

A Comparison Framework for Interleaved Persistence Modules

Shaun Harker, Miroslav Kramár, Rachel Levanger, Konstantin Mischaikow

► **To cite this version:**

Shaun Harker, Miroslav Kramár, Rachel Levanger, Konstantin Mischaikow. A Comparison Framework for Interleaved Persistence Modules. Journal of Applied and Computational Topology, Springer, 2019, 10.1007/s41468-019-00026-x . hal-01757092

HAL Id: hal-01757092

<https://hal.inria.fr/hal-01757092>

Submitted on 3 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Comparison Framework for Interleaved Persistence Modules

Shaun Harker^a, Miroslav Kramár^b, Rachel Levanger^c, Konstantin Mischaikow^a

^a*Department of Mathematics, Hill Center-Busch Campus, Rutgers University, 110 Frelingheusen Rd, Piscataway, NJ 08854-8019, USA*

^b*INRIA Saclay, 1 Rue Honor d'Estienne d'Orves, 91120 Palaiseau, France*

^c*Department of Electrical and Systems Engineering, University of Pennsylvania, 200 S. 33rd St., Philadelphia, PA 19104-6314, USA*

Abstract

We present a generalization of the induced matching theorem of [1] and use it to prove a generalization of the algebraic stability theorem for \mathbb{R} -indexed pointwise finite-dimensional persistence modules. Via numerous examples, we show how the generalized algebraic stability theorem enables the computation of rigorous error bounds in the space of persistence diagrams that go beyond the typical formulation in terms of bottleneck (or log bottleneck) distance.

1. Introduction

Persistent homology [14, 20, 17] is a key element in the rapidly-developing field of topological data analysis, where it is used both as a means of identifying geometric structures associated with data and as a data reduction tool. Any work with data involves approximations that arise from finite sampling, limits to measurement, and experimental or numerical errors. The results of this paper focus on obtaining rigorous bounds on the variations in persistence diagrams arising from these approximations.

To motivate this work, we begin with the observation that many problems in data analysis can be rephrased as a problem concerned with the analysis of the geometry induced by a scalar function $f: X \rightarrow \mathbb{R}$ defined on a topological space X . Two canonical examples are as follows. Assume that (X, ρ) is a metric space and let $\mathcal{X} \subseteq X$. Single-linkage hierarchical clustering problems based on \mathcal{X} are naturally associated with the function $f: X \rightarrow [0, \infty)$ given by

$$f(x) := \rho(x, \mathcal{X}) = \inf \{ \rho(x, \xi) : \xi \in \mathcal{X} \},$$

where clusters are derived from the connected components of the sublevel set

$$C(f, t) := \{x \in X : f(x) \leq t\}$$

Email addresses: sharker@math.rutgers.edu (Shaun Harker), miroslav.kramar@inria.fr (Miroslav Kramár), levanger@seas.upenn.edu (Rachel Levanger), mischaik@math.rutgers.edu (Konstantin Mischaikow)

for choices of $t \in [0, \infty)$. The collection $\{\mathbf{C}(f, t)\}_{t \in \mathbb{R}}$ is called the *sublevel set filtration* of X induced by f . Superlevel sets and superlevel set filtrations are defined similarly by considering the sets $\{x \in X : t \leq f(x)\}$ for every $t \in \mathbb{R}$.

Alternatively, assume that X is a topological domain and $f: X \rightarrow \mathbb{R}$ is a scalar value of a nonlinear physical model, e.g. the magnitude of vorticity or temperature field of a fluid, the chemical density in a reaction diffusion system, the magnitude of forces between particles in a granular system, etc. Patterns produced by these systems are often associated with sublevel or superlevel sets of f . In fact, the direct motivation for this work is to justify claims made in [16] concerning the time-evolution of patterns in convection models.

These examples are meant to motivate our interest in studying the geometry of the sets $\mathbf{C}(f, t)$. Homology provides a coarse but computable representation of this geometry. In particular, for each $t \in \mathbb{R}$, there is an assigned graded vector space

$$M(f)_t = H_\bullet(\mathbf{C}(f, t), \mathbf{k}),$$

where \mathbf{k} is a field. Because each $t \leq s$ implies $\mathbf{C}(f, t) \subseteq \mathbf{C}(f, s)$, the inclusion maps induce the following linear maps at the level of homology:

$$\varphi_{M(f)}(t, s): M(f)_t \rightarrow M(f)_s.$$

This homological information can be abstracted as follows.

Definition 1.1. A *persistence module* is a collection of vector spaces indexed by the real numbers, $\{V_t\}_{t \in \mathbb{R}}$, and linear maps $\{\varphi_V(s, t): V_s \rightarrow V_t\}_{s \leq t \in \mathbb{R}}$ satisfying the following conditions:

- (i) $\varphi_V(t, t) = \text{id}_{V_t}$ for every $t \in \mathbb{R}$, and
- (ii) $\varphi_V(s, t) \circ \varphi_V(r, s) = \varphi_V(r, t)$ for every $r \leq s \leq t$ in \mathbb{R} .

We write (V, φ_V) to denote the collection of vector spaces and compatible linear maps, and will sometimes just write V for the full persistence module when the maps are clear. We say that V is a *pointwise finite dimensional (PFD)* persistence module when every V_t is finite-dimensional.

As is described in Sections 2.1 and 4, a PFD persistence module gives rise to a *persistence diagram*, which is a set of points in $\overline{\mathbb{R}}^2 \times \mathbb{N}$, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. Given a PFD persistence module (V, φ_V) , we denote its associated persistence diagram by $\text{PD}(V)$.

Observe that we have outlined a procedure by which the sublevel sets of a scalar field f produce a persistence diagram PD . Returning to our examples, in the first case, it is reasonable to assume that the actual available data is $\mathcal{X}' \subseteq \mathcal{X} \subseteq X$, as opposed to \mathcal{X} , which represents the true set of objects upon which the clustering is to be based. In this case, collecting experimental or numerical data results in $f': X \rightarrow \mathbb{R}$, an approximation of the actual function of interest, f . Recent computational developments have led to the routine computation of PD' , the persistence diagram associated with \mathcal{X}' or f' . Thus, the natural question is this: how is PD' , the computed persistence diagram, related to PD , the persistence diagram of interest?




Figure 1: A schematic diagram illustrating the potential locations of persistence points of $\text{PD}(V)$ based on the computation of $\text{PD}(V^{\mathbb{Z}})$. The persistence point $(2,6)$ of $\text{PD}(V^{\mathbb{Z}})$ is matched with a persistence point of $\text{PD}(V)$ which must lie in the light gray region. If a persistence point of $\text{PD}(V^{\mathbb{Z}})$ occurs at one of the open circles, then it is possible that it is a computational artifact, i.e. that it is not matched with any persistence point in $\text{PD}(V)$. The dark gray region indicates the potential location of persistence points of $\text{PD}(V)$ that cannot be detected because of the integer valued approximation used to compute $\text{PD}(V^{\mathbb{Z}})$. Finally, if W is an arbitrary PFD persistence module and the bottleneck distance between $\text{PD}(V^{\mathbb{Z}})$ and $\text{PD}(W)$ is one, then $\text{PD}(W)$ may have a single point in the region indicated by the dashed square, and arbitrarily many persistence points in the region below the dashed line.

A fundamental result [14] in the theory of persistent homology is that a variety of metrics can be imposed on the space of persistence diagrams such that PD changes continuously with respect to L^∞ changes in f . However, these metrics provide limited control on the variation of individual points in the persistence diagram. Recent developments by Bauer and Lesnick [1] allow for comparisons of persistence modules through a matching of the associated persistence points. The primary theoretical results of this paper, Theorem 3.2 and Theorem 4.1, are extensions of Bauer and Lesnick's Induced Matching Theorem and Algebraic Stability Theorem, respectively.

As indicated above, the applications of these extensions provided the motivation for this paper. To give a particular example, consider the persistence module $V = (M(f), \varphi_{M(f)})$ associated with the scalar function $f: X \rightarrow \mathbb{R}$. However, assume that we are only able to sample the sublevel sets of f at the integers $\mathbb{Z} \subset \mathbb{R}$. As explained in Section 5.1, this sampling gives rise to a persistence module $V^{\mathbb{Z}}$. Figure 1 indicates the type of result that we obtain. What is shown is the persistence diagram $\text{PD}(V^{\mathbb{Z}})$ (which can be computed) which, for the region shown, is assumed to have a single persistence point at $(2,6)$. As a consequence of Proposition 5.2, we can conclude that the persistence

diagram of interest, $\text{PD}(V)$, contains a single persistence point in the light gray region. It is also possible that $\text{PD}(V)$ contains persistence points in the dark gray regions. This would correspond to geometrical features of $f: X \rightarrow \mathbb{R}$ that take place on a scale that is too fine to be detected by the integer-valued sampling. Finally, if a persistence point for $\text{PD}(V^{\mathbb{Z}})$ occurred at one of the open circles centered at $(n, n + 1)$, then this persistence point could be a computational artifact, i.e. it is not necessarily associated with any persistence point of $\text{PD}(V)$.

Figure 1 also indicates the advantage of comparing persistence diagrams using the matching theorems of this paper as opposed to the classical metrics such as the bottleneck distance [9]. In particular, if $\text{PD}(W)$ is an arbitrary persistence diagram whose bottleneck distance from $\text{PD}(V^{\mathbb{Z}})$ is one, then $\text{PD}(W)$ may have a single point in the region indicated by the dashed square and arbitrarily many persistence points in the region below the dashed line, versus a single point in the light gray box and arbitrarily many points in the dark gray region.

An outline of this paper is as follows. In Section 2 we review the essential concepts associated with persistence modules required for our results. This section defines the notions of persistence modules and their morphisms, interleavings, and induced matchings. Of particular note is the introduction of the concept of a non-constant translation pair that is used to extend the results of Bauer and Lesnick [1], where translation pairs are defined in terms of uniform translations. We also include a review of Galois connections, as we use these concepts for some proofs in Section 4.

Section 3 focuses on Theorem 3.2, which is an extension of the Induced Matching Theorem of [1]. The proof incorporates ideas from the theory of generalized interleavings of Bubenik, de Silva, and Scott [4].

Section 4 begins with the proof of Theorem 4.1, which follows closely the proof of the Algebraic Stability Theorem of [1]. The remainder of the section provides results, corollaries, and re-interpretations of Theorem 4.1. In particular, under the assumption that the maps in the translation pair are invertible, Corollary 4.4 provides an easy-to-state version of Theorem 4.1 that clarifies how translation pairs relate to stability in the space of persistence diagrams. Proposition 4.8 and Corollary 4.9 indicate how Theorem 4.1 applies to specific points in the associated persistence diagrams.

Finally, Section 5 provides examples of applications of Theorem 4.1. As indicated above Section 5.1 considers the problem of bounds on the desired persistence diagram under the assumption that values of the function $f: X \rightarrow \mathbb{R}$ can only be sampled discretely.

In Section 5.2, we consider the following problem associated with the first example of this introduction. Assume that one is given a large finite point cloud $\mathcal{X} \subset X$ for which one wishes to compute the persistence diagram $\text{PD}(V)$ for the persistence module $V = (M(f), \varphi_{M(f)})$. However, because of the size of \mathcal{X} , the computational cost of computing $\text{PD}(V)$ is prohibitive. At the time of this writing, this is a reasonable concern since the standard approach is to use a Vietoris-Rips complex (this is discussed at the beginning of Section 5.2) to compute $\text{PD}(V)$, and the size of this complex grows extremely fast as a function of the size of \mathcal{X} and the magnitude of f . This suggests that once the magnitude of f is too large, then one should subsample and compute an approximate persistence diagram $\text{PD}(V')$ based on $\mathcal{X}' \subset \mathcal{X}$. Proposition 5.6 provides a simple result bounding the locations of the persistence points in $\text{PD}(V)$ based on $\text{PD}(V')$. This result immediately suggests that if one could make use of a sequence of subsamples associated

with a sequence of values of f , then one could get a better approximation than just making use of a single subsampling. To obtain this result, we introduce in Section 5.2.2 the concept of stitching two persistence modules together to create a new persistence module. In Section 5.2.3, we outline how this can be used to obtain bounds on the persistence diagram of \mathcal{X} from a sequence of subsamples $\mathcal{X} = \mathcal{X}_0 \supset \mathcal{X}_1 \supset \cdots \supset \mathcal{X}_N$ and the associated persistence diagrams.

It can be argued that for applications, the most difficult task is the construction of the interleaving between the two persistence modules. However, as we hope the examples of Section 5 illustrate, once the interleavings are determined, working with our framework is straightforward. With this in mind, we include Table 1 in Section 5.3, providing an easily-referenced list of translation maps of generalized interleavings for common approximations to Vietoris-Rips and Čech filtrations.

2. Preliminaries

In this section, we summarize background material and establish notation for the work we present in this paper. In Section 2.1, we recall basic facts about persistence modules, their morphisms, and persistence diagrams. In Section 2.2 we provide a necessary background for interleavings of persistence modules. In Section 2.3 we give a treatment of monotone functions and Galois connections, and we define matchings. Section 2.4 introduces matchings between persistence diagrams induced by morphisms of persistence modules and recalls the results of Bauer and Lesnick [1] concerning these matchings.

2.1. Persistence Modules, Persistence Module Morphisms, and Persistence Diagrams

This section provides basic facts about persistence modules (Definition 1.1). For alternative treatments, see [1, 20, 4, 8].

Definition 2.1. A persistence module V is *trivial* if $V_t = 0$ for all $t \in \mathbb{R}$.

Definition 2.2. Let $J \subseteq \overline{\mathbb{R}}$ be a nonempty interval and let \mathbb{k} denote a field. The *interval persistence module* (I_J, φ_J) is defined by the vector spaces

$$(I_J)_t := \begin{cases} \mathbb{k} & \text{if } t \in J, \\ 0 & \text{otherwise,} \end{cases}$$

and transition maps

$$\varphi_J(s, t) := \begin{cases} \text{id}_{\mathbb{k}} & \text{if } s, t \in J, \\ 0 & \text{otherwise.} \end{cases}$$

Definition 2.3. Let (V, φ_V) and (W, φ_W) be persistence modules. A *persistence module morphism* $\phi: V \rightarrow W$ is a collection of linear maps $\{\phi_t: V_t \rightarrow W_t\}_{t \in \mathbb{R}}$ such that the following diagram commutes for all $s, t \in \mathbb{R}$ with $s \leq t$.

$$\begin{array}{ccc} V_s & \xrightarrow{\varphi_V(s,t)} & V_t \\ \downarrow \phi_s & & \downarrow \phi_t \\ W_s & \xrightarrow{\varphi_W(s,t)} & W_t \end{array}$$

5

If ϕ_t is injective (surjective) for every $t \in \mathbb{R}$, then we say that ϕ is a monomorphism (epimorphism). A persistence module morphism that is both a monomorphism and an epimorphism is an isomorphism.

Persistence modules and their morphisms form an abelian category [5]. Thus, it makes sense to talk about submodules, quotients, and direct sums of persistence modules. Moreover, the kernel and image of a persistence module morphism are submodules, and the cokernel of a persistence module morphism is a quotient persistence module. The following fundamental result (see [8, 10]) guarantees that nontrivial PFD persistence modules are direct sums of interval persistence modules.

Theorem 2.4. *Every non-trivial PFD persistence module V is a direct sum of interval persistence modules. Moreover, the direct sum decomposition of V into interval persistence modules is unique up to a reindexing of these interval persistence modules.*

This direct sum decomposition is called the *interval decomposition* of V , which we represent using the definitions that follow.

Definition 2.5. The set \mathbb{E} of *decorated points* is defined by

$$\mathbb{E} := \mathbb{R} \times \{-, +\} \cup \{-\infty, \infty\}.$$

For $t \in \mathbb{R}$, define $t^- := (t, -)$ and $t^+ := (t, +)$. Consider the ordering $- < +$ on the set $\{-, +\}$. Then there is a natural ordering on \mathbb{E} induced by a lexicographical ordering of \mathbb{R} and $\{-, +\}$, in that order, with $\{-\infty\}$ the minimal element and $\{\infty\}$ the maximal element.

Definition 2.6. Let $a, b \in \overline{\mathbb{R}}$ such that $a \leq b$. Any nonempty interval J with endpoints a and b can be represented by an ordered pair $(\mathcal{B}(J), \mathcal{D}(J))$ of decorated points where:

$$\mathcal{B}(J) := \begin{cases} -\infty & \text{if } a = -\infty, \\ a^- & \text{if } J \text{ is left closed,} \\ a^+ & \text{if } J \text{ is left open,} \end{cases} \quad \text{and} \quad \mathcal{D}(J) := \begin{cases} \infty & \text{if } b = \infty, \\ b^- & \text{if } J \text{ is right open,} \\ b^+ & \text{if } J \text{ is right closed.} \end{cases}$$

For an ordered pair (d_1, d_2) of decorated points with $d_1 < d_2$, we denote the interval they represent by $\langle d_1, d_2 \rangle$.

Definition 2.7. Let V be a PFD persistence module and \mathcal{J}_V be a multiset of interval persistence modules in the interval decomposition of V . Suppose that the function $m: \mathcal{J}_V \rightarrow \mathbb{N}$ assigns to every interval persistence module $I_J \in \mathcal{J}_V$ its multiplicity in \mathcal{J}_V . The *persistence diagram* of V is defined as the set

$$\mathbf{PD}(V) := \bigcup_{I_J \in \mathcal{J}_V} \{[\mathcal{B}(J), \mathcal{D}(J), 1], \dots, [\mathcal{B}(J), \mathcal{D}(J), m(I_J)]\} \subset \mathbb{E} \times \mathbb{E} \times \mathbb{N}.$$

Note that for every interval persistence module present in the interval decomposition of V , there is exactly one point in the persistence diagram. These points can be totally ordered as in the following definition.

Definition 2.8. Let \mathbf{PD} be a persistence diagram. The *left-handed ordering* of the points $[b, d, i] \in \mathbf{PD}$ is given by a lexicographical ordering applied to $(b, -d, i)$, where the minus sign indicates reversing the ordering for the second coordinate. The *right-handed ordering* of \mathbf{PD} is given by a lexicographical ordering applied to (d, b, i) .

2.2. Persistence Module Interleavings

In this section we review the notion of persistence module interleavings, introduced by Chazal, et al. in [7] and generalized by Bubenik, et al. in [4]. Interleavings provide a measure of similarity between persistence modules.

Definition 2.9. A function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is *monotone* if $x \leq y$ implies that $\sigma(x) \leq \sigma(y)$. If, in addition, $x \leq \sigma(x)$ for all $x \in \mathbb{R}$, then σ is called a *translation map*.

Definition 2.10. A pair (τ, σ) of monotone functions is a *translation pair* if $\tau \circ \sigma$ and $\sigma \circ \tau$ are translation maps.

Definition 2.11. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be monotone and let (V, φ_V) be a persistence module. The σ -shifted persistence module $(V(\sigma), \varphi_{V(\sigma)})$ is defined by the vector spaces

$$V(\sigma)_t := V_{\sigma(t)}$$

for $t \in \mathbb{R}$ and transition maps

$$\varphi_{V(\sigma)}(s, t) := \varphi_V(\sigma(s), \sigma(t))$$

for every $s \leq t \in \mathbb{R}$.

Definition 2.12. Let (V, φ_V) and (W, φ_W) be persistence modules, $\phi : V \rightarrow W$ a persistence module morphism, and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ a monotone function. The σ -shifted persistence module morphism $\phi(\sigma) : V(\sigma) \rightarrow W(\sigma)$ is defined by

$$\phi(\sigma)_t := \phi_{\sigma(t)}$$

for every $t \in \mathbb{R}$.

Definition 2.13. Let (V, φ_V) and (W, φ_W) be persistence modules and let (τ, σ) be a translation pair. The ordered pair of persistence modules (V, W) is (τ, σ) -interleaved if there exist persistence module morphisms $\phi : V \rightarrow W(\tau)$ and $\psi : W \rightarrow V(\sigma)$ such that

$$\psi(\tau)_t \circ \phi_t = \varphi_V[t, (\sigma \circ \tau)(t)]$$

and

$$\phi(\sigma)_t \circ \psi_t = \varphi_W[t, (\tau \circ \sigma)(t)]$$

for all $t \in \mathbb{R}$. We refer to these last two conditions as the *commutativity constraint* of the interleaving. The persistence module morphisms ϕ and ψ are called *interleaving morphisms*.

Definition 2.14. Given a persistence module V and a translation map σ , define a persistence module morphism $\phi_{\{V, \sigma\}} : V \rightarrow V(\sigma)$ by $(\phi_{\{V, \sigma\}})_t := \varphi_V(t, \sigma(t))$ for all $t \in \mathbb{R}$.

Remark 2.15. The notion of δ -interleaved persistence modules, presented in [1, 7, 8], is equivalent to the notion of (τ, σ) -interleaved persistence modules with $\tau(t) = t + \delta = \sigma(t)$.

Remark 2.16. Two persistence modules that are 0-interleaved are isomorphic as persistence modules.

Recall that the transition maps of the trivial persistence module are trivial. The following definition provides a way of quantifying the similarity between a persistence module V and the trivial persistence module in terms of a translation map.

Definition 2.17. Let σ be a translation map. A persistence module (V, φ_V) is σ -trivial if $\varphi_V(t, \sigma(t)) = 0$ for all $t \in \mathbb{R}$.

The following proposition provides information about the kernel and cokernel of the interleaving morphisms of two interleaved persistence modules.

Proposition 2.18. Let V and W be persistence modules such that (V, W) are (τ, σ) -interleaved via the morphisms $\phi : V \rightarrow W(\tau)$ and $\psi : W \rightarrow V(\sigma)$. Then

(i) $\ker \phi$ and $\operatorname{coker} \phi$ are $(\sigma \circ \tau)$ -trivial, and

(ii) $\ker \psi$ and $\operatorname{coker} \psi$ are $(\tau \circ \sigma)$ -trivial.

Proof. (i) The persistence module $\ker \phi$ is $(\sigma \circ \tau)$ -trivial if and only if $\varphi_{\ker \phi}(t, \sigma \circ \tau(t)) = 0$ for all $t \in \mathbb{R}$. By the commutativity constraint of a (τ, σ) -interleaving, we know that $\varphi_V(t, \sigma \circ \tau(t)) = \psi(\tau)_t \circ \phi_t$. Thus,

$$\varphi_V(t, \sigma \circ \tau(t))|_{\ker \phi_t} = \psi(\tau)_t \circ \phi_t|_{\ker \phi_t} = 0.$$

By the definition of the persistence module $\ker \phi$, we have

$$\varphi_{\ker \phi}(t, \sigma \circ \tau(t)) = \varphi_V(t, \sigma \circ \tau(t))|_{\ker \phi_t} = 0,$$

and so we are done.

The persistence module $\operatorname{coker} \phi$ is $(\sigma \circ \tau)$ -trivial if and only if $\varphi_{\operatorname{coker} \phi}(t, \sigma \circ \tau(t)) = 0$ for all $t \in \mathbb{R}$. Recall that the transition maps of the persistence module $\operatorname{coker} \phi$ are defined to be the unique linear maps $\varphi_{\operatorname{coker} \phi}(r, s)$ such that

$$\varphi_{\operatorname{coker} \phi}(r, s) \circ q_r = q_s \circ \varphi_{W(\tau)}(r, s)$$

for every $r \leq s \in \mathbb{R}$, where $q_r := (\alpha \mapsto \alpha + \operatorname{im} \phi_r)$ for every $\alpha \in W(\tau)_r$ is the quotient map. Thus, it suffices to show that $\operatorname{im} \varphi_{W(\tau)}(t, \sigma \circ \tau(t)) \subseteq \operatorname{im} \phi(\sigma \circ \tau)_t$ for every $t \in \mathbb{R}$. For $t \in \mathbb{R}$ we have

$$\begin{aligned} \varphi_{W(\tau)}(t, \sigma \circ \tau(t)) &= \varphi_W(\tau(t), \tau \circ \sigma \circ \tau(t)) \\ &= \phi(\sigma)_{\tau(t)} \circ \psi_{\tau(t)} \\ &= \phi_{\sigma \circ \tau(t)} \circ \psi_{\tau(t)}, \end{aligned}$$

where the first equality follows from the definition of the maps $\varphi_{W(\tau)}$, the second equality follows from the commutativity constraint of the interleaving morphisms ϕ and ψ , and the last equality follows from the definition of $\phi(\sigma)$. Hence, we have shown that

$$\operatorname{im} \varphi_{W(\tau)}(t, (\sigma \circ \tau)(t)) \subseteq \operatorname{im} \phi_{\sigma \circ \tau(t)}$$

for every $t \in \mathbb{R}$.

Part (ii) follows from (i) by reversing the roles of ϕ and ψ , creating a (σ, τ) -interleaving of W and V ; it follows directly that $\ker \psi$ and $\operatorname{coker} \psi$ are $(\tau \circ \sigma)$ -trivial. \square

We close this section by recalling a result [4, Proposition 2.2.11] that allows us to compose interleavings.

Proposition 2.19. *Let (U, φ_U) , (V, φ_V) , and (W, φ_W) be persistence modules such that (U, V) are (τ, σ) -interleaved and (V, W) are (τ', σ') -interleaved. Then the persistence modules (U, V) are $(\tau' \circ \tau, \sigma \circ \sigma')$ -interleaved.*

2.3. Galois Connections

In this section we provide a brief review of Galois connections (see [11]) and establish some Galois connections that are used in Section 4.

Definition 2.20. Let P and Q be posets and suppose $f : P \rightarrow Q$ and $g : Q \rightarrow P$ are monotone functions. The pair (f, g) is a Galois connection if for all $x \in P$ and all $y \in Q$

$$f(x) \leq y \text{ if and only if } x \leq g(y).$$

Proposition 2.21. *Suppose P , Q , and R are posets and $f : P \rightarrow Q$, $g : Q \rightarrow P$, $f' : Q \rightarrow R$, and $g' : R \rightarrow Q$ are monotone functions. Suppose further that (f, g) and (f', g') are Galois connections. Then $(f' \circ f, g \circ g')$ is a Galois connection.*

Proof. For all $x \in P$, $y \in R$, we have $(f' \circ f)(x) \leq y \Leftrightarrow f(x) \leq g'(y) \Leftrightarrow x \leq (g \circ g')(y)$. \square

We make use of Galois connections whose definition requires the poset \mathbb{R}_L of lower sets of \mathbb{R} (i.e. intervals $\langle -\infty, e \rangle$ for $e \in \mathbb{E}$) and the poset \mathbb{R}_U of upper sets of \mathbb{R} (i.e. intervals $\langle e, \infty \rangle$ for $e \in \mathbb{E}$). In both cases, the ordering is given by inclusion. Define the order isomorphisms $|\cdot\rangle : \mathbb{E} \rightarrow \mathbb{R}_L$ and $\langle \cdot | : \mathbb{E} \rightarrow \mathbb{R}_U$ as

$$|e\rangle := \langle -\infty, e \rangle \quad \text{and} \quad \langle e | := \langle e, \infty \rangle.$$

Moreover, for any set $S \subseteq \mathbb{R}$, define

$$\downarrow S = \{x \in \mathbb{R} : \exists y \in S \text{ s.t. } x \leq y\} \in \mathbb{R}_L,$$

$$\uparrow S = \{x \in \mathbb{R} : \exists y \in S \text{ s.t. } y \leq x\} \in \mathbb{R}_U.$$

Definition 2.22. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a monotone function. We define $\sigma^\downarrow : \mathbb{E} \rightarrow \mathbb{E}$, $\sigma^\uparrow : \mathbb{E} \rightarrow \mathbb{E}$, and $\sigma^* : \mathbb{E} \rightarrow \mathbb{E}$ by requiring that the following sets are equal:

$$\begin{aligned} |\sigma^\downarrow(e)\rangle &= \downarrow\{\sigma(x) : x \in |e\rangle\}, \\ \langle \sigma^\uparrow(e) | &= \uparrow\{\sigma(x) : x \in \langle e | \}, \text{ and} \\ |\sigma^*(e)\rangle &= \sigma^{-1}(|e\rangle), \text{ or, equivalently, } \langle \sigma^*(e) | = \sigma^{-1}(\langle e |). \end{aligned}$$

for all $e \in \mathbb{E}$. Note that these functions are defined since $|\cdot\rangle$ and $\langle \cdot |$ are order isomorphisms and σ is monotone.

Proposition 2.23. *Let $\sigma, \tau : \mathbb{R} \rightarrow \mathbb{R}$ be monotone functions. Then $(\sigma \circ \tau)^\uparrow = \sigma^\uparrow \circ \tau^\uparrow$, $(\sigma \circ \tau)^\downarrow = \sigma^\downarrow \circ \tau^\downarrow$, and $(\sigma \circ \tau)^* = \tau^* \circ \sigma^*$.*

Proof. It is easy to verify that $\uparrow(\sigma \circ \tau)(S) = \uparrow\sigma(\uparrow\tau(S))$, $\downarrow(\sigma \circ \tau)(S) = \downarrow\sigma(\downarrow\tau(S))$, and $(\sigma \circ \tau)^{-1}(S) = \tau^{-1}(\sigma^{-1}(S))$ for any $S \subseteq \mathbb{R}$. Now the result follows from Definition 2.22 and the above equalities applied to $S = \langle e |$ or $S = |e\rangle$ for $e \in \mathbb{E}$. \square

Proposition 2.24. *Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a monotone function. Then both $(\sigma^\downarrow, \sigma^\star)$ and $(\sigma^\star, \sigma^\uparrow)$ are Galois connections.*

Proof. We show $(\sigma^\downarrow, \sigma^\star)$ is a Galois connection. Suppose first that $\sigma^\downarrow(x) \leq y$ for some $x, y \in \mathbb{E}$. We show $x \leq \sigma^\star(y)$. Since $|\cdot\rangle$ is an order isomorphism, $\sigma^\downarrow(x) \leq y$ is equivalent to $|\sigma^\downarrow(x)\rangle \subseteq |y\rangle$. By the definition of σ^\downarrow , this is equivalent to $\downarrow\sigma(|x\rangle) \subseteq |y\rangle$. Taking the preimage of both sides yields $\sigma^{-1}(\downarrow\sigma(|x\rangle)) \subseteq \sigma^{-1}(|y\rangle)$. Since $|x\rangle \subseteq \sigma^{-1}(\sigma(|x\rangle)) \subseteq \sigma^{-1}(\downarrow\sigma(|x\rangle))$, we obtain $|x\rangle \subseteq \sigma^{-1}(|y\rangle)$. Since $|\cdot\rangle$ is an order isomorphism, we conclude that $x \leq \sigma^\star(y)$.

We now prove the converse. That is, we suppose that $x \leq \sigma^\star(y)$ and show $\sigma^\downarrow(x) \leq y$. From $x \leq \sigma^\star(y)$, we obtain $|x\rangle \subseteq \sigma^{-1}(|y\rangle)$. Applying σ to both sides and taking the downward closure gives $\downarrow\sigma(|x\rangle) \subseteq \downarrow\sigma(\sigma^{-1}(|y\rangle))$. See that $\downarrow\sigma(\sigma^{-1}(|y\rangle)) = |y\rangle$, hence $\downarrow\sigma(|x\rangle) \subseteq |y\rangle$, or equivalently, $\sigma^\downarrow(x) \leq y$, as desired. Hence, the pair $(\sigma^\downarrow, \sigma^\star)$ is a Galois connection. To show the pair $(\sigma^\star, \sigma^\uparrow)$ is a Galois connection, one proceeds similarly. \square

The following maps are used to move between points in $\overline{\mathbb{R}}$ and decorated points in \mathbb{E} .

Definition 2.25. The maps $\pi : \mathbb{E} \rightarrow \overline{\mathbb{R}}$, $i^- : \overline{\mathbb{R}} \rightarrow \mathbb{E}$ and $i^+ : \overline{\mathbb{R}} \rightarrow \mathbb{E}$ are defined by:

$$\pi(t^\pm) = t, \quad i^\pm(t) = t^\pm,$$

for $t \in \mathbb{R}$, and

$$\pi(\pm\infty) = \pm\infty, \quad i^\pm(\pm\infty) = \pm\infty.$$

Definition 2.26. Let $f : \mathbb{E} \rightarrow \mathbb{E}$ be a monotone function. Define functions $f_+ : \overline{\mathbb{R}} \rightarrow \overline{\mathbb{R}}$ and $f_- : \overline{\mathbb{R}} \rightarrow \overline{\mathbb{R}}$ via

$$f_+ := \pi \circ f \circ i^+ \quad \text{and} \quad f_- := \pi \circ f \circ i^-.$$

We close this section by establishing some Galois connections that will be needed later.

Proposition 2.27. *Both (i^-, π) and (π, i^+) are Galois connections.*

Proof. First we show that (π, i^+) is a Galois connection. Using the easily-verified relations $\pi \circ i^+ = \text{id}$ and $\text{id} \leq i^+ \circ \pi$, we have, for all $x \in \mathbb{E}$ and $y \in \overline{\mathbb{R}}$, the circle of implications $(\pi(x) \leq y) \Rightarrow ((i^+ \circ \pi)(x) \leq i^+(y)) \Rightarrow (x \leq i^+(y)) \Rightarrow (\pi(x) \leq (\pi \circ i^+)(y)) \Rightarrow (\pi(x) \leq y)$. Thus, $(\pi(x) \leq y) \Leftrightarrow (x \leq i^+(y))$. That is, the pair (π, i^+) is a Galois connection.

Showing that the pair (i^-, π) is a Galois connection proceeds similarly. Using the easily-verified relations $\pi \circ i^- = \text{id}$ and $i^- \circ \pi \leq \text{id}$, we have, for all $x \in \overline{\mathbb{R}}$ and $y \in \mathbb{E}$, the circle of implications $(x \leq \pi(y)) \Rightarrow (i^-(x) \leq (i^- \circ \pi)(y)) \Rightarrow (i^-(x) \leq y) \Rightarrow ((\pi \circ i^-)(x) \leq \pi(y)) \Rightarrow (x \leq \pi(y))$. Thus, $(x \leq \pi(y)) \Leftrightarrow (i^-(x) \leq y)$. That is, the pair (i^-, π) is a Galois connection. \square

Proposition 2.28. *Suppose $f, g : \mathbb{E} \rightarrow \mathbb{E}$ are monotone functions such that the pair (f, g) is a Galois connection. Then the pair (f_-, g_+) is a Galois connection.*

Proof. By Definition 2.26, we have $(f_-, g_+) = (\pi \circ f \circ i^-, \pi \circ g \circ i^+)$. The result follows from Proposition 2.21 and Proposition 2.27. \square

2.4. Induced Matchings on Persistence Diagrams

In this section, we summarize the work of Bauer and Lesnick [1, 2] on matchings of persistence diagrams of PDF persistence modules V and W induced by a morphism $\phi : V \rightarrow W$.

Definition 2.29. Let \mathcal{X} be a relation between sets S and T (i.e. $\mathcal{X} \subseteq S \times T$). We say that \mathcal{X} is a *matching* $\mathcal{X} : S \rightarrow T$ if \mathcal{X} is the graph of an injective function $\mathcal{X}' : S' \rightarrow T'$, where $S' \subseteq S$ and $T' \subseteq T$. We define the domain and image of a matching via $\text{dom } \mathcal{X} := \text{dom } \mathcal{X}'$ and $\text{im } \mathcal{X} := \text{im } \mathcal{X}'$, and we will use the notation $\mathcal{X}(s) = t$ to denote $(s, t) \in \mathcal{X}$.

We use the following notation to define matchings induced by morphisms.

Definition 2.30. Let (V, φ_V) be a PDF persistence module. For $b, d \in \mathbb{E}$, we define two subsets of the persistence diagram $\mathbf{PD}(V)$ by:

$$\begin{aligned}\mathbf{PD}_b(V) &:= \{[b, d', i] : [b, d', i] \in \mathbf{PD}(V)\}, \\ \mathbf{PD}^d(V) &:= \{[b', d, i] : [b', d, i] \in \mathbf{PD}(V)\}.\end{aligned}$$

If V is a PFD persistence module, then the sets $\mathbf{PD}_b(V)$ and $\mathbf{PD}^d(V)$ are countable for every $b, d \in \mathbb{E}$. The left-handed (right-handed) ordering on $\mathbf{PD}(V)$ induces a total ordering on $\mathbf{PD}_b(V)$ ($\mathbf{PD}^d(V)$). We will always consider these sets together with these induced orderings. Therefore, if we talk about the first n points in $\mathbf{PD}_b(V)$ or $\mathbf{PD}^d(V)$, we mean the n smallest points with respect to the induced ordering. The following proposition allows us to define matchings between the PFD persistence modules as introduced in [1].

Proposition 2.31 (Theorem 4.2 [1]). *Let V and W be PFD persistence modules, and let the symbol $|\cdot|$ denote the cardinality of a set.*

- (i) *If there exists a monomorphism $V \hookrightarrow W$, then $|\mathbf{PD}^d(V)| \leq |\mathbf{PD}^d(W)|$ for $d \in \mathbb{E}$.*
- (ii) *If there exists an epimorphism $V \twoheadrightarrow W$, then $|\mathbf{PD}_b(W)| \leq |\mathbf{PD}_b(V)|$ for $b \in \mathbb{E}$.*

The next two propositions establish the matchings induced by monomorphisms and epimorphisms. The proof of parts (i)-(iii) of each proposition is a simple consequence of the previous proposition, while (iv) follows from [1, Theorem 4.2].

Proposition 2.32. *Let V, W be PFD persistence modules. If there exists a monomorphism from V to W , then there exists a unique matching $\mathcal{X}_{i(V,W)} : \mathbf{PD}(V) \rightarrow \mathbf{PD}(W)$ which satisfies:*

- (i) *the domain of $\mathcal{X}_{i(V,W)}$ is $\mathbf{PD}(V)$,*
- (ii) *$\mathcal{X}_{i(V,W)}$ preserves the right-handed ordering,*
- (iii) *$\mathcal{X}_{i(V,W)}$ maps the points in $\mathbf{PD}^d(V)$ to the smallest $|\mathbf{PD}^d(V)|$ points in $\mathbf{PD}^d(W)$,*
- (iv) *if $\mathcal{X}_{i(V,W)}([b, d, i]) = [b', d', i']$, then $d = d'$ and $b' \leq b$.*

Proposition 2.33. *Let V, W be PFD persistence modules. If there exists an epimorphism from V to W , then there exist a unique matching $\mathcal{X}_{s(V,W)} : \mathbf{PD}(V) \rightarrow \mathbf{PD}(W)$ that satisfies:*

- (i) the image of $\mathcal{X}_{s(V,W)}$ is $\mathbf{PD}(W)$,
- (ii) the inverse relation $\mathcal{X}_{s(V,W)}^{-1}$ preserves the left-handed ordering,
- (iii) $\mathcal{X}_{s(V,W)}^{-1}$ maps the points in $\mathbf{PD}_b(W)$ to the smallest $|\mathbf{PD}_b(W)|$ points in $\mathbf{PD}_b(V)$,
- (iv) if $\mathcal{X}_{s(V,W)}([b, d, i]) = [b', d', i']$ then $b = b'$ and $d' \leq d$.

Every persistence module morphism $\phi: V \rightarrow W$ can be factored as the composition of an epimorphism and monomorphism as follows:

$$V \rightarrow \text{im } \phi \hookrightarrow W.$$

Therefore, we can define a matching $\mathcal{X}_\phi: \mathbf{PD}(V) \rightarrow \mathbf{PD}(W)$ via the composition of the following relations:

$$\mathcal{X}_\phi := \mathcal{X}_{s(V, \text{im } \phi)} \circ \mathcal{X}_{i(\text{im } \phi, W)}.$$

In general, it is not true that if $\phi: U \rightarrow V$ and $\psi: V \rightarrow W$ are PFD persistence module morphisms then $\mathcal{X}_{\psi \circ \phi} = \mathcal{X}_\psi \circ \mathcal{X}_\phi$. However, the following result [1, Proposition 5.7] provides hypotheses under which this is true.

Proposition 2.34. *Let $\phi: U \rightarrow V$ and $\psi: V \rightarrow W$ be PFD persistence module morphisms. If ϕ and ψ are either both injective or both surjective, then $\mathcal{X}_{\psi \circ \phi} = \mathcal{X}_\psi \circ \mathcal{X}_\phi$.*

Definition 2.35. Let $A, B \subseteq \mathbb{R}$. We say that A bounds B below, written $A \triangleleft B$, if for all $y \in B$, there exists some $x \in A$ with $x \leq y$. We say that B bounds A above, written $A \blacktriangleleft B$, if for all $x \in A$, there exists some $y \in B$ such that $x \leq y$. We say that B overlaps A above, written $A \blacktriangleleft\blacktriangleleft B$, if and only if each of the following conditions hold:

$$A \triangleleft B, \quad A \blacktriangleleft B, \quad \text{and} \quad A \cap B \neq \emptyset.$$

Proposition 2.36. [1, Proposition 5.3] *Let $\phi: V \rightarrow W$ be a PFD persistence module morphism. If $\mathcal{X}_\phi([b, d, i]) = [b', d', i']$, then $\langle b', d' \rangle \blacktriangleleft\blacktriangleleft \langle b, d \rangle$.*

3. Generalized Induced Matching Theorem

In this section we present a generalization of the Induced Matching Theorem of [1].

Definition 3.1. Let σ be a translation map and let $b, d \in \mathbb{E}$. An interval $\langle b, d \rangle$ is σ -trivial if $\langle b, d \rangle \cap \sigma(\langle b, d \rangle) = \emptyset$. A point $[b, d, i] \in \mathbb{E} \times \mathbb{E} \times \mathbb{N}$ is σ -trivial if $\langle b, d \rangle$ is σ -trivial. A point in $\mathbb{E} \times \mathbb{E} \times \mathbb{N}$ that is not σ -trivial is called σ -nontrivial.

Theorem 3.2. *Let $\phi: V \rightarrow W$ be a PFD persistence module morphism and σ a translation map. Suppose that $\mathcal{X}_\phi([b, d, i]) = [b', d', i']$.*

- (i) If $\text{coker } \phi$ is σ -trivial, then

$$\langle b, d \rangle \triangleleft \sigma(\langle b', d' \rangle) \text{ and } \text{im } \mathcal{X}_\phi \text{ contains all } \sigma\text{-nontrivial points in } \mathbf{PD}(W).$$

- (ii) If $\text{ker } \phi$ is σ -trivial, then

$$\sigma^{-1}(\langle b, d \rangle) \blacktriangleleft \langle b', d' \rangle \text{ and } \text{dom } \mathcal{X}_\phi \text{ contains all } \sigma\text{-nontrivial points in } \mathbf{PD}(V).$$

Note that the Induced Matching Theorem of [1] follows from Theorem 3.2 by setting $\sigma(t) = t + \delta$ for $\delta \geq 0$.

The remainder of this section is devoted to the proof of Theorem 3.2.

Definition 3.3. Let V be a persistence module and σ a translation map. Define vector spaces of the persistence module V^σ by

$$V_t^\sigma := \bigcup_{\{s: \sigma(s) \leq t\}} \text{im } \varphi_V(s, t)$$

for $t \in \mathbb{R}$. The linear maps φ_{V^σ} are given by restriction of the maps φ_V to V^σ .

The following lemma shows that V^σ is a persistence module.

Lemma 3.4. *Let V be a persistence module and σ a translation map. Then, V^σ is a persistence submodule of V .*

Proof. By definition, V_t^σ is a subspace of V_t for $t \in \mathbb{R}$. To see that V^σ is a persistence submodule of V , we must show that $\text{im } \varphi_V(s, t)|_{V_s^\sigma} \subseteq V_t^\sigma$ for all $s \leq t$. To do this, we consider $y \in V_s^\sigma$ and show that $\varphi(s, t)(y) \in V_t^\sigma$. By definition, there exists $x \in V_{t'}$ for some $t' \in \mathbb{R}$ such that $\sigma(t') \leq s$ and $\varphi_V(t', s)(x) = y$. Thus,

$$\varphi_V(s, t)(y) = \varphi_V(s, t)[\varphi_V(t', s)(x)] = \varphi_V(t', t)(x).$$

Since σ is a translation map, $t' \leq \sigma(t') \leq s \leq t$, and so $\varphi(s, t)(y) \in \text{im } \varphi_V(t', t) \subseteq V_t^\sigma$. \square

Lemma 3.5. *Let $\phi : V \rightarrow W$ be a persistence module morphism and σ a translation map.*

- (i) *If $\text{im } \phi$ is σ -trivial, then $W_t^\sigma \subseteq \text{im } \phi_t \subseteq W_t$ for every $t \in \mathbb{R}$, and*
- (ii) *if $\ker \phi$ is σ -trivial, then $\ker \phi_t \subseteq (\ker \phi_{\{V, \sigma\}})_t \subseteq V_t$ for every $t \in \mathbb{R}$.*

Proof. (i) By definition, given a morphism $\phi : V \rightarrow W$, the persistence module $\text{coker } \phi$ is σ -trivial if and only if

$$\varphi_{\text{coker } \phi}(t, \sigma(t)) = 0 \quad \text{for all } t \in \mathbb{R},$$

which is true if and only if

$$\text{im } \varphi_W(t, \sigma(t)) \subseteq \text{im } \phi(\sigma)_t \quad \text{for all } t \in \mathbb{R},$$

which again is true if and only if for each $t \in \mathbb{R}$ and each $x \in W_t$, there exists some $y \in V_{\sigma(t)}$ such that

$$\varphi_W(t, \sigma(t))(x) = \phi(\sigma)_t(y).$$

So, to prove that $W_t^\sigma \subseteq \text{im } \phi_t$, it is enough to show that $\text{im } \varphi_W(t', t) \subseteq \text{im } \phi_{t'}$ for every $t' \in \mathbb{R}$ such that $t' \leq \sigma(t)$. By commutativity of the diagram

$$\begin{array}{ccccc} V_{\sigma(t')} & \xrightarrow{\varphi_V(\sigma(t'), t)} & V_t & & \\ & & \downarrow \phi(\sigma)_{t'} & & \downarrow \phi_t \\ W_{t'} & \xrightarrow{\varphi_W(t', \sigma(t'))} & W_{\sigma(t')} & \xrightarrow{\varphi_W(\sigma(t'), t)} & W_t \end{array}$$

we have

$$\varphi_W(t', t)(x) = \varphi_W(\sigma(t'), t)[\varphi_W(t', \sigma(t'))(x)] = \phi_t[\varphi_V(\sigma(t'), t)(y)],$$

and so $\text{im } \varphi_W(t', t) \subseteq \text{im } \phi_t$.

To prove (ii), we show that $\ker \phi_t \subseteq (\ker \varphi_{\{V, \sigma\}})_t$ for all $t \in \mathbb{R}$ whenever ϕ has a σ -trivial kernel. By definition, $\ker \phi$ is σ -trivial if and only if

$$\varphi_V(t, \sigma(t))|_{\ker \phi_t} = \varphi_{\ker \phi}(t, \sigma(t)) = 0$$

for all $t \in \mathbb{R}$. Hence, $\ker \phi_t \subseteq \ker \varphi_V(t, \sigma(t)) = (\ker \varphi_{\{V, \sigma\}})_t$ for all $t \in \mathbb{R}$ if and only if $\ker \phi$ is σ -trivial. \square

We now study the relationship between the persistence module V and the persistence modules V^σ and $V/\ker \phi_{\{V, \sigma\}}$. We start by considering an interval persistence module.

Lemma 3.6. *Let I_J be an interval persistence module and σ a translation map. If $J \cap \sigma(J) \neq \emptyset$, then*

- (i) $I_J^\sigma \cong I_{J \cap \text{Conv}(\sigma(J))}$, where $\text{Conv}(\sigma(J))$ is the convex hull of $\sigma(J)$,
- (ii) $V/\ker \phi_{\{V, \sigma\}} \cong I_{J \cap \sigma^{-1}(J)}$,

If $J \cap \sigma(J) = \emptyset$, then both persistence modules are trivial.

Proof. (i) We first show that $(I_J^\sigma)_t \neq \emptyset$ for $t \in J \cap \text{Conv}(\sigma(J))$. Since σ is a translation map, for every $t \in \text{Conv}(\sigma(J))$, there exist $s \in J$ such that $s \leq \sigma(s) \leq t$. The fact that $t \in J$ implies $\text{im } \varphi_{I_J}(s, t) = \text{id}_k$ and hence $\emptyset \neq \text{im } \varphi_{I_J}(s, t) \subseteq (I_J^\sigma)_t$.

To finish the proof of (i), we need to show that $(I_J^\sigma)_t = 0$ if $t \notin J \cap \text{Conv}(\sigma(J))$. Suppose that $t \notin J$. Then $(I_J^\sigma)_t \subseteq (I_J)_t = 0$. On the other hand, if $t \notin \text{Conv}(\sigma(J))$ and $\sigma(s) \leq t$, then $s \notin J$, and so $\varphi_{I_J}(s, t) = 0$ for all s such that $\sigma(s) \leq t$.

If $J \cap \sigma(J) = \emptyset$, then $J \cap \text{Conv}(\sigma(J)) = \emptyset$ and we showed above that $\varphi_{I_J}(s, t) = 0$ for all $s \leq t$. It follows that I_J^σ is trivial. Similar arguments can be used to prove (ii), and we leave it to the reader. \square

For the following two definitions, for an interval $J \subseteq \mathbb{R}$, we recall the symbols $\mathcal{B}(J)$ and $\mathcal{D}(J)$ (Definition 2.6) give the left and right (decorated) endpoints of J , respectively.

Proposition 3.7. *Let V be a PFD persistence module and σ a translation map. Suppose that $[b, d, i] \in \mathbf{PD}(V)$ and the interval $J = \langle b, d \rangle \cap \text{Conv}(\sigma \langle b, d \rangle)$. Then $[\mathcal{B}(J), d, i] \in \mathbf{PD}(V^\sigma)$ if and only if $J \neq \emptyset$. Moreover,*

$$\mathcal{X}_{i(V^\sigma, V)}([\mathcal{B}(J), d, i]) = [b, d, i].$$

Proof. Let $[b, d, i] \in \mathbf{PD}(V)$. Since σ a translation map, $J = \langle \mathcal{B}(J), d \rangle$. By Lemma 3.6, the interval persistence module $I_{\langle b, d \rangle}^\sigma$ is nontrivial if and only if $J \neq \emptyset$. It follows from Theorem 2.4 that the interval persistence module $I_{\langle b, d \rangle}^\sigma (\cong I_{\langle \mathcal{B}(J), d \rangle})$ belongs to the interval decomposition of V^σ if and only if $J \neq \emptyset$. Thus, $[\mathcal{B}(J), d, i] \in \mathbf{PD}(V^\sigma)$ for every $[b, d, i] \in \mathbf{PD}(V)$ such that $\langle b, d \rangle \cap \text{Conv}(\sigma \langle b, d \rangle) \neq \emptyset$. Now, $\mathcal{X}_{i(V^\sigma, V)}([\mathcal{B}(J), d, i]) = [b, d, i]$ is a simple consequence of Proposition 2.32. \square

By using Lemma 3.6(ii) and similar reasoning as above, one can prove the following about the matching $\mathcal{X}_{s(V, V/\ker \phi_{\{V, \sigma\}})}$.

Proposition 3.8. *Let V be a PFD persistence module and σ be a translation map. Suppose that $[b, d, i] \in \mathbf{PD}(V)$. Then $[b, \mathcal{D}(J), i] \in \mathbf{PD}(V/\ker \phi_{\{V, \sigma\}})$ if and only if $J := \langle b, d \rangle \cap \sigma^{-1}(\langle b, d \rangle) \neq \emptyset$. Moreover,*

$$\mathcal{X}_{s(V, V/\ker \phi_{\{V, \sigma\}})}([b, d, i]) = [b, \mathcal{D}(J), i].$$

Proof of Theorem 3.2. Our proof closely follows the proof of the Induced Matching Theorem in [1]. To prove (i), we start by establishing the existence of certain matchings. By Lemma 3.5(i), W^σ is a submodule of $\text{im } \phi$ and so the matching $\mathcal{X}_{i(W, \text{im } \phi)}$ is defined. It follows from Lemma 3.4 that W^σ is a submodule of W and thus $\mathcal{X}_{i(W^\sigma, W)}$ is defined. Proposition 2.34 implies that the following diagram commutes.

$$\begin{array}{ccccc} & & \mathbf{PD}(W) & & \\ & \swarrow \mathcal{X}_{i(W^\sigma, W)} & \uparrow & \nwarrow \mathcal{X}_\phi & \\ \mathbf{PD}(W^\sigma) & & & & \mathbf{PD}(V) \\ & \swarrow \mathcal{X}_{i(W^\sigma, \text{im } \phi)} & \uparrow \mathcal{X}_{i(\text{im } \phi, W)} & \nwarrow \mathcal{X}_{s(V, \text{im } \phi)} & \\ & & \mathbf{PD}(\text{im } \phi) & & \end{array}$$

It follows from Proposition 2.33(i) that $\text{im } \mathcal{X}_\phi = \text{im } \mathcal{X}_{i(\text{im } \phi, W)}$. By commutativity of the left triangle, $\text{im } \mathcal{X}_{i(W^\sigma, W)} \subseteq \text{im } \mathcal{X}_{i(\text{im } \phi, W)}$. Now it follows from Proposition 3.7 that $\text{im } \mathcal{X}_\phi$ contains all σ -nontrivial points in $\mathbf{PD}(W)$.

To finish the proof of (i), we must show that if $\mathcal{X}_\phi([b, d, i]) = [b', d', i']$, then $\langle b, d \rangle \triangleleft \sigma(\langle b', d' \rangle)$. First we suppose that $[b', d', i']$ is σ -nontrivial. In this case for $J := \langle b', d' \rangle \cap \text{Conv}(\sigma(\langle b', d' \rangle))$, the point $([\mathcal{B}(J), d', i'] \in \mathbf{PD}(W^\sigma)$ and $\mathcal{X}_{i(W^\sigma, W)}([\mathcal{B}(J), d', i']) = [b', d', i']$. Since σ is a translation map, we have $\langle \mathcal{B}(J), d' \rangle \triangleleft \sigma(\langle b', d' \rangle)$. Due to commutativity of the above diagram, $\mathcal{X}_{i(W^\sigma, \text{im } \phi)}([\mathcal{B}(J), d', i']) = \mathcal{X}_{s(V, \text{im } \phi)}([b, d, i])$. It follows from Proposition 2.32(iv) and Proposition 2.33(iv) that $b \leq \mathcal{B}(J)$. Therefore, $\langle b, d \rangle \triangleleft \sigma(\langle b', d' \rangle)$.

On the other hand, if $[b', d', j]$ is σ -trivial, then $\langle b', d' \rangle \cap \sigma(\langle b', d' \rangle) = \emptyset$ and every point in $\sigma(\langle b', d' \rangle)$ is larger than d' . By Proposition 2.36, $b < d'$ and we get $\langle b, d \rangle \triangleleft \sigma(\langle b', d' \rangle)$.

The proof of (ii) is based on similar ideas combined with the commutativity of the diagram

$$\begin{array}{ccccc} & & \mathbf{PD}(V) & & \\ & \swarrow \mathcal{X}_{s(V, V/\ker \phi_{\{V, \sigma\}})} & \downarrow & \searrow \mathcal{X}_\phi & \\ \mathbf{PD}(V/\ker \phi_{\{V, \sigma\}}) & & & & \mathbf{PD}(W) \\ & \swarrow \mathcal{X}_{s(\text{im } \phi, V/\ker \phi_{\{V, \sigma\}})} & \downarrow \mathcal{X}_{s(V, \text{im } \phi)} & \nwarrow \mathcal{X}_{i(\text{im } \phi, W)} & \\ & & \mathbf{PD}(\text{im } \phi) & & \end{array}$$

and is left to the reader. □

4. Algebraic Stability Theorem for Generalized Interleavings

In this section we provide a generalization of the Algebraic Stability Theorem of [1].

Theorem 4.1. *Let (V, φ_V) and (W, φ_W) be PFD persistence modules such that (V, W) are (τ, σ) -interleaved via the morphisms $\phi : V \rightarrow W(\tau)$ and $\psi : W \rightarrow V(\sigma)$. There exists a matching $\mathcal{X} : \mathbf{PD}(V) \rightarrow \mathbf{PD}(W)$ such that $\mathcal{X}([b, d, i]) = [b', d', i']$ implies*

$$(\sigma \circ \tau)^{-1}(\langle b, d \rangle) \blacktriangleleft \tau^{-1}(\langle b', d' \rangle) \blacktriangleleft \langle b, d \rangle \triangleleft \sigma \circ \tau \circ \tau^{-1}(\langle b', d' \rangle). \quad (4.2)$$

Moreover, if $[b, d, i] \in \mathbf{PD}(V)$ is unmatched, then it is $(\sigma \circ \tau)$ -trivial, and if $[b', d', i'] \in \mathbf{PD}(W)$ is unmatched, then either $\langle b', d' \rangle \cap \text{im } \tau = \emptyset$ or $\tau^{-1}(\langle b', d' \rangle)$ is $(\sigma \circ \tau)$ -trivial.

As in the case of the Induced Matching Theorem, the result of [1] follows from setting $\tau(t) = \sigma(t) = t + \delta$ for $\delta \geq 0$.

Proof. It follows from Proposition 2.18(i) that $\ker \phi$ and $\text{coker } \phi$ are $(\sigma \circ \tau)$ -trivial. By Theorem 3.2, the domain of $\mathcal{X}_\phi : \mathbf{PD}(V) \rightarrow \mathbf{PD}(W(\tau))$ contains all σ -nontrivial points in $\mathbf{PD}(V)$, and its image contains all σ -nontrivial points in $\mathbf{PD}(W(\tau))$. Now suppose that $\mathcal{X}_\phi([b, d, i]) = [x, y, j]$. Then

$$(\sigma \circ \tau)^{-1}(\langle b, d \rangle) \blacktriangleleft \langle x, y \rangle \blacktriangleleft \langle b, d \rangle \triangleleft \sigma \circ \tau \langle x, y \rangle, \quad (4.3)$$

where the first and the last relations follow from Theorem 3.2 and the middle one from Proposition 2.36.

To finish the proof, we build an appropriate matching $\mathcal{X}' : \mathbf{PD}(W(\tau)) \rightarrow \mathbf{PD}(W)$. It follows from the definition of $W(\tau)$ that there is a one-to-one correspondence between the points in $\mathbf{PD}(W(\tau))$ and the set $\{[b', d', i'] \in \mathbf{PD}(W) : \langle b', d' \rangle \cap \text{im } \tau \neq \emptyset\}$. This correspondence can be realized by a matching $\mathcal{X}' : \mathbf{PD}(W(\tau)) \rightarrow \mathbf{PD}(W)$ such that $\mathcal{X}'([x, y, j]) = [b', d', i']$ implies $\langle x, y \rangle = \tau^{-1}(\langle b', d' \rangle)$. Thus, the desired matching \mathcal{X} is obtained by the composition $\mathcal{X}' \circ \mathcal{X}_\phi$. Condition (4.2) follows from (4.3) and the fact that $\langle x, y \rangle = \tau^{-1}(\langle b', d' \rangle)$. \square

Statement (4.2) in Theorem 4.1 may seem impractical. However, it can be rewritten as a set of inequalities concerning the endpoints of the intervals, and if the translation maps τ and σ are bijective, then the inequalities can be simplified considerably.

Corollary 4.4. *Let (V, φ_V) and (W, φ_W) be persistence modules such that (V, W) are (τ, σ) -interleaved. Suppose that $\mathcal{X}([b, d, i]) = [b', d', i']$, where \mathcal{X} is the matching given by Theorem 4.1. Condition (4.2) implies that the inequalities:*

$$\begin{aligned} \tau^*(b') &\leq b, & \tau^*(d') &\leq d, \\ \sigma^*(b) &\leq \tau^\dagger \circ \tau^*(b'), & \sigma^*(d) &\leq \tau^\dagger \circ \tau^*(d') \end{aligned}$$

hold. Moreover, if the maps τ and σ are bijections, then the above inequalities reduce to

$$\begin{aligned} \tau^*(b') &\leq b, & \tau^*(d') &\leq d, \\ \sigma^*(b') &\leq b, & \sigma^*(d') &\leq d, \end{aligned}$$

and we have that if $[b, d, i] \in \mathbf{PD}(V)$ is unmatched, then it is $(\sigma \circ \tau)$ -trivial, and if $[b', d', i'] \in \mathbf{PD}(W)$ is unmatched, then it is $(\tau \circ \sigma)$ -trivial.

Proof. We start by showing that the first set of inequalities follows from Condition (4.2). The relation $\tau^{-1}(\langle b', d' \rangle) \triangleleft \langle b, d \rangle$ together with the monotonicity of τ implies that $\langle b', d' \rangle \triangleleft \tau(\langle b, d \rangle)$, and by definition we thus have

$$\langle b', d' \rangle \triangleleft \tau(\langle b, d \rangle) \text{ and } \langle b', d' \rangle \triangleleft \tau(\langle b, d \rangle).$$

The inequality $\tau^*(b') \leq b$ ($\tau^*(d') \leq d$) follows from definition of τ^* , and the first (second) relation stated above. Now we prove that $\sigma^*(b) \leq \tau^\dagger \circ \tau^*(b')$. Starting from the relation $\langle b, d \rangle \triangleleft \sigma \circ \tau \circ \tau^{-1}(\langle b', d' \rangle)$, we obtain that

$$\langle b, d \rangle \triangleleft \uparrow \{ \sigma \circ \tau \circ \tau^{-1}(\langle b', d' \rangle) \} = \langle (\sigma \circ \tau)^\dagger \circ \tau^*(b'), \infty \rangle.$$

Hence $b \leq (\sigma \circ \tau)^\dagger \circ \tau^*(b')$, and the desired inequality follows from the definition of σ^* . The last inequality follows again from the definitions of σ^* , τ^\dagger , and τ^* , and the relation $(\sigma \circ \tau)^{-1}(\langle b, d \rangle) \triangleleft \tau^{-1}(\langle b', d' \rangle)$ is guaranteed by Condition (4.2).

Now the moreover part. We suppose that τ and σ are invertible. Combining the invertibility of τ with the definitions of τ^\dagger and τ^* (i.e. Definition 2.22), it is routine to verify $\tau^\dagger \circ \tau^* = \text{id}$. Similarly, $\sigma^\dagger \circ \sigma^* = \text{id}$. The second set of inequalities follows. The statement that if $[b, d, i] \in \mathbf{PD}(V)$ is unmatched, then it is $(\sigma \circ \tau)$ -trivial carries over from Theorem 4.1. The statement that if $[b', d', i'] \in \mathbf{PD}(W)$ is unmatched, then it is $(\tau \circ \sigma)$ -trivial follows from Theorem 4.1 as well, observing that the $\langle b', d' \rangle \cap \text{im } \tau = \emptyset$ case cannot happen when τ is surjective, and that the condition that $\tau^{-1}(\langle b', d' \rangle)$ is $(\sigma \circ \tau)$ -trivial is equivalent to $\langle b', d' \rangle$ being $(\tau \circ \sigma)$ -trivial when τ is injective. \square

Most of the algorithms for computing persistence diagrams do not store information about decorations of the endpoints, and they produce undecorated persistence diagrams. To define undecorated persistence diagrams we will use the maps introduced in Definition 2.25.

Definition 4.5. Let V be a PDF persistence module and $\mathbf{PD}(V)$ its persistence diagram. The *undecorated persistence diagram* $\text{PD}(V)$ of V is a subset of $\mathbb{R}^2 \times \mathbb{N}_{>0}$ such that there exists a bijection $\mathcal{X}_V: \mathbf{PD}(V) \rightarrow \text{PD}(V)$ with the following property: if $\mathcal{X}_V([b, d, i]) = [b', d', i']$, then $b' = \pi(b)$ and $d' = \pi(d)$.

To formulate an analog of Theorem 4.1 for undecorated persistence diagrams we make use of the following functions.

Definition 4.6. Let $\tau: \mathbb{R} \rightarrow \mathbb{R}$ be a monotone function. We define monotone functions $\tau_L, \tau_L^\dagger, \tau_R, \tau_R^\dagger: \overline{\mathbb{R}} \rightarrow \overline{\mathbb{R}}$ as:

1. $\tau_L(\pm\infty) = \lim_{x \rightarrow \pm\infty} \tau(x)$ and $\tau_L(x) := \lim_{y \rightarrow x^-} \tau(y)$ for $x \in \mathbb{R}$.
2. $\tau_R(\pm\infty) = \lim_{x \rightarrow \pm\infty} \tau(x)$ and $\tau_R(x) := \lim_{y \rightarrow x^+} \tau(y)$ for $x \in \mathbb{R}$.
3. $\tau_L^\dagger(\pm\infty) := \pm\infty$ and $\tau_L^\dagger(x) = \inf \{y: \tau(y) > x\}$ for $x \in \mathbb{R}$.
4. $\tau_R^\dagger(\pm\infty) := \pm\infty$ and $\tau_R^\dagger(x) := \inf \{y: \tau(y) \geq x\}$ for $x \in \mathbb{R}$.

Proposition 4.7. Let $\tau: \mathbb{R} \rightarrow \mathbb{R}$ be a monotone function. We have that $\tau_L = (\tau^\dagger)_-$, $\tau_R = (\tau^\dagger)_+$, $\tau_L^\dagger = (\tau^*)_+$, and $\tau_R^\dagger = (\tau^*)_-$. Moreover, the pair (τ_R^\dagger, τ_R) is a Galois connection and the pair (τ_L, τ_L^\dagger) is a Galois connection.

Proof. To establish the first part of the result, one performs a routine verification (which we omit) that the functions defined in Definition 4.6 could have been alternatively defined using the concepts in Definition 2.22 and Definition 2.26 according to the formulas given. Now the moreover part. Since $(\tau_L, \tau_L^\dagger) = ((\tau^\downarrow)_-, (\tau^*)_+)$ and $(\tau_R^\dagger, \tau_R) = ((\tau^*)_-, (\tau^\uparrow)_+)$, the result follows from Proposition 2.24 and Proposition 2.28. \square

Proposition 4.8. *Let (V, φ_V) and (W, φ_W) be persistence modules such that (V, W) are (τ, σ) -interleaved. Then there exists a matching $\mathcal{X}: \text{PD}(V) \rightarrow \text{PD}(W)$ with the following properties. If $\mathcal{X}([b, d, i]) = [b', d', i']$, then*

$$\begin{aligned} \tau_R^\dagger(b') &\leq b \leq (\sigma_R \circ \tau_R \circ \tau_L^\dagger)(b'), \\ \tau_R^\dagger(d') &\leq d \leq (\sigma_R \circ \tau_R \circ \tau_L^\dagger)(d'), \\ (\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger)(b) &\leq b' \leq \tau_R(b), \\ (\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger)(d) &\leq d' \leq \tau_R(d). \end{aligned}$$

Moreover, if $[b, d, i] \in \text{PD}(V)$ is unmatched, then

$$d \leq (\sigma_R \circ \tau_R)(b),$$

and if $[b', d', i'] \in \text{PD}(W)$ is unmatched, then

$$d' \leq (\tau_R \circ \sigma_R \circ \tau_R \circ \tau_L^\dagger)(b').$$

Proof. We consider the matching $\mathcal{X}: \text{PD}(V) \rightarrow \text{PD}(W)$ defined by

$$\mathcal{X} := \mathcal{X}_W \circ \mathcal{X}' \circ \mathcal{X}_V^{-1},$$

where $\mathcal{X}_V, \mathcal{X}_W$ are bijections given by Definition 4.5 and \mathcal{X}' is the matching from Theorem 4.1. We start by proving the inequalities for the end points. We only need to show that $\tau_R^\dagger(b') \leq b$, $\tau_R^\dagger(d') \leq d$, $(\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger)(b) \leq b'$, and $(\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger)(d) \leq d'$ since by Proposition 4.7 the other inequalities can be recovered using Galois connections (e.g. $\tau_R^\dagger(b') \leq b$ if and only if $b' \leq \tau_R(b)$).

We only prove $\tau_R^\dagger(b') \leq b$ since the rest can be obtain by using similar arguments. Let $[c, e, j] = \mathcal{X}_V^{-1}([b, d, i])$ and $[c', e', j'] = \mathcal{X}_W^{-1}([b', d', i'])$. Then $\mathcal{X}'([c, e, j]) = [c', e', j']$. It follows from Corollary 4.4 that $\tau^*(c') \leq c$, and so $\pi \circ \tau^*(c') \leq \pi(c)$. Since $i^- \circ \pi(x) \leq \text{id}(x)$ for all $x \in \mathbb{E}$, we get $\pi \circ \tau^* \circ i^- \circ \pi(c') \leq \pi(c)$. By Proposition 4.7, we get $\tau_R^\dagger = \pi \circ \tau^* \circ i^-$, which implies that $\tau_R^\dagger(\pi(c')) \leq \pi(c)$. It follows from the definition of \mathcal{X}_V and \mathcal{X}_W that $b' = \pi(c')$ and $b = \pi(c)$. Combining this with the previous inequality yields $\tau_R^\dagger(b') \leq b$.

Now we assume that $[b, d, i] \in \text{PD}(V)$ is not in $\text{dom } \mathcal{X}$. Again, let $[c, e, j] = \mathcal{X}_V^{-1}([b, d, i])$. By Theorem 4.1, the point $[c, e, j]$ is $(\sigma \circ \tau)$ -trivial, i.e. $e \leq \sigma^\uparrow \circ \tau^\uparrow(c)$. Since $\text{id}(x) \leq i^+ \circ \pi(x)$ for all $x \in \mathbb{E}$, we have $\pi(e) \leq \pi \circ \sigma^\uparrow \circ (i^+ \circ \pi) \circ \tau^\uparrow \circ (i^+ \circ \pi)(c)$. Proposition 4.7 implies that $\pi(e) \leq \sigma_R \circ \tau_R(\pi(c))$. By using $\pi(c) = b$ and $\pi(c') = b'$, we obtain that $d \leq \sigma_R \circ \tau_R(b)$. The inequality for the points $[b', d', i'] \in \text{PD}(W)$ that are not in $\text{im } \mathcal{X}$ can be achieved by using similar methods as above and is left to the reader. \square

Corollary 4.9. *Let (V, φ_V) and (W, φ_W) be persistence modules such that (V, W) are (τ, σ) -interleaved. Suppose that the maps τ and σ are bijections. If $\mathcal{X}: \text{PD}(V) \rightarrow \text{PD}(W)$ is the matching given by Proposition 4.8 and $\mathcal{X}([b, d, i]) = [b', d', i']$, then*

$$\sigma^{-1}(b) \leq b' \leq \tau(b) \quad \text{and} \quad \sigma^{-1}(d) \leq d' \leq \tau(d).$$

If $[b, d, i] \in \text{PD}(V)$ is unmatched, then $d \leq (\sigma \circ \tau)(b)$, and if $[b', d', i'] \in \text{PD}(W)$ is not in $\text{im } \mathcal{X}$, then $d' \leq (\tau \circ \sigma)(b')$.

Proof. If τ and σ are invertible, then $\tau_L = \tau_R = \tau$, $\sigma_L = \sigma_R = \sigma$, $\tau_L^\dagger = \tau_R^\dagger = \tau^{-1}$, and $\sigma_L^\dagger = \sigma_R^\dagger = \sigma^{-1}$. The proof is obtained by evaluating expressions in Proposition 4.8. \square

5. Applications

We illustrate the use of results obtained in Section 4 through a series of applications. Our first example examines the relationship between \mathbb{Z} -indexed and \mathbb{R} -indexed persistence modules. The second example focuses on obtaining bounds on errors that arise from computational limitations to obtaining the true persistence diagrams for large point clouds. We conclude with a table indicating how to apply Theorem 4.1 for a variety of approximations that are commonly used.

5.1. Discretizing a persistence module

The \mathbb{R} -indexed persistence module V derived by considering the sublevel set filtration of a function $f: X \rightarrow \mathbb{R}$ provides a characterization of the topography of f . However, in practice only a finite number of calculations can be performed. A simple idealization is to assume that calculations are performed only at integer values of f . This leads to the following definition.

Definition 5.1. The \mathbb{Z} -discretized persistence module $V^{\mathbb{Z}}$ is defined as follows. Set

$$V_t^{\mathbb{Z}} := V_{\lfloor t \rfloor} \quad \text{and} \quad \varphi_{V^{\mathbb{Z}}}(s, t) := \varphi_V(\lfloor s \rfloor, \lfloor t \rfloor),$$

where $\lfloor \cdot \rfloor$ is the floor function.

The following proposition provides an answer to the following question: given the \mathbb{Z} -discretized persistence module $V^{\mathbb{Z}}$, what are the constraints on the persistence diagram associated to the persistence module V ?

Proposition 5.2. *If V is an \mathbb{R} -indexed PFD persistence module and $V^{\mathbb{Z}}$ is the associated \mathbb{Z} -discretized PFD persistence module, then the following are true:*

- (i) $(V, V^{\mathbb{Z}})$ are (τ, σ) -interleaved, where $\tau(t) = t$ and $\sigma(t) = \lceil t \rceil$; and
- (ii) there exists a matching $\mathcal{X}: \text{PD}(V^{\mathbb{Z}}) \leftrightarrow \text{PD}(V)$ such that if $\mathcal{X}([b, d, i]) = [b', d', i']$, then

$$b - 1 \leq b' \leq b \quad \text{and} \quad d - 1 \leq d' \leq d.$$

Additionally, any unmatched points $[b, d, i] \in \text{PD}(V^{\mathbb{Z}})$ satisfy $d = b + 1$, and unmatched points in $[b', d', i'] \in \text{PD}(V)$ satisfy $d' \leq \lfloor b' + 1 \rfloor$.

Proof. (i) Define persistence module morphisms $\phi : V^{\mathbb{Z}} \rightarrow V(\tau) = V$ by $\phi_t := \varphi_V(\lfloor t \rfloor, t)$ and $\psi : V \rightarrow V^{\mathbb{Z}}(\sigma)$ by $\psi_t := \varphi_V(t, \lceil t \rceil)$. Observe that

$$\begin{aligned} \psi(\tau)_t \circ \phi_t &= \psi_{\tau(t)} \circ \varphi_V(\lfloor t \rfloor, t) \\ &= \varphi_V(t, \lceil t \rceil) \circ \varphi_V(\lfloor t \rfloor, t) \\ &= \varphi_V(\lfloor t \rfloor, \lceil t \rceil) \\ &= \varphi_{V^{\mathbb{Z}}}(t, \sigma \circ \tau(t)). \end{aligned}$$

It is left to the reader to check that $\varphi(\sigma)_t \circ \psi_t = \varphi_V[t, (\tau \circ \sigma)(t)]$, and therefore that $(V, V^{\mathbb{Z}})$ are (τ, σ) -interleaved.

(ii) Observe that if $[b, d, i] \in \text{PD}(V^{\mathbb{Z}})$, then $b, d \in \mathbb{Z}$. The proof now follows from Proposition 4.8 and the fact that for $t \in \mathbb{R}$

$$\tau_R(t) = \tau_L(t) = \tau_R^\dagger(t) = \tau_L^\dagger(t) = t$$

and

$$\sigma_R^\dagger(t) = \lceil t - 1 \rceil.$$

□

See Figure 1 for an illustration of an estimate of $\text{PD}(V)$ from $\text{PD}(V^{\mathbb{Z}})$.

5.2. Computing persistence diagrams for large point clouds

We now turn to the question of computing persistence diagrams for large point clouds. For point clouds in arbitrary metric spaces, a standard approach makes use of a filtration of the associated Vietoris-Rips complex, which we define next.

Definition 5.3. Let (X, d) be a finite metric space with metric d . The *Vietoris-Rips complex of X at scale t* , denoted by $\mathcal{R}(X, t)$, is the simplicial complex with vertices given by X and containing the N -simplex $[x_{i_0}, \dots, x_{i_N}]$ if and only if $d(x_{i_j}, x_{i_k}) \leq 2t$ for all $j, k = 0, \dots, N$.

The collection $\{\mathcal{R}(X, t)\}_{t \in \mathbb{R}}$ is called the *Vietoris-Rips filtration associated to X* .

Definition 5.4. Let (X, d) be a finite metric space. Fix a field \mathbf{k} . The *persistence module induced by the Vietoris-Rips filtration associated to X* , denoted by $M^{\mathcal{R}}(X)$, is defined via simplicial homology as follows:

$$M^{\mathcal{R}}(X)_t := H_*(\mathcal{R}(X, t), \mathbf{k}), \quad t \in \mathbb{R}$$

and the transition maps $\varphi_{M^{\mathcal{R}}(X)}(t, s)$ are the associated linear maps on homology induced by the inclusion maps $j_{X;t,s} : \mathcal{R}(X, t) \rightarrow \mathcal{R}(X, s)$.

We remark that given a finite metric space X , the induced persistence module $M^{\mathcal{R}}(X)$ is a PFD persistence module.

Observe that for large X , the computational cost of determining $H_*(\mathcal{R}(X, t), \mathbf{k})$ grows rapidly as a function of t . If $Y \subset X$, then one expects that it is cheaper to compute $H_*(\mathcal{R}(Y, t), \mathbf{k})$. The goal of this section is twofold: first, to quantify the difference between $M^{\mathcal{R}}(X)$ and $M^{\mathcal{R}}(Y)$; and second, to suggest an iterative procedure, motivated by [12], for obtaining reasonable approximations of $M^{\mathcal{R}}(X)$.




Figure 2: (left) Schematic diagram illustrating the quality of the matching from Proposition 5.6. The persistence diagram $\text{PD}(M^{\mathcal{R}}(Y))$ is shown. The dark gray region indicates the possible locations of the unmatched points for the persistence diagrams $\text{PD}(M^{\mathcal{R}}(Y))$ and $\text{PD}(M^{\mathcal{R}}(X))$. The light gray region indicates the possible location of the point of $\text{PD}(M^{\mathcal{R}}(X))$ that is matched to the point of $\text{PD}(M^{\mathcal{R}}(Y))$ that is shown. (right) Schematic diagram illustrating the quality of the matching from Corollary 5.15. The persistence diagram $\text{PD}(M^{\mathcal{R}}(U))$ is shown. Persistence points for the persistence diagrams $\text{PD}(M^{\mathcal{R}}(U))$ and $\text{PD}(M^{\mathcal{R}}(X))$ agree in the region $[0, t_0] \times [0, t_0]$. Unmatched points for the persistence diagrams $\text{PD}(M^{\mathcal{R}}(U))$ and $\text{PD}(M^{\mathcal{R}}(X))$ will lie in the dark gray region. The light gray region indicates the possible location of the point of $\text{PD}(M^{\mathcal{R}}(X))$ that is matched to the point of $\text{PD}(M^{\mathcal{R}}(U))$ that is shown.

5.2.1. Subsampling a large point cloud

Definition 5.5. Let (X, d) be a finite metric space. A subset $Y \subset X$ is a δ -approximation of X if for every $x \in X$ there exists a $y \in Y$ such that $d(x, y) \leq \delta$.

Proposition 5.6. If (X, d) is a finite metric space and Y is a δ -approximation of X , then the following approximations hold.

- (i) The persistence modules $(M^{\mathcal{R}}(Y), M^{\mathcal{R}}(X))$ are (τ, σ) -interleaved, where $\tau(t) = t$ and $\sigma(t) = t + \delta$.
- (ii) There exists a matching $\mathcal{X} : \text{PD}(M^{\mathcal{R}}(Y)) \leftrightarrow \text{PD}(M^{\mathcal{R}}(X))$ such that if $\text{PD}_{\mathcal{X}}(b, d, i) = (b', d', i')$, then $b - \delta \leq b' \leq b$ and $d - \delta \leq d' \leq d$. Moreover, all unmatched points in $\text{PD}(M^{\mathcal{R}}(Y))$ and $\text{PD}(M^{\mathcal{R}}(X))$ are at most δ above the diagonal.

Figure 2(left) provides an illustration of Proposition 5.6(ii). Observe that since τ and σ are invertible, Proposition 5.6(ii) follows from Proposition 5.6(i) and Corollary 4.9. The proof of Proposition 5.6(i) occupies the remainder of this section. We begin with some preliminary arguments.

Lemma 5.7. Let $Y, Y' \subseteq X$ and $\delta \geq 0$. If $\gamma : Y \rightarrow Y'$ satisfies $d(x, \gamma(x)) \leq \delta$ for all $x \in Y$, then $\tilde{\gamma}_t : \mathcal{R}(Y, t) \rightarrow \mathcal{R}(Y', t + \delta)$, defined by

$$\tilde{\gamma}_t[x_0, \dots, x_k] = [\gamma(x_0), \dots, \gamma(x_k)] \quad \text{for any simplex } [x_0, \dots, x_k] \in \mathcal{R}(Y, t),$$

is a simplicial map.

Proof. To prove that $\tilde{\gamma}$ is a simplicial map, we need to show that for every k -simplex $[x_0, \dots, x_k] \in \mathcal{R}(Y, t)$, the k -simplex $[\gamma(x_0), \dots, \gamma(x_k)]$ is a simplex in $\mathcal{R}(Y', t + \delta)$. Since the simplices in a Vietoris-Rips complex are fully determined by its 1-skeleton, we only need to show that the 1-skeleton of $\mathcal{R}(Y, t)$ is mapped to the 1-skeleton of $\mathcal{R}(Y', t + \delta)$. Recall that $[x, y]$ is an edge in $\mathcal{R}(Y, t)$ if and only if $d(x, y) \leq 2t$. Thus, we have

$$\begin{aligned} d(\gamma(x), \gamma(y)) &\leq d(\gamma(x), x) + d(x, y) + d(y, \gamma(y)) \\ &\leq \delta + 2t + \delta \\ &= 2(t + \delta), \end{aligned}$$

and so $[\gamma(x), \gamma(y)]$ is either a 1-simplex or a 0-simplex in $\mathcal{R}(Y', t + \delta)$. \square

Let $Y \subset X$ be a δ -approximation and let $\iota_t: \mathcal{R}(Y, t) \rightarrow \mathcal{R}(X, t)$ denote the inclusion map. Set

$$\phi_t := \iota_{t*}: H_*(\mathcal{R}(Y, t)) \rightarrow H_*(\mathcal{R}(X, t)).$$

Since Y is a δ -approximation, there exists $\gamma: X \rightarrow Y$ such that $d(x, \gamma(x)) \leq \delta$ for all $x \in X$, and $\gamma(y) = y$ for all $y \in Y$. By Lemma 5.7, $\tilde{\gamma}_t: \mathcal{R}(X, t) \rightarrow \mathcal{R}(Y, t + \delta)$ is a simplicial map and hence we can define

$$\psi_t := \tilde{\gamma}_{t*}: H_*(\mathcal{R}(X, t)) \rightarrow H_*(\mathcal{R}(Y, t + \delta)).$$

Our goal is to show that $\phi: M^{\mathcal{R}}(Y) \rightarrow M^{\mathcal{R}}(X)$ and $\psi: M^{\mathcal{R}}(X) \rightarrow M^{\mathcal{R}}(Y)(\delta)$ are persistence module morphisms that guarantee that the persistence modules $(M^{\mathcal{R}}(Y), M^{\mathcal{R}}(X))$ are (τ, σ) -interleaved, and therefore, provide a proof of Proposition 5.6(i).

Observe that $\gamma_t \circ \iota_t = j_{Y; t, t + \delta}$ and hence

$$\psi(\tau)_t \circ \phi_t = \psi_t \circ \phi_t = \varphi_{M^{\mathcal{R}}(Y)}[t, t + \delta] = \varphi_{M^{\mathcal{R}}(Y)}[t, (\sigma \circ \tau)t]. \quad (5.8)$$

The challenge is to show that the middle equality holds for

$$\phi(\sigma)_t \circ \psi_t = \phi_{t + \delta} \circ \psi_t = \varphi_{M^{\mathcal{R}}(X)}[t, t + \delta] = \varphi_{M^{\mathcal{R}}(X)}[t, (\tau \circ \sigma)t]. \quad (5.9)$$

For purposes of the next section, we prove a more general result than necessary.

Lemma 5.10. *Consider $Y, Y' \subseteq X$ and $\delta \geq 0$. Let $\iota'_t: \mathcal{R}(Y', t) \rightarrow \mathcal{R}(X, t + \delta)$ and $\iota_t: \mathcal{R}(Y, t) \rightarrow \mathcal{R}(X, t)$ be the simplicial maps induced by inclusion. If $\gamma: Y' \rightarrow Y$ satisfies $d(y, \gamma(y)) \leq \delta$ for all $y \in Y'$ and $\tilde{\gamma}_t: \mathcal{R}(Y', t) \rightarrow \mathcal{R}(Y, t + \delta)$ is the simplicial map as defined in Lemma 5.7, then $\iota_{(t + \delta)*} \circ \tilde{\gamma}_t$ and ι'_t are homotopic and hence*

$$\iota_{(t + \delta)*} \circ \tilde{\gamma}_t = \iota'_{t*}.$$

Proof. To prove this we make use of the theory of simplicial sets [19, 15] and begin with the remark that by [19, Lemma 8.3.13, Theorem 8.3.8] it is sufficient to prove that $\iota_{(t + \delta)*} \circ \tilde{\gamma}_t$ and ι'_t are homotopic.

Given a simplicial complex \mathcal{K} let $\bar{\mathcal{K}}$ denote the associated simplicial set. To establish notation let $\bar{\mathcal{K}}_k$ denote the k -dimensional simplices in \mathcal{K} and let $d_i: \bar{\mathcal{K}}_k \rightarrow \bar{\mathcal{K}}_{k-1}$ and $s_i: \bar{\mathcal{K}}_k \rightarrow \bar{\mathcal{K}}_{k+1}$ be the delete and duplicate i -th vertex operations defined by

$$d_i[v_0, \dots, v_k] := [v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k]$$

and

$$s_i[v_0, \dots, v_k] := [v_0, \dots, v_i, v_i, \dots, v_k].$$

We claim that the functions $h_i: \bar{\mathcal{R}}(Y', t)_k \rightarrow \bar{\mathcal{R}}(X, t + \delta)_{k+1}$, $i = 0, \dots, k$, defined by

$$h_i([x_0, \dots, x_k]) = [x_0, \dots, x_i, \gamma(x_i), \dots, \gamma(x_k)]$$

provide a simplicial homotopy between $\iota_{(t+\delta)} \circ \tilde{\gamma}_t$ and ι'_t . Recall that to justify this claim, it is sufficient to verify the following equalities:

$$d_0 h_0 = \iota_{(t+\delta)} \circ \tilde{\gamma}_t \quad \text{and} \quad d_{k+1} h_k = \iota'_t$$

$$d_i h_j = \begin{cases} h_{j-1} d_i & \text{if } i < j \\ d_i h_{i-1} & \text{if } i = j \neq 0 \\ h_j d_{i-1} & \text{if } i > j + 1 \end{cases} \quad \text{and} \quad s_i h_j = \begin{cases} h_{j+1} s_i & \text{if } i \leq j \\ h_j s_{i-1} & \text{if } i > j. \end{cases}$$

We leave these calculations to the reader. \square

Proof of Proposition 5.6(i). As indicated above, the proof of Proposition 5.6(i) follows from (5.8), which has already been justified, and (5.9), which follows from Lemma 5.10 under the assumption that $Y' = X$. \square

5.2.2. Stitching persistence modules

Proposition 5.6 demonstrates that given a finite metric space (X, d) , bounds on the persistence diagram $\text{PD}(M^{\mathcal{R}}(X))$ can be determined from $\text{PD}(M^{\mathcal{R}}(Y))$ under the assumption that Y is a δ -approximation of X . The motivation for using Y is that for large t it may not be feasible to compute $H_*(\mathcal{R}(X, t); \mathbf{k})$. However, for small t , the size of the complex $\mathcal{R}(X, t)$ is on the order of the size of X . Furthermore, if $t < \delta$, then $\mathcal{R}(Y, t)$ will fail to capture the fine geometric structure of X , and hence, as indicated by Figure 2(left), points in this region of $M^{\mathcal{R}}(X)$ will be unmatched with respect to $M^{\mathcal{R}}(Y)$. For this reason, we would like to construct a persistence module determined by $M^{\mathcal{R}}(X)$ in the range $t < \delta$ and $M^{\mathcal{R}}(Y)$ in the range $t \geq \delta$. This will provide finer information, as illustrated in Figure 2(right). Observe that this suggests the need to be able to stitch together persistence modules and motivates the following definition.

Definition 5.11. Let (V, φ_V) , $(V', \varphi_{V'})$ and (W, φ_W) be persistence modules such that (V, W) are (τ, σ) -interleaved via the morphisms $\phi: V \rightarrow W(\tau)$ and $\psi: W \rightarrow V(\sigma)$, and that (V', W) are (τ', σ') -interleaved via $\phi': V' \rightarrow W(\tau')$ and $\psi': W \rightarrow V'(\sigma')$. A point $t_0 \in \mathbb{R}$ is a (V, V') *stitch point* if

$$t_0 \leq (\sigma' \circ \tau)(t_0).$$

The (V, V') *persistence module stitched through W at stitch point t_0* consists of vector spaces

$$U_t = \begin{cases} V_t & \text{if } t < t_0 \\ V_{t_0} & \text{if } t_0 \leq t < (\sigma' \circ \tau)(t_0) \\ V'_t & \text{if } (\sigma' \circ \tau)(t_0) \leq t \end{cases}$$

and linear maps

$$\varphi_U(s, t) = \begin{cases} \varphi_V(\min\{s, t_0\}, \min\{t, t_0\}) & \text{if } s \leq t < (\sigma' \circ \tau)(t_0) \\ \varphi_{V'}(t, (\sigma' \circ \tau)(t_0)) \circ \psi'_{\tau(t_0)} \circ \phi_{t_0} \circ \varphi_V(\min\{s, t_0\}, t_0) & \text{if } s < (\sigma' \circ \tau)(t_0) < t \\ \varphi_{V'}(s, t) & \text{if } (\sigma' \circ \tau)(t_0) \leq t. \end{cases}$$

We denote this persistence module by $U = U(W; V, t_0, V')$ and for the sake of simplicity refer to it as the *stitched persistence module*.

The following diagram (with unlabeled arrows assumed to be the appropriate transition maps) shows the idea behind the vector spaces of U and the transition maps $\varphi_U(s, t)$ for $s < (\sigma' \circ \tau)(t_0) \leq t$.

$$\begin{array}{ccccccc} V_s & \longrightarrow & V_{t_0} & \xrightarrow{\psi'_{\tau(t_0)} \circ \phi_{t_0}} & V'_{(\sigma' \circ \tau)(t_0)} & \longrightarrow & V'_t \\ & & \searrow \phi_{t_0} & & \nearrow \psi'_{\tau(t_0)} & & \\ W_s & \longrightarrow & W_{\tau(t_0)} & \longrightarrow & & \longrightarrow & W_t \end{array}$$

Proposition 5.12. *Let (V, φ_V) , $(V', \varphi_{V'})$, and (W, φ_W) be PFD persistence modules such that (V, W) are (τ, σ) -interleaved via the morphisms $\phi : V \rightarrow W(\tau)$ and $\psi : W \rightarrow V(\sigma)$, and that (V', W) are (τ', σ') -interleaved via $\phi' : V' \rightarrow W(\tau')$ and $\psi' : W \rightarrow V'(\sigma')$. Assume that t_0 is a stitch point. If $U = U(W; V, t_0, V')$ is the stitched persistence module, then the following statements are true:*

- (i) U is a PFD persistence module;
- (ii) (U, W) are (η, ρ) -interleaved where

$$\eta(t) = \begin{cases} \tau(t) & \text{if } t \leq t_0 \\ \tau(t_0) & \text{if } t_0 < t < (\sigma' \circ \tau)(t_0) \\ \tau'(t) & \text{if } (\sigma' \circ \tau)(t_0) \leq t; \end{cases}$$

$$\rho(t) = \begin{cases} \sigma(t) & \text{if } t < \sigma^{-1}(t_0) \\ (\sigma' \circ \tau)(t_0) & \text{if } \sigma^{-1}(t_0) \leq t < \tau(t_0) \\ \sigma'(t) & \text{if } \tau(t_0) \leq t. \end{cases}$$

Proof. (i) Since V and V' are PFD persistence modules, U_t is finite dimensional for each $t \in \mathbb{R}$. It is left to the reader to check that $\varphi_U(t, t) = \text{id}_{U_t}$ and $\varphi_U(s, t) \circ \varphi_U(r, s) = \varphi_U(r, t)$ for every $r \leq s \leq t$ in \mathbb{R} .

(ii) To show that (U, W) are (η, ρ) -interleaved we will show that the morphisms $\bar{\phi} : U \rightarrow W(\eta)$ and $\bar{\psi} : W \rightarrow U(\rho)$, where

$$\bar{\phi}_t = \begin{cases} \phi_t & \text{if } t \leq t_0 \\ \phi_{t_0} & \text{if } t_0 < t < (\sigma' \circ \tau)(t_0) \\ \phi'_t & \text{if } (\sigma' \circ \tau)(t_0) \leq t, \end{cases}$$

$$\bar{\psi}_t = \begin{cases} \psi_t & \text{if } t < \sigma^{-1}(t_0) \\ \psi'_{\tau(t_0)} \circ \varphi_W(t, \tau(t_0)) & \text{if } \sigma^{-1}(t_0) \leq t < \tau(t_0) \\ \psi'_t & \text{if } \tau(t_0) \leq t, \end{cases}$$

give the desired interleaving of U and W . To show that $\bar{\phi}$ and $\bar{\psi}$ are persistence module morphisms, we first note that the monotone functions η and ρ line up with the indices of the shifts of $\bar{\phi}$ and $\bar{\psi}$ by inspection. For what follows, let $s \leq t \in \mathbb{R}$.

We will now show that $\bar{\phi}: U \rightarrow W(\eta)$ is a persistence module morphism. If $s \leq t < (\sigma' \circ \tau)(t_0)$, then $\bar{\phi}_s = \phi_{\min\{s, t_0\}}$ and $\bar{\phi}_t = \phi_{\min\{t, t_0\}}$, and if $(\sigma' \circ \tau)(t_0) \leq s \leq t$, then $\bar{\phi}_s = \phi'_s$ and $\bar{\phi}_t = \phi'_t$, and so these cases hold. Now suppose that $s < (\sigma' \circ \tau)(t_0) \leq t$. The diagram

$$\begin{array}{ccccccc}
V_s & \longrightarrow & V_{t_0} & \overset{\psi'_{\tau(t_0)} \circ \phi_{t_0}}{\dashrightarrow} & V'_{(\sigma' \circ \tau)(t_0)} & \longrightarrow & V'_t \\
& \searrow \phi_s & \searrow \phi_{t_0} & \nearrow \psi'_{\tau(t_0)} & \searrow \phi'_{(\sigma' \circ \tau)(t_0)} & \searrow \phi'_t & \\
& & W_{\tau(s)} & \longrightarrow & W_{\tau(t_0)} & \longrightarrow & W_{(\tau' \circ \sigma' \circ \tau)(t_0)} & \longrightarrow & W_{\tau'(t)}
\end{array}$$

where unlabeled arrows are transition maps, commutes since both ϕ and ϕ' are persistence module morphisms and (V', W) are (τ', σ') -interleaved. It follows that $\bar{\phi}_t \circ \varphi_U(s, t) = \varphi_{W(\eta)}(s, t) \circ \bar{\phi}_s$, and thus $\bar{\phi}: U \rightarrow W(\eta)$ is a persistence module morphism.

Now we will show that $\bar{\psi}: W \rightarrow U(\rho)$ is a persistence module morphism. If $s \leq t < \sigma^{-1}(t_0)$, then $\bar{\psi}_s = \psi_s$ and $\bar{\psi}_t = \psi_t$, and if $\tau(t_0) \leq s \leq t$, then $\bar{\psi}_s = \psi'_s$ and $\bar{\psi}_t = \psi'_t$, so these cases are clear. Suppose that $s < \sigma^{-1}(t_0) \leq t \leq \tau(t_0)$. This choice of s and t yield $\bar{\psi}_s = \psi_s$ and $\bar{\psi}_t = \psi'_{\tau(t_0)} \circ \varphi_W(t, \tau(t_0))$ by definition. Since $s < \sigma^{-1}(t_0)$, it follows that $s \leq (\tau \circ \sigma)(s) \leq \tau(t_0)$, and so the following diagram

$$\begin{array}{ccccc}
& & V_{\sigma(s)} & \longrightarrow & V_{t_0} & \overset{\psi'_{\tau(t_0)} \circ \phi_{t_0}}{\dashrightarrow} & V'_{(\sigma' \circ \tau)(t_0)} \\
& \nearrow \psi_s & & \searrow \phi_s & \searrow \phi_{t_0} & \nearrow \psi'_{\tau(t_0)} & \\
W_s & \longrightarrow & W_{(\tau \circ \sigma)(s)} & \longrightarrow & W_{\tau(t_0)} & \longrightarrow & \\
& \searrow & & & & & \\
& & & & W_t & \xrightarrow{\bar{\psi}_t} &
\end{array}$$

where unlabeled arrows are transition maps, commutes. Hence, $\bar{\psi}_t \circ \varphi_W(s, t) = \varphi_{U(\rho)}(s, t) \circ \bar{\psi}_s$ in this case as well. For $\sigma^{-1}(t_0) \leq s \leq t \leq \tau(t_0)$, then

$$\begin{aligned}
\bar{\psi}_t \circ \varphi_W(s, t) &= \psi'_{\tau(t_0)} \circ \varphi_W(t, \tau(t_0)) \circ \varphi_W(s, t) \\
&= \varphi_U((\sigma' \circ \tau)(t_0), (\sigma' \circ \tau)(t_0)) \circ \psi'_{\tau(t_0)} \circ \varphi_W(s, \tau(t_0)) \\
&= \varphi_{U(\rho)}(s, t) \circ \bar{\psi}_s,
\end{aligned}$$

and so this case also holds. If $s < \sigma^{-1}(t_0) \leq \tau(t_0) < t$, the commutativity of the diagram

$$\begin{array}{ccccccc}
& & V_{\sigma(s)} & \longrightarrow & V_{t_0} & \overset{\psi'_{\tau(t_0)} \circ \phi_{t_0}}{\dashrightarrow} & V'_{(\sigma' \circ \tau)(t_0)} & \longrightarrow & V'_{\sigma'(t)} \\
& \nearrow \psi_s & & \searrow \phi_s & \searrow \phi_{t_0} & & \nearrow \psi'_{\tau(t_0)} & & \nearrow \psi'_t \\
W_s & \longrightarrow & W_{(\tau \circ \sigma)(s)} & \longrightarrow & W_{\tau(t_0)} & \longrightarrow & W_t & &
\end{array}$$

ensures that $\bar{\psi}_t \circ \varphi_W(s, t) = \varphi_{U(\rho)}(s, t) \circ \bar{\psi}_s$. Finally, it remains to check the case when $\sigma^{-1}(t_0) \leq s \leq \tau(t_0) < t$. This follows from extending the case when $\sigma^{-1}(t_0) \leq s \leq t \leq \tau(t_0)$ by the right-most commutative square in the previous diagram. Thus, we have shown that $\bar{\phi}: U \rightarrow W(\eta)$ and $\bar{\psi}: U \rightarrow W(\eta)$ are persistence module morphisms.

We must now show that $\bar{\phi}$ and $\bar{\psi}$ satisfy the commutativity constraints of Definition 2.13. That is, for every $t \in \mathbb{R}$, we must show that

$$\bar{\psi}(\eta)_t \circ \bar{\phi}_t = \varphi_U(t, (\rho \circ \eta)(t)) \quad (5.13)$$

and

$$\bar{\phi}(\rho)_t \circ \bar{\psi}_t = \varphi_W(t, (\rho \circ \eta)(t)) \quad (5.14)$$

for all $t \in \mathbb{R}$. First we prove (5.13) for all $t \in \mathbb{R}$. When $\tau(t) < \sigma^{-1}(t_0)$, then (5.13) follows from the fact that V and W are (τ, σ) -interleaved. Similarly, when $(\sigma' \circ \tau)(t_0) \leq t$, (5.13) holds since V' and W are (τ', σ') -interleaved. If $t_0 \leq t < (\sigma' \circ \tau)(t_0)$, then (5.13) holds by the definition of (U, φ_U) . Now suppose that $t \leq t_0$ such that $\sigma^{-1}(t_0) \leq \tau(t)$. Then (5.13) holds from the commutativity of the following diagram.

$$\begin{array}{ccccc}
V_t & \longrightarrow & V_{t_0} & \overset{\psi'_{\tau(t_0)} \circ \phi_{t_0}}{\dashrightarrow} & V'_{(\sigma' \circ \tau)(t_0)} \\
& \searrow \phi_t & & \searrow \phi_{t_0} & \nearrow \psi'_{\tau(t_0)} \\
& & W_{\tau(t)} & \longrightarrow & W_{\tau(t_0)} \\
& & & \searrow & \nearrow \\
& & & & \bar{\psi}_{\tau(t)}
\end{array}$$

Thus we have shown that (5.13) holds for every $t \in \mathbb{R}$.

We now verify that (5.14) holds for every $t \in \mathbb{R}$. When $t < \sigma^{-1}(t_0)$, then (5.14) follows since V and W are (τ, σ) -interleaved. Similarly, if $\tau(t_0) \leq t$, then (5.14) follows since V' and W are (τ', σ') -interleaved. For $\sigma^{-1}(t_0) \leq t < \tau(t_0)$, then (5.14) follows by the commutativity of the following diagram.

$$\begin{array}{ccccc}
& & V'_{(\sigma' \circ \tau)(t_0)} & & \\
& \nearrow \bar{\psi}_t & & \searrow \bar{\phi}_{(\sigma' \circ \tau)(t_0)} & \\
W_t & \longrightarrow & W_{\tau(t_0)} & \longrightarrow & W_{(\tau' \circ \sigma' \circ \tau)(t_0)} \\
& & \nearrow \psi'_{\tau(t_0)} & &
\end{array}$$




Figure 3: Functions η and ρ that provide an interleaving between the stitched persistence module $U = (M^{\mathcal{R}}(X); M^{\mathcal{R}}(X), t_0, M^{\mathcal{R}}(Y))$ and $(M^{\mathcal{R}}(X))$ where Y is a δ -approximation of a finite metric space X .

Thus, we have also shown that (5.14) holds for every $t \in \mathbb{R}$. It follows that (η, ρ) is a translation pair. \square

We conclude this section with a simple, but hopefully illustrative, application of Proposition 5.12.

Corollary 5.15. *Let (X, d) be a finite metric space and let $Y \subset X$ be a δ -approximation of X with $\delta > 0$. Let $U = U(M^{\mathcal{R}}(X); M^{\mathcal{R}}(X), t_0, M^{\mathcal{R}}(Y))$ be the stitched persistence module and $\text{PD}(U)$ its persistence diagram. If $[b, d, i] \in \text{PD}(U)$, then neither b or d is in the interval $(t_0, t_0 + \delta)$. Moreover, there exists a matching $\mathcal{X} : \text{PD}(U) \rightarrow \text{PD}(M^{\mathcal{R}}(X))$ such that if $\mathcal{X}([b, d, i]) = [b', d', i']$ and*

$$\begin{array}{llll}
 \text{if } & b < d \leq t_0 & \text{then} & b' = b & \text{and} & d' = d; \\
 \text{if } & b \leq t_0 < t_0 + \delta \leq d & \text{then} & b' = b & \text{and} & \max(t_0, d - \delta) \leq d' \leq d; \\
 \text{if } & b = t_0 + \delta < d & \text{then} & t_0 \leq b' \leq t_0 + \delta & \text{and} & \max(t_0, d - \delta) \leq d' \leq d; \\
 \text{if } & t_0 + \delta < b < d & \text{then} & b - \delta \leq b' \leq b & \text{and} & d - \delta \leq d' \leq d.
 \end{array}$$

All unmatched points $[b, d, i] \in \text{PD}(U)$ satisfy

$$t_0 + \delta \leq b \quad \text{and} \quad d \leq b + \delta,$$

and all unmatched points $[b', d', i'] \in \text{PD}(M^{\mathcal{R}}(X))$ satisfy

$$t_0 < b' < d' \leq b' + \delta.$$

Proof. It follows from the definition of the Vietoris-Rips filtration that every interval in $\mathbf{PD}(M^{\mathcal{R}}(X))$ and $\mathbf{PD}(M^{\mathcal{R}}(Y))$ has a closed left-hand endpoint and an open right-hand endpoint, and so $b < d$ and $b' < d'$. Moreover, by the definition of U , we cannot have $t_0 < b < t_0 + \delta$ or $t_0 < d < t_0 + \delta$. Observe that the persistence modules $(M^{\mathcal{R}}(X), M^{\mathcal{R}}(X))$ are $(\tau(t) = t, \sigma(t) = t)$ -interleaved and that $(M^{\mathcal{R}}(Y), M^{\mathcal{R}}(X))$ are $(\tau'(t) = t, \sigma'(t) = t + \delta)$ -interleaved. The result follows by applying Proposition 5.12 and Proposition 4.8, with the additional observation that the identity map yielding $U_{t_0} = M^{\mathcal{R}}(X)_{t_0}$ and the definition of U forces that every point $[t_0, d, i] \in \mathbf{PD}(U)$ is matched to some point $[t_0, d', i'] \in \mathbf{PD}(M^{\mathcal{R}}(X))$ and vice versa. For the reader's benefit, we indicate the forms of η, ρ, η^\dagger and ρ^\dagger in Figure 3. \square

5.2.3. Iterated subsampling of a large point cloud

The goal of this section is to demonstrate that the techniques developed in Sections 5.2.2 and 5.2.3 can be used to obtain a multiscale approximation of the persistence diagram of a large point cloud X . Our aim is to highlight the method as opposed to presenting an optimal result, and thus we begin with a sequence of δ -approximations of X .

Definition 5.16. Let (X, d) be a finite metric space. A sequence $\mathcal{Y} = \{Y_i \subseteq X\}_{i=0}^m$ is a $\Delta = \{\delta_i > 0\}_{i=1}^m$ sampling of X if

- (i) $Y_0 = X, Y_{i+1} \subset Y_i$, and $\delta_i < \delta_{i+1}$ for all i , and
- (ii) for every $i > 0, Y_i$ is a δ_i approximation of Y_{i-1} .

Definition 5.17. Let $\mathcal{Y} = \{Y_i \subseteq X\}_{i=0}^m$ be a $\Delta = \{\delta_i > 0\}_{i=1}^m$ sampling of a finite metric space (X, d) . An *admissible stitching sequence* is a sequence $\mathcal{T} = \{t_i \geq 0\}_{i=1}^m$ satisfying

$$t_{i+1} \geq t_i + \delta_i, \quad i = 1, \dots, m-1.$$

The associated *stitched Vietoris-Rips persistence module*

$$U = U(M^{\mathcal{R}}(X); \mathcal{Y}, \Delta, \mathcal{T})$$

is defined inductively as follows.

Using the fact that the pair $(M^{\mathcal{R}}(X), M^{\mathcal{R}}(X))$ is (τ_0, σ_0) interleaved, where $\tau_0(t) = t$ and $\sigma_0(t) = t$, Proposition 5.6(i) guarantees that $(M^{\mathcal{R}}(Y_1), M^{\mathcal{R}}(X))$ is (τ_1, σ_1) interleaved, where $\tau_1(t) = t$ and $\sigma_1(t) = t + \delta_1$. Set

$$U_1 := U(M^{\mathcal{R}}(X); M^{\mathcal{R}}(X), t_1, M^{\mathcal{R}}(Y_1)).$$

By Proposition 5.12, $(U_1, M^{\mathcal{R}}(X))$ is (η_1, ρ_1) interleaved, where η_1 and ρ_1 are given as in Proposition 5.12(ii).

By Definition 5.16, Y_i is a δ_i -approximation of X , and hence by Proposition 5.6(i), the persistence modules $(M^{\mathcal{R}}(Y_i), M^{\mathcal{R}}(X))$ are (τ_i, σ_i) interleaved where $\tau_i(t) = t$ and $\sigma_i(t) = t + \delta_i$. Inductively, for $i = 1, \dots, m$, we use that the persistence modules $(U_{i-1}, M^{\mathcal{R}}(X))$ are (η_{i-1}, ρ_{i-1}) -interleaved and $(M^{\mathcal{R}}(Y_i), M^{\mathcal{R}}(X))$ are (τ_i, σ_i) -interleaved to define

$$U_i := U(M^{\mathcal{R}}(X); U_{i-1}, t_i, M^{\mathcal{R}}(Y_i))$$

and, again, use Proposition 5.12(ii) to define (η_i, ρ_i) that gives the interleaving of the pair $(U_i, M^{\mathcal{R}}(X))$.

Point clouds in \mathbb{R}^n			
<i>Approximate complex</i>	<i>Complex of Interest</i>	$\tau(t)$	$\sigma(t)$
Vietoris-Rips [14, Proof of Vietoris-Rips Lemma.]	Čech	$t\sqrt{\frac{2n}{n+1}}$	t
Net-tree [3, Proof of Proposition 20]	Čech	t	$(1 + \varepsilon)^2 t$
Graph induced complex [13, Proposition 2.8]	Vietoris-Rips	$t + 2\varepsilon$	t
Sparsified Vietoris-Rips [12, Claim 6.1]	Vietoris-Rips	t	$(1 + \varepsilon)t$
Point clouds in an arbitrary metric space			
<i>Approximate complex</i>	<i>Complex of Interest</i>	$\tau(t)$	$\sigma(t)$
Vietoris-Rips	Čech filtration	$2t$	t
Relaxed Vietoris-Rips [18, Lemma 4]	Vietoris-Rips	t	$\left(\frac{1}{1-2\varepsilon}\right) t$
Sparse weighted Rips [6, Lemma 6.13]	Vietoris-Rips	t	$\left(\frac{1+\sqrt{1+\delta^2\varepsilon}}{1-\varepsilon}\right) t$

Table 1: A table of approximations for Vietoris-Rips and Čech filtrations for point clouds. The first column gives the approximation and a reference to the construction of the approximation (i.e. the explicit construction of the interleavings), and the second column gives the complex that is being approximated. The third and fourth columns list the translation maps for (τ, σ) -interleavings of the associated persistence modules induced by taking homology of the associated filtrations. The values $\delta, \varepsilon \geq 0$ are parameters specified by the approximations where applicable.

Observe that having constructed $U = U(M^{\mathcal{R}}(X); \mathcal{Y}, \Delta, \mathcal{T})$, we have that $(U, M^{\mathcal{R}}(X))$ is (η_m, ρ_m) interleaved. This implies that an analogue of Corollary 5.15 provides a quantitative comparison of the associated persistence diagrams. An explicit description of this corollary is probably of limited interest. For many implementations, the computational bottleneck for obtaining a persistence diagram is the memory constraint associated with the Vietoris-Rips complex $\mathcal{R}(X, t)$. Thus, a desirable strategy is to: compute $M^{\mathcal{R}}(X)$ over an interval $[0, t_0]$; downsample to $Y_1 \subset X$; compute U_1 using $M^{\mathcal{R}}(Y_1)$ over an interval $[t_0, t_1]$; downsample to $Y_2 \subset Y_1$; and repeat the process. An open question is how to optimize the choice of the locations t_i of downsampling, and the $\delta_i \geq 0$ used to construct the downsampled sets $\{Y_i\}_{i=1}^m$.

5.3. A comparison of approximations of Vietoris-Rips and Čech filtrations

In applications, a persistence module is associated to a finite metric space (X, d) via the construction of a simplicial complex. There is typically a natural choice of complex for the problem of interest (e.g. Čech complex). However, the Vietoris-Rips complex is usually more manageable than the Čech complex. Table 1 provides a list of examples

of pairs of filtrations and their approximations that have appeared in the literature. Proposition 4.8 provides a general quantitative comparison of persistence diagrams given an interleaving between the associated persistence modules.

Table 1 explicitly defines interleavings and the interested reader can derive the bounds for the matching of persistence diagrams using Corollary 4.9 since all of the maps τ and σ in the table are bijections. Note that Proposition 2.19 enables one to keep track of errors even when multiple approximation steps have been used. For example, say that one desires to make a statement about the persistence diagram corresponding to the Čech filtration of a finite point cloud in \mathbb{R}^n via the persistence diagram corresponding to a filtration of the Sparsified Vietoris-Rips complex from [12] with parameter ε . Then the (η, ρ) -interleaving between the persistence module induced by the Čech filtration and the persistence module induced by the Sparsified Vietoris-Rips complex filtration is given by $\eta(t) = t\sqrt{\frac{2n}{n+1}}$ and $\rho(t) = (1 + \varepsilon)t$, where an intermediate approximation uses the Vietoris-Rips complex filtration (e.g. the translation pairs that one plugs into Proposition 2.19 come from the first and last rows of the first section of the table).

6. Acknowledgements

R. L. would like to thank Chuck Weibel, Michael Lesnick, and Ulrich Bauer for the many insightful discussions that led to the results presented in this paper.

M. K. was supported by ERC project GUDHI (Geometric Understanding in Higher Dimensions). R. L. was supported by DARPA contracts HR0011-17-1-0004 and HR0011-16-2-0033. S. H. and K. M. were partially supported by grants NSF-DMS-1125174, 1248071, 1521771, NIH 1R01GM126555-01 and DARPA contracts HR0011-16-2-0033, FA8750-17-C-0054.

References

References

- [1] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. *Proceedings of the Thirtieth Annual Symposium Computational Geometry*, page 355, 2014.
- [2] Ulrich Bauer and Michael Lesnick. Persistence diagrams as diagrams: A categorification of the stability theorem, 2016. arXiv:1610.10085.
- [3] Magnus Bakke Botnan and Gard Spreemann. Approximating persistent homology in euclidean space through collapses. *Applicable Algebra in Engineering, Communication and Computing*, 26(1-2):73–101, 2015.
- [4] P. Bubenik, J. Scott, and V. De Silva. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, page 31p., 2014.
- [5] Peter Bubenik and Jonathan A. Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51(3):600–627, Apr 2014.
- [6] Mickal Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. *Computational Geometry*, 58:70 – 96, 2016.
- [7] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry, SCG '09*, pages 237–246, New York, NY, USA, 2009. ACM.
- [8] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The Structure and Stability of Persistence Modules (Springer Briefs in Mathematics)*. Springer, 2016.

- [9] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, January 2007.
- [10] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and Its Applications*, 14(05):1550066, 2015.
- [11] B. A Davey and H. A Priestley. Introduction to lattices and order. *Cambridge University Press*, pages xii+298, 2002.
- [12] Tamal K. Dey, Fengtao Fan, and Yusu Wang. Computing topological persistence for simplicial maps. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry, SOCG'14*, pages 345:345–345:354, New York, NY, USA, 2014. ACM.
- [13] Tamal Krishna Dey, Fengtao Fan, and Yusu Wang. Graph induced complex on point data. In *Proceedings of the Twenty-ninth Annual Symposium on Computational Geometry, SoCG '13*, pages 107–116, New York, NY, USA, 2013. ACM.
- [14] Herbert Edelsbrunner and John L. Harer. *Computational topology : an introduction*. American Mathematical Society, Providence (R.I.), 2010.
- [15] Greg Friedman. Survey article: An elementary illustrated introduction to simplicial sets. *Rocky Mountain J. Math.*, 42(2):353–423, 04 2012.
- [16] Miroslav Kramár, Rachel Levanger, Jeffrey Tithof, Balachandra Suri, Mu Xu, Mark Paul, Michael F. Schatz, and Konstantin Mischaikow. Analysis of Kolmogorov flow and Rayleigh-Bénard convection using persistent homology. *Phys. D*, 334:82–98, 2016.
- [17] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.
- [18] Donald R. Sheehy. Linear-size approximations to the Vietoris-rips filtration. *Discrete and Computational Geometry*, 49(4):778–796, 2013.
- [19] C.A. Weibel. *An Introduction to Homological Algebra*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1995.
- [20] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2004.