HHS Public Access

Author manuscript

J Appl Comput Topol. Author manuscript; available in PMC 2020 June 01.

Published in final edited form as:

J Appl Comput Topol. 2019 June ; 3(1-2): 85–118. doi:10.1007/s41468-019-00026-x.

A comparison framework for interleaved persistence modules

Shaun Harker[1], **Miroslav Kramár**[2], **Rachel Levanger**[3], **Konstantin Mischaikow**[1]

[1]Department of Mathematics, Hill Center-Busch Campus, Rutgers University, 110 Frelingheusen Rd, Piscataway, NJ 08854-8019, USA

[2]INRIA Saclay, 1 Rue Honoré d'Estienne d'Orves, 91120 Palaiseau, France

[3]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 as reported by Bauer and Lesnick (in: Proceedings of the thirtieth annual symposium computational geometry 2014) 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.

Keywords

Persistence module; Persistence diagram; Persistent homology; Error bounds; Topological data analysis

Mathematics Subject Classification

55N35; 55U10; 65G99

1 Introduction

Persistent homology, see (Edelsbrunner and Harer 2010; Oudot 2015), or (Zomorodian and Carlsson 2004), 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.

Shaun Harker, sharker@math.rutgers.edu.

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 $f: X \to \mathbb{R}$ 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 $f: X \to [0, \infty)$ 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

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

for choices of $t \in [0, \infty)$. The collection $\{\mathrm{C}(f, t)\}_{t \in \mathbb{R}}$ 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 $f: X \to \mathbb{R}$ 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 Kramár et al. (2016) concerning the time-evolution of patterns in convection models. These examples are meant to motivate our interest in studying the geometry of the sets $\mathrm{C}(f, t)$. 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

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

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

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

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 $\{\varphi_V(s, t): V_s \to V_t\}_{s \leq t \in \mathbb{R}}$ satisfying the following conditions:

- **i.** $\varphi_V(t, t) = \mathrm{id}_{V_t}$ 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 Sects. 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 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 $f': X \to \mathbb{R}$, 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?

A fundamental result by Cohen-Steiner et al. (2007) 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. Recent developments by Bauer and Lesnick (2014) allow for comparisons of persistence modules through a matching of the associated persistence points. The primary theoretical results of this paper, Theorems 3.2 and 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 $f: X \to \mathbb{R}$. 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 Sect. 5.1, this sampling gives rise to a persistence module $V^{\mathbb{Z}}$. Assume that the persistence diagram $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$ has a single persistence point (2, 6) as shown in Fig. 1. As a consequence of Proposition 5.2, we can conclude that the persistence diagram of interest, PD(V), contains a single persistence point in the light gray region and possibly some other persistence points in the dark gray regions. This would correspond to geometrical features of $f: X \to \mathbb{R}$ that take place on a scale that is too fine to be detected by the integer-valued sampling. Finally, if a persistence point for $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$ 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 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, see (Cohen-Steiner et al. 2007). In particular, if PD(W) is an arbitrary persistence diagram whose bottleneck distance from $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$ is one, then 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 Sect. 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 (2014), 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 Sect. 4.

Section 3 focuses on Theorem 3.2, which is an extension of the Induced Matching Theorem of Bauer and Lesnick (2014). The proof incorporates ideas from the theory of generalized interleavings of Bubenik et al. (2014).

Section 4 begins with the proof of Theorem 4.1, which follows closely the proof of the Algebraic Stability Theorem of Bauer and Lesnick (2014). 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 4.2 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 4.6 and Corollary 4.7 indicate how Theorem 4.1 applies to specific points in the associated persistence diagrams.

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

In Sect. 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 PD(V) for the persistence module $V = (M(f), \phi_{M(f)})$. However, because of the size of $\mathcal{X}$, the computational cost of computing 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 Sect. 5.2) to compute 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 PD(V') based on $\mathcal{X}' \subset \mathcal{X}$. Proposition 5.6 provides a simple result bounding the locations of the persistence points in PD(V) based on 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 Sect. 5.2.2 the concept of stitching two persistence modules together to create a new persistence module. In Sect. 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 Sect. 5 illustrate, once the interleavings are determined, working with our framework is straightforward. With this in mind, we include Table 1 in Sect. 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 Sect. 2.1, we recall basic facts about persistence modules, their morphisms, and persistence diagrams. In Sect. 2.2 we provide a necessary background for interleavings of persistence modules. In Sect. 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 (2014) 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 (Bauer and Lesnick 2014; Bubenik et al. 2014; Chazal et al. 2016), or Zomorodian and Carlsson (2004).

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

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

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

and transition maps

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

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

$$\begin{tikzcd} V_s \arrow[r, "\varphi_V(s,t)"] \arrow[d, "\phi_s"] & V_t \arrow[d, "\phi_t"] \\ W_s \arrow[r, "\varphi_W(s,t)"] & W_t \end{tikzcd}$$

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, as shown in Bubenik and Scott (2014). 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 Chazal et al. 2016; Crawley-Boevey 2015) 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 $\overline{\mathbb{R}}$ 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 $a \le b$. Any nonempty interval J with endpoints a and b can be represented by an ordered pair $(\mathscr{B}(J), \mathscr{D}(J))$ of decorated points where:

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

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 $\mathscr{J}_V$ be a multiset of interval persistence modules in the interval decomposition of V. Suppose that the function $m: \mathscr{J}_V \to \mathbb{N}$ assigns to every interval persistence module $\mathrm{k}_J \in \mathscr{J}_V$ its multiplicity in $\mathscr{J}_V$. The *persistence diagram* of V is defined as the set

$$\mathbf{PD}(V) := \bigcup_{\mathrm{k}_J \in \mathscr{J}_V} \{[\mathscr{B}(J), \mathscr{D}(J), 1], \ldots, [\mathscr{B}(J), \mathscr{D}(J), m(\mathrm{k}_J)]\} \subset \mathbb{E} \times \mathbb{E} \times \mathbb{N}.$$

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 **PD** be a persistence diagram. The *left-handed ordering* of the points $[b, d, i] \in$ **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 **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. (2009) and generalized by Bubenik et al. (2014). Interleavings provide a measure of similarity between persistence modules.

Definition 2.9 A function $\sigma: \mathbb{R} \to \mathbb{R}$ is *monotone* if $x \le y$ implies that $\sigma(x) \le \sigma(y)$. If, in addition, $x \le \sigma(x)$ 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 $\sigma: \mathbb{R} \to \mathbb{R}$ 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, $\phi: V \to W$ a persistence module morphism, and $\sigma: \mathbb{R} \to \mathbb{R}$ a monotone function. The *σ-shifted persistence module morphism* $\phi(\sigma): V(\sigma) \to W(\sigma)$ 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 $\phi: V \to W(\tau)$ and $\psi: W \to V(\sigma)$ 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 $\phi_{\{V,\sigma\}}: V \to V(\sigma)$ 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 Bauer and Lesnick (2014), Chazal et al. (2009), and Chazal et al. (2016), 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 $\phi: V \to W(\tau)$ and $\psi: W \to V(\sigma)$. Then*

i. ker ϕ *and* coker ϕ *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 $r \le s \in \mathbb{R}$, 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. □

We close this section by recalling a result that allows us to compose interleavings. While the formulation of the definition of a (σ, τ)-interleaving above differs slightly from that in Bubenik et al. (2014), it is straightforward to show that the result still goes through with our more general definition (the proof ultimately relies on the notion of a translation pair, and does not explicitly require that each map is a translation map). Additionally, our definition of a (σ, τ)-interleaving is also given in the follow-up paper (Bubenik et al. 2017).

Proposition 2.19 (Bubenik et al. 2014, Proposition 2.2.11) *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 Davey and Priestley 2002) and establish some Galois connections that are used in Sect. 4.

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

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

Proposition 2.21 *Suppose P, Q, and R are posets and* $f: P \to Q$, $g: Q \to P$, $f': Q \to R$, *and* $g': R \to Q$ *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 $f' \circ f(x) \le y \Leftrightarrow f(x) \le g'(y) \Leftrightarrow x \le g \circ g'(y)$. □

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 $|\cdot\rangle: \mathbb{E} \to \mathbb{R}_L$ and $\langle \cdot | \mathbb{E} \to \mathbb{R}_U$ 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

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

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

$$\left|\sigma^{\downarrow}(e)\right\rangle = \downarrow \left\{\sigma(x) : x \in \left|e\right\rangle\right\},$$
$$\left\langle\sigma^{\uparrow}(e)\right| = \uparrow \left\{\sigma(x) : x \in \left\langle e\right|\right\}, \text{ and}$$
$$|\sigma^{\star}(e)\rangle = \sigma^{-1}(|e\rangle), \text{ or, equivalently, } \langle\sigma^{\star}(e)| = \sigma^{-1}(\langle e|).$$

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

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 $e \in \mathbb{E}$. ☐

Proposition 2.24 *Let $\sigma: \mathbb{R} \to \mathbb{R}$ 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 $\sigma^{\downarrow}(x) \le y$ for some $x, y \in \mathbb{E}$. We show $x \le \sigma^{\star}(y)$. Since $|\cdot\rangle$ is an order isomorphism, $\sigma^{\downarrow}(x) \le y$ is equivalent to $|\sigma^{\downarrow}(x)\rangle \subseteq |y\rangle$. By the definition of $\sigma^{\downarrow}$, 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 $x \le \sigma^{\star}(y)$.

We now prove the converse. That is, we suppose that $x \le \sigma^{\star}(y)$ and show $\sigma^{\downarrow}(x) \le y$. From $x \le \sigma^{\star}(y)$, 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, $\sigma^{\downarrow}(x) \le y$, 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. ☐

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

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

$$\pi\left(t^{\pm}\right) = t,\, i^{\pm}(t) = t^{\pm},$$

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

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

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

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

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 $\pi \circ i^{+} = \mathrm{id}$ and $\mathrm{id} \le i^{+} \circ \pi$, we have, for all $x \in \mathbb{E}$ and $y \in \overline{\mathbb{R}}$, the circle of implications $(\pi(x) \le y) \Rightarrow (i^{+} \circ \pi(x) \le i^{+}(y)) \Rightarrow x \le i^{+}(y)) \Rightarrow (\pi(x) \le \pi \circ i^{+}(y)) \Rightarrow (\pi(x) \le y)$. Thus, $(\pi(x) \le y) \Leftrightarrow (x \le i^{+}(y))$. 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 $\pi \circ i^{-} = \mathrm{id}$ and $i^{-} \circ \pi \le \mathrm{id}$, we have, for all $x \in \overline{\mathbb{R}}$ and $y \in \mathbb{E}$, the circle of implications $(x \le \pi(y)) \Rightarrow (i^-(x) \le i^- \circ \pi(y)) \Rightarrow (i^-(x) \le y) \Rightarrow (\pi \circ i^-(x) \le \pi(y)) \Rightarrow (x \le \pi(y))$. Thus, $(x \le \pi(y)) \Leftrightarrow (i^-(x) \le y)$. That is, the pair (i^-, π) is a Galois connecion.

Proposition 2.28 *Suppose* $f, g: \mathbb{E} \to \mathbb{E}$ *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 Propositions 2.21 and 2.27. □

2.4 Induced matchings on persistence diagrams

In this section, we summarize the work of Bauer and Lesnick (2014) and Bauer and Lesnick (2016) on matchings of persistence diagrams of PDF persistence modules V and W induced by a morphism $\phi: V \to W$.

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* $\mathcal{X}: S \nrightarrow T$ if $\mathcal{X}$ is the graph of an injective function $\mathcal{X}': S' \to T'$, where $S' \subseteq S$ and $T' \subseteq T$. We define the domain and image of a matching via $\mathrm{dom}\, \mathcal{X} := \mathrm{dom}\, \mathcal{X}'$ and $\mathrm{im}\, \mathcal{X} := \mathrm{im}\, \mathcal{X}'$, 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:

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

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 by Bauer and Lesnick (2014).

Proposition 2.31 (Theorem 4.2, Bauer and Lesnick 2014) *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* $|\mathbf{PD}^d(V)| \le |\mathbf{PD}^d(W)|$ *for* $d \in \mathbb{E}$.
- **ii.** *If there exists an epimorphism* $V \twoheadrightarrow W$, *then* $|\mathbf{PD}_b(W)| \le |\mathbf{PD}_b(V)|$ *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 (Bauer and Lesnick 2014, 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* $\mathcal{X}_{i(V,W)}: \mathbf{PD}(V) \nrightarrow \mathbf{PD}(W)$ *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 epiomorphism from V to W, then there exist a unique matching* $\mathcal{X}_{s(V,W)}: \mathbf{PD}(V) \nrightarrow \mathbf{PD}(W)$ *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 $\phi: V \to W$ can be factored as the composition of an epimorphism and monomorphism as follows:

$$V \twoheadrightarrow \operatorname{im}\phi \hookrightarrow W$$

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

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

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

Proposition 2.34 (Proposition 5.7, Bauer and Lesnick 2014) *Let* $\phi: U \to V$ *and* $\psi: V \to W$ *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*, if for all $y \in B$, there exists some $x \in A$ with $x \leq y$. We say that B *bounds* A *above*, if for all $x \in A$, there exists some $y \in B$ such that $x \leq y$. We say that B *overlaps* A *above*, if and only if each of the following conditions hold: A bounds B below, B bounds A above, and $A \cap B \neq \varnothing$.

Proposition 2.36 (Proposition 5.3, Bauer and Lesnick 2014) *Let* $\phi: V \to W$ *be a PFD persistence module morphism. If* $\mathcal{X}_\phi([b,d,i]) = [b',d',i']$, *then* $\langle b, d\rangle$ *overlaps* $\langle b', d'\rangle$ *above.*

3 Generalized induced matching theorem

In this section we present a generalization of the Induced Matching Theorem of Bauer and Lesnick (2014).

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

- **i.** *If* coker ϕ *is σ-trivial, then $\langle b, d\rangle$ bounds $\sigma(\langle b', d'\rangle)$ below and* $\operatorname{im} \mathcal{X}_\phi$ *contains all σ-nontrivial points in* $\mathbf{PD}(W)$.
- **ii.** *If* ker ϕ *is σ-trivial, then $\langle b', d'\rangle$ bounds $\sigma^{-1}(\langle b, d\rangle)$ above and* $\operatorname{dom} \mathcal{X}_\phi$ *contains all σ-nontrivial points in* $\mathbf{PD}(V)$.

Note that the Induced Matching Theorem of Bauer and Lesnick (2014) 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

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

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 $\operatorname{im} \varphi_V(s, t)|_{V_s^\sigma} \subseteq V_t^\sigma$ 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, we have that $t' \leq \sigma(t') \leq s \leq t$, and so $\varphi(s, t)(y) \in \operatorname{im} \varphi_V(t', t) \subseteq V_t^\sigma$.

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

- **i.** *If* coker ϕ *is σ-trivial, then* $W_t^\sigma \subseteq \operatorname{im} \phi_t \subseteq W_t$ *for every* $t \in \mathbb{R}$, *and*
- **ii.** *if* ker ϕ *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 $\phi : V \to W$, the persistence module coker ϕ is σ-trivial if and only if

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

which is true if and only if

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

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

$$\begin{tikzcd} & V_{\sigma(t')} \arrow[r, "\varphi_V(\sigma(t'),t)"] \arrow[d, "\phi(\sigma)_{t'}"] & V_t \arrow[d, "\phi_t"] \\ W_{t'} \arrow[r, "\varphi_W(t',\sigma(t'))"] & W_{\sigma(t')} \arrow[r, "\varphi_W(\sigma(t'),t)"] & W_t \end{tikzcd}$$

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

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 k_J be an interval persistence module and σ a translation map. If $J \cap \sigma(J) \neq \varnothing$, then*

- **i.** $k_J^{\sigma} \cong k_{J \cap Conv(\sigma(J))}$, *where $Conv(\sigma(J))$ is the convex hull of $\sigma(J)$,*
- **ii.** $k_J/\ker \phi_{\{k_J, \sigma\}} \cong k_{J \cap \sigma^{-1}(J)}$,

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

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

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

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

For the following two definitions, for an interval $J \subseteq \mathbb{R}$, we recall the symbols $\mathscr{B}(J)$ and $\mathscr{D}(J)$ (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 $J = \langle b, d\rangle \cap \mathrm{Conv}(\sigma\langle b, d\rangle)$. Then $[\mathscr{B}(J), d, i] \in \mathbf{PD}(V^\sigma)$ if and only if $J \neq \varnothing$. Moreover, in that case,*

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

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

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

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 $[b, \mathscr{D}(J), i] \in \mathbf{PD}(V/\ker \phi_{\{V, \sigma\}})$ if and only if $J := \langle b, d\rangle \cap \sigma^{-1}(\langle b, d\rangle) \neq \varnothing$. Moreover, in that case,*

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

Proof of Theorem 3.2 Our proof closely follows the proof of the Induced Matching Theorem in Bauer and Lesnick (2014). To prove (i), we start by establishing the existence of certain matchings. By Lemma 3.5(i), W^σ is a submodule of im ϕ and so the matching

$\mathcal{X}_{i(W,\ \mathrm{im}\ \phi)}$ is defined. It follows from Lemma 3.4 that W^{σ} is a submodule of W and thus $\mathcal{X}_{i\left(W^{\sigma},W\right)}$ is defined. Proposition 2.34 implies that the following diagram commutes.

$$\begin{tikzcd}
 & \mathbf{PD}(W) & \\
\mathbf{PD}(W^{\sigma}) \arrow[ur, hook, "\mathcal{X}_{i(W^{\sigma},W)}"] \arrow[dr, hook, "\mathcal{X}_{i(W^{\sigma},\mathrm{im}\ \phi)}"'] & & \mathbf{PD}(V) \arrow[ul, "\mathcal{X}_{\phi}"'] \arrow[dl, two heads, "\mathcal{X}_{s(V,\mathrm{im}\ \phi)}"] \\
 & \mathbf{PD}(\mathrm{im}\ \phi) \arrow[uu, hook, "\mathcal{X}_{i(\mathrm{im}\ \phi,W)}"'] &
\end{tikzcd}$$

It follows from Proposition 2.33(i) that $\mathrm{im}\ \mathcal{X}_{\phi} = \mathrm{im}\ \mathcal{X}_{i(\mathrm{im}\ \phi,\,W)}$. By commutativity of the left triangle, $\mathrm{im}\ \mathcal{X}_{i\left(W^{\sigma},W\right)} \subseteq \mathrm{im}\ \mathcal{X}_{i(\mathrm{im}\ \phi,\,W)}$. Now it follows from Proposition 3.7 that $\mathrm{im}\ \mathcal{X}_{\phi}$ 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 $\langle b, d\rangle$ bounds $\sigma(\langle b', d'\rangle)$ below. First we suppose that $[b', d', i']$ is σ-nontrivial. In this case for $J := \langle b', d'\rangle \cap \mathrm{Conv}(\sigma\langle b', d'\rangle)$, the point $[\mathscr{B}(J), d', i'] \in \mathbf{PD}\left(W^{\sigma}\right)$ and $\mathcal{X}_{i\left(W^{\sigma},W\right)}([\mathscr{B}(J),d',i']) = [b',d',i']$. Since σ a translation map, we have $\langle \mathscr{B}(J), d'\rangle$ bounds $\sigma(\langle b', d'\rangle)$ below. Due to commutativity of the above diagram, $\mathcal{X}_{i\left(W^{\sigma},\mathrm{im}\ \phi\right)}([\mathscr{B}(J),d',i']) = \mathcal{X}_{s(V,\,\mathrm{im}\ \phi)}([b,d,i])$. It follows from Proposition 2.32(iv) and Proposition 2.33(iv) that $b \le \mathscr{B}(J)$. Therefore, $\langle b, d\rangle$ bounds $\sigma(\langle b', d'\rangle)$ below. On the other hand, if $[b', d', j]$ is σ-trivial, then $\langle b', d'\rangle \cap \sigma(\langle b', d'\rangle) = \varnothing$ and every point in $\sigma(\langle b', d'\rangle)$ is larger than d'. By Proposition 2.36, $b < d'$ and we get $\langle b, d\rangle$ bounds $\sigma(\langle \mathrm{b}', \mathrm{d}'\rangle)$ below.

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

$$\begin{tikzcd}
 & \mathbf{PD}(V) \arrow[dl, two heads, "\mathcal{X}_{s(V,V/\ker\phi_{\{V,\sigma\}})}"'] \arrow[dr, "\mathcal{X}_{\phi}"] \arrow[dd, two heads, "\mathcal{X}_{s(V,\mathrm{im}\ \phi)}"] & \\
\mathbf{PD}(V/\ker\phi_{\{V,\sigma\}}) & & \mathbf{PD}(W) \\
 & \mathbf{PD}(\mathrm{im}\ \phi) \arrow[ul, two heads, "\mathcal{X}_{s(\mathrm{im}\ \phi,V/\ker\phi_{\{V,\sigma\}})}"] \arrow[ur, hook, "\mathcal{X}_{i(\mathrm{im}\ \phi,W)}"'] &
\end{tikzcd}$$

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 Bauer and Lesnick (2014).

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 $\phi: V \to W(\tau)$ and $\psi: W \to V(\sigma)$. There exists a matching $\mathcal{X}: \mathbf{PD}(V) \nrightarrow \mathbf{PD}(W)$ such that $\mathcal{X}([b,d,i]) = [b',d',i']$ implies*

- **i.** $\tau^{-1}(\langle b', d'\rangle)$ *bounds* $(\sigma\circ\tau)^{-1}$, $(\langle b, d\rangle)$ *above,*

ii. $\langle b, d\rangle$ *overlaps* $\tau^{-1}(\langle b', d'\rangle)$ *above, and*

iii. $\langle b, d\rangle$ *bounds* $\sigma \circ \tau \circ \tau^{-1}(b', d'\rangle)$ *below.*

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* $\langle b', d'\rangle \cap \operatorname{im} \tau = \varnothing$ *or* $\tau^{-1}(\langle b', d'\rangle)$ *is* $(\sigma \circ \tau)$*-trivial.*

As in the case of the Induced Matching Theorem, the result of Bauer and Lesnick (2014) follows from setting $\tau(t) = \sigma(t)$ $t + \delta$ for $\delta \geq 0$. It is worth pointing out that Theorem 4.1 is not self-dual, in contrast to the result of Bauer and Lesnick (2014). This is due to the fact that the maps τ and σ from the interleavings might not be bijections.

Proof It follows from Propostion 2.18(i) that ker ϕ and coker ϕ are $(\sigma \circ \tau)$-trivial. By Theorem 3.2, the domain of $\mathcal{X}_\phi : \mathbf{PD}(V) \nrightarrow \mathbf{PD}(W(\tau))$ 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

a. $\langle x, y\rangle$ bounds $(\sigma \circ \tau)^{-1}(\langle b, d\rangle)$ above by Theorem 3.2,

b. $\langle b, d\rangle$ overlaps $\langle x, y\rangle$ above by Proposition 2.36, and

c. $\langle b, d\rangle$ bounds $\sigma \circ \tau \langle x, y\rangle$ below by Theorem 3.2.

To finish the proof, we build an appropriate matching $\mathcal{X}' : \mathbf{PD}(W(\tau)) \nrightarrow \mathbf{PD}(W)$. 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

$$\{[b', d', i'] \in \mathbf{PD}(W) : \langle b', d'\rangle \cap \operatorname{im} \tau \neq \varnothing\}.$$

This correspondence can be realized by a matching $\mathcal{X}' : \mathbf{PD}(W(\tau)) \nrightarrow \mathbf{PD}(W)$ 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 $\mathcal{X}' \circ \mathcal{X}_\phi$. Conditions (i–iii) then follow from (a–c) above and the fact that $\langle x, y\rangle = \tau^{-1}(\langle b', d'\rangle)$. □

Statements (i–iii) 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 4.2 *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, then the inequalities:*

$$\begin{aligned} &\tau^\star(b') \leq b, \quad \tau^\star(d') \leq d, \\ &\sigma^\star(b) \leq \tau^\uparrow \circ \tau^\star(b'), \quad \sigma^\star(d) \leq \tau^\uparrow \circ \tau^\star(d') \end{aligned}$$

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

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

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 Theorem 4.1(i–iii).

The fact that $\langle b, d\rangle$ overlaps $\tau^{-1}(\langle b', d'\rangle)$ above together with the monotonicity of τ implies that $\tau(\langle b, d\rangle)$ overlaps $\langle b', d'\rangle$ above, and by definition we thus have $\langle b', d'\rangle$ bounds $\tau(\langle b, d\rangle)$ below and $\tau(\langle b, d\rangle)$ bounds $\langle b', d'\rangle$ above.

The inequality $\tau^\star(b') \le b$ ($\tau^\star(d') \le d$) follows from definition of $\tau^\star$ and the fact that $\langle b', d'\rangle$ bounds $\tau(\langle b, d\rangle)$ below ($\tau(\langle b, d\rangle)$ bounds $\langle b', d'\rangle$ above). Now we prove that $\sigma^\star(b) \le \tau^\uparrow \circ \tau^\star(b')$. Starting from the fact that $\langle b, d\rangle$ bounds $\sigma \circ \tau \circ \tau^{-1}(\langle b', d'\rangle)$ below, we obtain $\langle b, d\rangle$ that bounds.

$$\uparrow \left\{\sigma \circ \tau \circ \tau^{-1}(\langle b', d'\rangle)\right\} = \left\langle(\sigma \circ \tau)^{\uparrow} \circ \tau^{\star}(b'), \infty\right\rangle$$

below. Hence $b \le (\sigma \circ \tau)^\uparrow \circ \tau^\star(b')$, and the desired inequality follows from the definition of $\sigma^\star$. The last inequality follows again from the definitions of $\sigma^\star$, $\tau^\uparrow$, and $\tau^\star$, and thus, by Theorem 4.1(iii), we have that $(\sigma \circ \tau)^{-1}(\langle b, d\rangle)$ bounds $\tau^{-1}(\langle b', d'\rangle)$ below.

Now the moreover part. We suppose that τ and σ are invertible. Combining the invertibility of τ with the definitions of $\tau^\uparrow$ and $\tau^\star$ (i.e. Definition 2.22), it is routine to verify $\tau^\uparrow \circ \tau^\star = \mathrm{id}$. Similarly, $\sigma^\uparrow \circ \sigma^\star = \mathrm{id}$. 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 $\langle b', d'\rangle \cap \mathrm{im}\, \tau = \varnothing$ 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. □

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 4.3 Let V be a PDF persistence module and $\mathbf{PD}(V)$ its persistence diagram. The *undecorated persistence diagram* $\mathrm{PD}(V)$ of V is a subset of $\overline{\mathbb{R}}^2 \times \mathbb{N}_{>0}$ with the following properties:

- **i.** if $[b, d, n] \in \mathrm{PD}(V)$, then $[b, d, m] \in \mathrm{PD}(V)$ for all $m < n$; and
- **ii.** there exists a bijection $\mathscr{X}_V : \mathbf{PD}(V) \twoheadrightarrow \mathrm{PD}(V)$ such that if $\mathscr{X}_V([b, d, i]) = [b', d', i']$, then $b' = \pi(b)$ and $d' = \pi(d)$.

It is clear that by condition (ii) the subset $\mathrm{PD}(V) \subseteq \overline{\mathbb{R}}^2 \times \mathbb{N}_{>0}$ exists, and by condition (i) it is unique. To formulate an analog of Theorem 4.1 for undecorated persistence diagrams we make use of the following functions.

Definition 4.4 Let $\tau:\mathbb{R}\to\mathbb{R}$ be a monotone function. We define monotone functions $\tau_L, \tau_L^\dagger, \tau_R, \tau_R^\dagger:\overline{\mathbb{R}}\to\overline{\mathbb{R}}$ as:

1. $\tau_L(\pm\infty) = \lim_{x\to\pm\infty} \tau(x)$ and $\tau_L(x) := \lim_{y\to x^-} \tau(y)$ for $x\in\mathbb{R}$.
2. $\tau_R(\pm\infty) = \lim_{x\to\pm\infty} \tau(x)$ and $\tau_R(x) := \lim_{y\to x^+} \tau(y)$ for $x\in\mathbb{R}$.
3. $\tau_L^\dagger(\pm\infty) := \pm\infty$ and $\tau_L^\dagger(x) = \inf\left\{y : x < \tau(y)\right\}$ for $x\in\mathbb{R}$.
4. $\tau_R^\dagger(\pm\infty) := \pm\infty$ and $\tau_R^\dagger(x) := \inf\left\{y : x \le \tau(y)\right\}$ for $x\in\mathbb{R}$.

Proposition 4.5 *Let $\tau:\mathbb{R}\to\mathbb{R}$ be a monotone function. We have that $\tau_L = (\tau^\downarrow)_-$, $\tau_R = (\tau^\uparrow)_+$, $\tau_L^\dagger = (\tau^\star)_+$, and $\tau_R^\dagger = (\tau^\star)_-$. Moreover, the pair $(\tau_R^\dagger, \tau_R)$ is a Galois connection and the pair $(\tau_L, \tau_L^\dagger)$ 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 4.4 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 $\left(\tau_L, \tau_L^\dagger\right) = \left(\left(\tau^\downarrow\right)_-, \left(\tau^\star\right)_+\right)$ and $\left(\tau_R^\dagger, \tau_R\right) = \left(\left(\tau^\star\right)_-, \left(\tau^\uparrow\right)_+\right)$, the result follows from Proposition 2.24 and Proposition 2.28. □

Proposition 4.6 *Let (V, φ_V) and (W, φ_W) be persistence modules such that (V, W) are (τ, σ)-interleaved. Then there exists a matching $\mathcal{X}: PD(V) \nrightarrow PD(W)$ with the following properties. If $\mathcal{X}([b, d, i]) = [b', d', i']$, then*

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

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

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

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

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

Proof We consider the matching $\mathcal{X}:\mathrm{PD}(V)\nrightarrow\mathrm{PD}(W)$ 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 4.3 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 $\tau_R^\dagger(b') \le b$, $\tau_R^\dagger(d') \le d$, $\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger(b) \le b'$ and $\left(\tau_L \circ \tau_R^\dagger \circ \sigma_R^\dagger\right)(d) \le d'$ since by Proposition 4.5 the other inequalities can be recovered using Galois connections (e.g. $\tau_R^\dagger(b') \le b$ if and only if $b' \le \boldsymbol{\tau}_R(b)$).

We only prove $\tau_R^\dagger(b') \le b$ 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 4.2 that $\boldsymbol{\tau}^\star(c') \le c$, and so $\boldsymbol{\pi} \circ \boldsymbol{\tau}^\star(c') \le \boldsymbol{\pi}(c)$. Since $i^- \circ \boldsymbol{\pi}(x) \le \mathrm{id}(x)$ for all $x \in \mathbb{E}$, we get $\boldsymbol{\pi} \circ \boldsymbol{\tau}^\star \circ i^- \circ \boldsymbol{\pi}(c') \le \boldsymbol{\pi}(c)$. By Proposition 4.5, we get $\tau_R^\dagger = \pi \circ \tau^\star \circ i^-$, which implies that $\tau_R^\dagger(\pi(c')) \le \pi(c)$. It follows from the definition of $\mathcal{X}_V$ and $\mathcal{X}_W$ that $b' = \boldsymbol{\pi}(c')$ and $b = \boldsymbol{\pi}(c)$. Combining this with the previous inequality yields $\tau_R^\dagger(b') \le b$.

Now we assume that $[b, d, i] \in \mathrm{PD}(V)$ is not in 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 $(\boldsymbol{\sigma} \circ \boldsymbol{\tau})$-trivial, i.e. $e \le \boldsymbol{\sigma}^\uparrow \circ \boldsymbol{\tau}^\uparrow(c)$. Since $\mathrm{id}(x) \le i^+ \circ \boldsymbol{\pi}(x)$ for all $x \in \mathbb{E}$. we have $\boldsymbol{\pi}(e) \le \boldsymbol{\pi} \circ \boldsymbol{\sigma}^\uparrow \circ (i^+ \circ \boldsymbol{\pi}) \circ \boldsymbol{\tau}^\uparrow \circ (i^+ \circ \boldsymbol{\pi})(c)$. Proposition 4.5 implies that $\boldsymbol{\pi}(e) \le \boldsymbol{\sigma}_R \circ \boldsymbol{\tau}_R(\boldsymbol{\pi}(c))$. By using $\boldsymbol{\pi}(c) = b$ and $\boldsymbol{\pi}(c') = b'$, we obtain that $d \le \boldsymbol{\sigma}_R \circ \boldsymbol{\tau}_R(b)$. The inequality for the points $[b', d', i'] \in \mathrm{PD(W)}$ that are not in im $\mathcal{X}$ can be achieved by using similar methods as above and is left to the reader. ☐

Corollary 4.7 *Let (V, φ_V) and (W, φ_W) be persistence modules such that (V, W) are $(\boldsymbol{\tau}, \boldsymbol{\sigma})$-interleaved. Suppose that the maps $\boldsymbol{\tau}$ and $\boldsymbol{\sigma}$ are bijections. If $\mathcal{X}: PD(V) \to PD(W)$ is the matching given by Proposition 4.6 and $\mathcal{X}([b, d, i]) = [b', d', i']$, then*

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

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

Proof If $\boldsymbol{\tau}$ and $\boldsymbol{\sigma}$ are invertible, then $\boldsymbol{\tau}_L = \boldsymbol{\tau}_R = \boldsymbol{\tau}$, $\boldsymbol{\sigma}_L = \boldsymbol{\sigma}_R = \boldsymbol{\sigma}$, $\tau_L^\dagger = \tau_R^\dagger = \tau^{-1}$, and $\sigma_L^\dagger = \sigma_R^\dagger = \sigma^{-1}$. The proof is obtained by evaluating expressions in Proposition 4.6. ☐

5 Applications

We illustrate the use of results obtained in Sect. 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 $f: X \to \mathbb{R}$ 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* $\mathcal{X}: PD(V^{\mathbb{Z}}) \nrightarrow PD(V)$ *such that if* $\mathcal{X}([b,d,i]) = [b', d', i']$, *then*

$$b - 1 \le b' \le b \text{ and } d - 1 \le d' \le d\,.$$

Additionally, any unmatched points $[b', d', i'] \in PD(V)$ *satisfy* $d' < \lfloor b' + 1 \rfloor$, *and all points in* $PD(V^{\mathbb{Z}})$ *are matched.*

Proof (i) Define persistence module morphisms $\phi: V^{\mathbb{Z}} \to V(\tau) = V$ by $\phi_t := \varphi_V(\lfloor t \rfloor, t)$ and $\psi: V \to V^{\mathbb{Z}}(\sigma)$ 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) First, note that that if $[b,d,i] \in \mathrm{PD}(V^{\mathbb{Z}})$ then $b, d \in \mathbb{Z}$. This, together with a direct application of Proposition 4.6 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

$$\sigma_R(t) = \lfloor t + 1 \rfloor \quad \text{and} \quad \sigma_R^{\dagger}(t) = \lceil t - 1 \rceil,$$

yields the bounds for the matched points. Since the matching given by Proposition 4.6 inherits the matching on $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$ and PD(V) from the matching given by Theorem 4.1 on **PD**$\left(V^{\mathbb{Z}}\right)$ and **PD**(V), we prove the statements about the unmatched points using this latter result. By the definition of the maps τ, σ and the interleaving map $\sigma: V^{\mathbb{Z}} \to V(\tau) = V$, we see that Theorem 4.1 implies that unmatched intervals in both **PD**$\left(V^{\mathbb{Z}}\right)$ and **PD**(V) are σ-trivial. By the definition of $V^{\mathbb{Z}}$, any interval $\langle c, e \rangle$ in **PD**$\left(V^{\mathbb{Z}}\right)$ necessarily has $\pi(c) \in \mathbb{Z}$ and so $\sigma \circ \pi(c) = \pi(c)$. Hence, no such interval is σ-trivial, and so every point in $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$ is matched. Finally, the only σ-trivial intervals $\langle c', e' \rangle$ in **PD**(V) are those that do not contain an integer, and thus any unmatched point $[b', d', n'] \in$ PD(V) necessarily has $d' < \lfloor b' + 1 \rfloor$. □

See Fig. 1 for an illustration of an estimate of PD(V) from $\mathrm{PD}\left(V^{\mathbb{Z}}\right)$.

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 $\mathscr{R}(X, t)$, is the simplicial complex with vertices given by X and containing the N-simplex $\left[x_{i_0}, \ldots, x_{i_N}\right]$ if and only if $d\left(x_{\mathrm{k}_J}, x_{i_k}\right) \leq 2t$ for all j, k = 0, …, N.

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

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

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

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

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

Observe that for large X, the computational cost of determining $H_*(\mathscr{R}(X, t), \mathrm{k})$ grows rapidly as a function of t. If $Y \subset X$, then one expects that it is cheaper to compute $H_*(\mathscr{R}(Y, t), \mathrm{k})$.

The goal of this section is twofold: first, to quantify the difference between $M^{\mathscr{R}}(X)$ and $M^{\mathscr{R}}(Y)$; and second, to suggest an iterative procedure, motivated by Dey et al. (2014), for obtaining reasonable approximations of $M^{\mathscr{R}}(X)$.

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 $(M^{\mathscr{R}}(Y), M^{\mathscr{R}}(X))$ are (τ, σ)-interleaved, where $\tau(t) = t$ and $\sigma(t) = t + \delta$.*

ii. *There exists a matching $\mathscr{X}: PD\left(M^{\mathscr{R}}(Y)\right) \nrightarrow PD\left(M^{\mathscr{R}}(X)\right)$ such that if $PD_{\mathscr{X}}(b, d, i) = (b', d', i')$, then $b - \delta \leq b' \leq b$ and $d - \delta \leq d' \leq d$. Moreover, all unmatched points in $PD\left(M^{\mathscr{R}}(Y)\right)$ and $PD\left(M^{\mathscr{R}}(X)\right)$ 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 4.7. 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 $\gamma: Y \to Y'$ satisfies $d(x, \gamma(x)) \leq \delta$ for all $x \in Y$, then $\tilde{\gamma}_t: \mathscr{R}(Y, t) \to \mathscr{R}(Y', t + \delta)$, defined by*

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

is a simplicial map.

Proof To prove that $\tilde{\gamma}$ is a simplicial map, we need to show that for every k-simplex $[x_0, \ldots, x_k] \in \mathscr{R}(Y, t)$, its image $[\gamma(x_0), \ldots, \gamma(x_k)]$ under $\tilde{\gamma}$ is a simplex in $\mathscr{R}(Y', t + \delta)$. 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 $\mathscr{R}(Y, t)$ is mapped to the 1-skeleton of $\mathscr{R}(Y', t + \delta)$. Recall that $[x, y]$ is an edge in $\mathscr{R}(Y, t)$ 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 $\mathscr{R}(Y', t + \delta)$. □

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

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

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

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

Our goal is to show that $\phi: M^{\mathscr{R}}(Y) \to M^{\mathscr{R}}(X)$ and $\psi: M^{\mathscr{R}}(X) \to M^{\mathscr{R}}(Y)(\delta)$ are persistence module morphisms that guarantee that the persistence modules $(M^{\mathscr{R}}(Y), M^{\mathscr{R}}(X))$ 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

$$\psi(\tau)_t \circ \phi_t = \psi_t \circ \phi_t = \varphi_{M^{\mathscr{R}}(Y)}[t, t+\delta] = \varphi_{M^{\mathscr{R}}(Y)}[t, (\sigma \circ \tau)t]. \tag{5.1}$$

The challenge is to show that the middle equality holds for

$$\phi(\sigma)_t \circ \psi_t = \phi_{t+\delta} \circ \psi_t = \varphi_{M^{\mathscr{R}}(X)}[t, t+\delta] = \varphi_{M^{R}(X)}[t, (\tau \circ \sigma)t]. \tag{5.2}$$

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

Lemma 5.8 *Consider Y, $Y' \subseteq X$ and $\delta \ge 0$. Let $\iota_t': \mathscr{R}(Y', t) \to \mathscr{R}(X, t+\delta)$ and $\iota_t: \mathscr{R}(Y, t) \to \mathscr{R}(X, t)$ be the simplicial maps induced by inclusion. If $\gamma: Y' \to Y$ satisfies $d(y, \gamma(y)) \le \delta$ for all $y \in Y'$ and $\tilde{\gamma}_t: \mathscr{R}(Y', t) \to \mathscr{R}(Y, t+\delta)$ 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 (see Weibel 1995 or Friedman 2012) and begin with the remark that by (Weibel 1995, 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 $\mathscr{K}$ let $\overline{\mathscr{K}}$ denote the associated simplicial set. To establish notation let $\overline{\mathscr{K}}_k$ denote the k-dimensional simplices in $\mathscr{K}$ and let $d_i: \overline{\mathscr{K}}_k \to \overline{\mathscr{K}}_{k-1}$ and $s_i: \overline{\mathscr{K}}_k \to \overline{\mathscr{K}}_{k+1}$ be the delete and duplicate i-th vertex operations defined by

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

and

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

We claim that the functions $h_i: \overline{\mathscr{R}}(Y', t)_k \to \overline{\mathscr{R}}(X, t+\delta)_{k+1}$, $i = 0, \ldots, k$, defined by

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

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. ☐

Proof of Proposition 5.6 (i) As indicated above, the proof of Proposition 5.6(i) follows from (5.1), which has already been justified, and (5.2), which follows from Lemma 5.8 under the assumption that $Y' = X$. ☐

5.2.2 Stitching persistence modules—

We begin this section by giving a useful construction and then motivate it with an application.

Definition 5.9 Let (V, φ_V) and (W, φ_W) be persistence modules and $s_0 \leq t_0$. Suppose that ϕ : $V_{s_0} \to W_{t_0}$ is a linear map. If $s_0 = t_0$ we require that the vector spaces V_{s_0} and W_{t_0} are the same, and that ϕ is the identity map. The *stitched persistence module* $U = U(V, W, \phi)$ is defined by the vector spaces

$$U_t = \begin{cases} V_t & \text{if } t \leq s_0 \\ V_{s_0} & \text{if } s_0 \leq t < t_0 \\ W_t & \text{if } t_0 \leq t \end{cases}$$

and linear maps

$$\varphi_U(s,t) = \begin{cases} \varphi_V(\min\{s, s_0\}, \min\{t, s_0\}) & \text{if } s \leq t < t_0 \\ \varphi_W(t_0, t) \circ \phi \circ \varphi_V(\min\{s, s_0\}, s_0) & \text{if } s < t_0 \leq t \\ \varphi_W(s,t) & \text{if } t_0 \leq s \leq t. \end{cases}$$

Remark 5.10 The stitched persistence module $U = U(V, W, \phi)$ is a PFD persistence module.

Proof Since V and W are PFD persistence modules, U_t is finite dimensional for each $t \in \mathbb{R}$. It is left to the reader to check that $\varphi_U(t, t)$ id_{U_t} for every $t \in \mathbb{R}$ and that $\varphi_U(s, t) \circ \varphi_U(r, s) = \varphi_U(r, t)$ for every $r \leq s \leq t$ in $\mathbb{R}$.

We showed that the persistence module $M^{\mathscr{R}}(X)$ of a finite metric space (X, d) can be approximated by the persistence module $M^{\mathscr{R}}(Y_\delta)$, where Y_δ is a δ-approximation of X. The quality of this approximation and the computational cost of constructing $M^{\mathscr{R}}(Y_\delta)$ decreases as δ increases. In particular, given fixed computational resources, the value of t for which the vector spaces $M^{\mathscr{R}}(Y_\delta)_t$ and the maps between them can be computed increases with δ. Therefore, to obtain a computable approximation of $M^{\mathscr{R}}(X)$ that is better than $M^{\mathscr{R}}(Y_\delta)$, we stitch together the persistence modules $M^{\mathscr{R}}(X)$ and $M^{\mathscr{R}}(Y_\delta)$ at t_0, chosen in such a way that the vector spaces $M^{\mathscr{R}}(Y_\delta)_t$ (and the maps) can be computed. Figure 2 demonstrates how the approximation given by $M^{\mathscr{R}}(Y_\delta)$ improves when it is stitched with $M^{\mathscr{R}}(X)$. Alternatively, we could consider stitching together two persistence modules $M^{\mathscr{R}}(Y_\delta)$ and $M^{\mathscr{R}}(Y_{\delta'})$ for $\delta < \delta'$, or even iterate the this procedure, as explained in the following section.

This motivates our interest in combining (or stitching together) two different approximations of a persistence module (W, φ_W). We recall that the quality of the approximation can be assessed by interleavings. Thus, in the rest of the sections that follow, by saying that (V, φ_V) and $(V', \varphi_{V'})$ are two different approximations of (W, φ_W), we mean that the persistence modules:

1. (V, W) are (τ, σ)-interleaved via $\phi : V \to W(\tau)$ and $\psi : W \to V(\sigma)$; and
2. (V', W) are (τ', σ')-interleaved via $\phi' : V' \to W(\tau')$ and $\psi' : W \to V'(\sigma')$.

In the following definition, we explain when two different approximations of W can be stitched together, and then we consider the best interleaving that is guaranteed to exist between W and the stitched persistence module.

Proposition 5.11 *Let (V, φ_V) and $(V', \varphi_{V'})$ be two different approximations of (W, φ_W) and suppose that $t_0 \in \mathbb{R}$ satisfies $t_0 \leq \sigma' \circ \tau(t_0)$. Let $U = U\left(V, V', \psi'_{\tau(t_0)} \circ \phi_{t_0}\right)$ be the stitched persistence module. Then (U, W) are (η, ρ)-interleaved where*

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

$$\rho(t) = \begin{cases} \sigma(t) & if\ \sigma(t) \leq t_0 \\ (\sigma' \circ \tau)(t_0) & if\ t_0 < \sigma(t)\ and\ t < \tau(t_0) \\ \sigma'(t) & if\ \tau(t_0) \leq t\,. \end{cases}$$

The following diagram (with unlabeled arrows assumed to be the appropriate transition maps) shows the idea behind how the vector spaces of U and the transition maps $\varphi_U(s, t)$ for $s < \sigma' \circ \tau(t_0) \leq t$ fit with Definition 5.9.

$$\begin{tikzcd}
V_s \arrow[r] & V_{t_0} \arrow[rr, dashed, "\psi'_{\tau(t_0)} \circ \phi_{t_0}"] \arrow[dr, "\phi_{t_0}"] & & V'_{(\sigma' \circ \tau)(t_0)} \arrow[r] & V'_t \\
W_s \arrow[rr] & & W_{\tau(t_0)} \arrow[ur, "\psi'_{\tau(t_0)}"] \arrow[rr] & & W_t
\end{tikzcd}$$

Proof We start by showing that (η, ρ) is a translation pair, i.e. both functions are monotone and the functions $\rho \circ \eta$ and $\eta \circ \rho$ are translation maps. First we prove that η is monotone. By its definition and the monotonicity of the functions τ, τ', we get that η is monotone for $t < \sigma$

$\circ\ \tau(t_0)$ as well as for $\sigma \circ \tau(t_0) \le s$. Thus, we only have to show that $\eta(s) \le \eta(t)$ for any $s < \sigma \circ \tau(t_0) \le t$. This is true because

$$\eta(s) \le \tau(t_0) \le \tau' \circ \sigma' \circ \tau(t_0) \le \eta(t),$$

where the first and last inequality follow from definition of η and the middle one holds because the map $\tau' \circ \sigma'$ is a translation map. The proof of the monotonicity of ρ is based on the inequality $t_0 \le \sigma' \circ \tau(t_0)$, which holds by assumption.

The map $\rho \circ \eta$ is defined by $\sigma \circ \tau$ for t such that $\sigma \circ \tau(t) \le t_0$ and by $\sigma' \circ \tau'$ for $\sigma' \circ \tau(t_0) \le t$. Thus, we see that $\rho \circ \eta$ is a translation map for these values of t. For all the other values of t we have $\rho \circ \eta(t) = \sigma' \circ \tau(t_0)$. However, in this case $t < \sigma' \circ \tau(t_0)$, which proves that the map is a translation map on $\mathbb{R}$. We leave it to the reader to show that $\eta \circ \rho$ is a translation map as well.

In the rest of the proof we will show that the morphisms $\bar{\phi}: U \to W(\eta)$ and $\bar{\psi}: W \to U(\rho)$ defined by

$$\bar{\phi}_t = \begin{cases} \phi_t & \text{if } t \le t_0 \\ \phi_{t_0} & \text{if } t_0 < t < \sigma' \circ \tau(t_0) \\ \phi'_t & \text{if } \sigma' \circ \tau(t_0) \le t, \end{cases}$$
$$\bar{\psi}_t = \begin{cases} \psi_t & \text{if } \sigma(t) \le t_0 \\ \psi'_{\tau(t_0)} \circ \varphi_W(t, \tau(t_0)) & \text{if } t_0 < \sigma(t) \text{ and } t < \tau(t_0) \\ \psi'_t & \text{if } \tau(t_0) \le t, \end{cases}$$

give the desired interleaving of U and W. We leave it to the reader to check that the shifts in indices of $\bar{\phi}$ and $\bar{\psi}$ are consistent with the shifts given by the maps η and ρ.

Showing that $\bar{\phi}: U \to W(\eta)$ and $\bar{\psi}: W \to U(\rho)$ are persistence module morphisms is done by using the fact that ϕ, ϕ', ψ and ψ' are persistence module morphisms and unwrapping the definitions of $\bar{\phi}$, $\bar{\psi}$, η, ρ and φ_U based on the values $s \le t \in R$. While each case is a straightforward computation, there are many of them. Thus, we only show that $\bar{\phi}_t \circ \varphi_U(s,t) = \varphi_{W(\eta)}(s,t) \circ \bar{\phi}_s$ for $s < t_0 < \sigma' \circ \tau(t_0) \le t$ and leave the rest to the reader. In this case,

$$\begin{aligned}
\bar{\phi}_t \circ \varphi_U(s,t) &= \phi'_t \circ \varphi_{V'}(\sigma' \circ \tau(t_0), t) \circ \psi'_{\tau(t_0)} \circ \phi_{t_0} \circ \varphi_V(s, t_0) \\
&= \phi'_t \circ \varphi_{V'}(\sigma' \circ \tau(t_0), t) \circ \psi'_{\tau(t_0)} \circ \varphi_W(\tau(s), \tau(t_0)) \circ \phi_S \\
&= \varphi_W(\tau' \circ \sigma' \circ \tau(t_0), \tau'(t)) \circ \phi'_{(\sigma' \circ \tau)(t_0)} \circ \psi'_{\tau(t_0)} \circ \varphi_W(\tau(s), \tau(t_0)) \circ \phi_S \\
&= \varphi_W(\tau' \circ \sigma' \circ \tau(t_0), \tau'(t)) \circ \varphi_W(\tau(t_0), \tau' \circ \sigma' \circ \tau(t_0)) \circ \varphi_W(\tau(s), \tau(t_0)) \circ \phi_S \\
&= \varphi_W(\tau(s), \tau'(t)) \circ \phi_s \\
&= \varphi_{W(\eta)}(s,t) \circ \bar{\phi}_S\,.
\end{aligned}$$

□

We conclude this section with a simple, but hopefully illustrative, application of Proposition 5.11.

Corollary 5.12 *Let* (X, d) *be a finite metric space and let* $Y \subset X$ *be a* δ*-approximation of* X *with* $\delta > 0$. *Let* $U = U\left(M^{\mathscr{R}}(X), M^{\mathscr{R}}(Y), \tilde{\gamma}_{t_0*}\right)$ *be the stitched persistence module with* $\tilde{\gamma}$ *defined as in Lemma* 5.7*, and* $PD(U)$ *its persistence diagram. If* $[b, d, i] \in PD(U)$*, then neither* b *or* d *is in the interval* $(t_0, t_0 + \delta)$. *Moreover, there exists a matching* $\mathscr{X}: PD(U) \nrightarrow PD\left(M^{\mathscr{R}}(X)\right)$ *such that if* $\mathscr{X}([b, d, i]) = [b', d', i']$ *and*

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

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

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

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

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

Proof It follows form the definition of the Vietoris–Rips filtration that every interval in $\mathbf{PD}\left(M^{\mathscr{R}}(X)\right)$ and $\mathbf{PD}\left(M^{\mathscr{R}}(Y)\right)$ 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 $(M^{\mathscr{R}}(X), M^{\mathscr{R}}(X))$ are $(\tau(t) = t, \sigma(t) = t)$-interleaved and that $(M^{\mathscr{R}}(Y), M^{\mathscr{R}}(X))$ are $(\tau'(t) = t, \sigma'(t) = t + \delta)$-interleaved. The result follows by applying Propositions 5.11 and 4.6, with the additional observation that the identity map yielding $U_{t_0} = M^{\mathscr{R}}(X)_{t_0}$ and the definition of U forces that every point $[t_0, d, i] \in \mathrm{PD}(U)$ is matched to some point $[t_0, d', i'] \in \mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$ and vice versa. For the reader's benefit, we indicate the forms of η, ρ, $\eta_\star^\dagger$ and $\rho_\star^\dagger$ in Fig. 3. □

5.2.3 Iterated subsampling of a large point cloud

The goal of this section is to demonstrate that the techniques developed in Sects. 5.2.2 and 5.2.1can 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 5.13 Let (X, d) be a finite metric space. A sequence $\mathscr{Y} = \{Y_i \subseteq X\}_{i=0}^m$ 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 X.

Definition 5.14 Let $\mathscr{Y} = \{Y_i \subseteq X\}_{i=0}^{m}$ be a $\Delta = \{\delta_i > 0\}_{i=1}^{m}$ sampling a finite metric space (X, d). An *admissible stitching sequence* is a sequence $\mathscr{T} = \{t_i \geq 0\}_{i=1}^{m}$ satisfying

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

The associated *stitched Vietoris–Rips* persistence module

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

is defined inductively as follows:

$$\begin{aligned} &U_0 := M^{\mathscr{R}}(X), \\ &U_i := U\left(U_{i-1}, M^{\mathscr{R}}(Y_i), (\tilde{\gamma}_i)_{t_i *}\right), \text{ and} \\ &U := U_m, \end{aligned}$$

where $(\tilde{\gamma}_i)_{t_i *}$ is the map from $M^{\mathscr{R}}(X)_{t_i}$ to $M^{\mathscr{R}}(Y_i)_{t_i + \delta_i}$ induced at the level of homology.

Proposition 5.6(i) guarantees that $(M^{\mathscr{R}}(Y_i), M^{\mathscr{R}}(X))$ are $(t, t + \delta_i)$ interleaved. By repeated application of Proposition 5.11, we could obtain the precise interleavings between $U\left(\mathscr{Y}, \Delta, \mathscr{T}; M^{\mathscr{R}}(X)\right)$ and $M^{\mathscr{R}}(X)$ and provide a quantitative comparison of their persistence diagrams as we did in Corollary 5.12. The detailed bounds on the relationship between the persistence diagrams are rather complicated to state precisely. Thus, we limit ourselves to remarking that $U\left(\mathscr{Y}, \Delta, \mathscr{T}; M^{\mathscr{R}}(X)\right)$ and $M^{\mathscr{R}}(X)$ are roughly $(t, t + \delta_i)$ interleaved on the interval $(t_i, t_{i+1}]$. The only exceptions occur on the intervals $(t_{i+1} - \delta_i, t_{i+1} + \delta_{i+1})$ that overlap the stitching points, and in this case, the bounds on the errors are no worse than $\delta_i + \delta_{i+1}$.

In practice, the main constraint in computing persistence diagrams of a Vietoris–Rips filtration is the memory constraint associated with storing the Vietoris–Rips complex $\mathscr{R}(X, t)$. Thus, a desirable strategy is to: compute $M^{\mathscr{R}}(X)$ over a longest possible interval $[0, t_0]$; downsample to $Y_1 \subset X$; compute U_1 using $M^{\mathscr{R}}(Y_1)$ over a longest possible 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 Cech 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 4.6

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 4.7 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 Dey et al. (2014) 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).

Acknowledgements

R. L. would like to thank Charles Weibel, Michael Lesnick, and Ulrich Bauer for the many insightful discussions that led to the results presented in this paper. The authors also thank the anonymous reviewers for their suggested corrections. On behalf of all authors, the corresponding author states that there is no conflict of interest.

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. M. K. was supported by ERC Gudhi (ERC-2013-ADG-339025). R. L. was supported by DARPA contracts HR0011-17-1-0004 and HR0011-16-2-0033.

References

Bauer U, Lesnick M: Induced matchings of barcodes and the algebraic stability of persistence. In: Proceedings of the Thirtieth Annual Symposium Computational Geometry p. 355 (2014)

Bauer U, Lesnick M: Persistence diagrams as diagrams: A categorification of the stability theorem. (2016) arXiv:1610.10085

Botnan M, Spreemann G: Approximating persistent homology in euclidean space through collapses. Appl. Algebra Eng. Commun. Comput 26(1–2), 73–101 (2015). 10.1007/s00200-014-0247-y

Bubenik P, Scott JA: Categorification of persistent homology. Discrete Comput. Geom 51(3), 600–627 (2014). 10.1007/s00454-014-9573-x

Bubenik P, de Silva V, Scott J: Metrics for generalized persistence modules. Found. Comput. Math 15(6), 1501–1531 (2015). 10.1007/s10208-014-9229-5

Bubenik P, de Silva V, Scott J: Categorification of gromov-hausdorff distance and interleaving of functors (2017). arXiv:1707.06288

Buchet M, Chazal F, Oudot SY, Sheehy DR: Efficient and robust persistent homology for measures. Comput. Geom 58, 70–96 (2016). 10.1016/j.comgeo.2016.07.001

Chazal F, Cohen-Steiner D, Glisse M, Guibas LJ, Oudot SY: Proximity of persistence modules and their diagrams In: Proceedings of the twenty-fifth annual symposium on computational geometry, ACM, New York, NY, USA, SCG ‘09, pp. 237–246, (2009). 10.1145/1542362.1542407

Chazal F, de Silva V, Glisse M, Oudot S: The Structure and Stability of Persistence Modules (Springer-Briefs in Mathematics). Springer, Berlin (2016)

Cohen-Steiner D, Edelsbrunner H, Harer J: Stability of persistence diagrams. Discrete Comput. Geom 37(1), 103–120 (2007). 10.1007/s00454-006-1276-5

Crawley-Boevey W: Decomposition of pointwise finite-dimensional persistence modules. J. Algebra Appl. 14(05), 1550066 (2015). 10.1142/S0219498815500668

Davey BA, Priestley HA: Introduction to Lattices and Order, p. xii+298 Cambridge University Press, Cambridge (2002)

Dey TK, Fan F, Wang Y: Graph induced complex on point data In: Proceedings of the twenty-ninth annual symposium on computational geometry, ACM, New York, NY, USA, SoCG ‘13, pp 107–116, 10.1145/2462356.2462387 (2013)

Dey TK, Fan F, Wang Y: Computing topological persistence for simplicial maps In: Proceedings of the thirtieth annual symposium on computational geometry, ACM, New York, NY, USA, SOCG’14, pp 345:345–345:354, (2014) 10.1145/2582112.2582165

Edelsbrunner H, Harer JL: Computational Topology : an Introduction. American Mathematical Society, Providence (2010)

Friedman G: Survey article: an elementary illustrated introduction to simplicial sets. Rocky Mountain J. Math 42(2), 353–423 (2012). 10.1216/RMJ-2012-42-2-353

Kramár M, Levanger R, Tithof J, Suri B, Xu M, Paul M, Schatz MF, Mischaikow K: Analysis of Kolmogorov flow and Rayleigh–Bénard convection using persistent homology. Phys D 334, 82–98 (2016)

Oudot SY: Persistence Theory: from Quiver Representations to Data Analysis, Mathematical Surveys and Monographs, vol. 209 American Mathematical Society, Providence (2015)

Sheehy D: Linear-size approximations to the vietoris? rips filtration. Discrete Comput. Geom 49(4),778–796 (2013). 10.1007/s00454-013-9513-1

Weibel C: An Introduction to Homological Algebra Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge (1995)

Zomorodian A, Carlsson G: Computing persistent homology. Discrete Comput. Geom 33(2), 249–274 (2004). 10.1007/s00454-004-1146-y

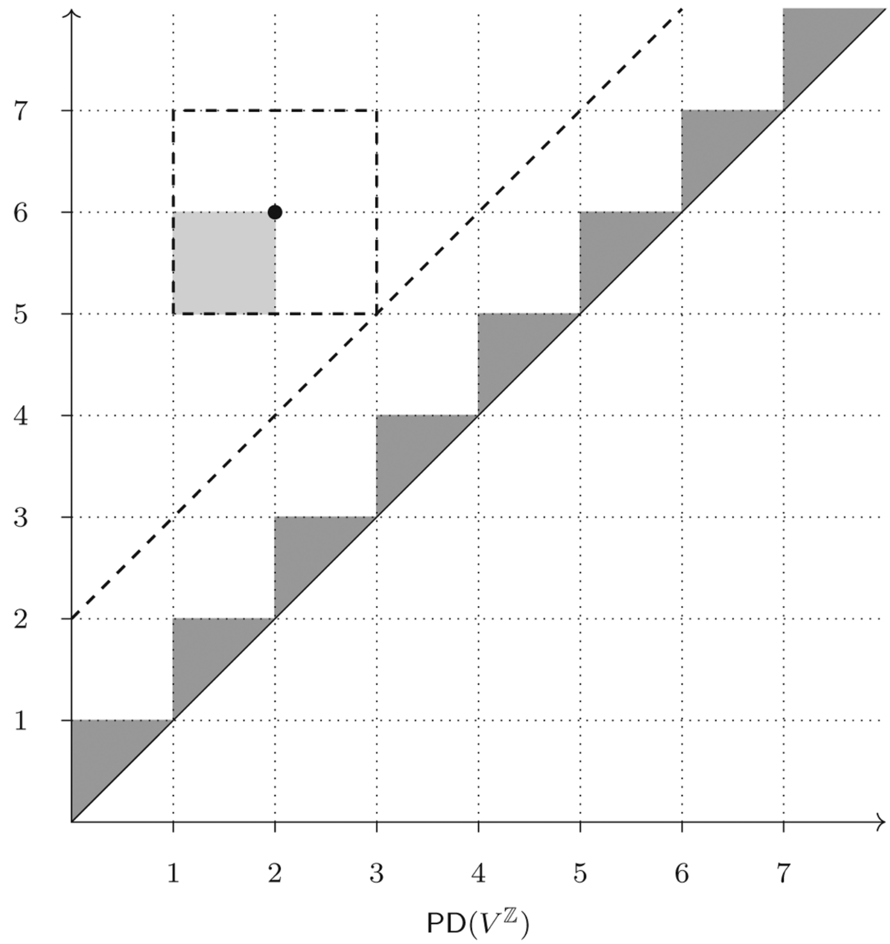

Fig. 1.

A schematic diagram illustrating the potential locations of persistence points of PD(V) based on the computation of PD$\left(V^{\mathbb{Z}}\right)$. The persistence point (2, 6) of PD$\left(V^{\mathbb{Z}}\right)$ is matched with a persistence point of PD(V) which must lie strictly within the light gray region. The dark gray region indicates the potential location of persistence points of PD(V) that cannot be detected because of the integer-valued approximation used to compute PD$\left(V^{\mathbb{Z}}\right)$. Finally, if W is an arbitrary PFD persistence module and the bottleneck distance between PD$\left(V^{\mathbb{Z}}\right)$ and PD(W) is one, then 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

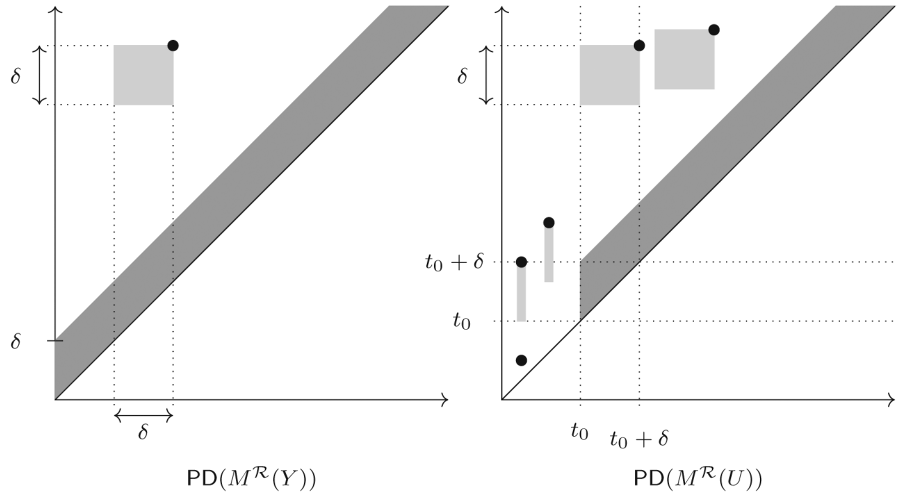

Fig. 2.

(left) Schematic diagram illustrating the quality of the matching from Proposition 5.6. The persistence diagram $\mathrm{PD}\left(M^{\mathscr{R}}(Y)\right)$ is shown. The dark gray region indicates the possible locations of the unmatched points for the persistence diagrams $\mathrm{PD}\left(M^{\mathscr{R}}(Y)\right)$ and $\mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$. The light gray region indicates the possible location of the point of $\mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$ that is matched to the point of $\mathrm{PD}\left(M^{\mathscr{R}}(Y)\right)$ that is shown. (right) Schematic diagram illustrating the quality of the matching from Corollary 5.12. The persistence diagram $\mathrm{PD}\left(M^{\mathscr{R}}(U)\right)$ is shown. Persistence points for the persistence diagrams $\mathrm{PD}\left(M^{\mathscr{R}}(U)\right)$ and $\mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$ agree in the region $[0, t_0) \times [0, t_0)$. Unmatched points for the persistence diagrams $\mathrm{PD}\left(M^{\mathscr{R}}(U)\right)$ and $\mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$ will lie in the dark gray region. The light gray region indicates the possible location of the point of $\mathrm{PD}\left(M^{\mathscr{R}}(X)\right)$ that is matched to the point of $\mathrm{PD}\left(M^{\mathscr{R}}(U)\right)$ that is shown

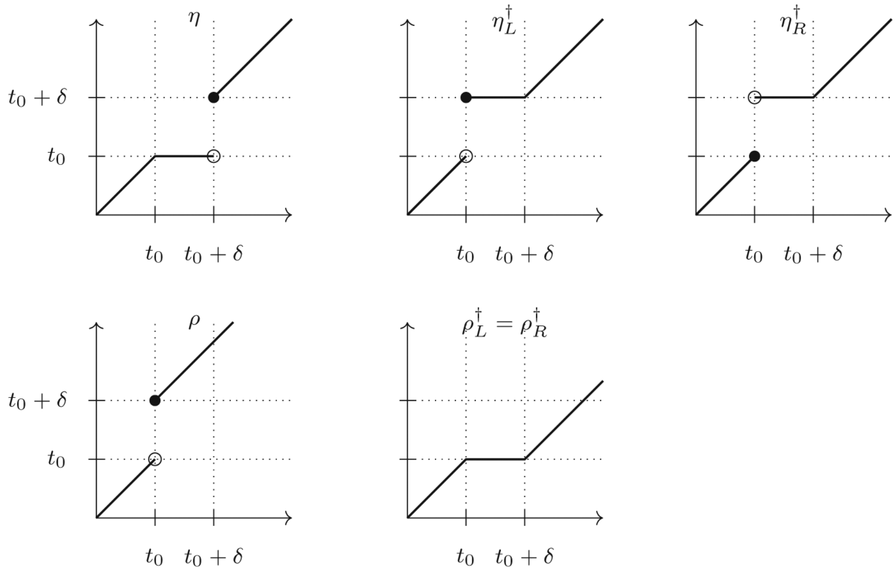

Fig. 3.

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

Table 1

A table of approximations for Vietoris–Rips and Čech filtrations for point clouds

Approximate complex	Complex of interest	$\tau(t)$	$\sigma(t)$
Point clouds in $\mathbb{R}^n$			
Vietoris-Rips (Edelsbrunner and Harer 2010, Proof of Vietoris-Rips Lemma.)	Čech	$t\sqrt{\frac{2n}{n+1}}$	t
Net-tree (Botnan and Spreemann 2015, Proof of Proposition 20)	Čech	t	$(1+\varepsilon)^2 t$
Graph induced complex (Dey et al. 2013, Proposition 2.8)	Vietoris-Rips	$t+2\varepsilon$	t
Sparsified Vietoris-Rips (Dey et al. 2014, Claim 6.1)	Vietoris-Rips	t	$(1+\varepsilon)t$
Point clouds in an arbitrary metric space			
Vietoris-Rips	Čech filtration	$2t$	t
Relaxed Vietoris-Rips (Sheehy 2013, Lemma 4)	Vietoris-Rips	t	$\left(\frac{1}{1-2\varepsilon}\right)t$
Sparse weighted Rips (Buchet et al. 2016, Lemma 6.13)	Vietoris-Rips	t	$\left(\frac{1+\sqrt{1+\delta^2\varepsilon}}{1-\varepsilon}\right)t$

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