherwise.} \end{cases}$$

See Figure 7. The bottleneck distance, d_B , between σ and ν is $d_B(\sigma, \nu) = 1$. We now compute the edit distance, d_{Fnc} , between σ and ν . Consider a third integral function $\tau : \bar{Q} \rightarrow \mathbb{Z}$ where $Q = \{0 < 0.5 < 1 < 1.5 < 2 < 2.5 < 3\}$ is a finite, totally ordered metric lattice where the distance between any two elements is the absolute value of the difference and

$$\nu[a, b] := \begin{cases} 1 & \text{if } [a, b] = [0, 1.5], [1.5, 3] \\ 0 & \text{otherwise.} \end{cases}$$

Let $\alpha : P \rightarrow Q$ be the bounded lattice function defined as follows

$$\alpha(0) := 0 \quad \alpha(1) := 1.5 \quad \alpha(2) := 1.5 \quad \alpha(3) := 3.$$

We now have a pair of charge-preserving morphisms $(\sigma, \tau, \bar{\alpha})$ and $(\nu, \tau, \bar{\alpha})$. Thus $d_{\text{Fnc}}(\sigma, \nu) \leq 2\|\bar{\alpha}\| = 2\|\alpha\| = 2(0.5) = 1$. Further, this is a shortest path between σ and ν in Fnc . Therefore $d_{\text{Fnc}}(\sigma, \nu) = 1$.

Example 9.12: Let P be the metric lattice defined in Example 9.11 and $\sigma, \tau : \bar{P} \rightarrow \mathbb{Z}$ be defined as

$$\sigma[a, b] := \begin{cases} 1 & \text{if } [a, b] = [1, 2] \\ 0 & \text{otherwise} \end{cases} \quad \nu[a, b] := \begin{cases} 1 & \text{if } [a, b] = [0, 3] \\ 0 & \text{otherwise.} \end{cases}$$




Figure 7: Above are three integral functions $\sigma, v : \bar{P} \rightarrow \mathbb{Z}$ and $\tau : \bar{Q} \rightarrow \mathbb{Z}$ drawn as barcodes and two charge-preserving morphisms $(\sigma, \tau, \bar{\alpha})$ and $(v, \tau, \bar{\alpha})$.




Figure 8: Above are two integral functions $\sigma, \tau : \bar{P} \rightarrow \mathbb{Z}$ drawn as barcodes and a charge-preserving morphism $(\sigma, \tau, \bar{\alpha})$.

See Figure 8. The bottleneck distance between σ and τ is 1. We now compute the edit distance $d_{\text{Fnc}}(\sigma, \tau)$. Let $\alpha : P \rightarrow P$ be the bounded lattice function defined by

$$\alpha(0) := 0 \quad \alpha(1) := 0 \quad \alpha(2) := 3 \quad \alpha(3) := 3.$$

The lattice map α induces a charge-preserving morphism $(\sigma, \tau, \bar{\alpha})$ with distortion 2. This is the shortest path between σ and τ in Fnc so $d_{\text{Fnc}}(\sigma, \tau) = 2$.

References

- [1] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017.
- [2] M. Barnabei, A. Brini, and G.-C. Rota. The theory of Möbius functions. *Russian Mathematical Surveys*, 41(3):135–188, 1986.
- [3] Ulrich Bauer, Claudia Landi, and Facundo Mémoli. The Reeb graph edit distance is universal. *Foundations of Computational Mathematics*, December 2020.
- [4] Leo Betthausen, Peter Bubenik, and Parker B. Edwards. Graded persistence diagrams and persistence landscapes. *arXiv 1904.12807*, 2019.
- [5] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 16(3):77–102, 2015.
- [6] Peter Bubenik. The persistence landscape and some of its properties. In Nils A. Baas, Gunnar E. Carlsson, Gereon Quick, Markus Szymik, and Marius Thauale, editors, *Topological Data Analysis*, pages 97–117. Springer International Publishing, 2020.
- [7] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2001.
- [8] G. Carlsson, F. Mémoli, A. Ribeiro, and S. Segarra. Axiomatic construction of hierarchical clustering in asymmetric networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5219–5223, 2013.
- [9] Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.
- [10] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- [11] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the Twenty-Second Annual Symposium on Computational Geometry*, SCG ’06, pages 119–126, New York, NY, USA, 2006. Association for Computing Machinery.
- [12] Justin Curry and Amit Patel. Classification of constructible cosheaves. *Theory and Applications of Categories*, 35(27):1012–1047, 2020.
- [13] Barbara Di Fabio and Claudia Landi. The edit distance for Reeb graphs of surfaces. *Discrete & Computational Geometry*, 55(2):423–461, 2016.

- [14] Patrizio Frosini and Claudia Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9:596–603, 11 2001.
- [15] Gregory Henselman and Robert Ghrist. Matroid Filtrations and Computational Persistent Homology. *arXiv e-prints*, page arXiv:1606.00199, Jun 2016.
- [16] Gregory Henselman-Petrusek. Matroids and Canonical Forms: Theory and Applications. *arXiv e-prints*, page arXiv:1710.06084, Oct 2017.
- [17] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl. Deep learning with topological signatures. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, pages 1633–1643, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [18] Woojin Kim and Facundo Mémoli. Generalized persistence diagrams for persistence modules over posets. *arXiv 1810.11517*, 2018.
- [19] Claudia Landi and Patrizio Frosini. New pseudodistances for the size function space. In Robert A. Melter, Angela Y. Wu, and Longin Jan Latecki, editors, *Vision Geometry VI*, volume 3168, pages 52 – 60. International Society for Optics and Photonics, SPIE, 1997.
- [20] Alex McCleary and Amit Patel. Bottleneck stability for generalized persistence diagrams. *Proceedings of the American Mathematical Society*, 148:3149–3161, 2020.
- [21] Amit Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, 1(3):397–419, Jun 2018.
- [22] J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt. A stable multi-scale kernel for topological machine learning. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4741–4748, 2015.
- [23] Gian Carlo Rota. On the foundations of combinatorial theory I. Theory of Möbius functions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2(4):340–368, 1964.