



Persistent extensions and analogous bars: data-induced relations between persistence barcodes

Hee Rhang Yoon¹ · Robert Ghrist² · Chad Giusti³

Received: 5 February 2022 / Revised: 9 October 2022 / Accepted: 13 February 2023 /

Published online: 15 April 2023

© The Author(s) 2023

Abstract

A central challenge in topological data analysis is the interpretation of barcodes. The classical algebraic-topological approach to interpreting homology classes is to build maps to spaces whose homology carries semantics we understand and then to appeal to functoriality. However, we often lack such maps in real data; instead, we must rely on a cross-dissimilarity measure between our observations of a system and a reference. In this paper, we develop a pair of computational homological algebra approaches for relating persistent homology classes and barcodes: *persistent extension*, which enumerates potential relations between homology classes from two complexes built on the same vertex set, and the method of *analogous bars*, which utilizes persistent extension and the witness complex built from a cross-dissimilarity measure to provide relations across systems. We provide an implementation of these methods and demonstrate their use in comparing homology classes between two samples from the same metric space and determining whether topology is maintained or destroyed under clustering and dimensionality reduction.

Keywords Persistent homological algebra · Persistence modules · Witness complex · Homology class relations

✉ Hee Rhang Yoon
irishryoon@gmail.com

Robert Ghrist
ghrist@math.upenn.edu

Chad Giusti
chad.giusti@oregonstate.edu

¹ Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK

² Department of Mathematics and Electrical and Systems Engineering, University of Pennsylvania, 3451 Walnut Street, Philadelphia, PA 19104, USA

³ Department of Mathematics, Oregon State University, Kidder Hall 368, 2000 SW Campus Way, Corvallis, Oregon 97331–4605, USA

Mathematics Subject Classification 55N31 · 68T09

1 Introduction

Persistent homology (Edelsbrunner et al. 2000) measures how structure (encoded as *homology classes*) varies with the parameter in a parameterized space. In principle, to apply this tool to data, we design parameterized combinatorial encodings in which the presence or absence of homology classes describes features of interest. In practice, it is prohibitively difficult and time-consuming to construct novel combinatorial structures which would be easy to interpret in terms of our data sets. Rather, we employ standard encodings for which theory and computational tools are readily available.

Currently, the most common approach is to take as input a pairwise dissimilarity measure M_P for system constituents $P = \{p_1, \dots, p_n\}$ and to construct from this data a *weighted clique complex* $X_P^\bullet = X(M_P)^\bullet$; vertices correspond to the elements in the system and simplices are weighted by the maximum pairwise dissimilarity of their vertices. When M_P is a metric, this produces the usual weighted *Vietoris-Rips complex* for a point cloud. We then employ computational homological algebra to determine the complex's k th persistent reduced¹ homology, $P\tilde{H}_k(X_P^\bullet)$. Most current software returns this information in the form of a *barcode*, $\text{BC}_k(X_P^\bullet)$, which encodes the homology classes and their *birth* and *death* parameters.

We are then left with the problem of interpreting the barcode. This is difficult in part due to the intricate structure of the weighted clique complexes, including trade-offs between fidelity of representations and their combinatorial and computational complexity, and in part due to the abstract nature of persistent homology classes, made up of sequences of affine subspaces of quotient vector spaces. As such, most modern applications of persistent homology involve vectorization of the barcode as a statistic for differentiating classes of data, or *ad hoc* interpretations of classes.

Success stories from classical algebraic topology suggest that to understand homology classes we should appeal to functoriality. By constructing maps to or from reference objects which are better understood, topologists have developed *semantics* for (co)homology classes in a range of contexts. Perhaps the best known classical examples are classifying maps and characteristic classes for vector bundles (Milnor 1956), which have been adapted to provide nonlinear coordinate systems for data sets (Perea 2020).

In the context of applied topology, our reference often comes in the form of more data: observations of some collection $Q = \{q_1, \dots, q_m\}$ of inputs to, outputs from, or known correlates for activity in the system P . From a dissimilarity matrix M_Q derived from these data Q , we can obtain another combinatorial space and compute a barcode $\text{BC}_k(X_Q^\bullet)$. When this second data set is accessible, experimentally or theoretically, we are better equipped to assign semantics to classes in $P\tilde{H}_k(X_P^\bullet)$. However, *a priori* there is insufficient information to construct the maps we need to apply functoriality, and

¹ All complexes in this paper are non-empty. To simplify statements of results, we will always use reduced homology, which we denote \tilde{H}_k , and we will omit the word *reduced* to avoid clutter.

we are left with the problem of determining whether and how classes from $P\tilde{H}_k(X_Q^\bullet)$ correspond to classes in $P\tilde{H}_k(X_P^\bullet)$.

To do so, we require some notion of how P and Q are related. In scientific and engineering applications, a common and effectively minimal way to satisfy this requirement is through a measure $M_{P,Q}$ of cross-dissimilarity between P and Q . For parameterized clique complexes, $M_{P,Q}$ measures dissimilarity between the vertices of X_P^\bullet and X_Q^\bullet . It is tempting to construct a Vietoris-Rips complex on the joint dissimilarity measure on $P \cup Q$ given by combining M_P , M_Q , and $M_{P,Q}$ and then to apply recently developed methods, including induced matching (Bauer and Lesnick 2014, 2020), cycle registration (Reani and Bobrowski 2021), or basis-independent partial matching (Gonzalez-Díaz and Soriano-Trigueros 2020), to match classes via the zig-zag of induced inclusion maps. However, as we will discuss in Sect. 3, even in simple cases involving pairs of point clouds in the same metric space this approach can fail to produce matches which satisfy our intuition.

The alternative we develop here is to observe that $M_{P,Q}$ is precisely the data of a weighted *witness complex*² (Dowker 1952; de Silva and Carlsson 2004). The witness complex $W_{P,Q}^\bullet$ has vertices P and faces indexed by Q : the face q_c at parameter ℓ has vertices $\{p_r \in P : (M_{P,Q})_{r,c} \leq \ell\}$. The Functorial Dowker Theorem (Dowker 1952; Chowdhury and Mémoli 2018) provides an explicit isomorphism $P\tilde{H}_k(W_{P,Q}^\bullet) \cong P\tilde{H}_k(W_{Q,P}^\bullet)$, and thus our cross-dissimilarity matrix provides a bridge between persistent homology classes in simplicial complexes with vertices P and Q .

To apply this information, in Sect. 4 we introduce *persistent extension*, a general method for comparing persistent homology classes between complexes Z^\bullet and Y^\bullet supported on the same vertex set. Suppose Y^ϵ is contractible for large ϵ , and fix some parameter ψ of Z^\bullet . Persistent extension takes as input a class $[\tau] \in \tilde{H}_k(Z^\psi)$ and outputs an enumeration of persistent homology classes in $P\tilde{H}_k(Y^\bullet)$ that contain representatives of $[\tau]$. To do so, we observe that classes of infinite persistence in the *auxiliary* filtered complex $Z^\psi \cap Y^\bullet$ are precisely the elements of $\tilde{H}_k(Z^\psi)$. We then leverage a zig-zag through the auxiliary persistence module to set up systems of linear equations that describe the image in $P\tilde{H}_k(Y^\bullet)$ of all elements of $P\tilde{H}_k(Z^\psi \cap Y^\ell)$ which map to $[\tau]$ as $\ell \rightarrow \infty$. As it is common to speak of persistent homology in terms of the barcode, in Sect. 4.2 we develop the necessary tools to translate between the persistence module-level information and the barcodes.

In Sect. 5, we combine the isomorphism from Dowker’s theorem with persistent extension to obtain a pair of processes for relating persistent homology classes in $P\tilde{H}_k(X_Q^\bullet)$ to those in $P\tilde{H}_k(X_P^\bullet)$. The first, the *feature-centric analogous bars* method, focuses on an individual bar $\tau \in \text{BC}_k(X_Q^\bullet)$ and enumerates subsets of $\text{BC}_k(X_P^\bullet)$ which potentially correspond to τ given the structure of $M_{P,Q}$. The *similarity-centric analogous bars* method, on the other hand, focuses on significant features in $\text{BC}_k(W_{P,Q}^\bullet)$, which intuitively describe very strong relationships between subsets of bars in $\text{BC}_k(X_Q^\bullet)$ and $\text{BC}_k(X_P^\bullet)$ in the same way long lifetimes in Vietoris-Rips

² We use a slightly different notion of witness complex than the one introduced in de Silva and Carlsson (2004). Our witness complex would be identical to that of de Silva and Carlsson (2004) if we assign a birth time of 0 to all vertices and take $v = 0$ in their definition. Our witness complex is a filtration of the Dowker complex.

complexes of point clouds intuitively correspond to significant features of the support of the underlying distribution.

In this paper, we assume that all barcodes have unique birth-death pairs. This assumption allows us to efficiently enumerate all possible basis of persistent homology. Such enumeration enables the persistent extension and analogous bars methods to be independent of the choice of basis. As discussed in Sect. 4.2.3, the additional assumptions of unique death parameters and \mathbb{F}_2 coefficients substantially simplifies the algorithm by removing one of the enumeration processes.

This paper is organized as follows. In Sect. 2, we recall the relevant algebraic and topological notions and set the notations and assumptions, and in Sect. 3, we discuss specific difficulties that arise when attempting to match bars between simplicial complexes, including a discussion of existing methods. Having established these preliminaries, in Sects. 4 and 5, we provide the details of the methods of persistent extension and analogous bars and demonstrate their use on simple examples. In Sect. 6 we demonstrate some simple applications: using analogous bars to compare the barcodes of two samples from the same distribution on a metric space, and using persistent extension to determine if topological structures are preserved under clustering and dimensionality reduction. Finally, in Sect. 7 we discuss context and future directions. To improve readability, we defer technical proofs to appendices.

1.1 Contributions

The principal contributions of this paper are:

- the *persistent extension* method, which compares persistent homology classes and barcodes between filtered simplicial complexes built on a common vertex set,
- the *analogous bars* method, which compares persistent homology classes and barcodes between two distinct clique complexes using a cross-dissimilarity measure,
- implementation of these methods built using the Eirene persistent homology package (Henselman and Ghrist 2016), available at https://github.com/UDATG/analogous_bars, and
- demonstration on some toy examples, along with a discussion of potential applications.

Remark 1 While the authors were writing this paper, independent work (Jacquard et al. 2022) investigating the structure of the set of bases for persistence barcodes appeared. While the fundamental aims of the two papers are different, Jacquard et al. (2022) establishes results that generalize our Lemmas 2.2.5, 2.2.6, and 4.2.3, in the course of developing a more complete picture of the relationship between persistence modules and barcodes. We retain our versions of these results in this manuscript for completeness and to save the reader effort in translating their work to our language, notation, and perspective.

2 Preliminaries

We begin by recalling relevant definitions and results from persistent homology. Along the way we will introduce some new terminology and results about change of basis of a persistence module. We assume readers are familiar with the general theory of persistent homology. A more complete development of this material can be found in, for example, Edelsbrunner et al. (2000); Zomorodian and Carlsson (2005); Carlsson (2009); Ghrist (2008). Fix a field \mathbb{F} throughout.

2.1 Primer on persistent homology

Definition 2.1.1 A persistence module is a \mathbb{Z} -graded \mathbb{F} -vector space $V^\bullet = \bigoplus_{\ell=-\infty}^{\infty} V^\ell$ equipped with linear *structure maps* $\{\phi^\ell : V^\ell \rightarrow V^{\ell+1}\}$, thought of as a $\mathbb{F}[x]$ -module in which x acts on elements of V^ℓ via ϕ^ℓ . We say a persistence module has *finite support* if $V^\ell = 0$ for all ℓ outside some compact interval.

We will omit the field from the notation and simply say “persistence module” throughout. Similarly, we will usually refer to a persistence module by the name of its underlying graded vector space, leaving the structure maps implicit. Finally, as all of our persistence modules have finite support, we will omit vector spaces and maps outside of the support from diagrams without further comment.

The prototypical examples of persistence modules are the *interval modules*.

Definition 2.1.2 For $\beta, \delta \in \mathbb{Z}$, the *interval module* $I[\beta, \delta)^\bullet$ is the persistence module $\bigoplus_{\ell=-\infty}^{\infty} I[\beta, \delta)^\ell$, where

$$I[\beta, \delta)^\ell \cong \begin{cases} \mathbb{F} & \beta \leq \ell < \delta \\ 0 & \text{else} \end{cases},$$

along with structure maps $\phi^\ell : I[\beta, \delta)^\ell \rightarrow I[\beta, \delta)^{\ell+1}$ for all $\ell \in \mathbb{Z}$ given by

$$\phi^\ell(f) = \begin{cases} f & \ell \neq \beta - 1, \delta \\ 0 & \text{else} \end{cases}.$$

Our primary example of interest arises in the context of filtered topological spaces.

Definition 2.1.3 Let

$$Z^\bullet = Z^1 \xrightarrow{t^1} Z^2 \xrightarrow{t^2} \dots \xrightarrow{t^{N-2}} Z^{N-1} \xrightarrow{t^{N-1}} Z^N$$

be a filtered topological space. The *degree- k reduced persistent homology* of Z^\bullet is the persistence module $P\tilde{H}_k(Z^\bullet)$ with

$$(P\tilde{H}_k(Z^\bullet))^\ell = \begin{cases} \tilde{H}_k(Z^\ell; \mathbb{F}) & 1 \leq \ell \leq N \\ 0 & \text{else} \end{cases}$$

with structure maps given by the induced maps on reduced homology ι_*^ℓ where applicable and zero maps elsewhere.

A fundamental result in the study of persistent homology tells us that, for finite data, persistent homology can be decomposed as a collection interval modules.

Theorem 2.1.4 (Interval decomposition for persistent homology (Zomorodian and Carlsson 2005)) *Let Z^\bullet be a filtered finite simplicial complex and k a non-negative integer. Then there exists a triple*

$$(\text{BC}_k(Z^\bullet), \beta, \delta)$$

called the barcode of the persistence module $P\tilde{H}_k(Z^\bullet)$, where

1. $\text{BC}_k(Z^\bullet)$ is a finite set whose elements we call bars,
2. $\beta, \delta : \text{BC}_k(Z^\bullet) \rightarrow \mathbb{Z}$ are functions that respectively record the birth and death parameters of each bar $\tau \in \text{BC}_k(Z^\bullet)$, so that
3. $\beta(\tau) < \delta(\tau)$ for all $\tau \in \text{BC}_k(Z^\bullet)$,

that is unique up to isomorphism,³ along with an isomorphism of persistence modules

$$\mathcal{B} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \xrightarrow{\cong} P\tilde{H}_k(Z^\bullet),$$

where $\mathbb{I}_{\text{BC}_k(Z^\bullet)} = \bigoplus_{\tau \in \text{BC}_k(Z^\bullet)} I[\beta(\tau), \delta(\tau)]$ is the barcode module for Z^\bullet in dimension k . This isomorphism is called an interval decomposition of $P\tilde{H}_k(Z^\bullet)$.

By abuse, we will usually refer to the barcode $(\text{BC}_k(Z^\bullet), \beta, \delta)$ using only the name of its underlying set, $\text{BC}_k(Z^\bullet)$. Further, we will often refer to the collection of bars which are *alive* at a particular parameter $\ell \in \mathbb{Z}$ by

$$\text{BC}_k(Z^\ell) = \{\tau \in \text{BC}_k(Z^\bullet) \mid \beta(\tau) \leq \ell < \delta(\tau)\}.$$


The *length* or *lifetime* of a bar $\tau \in \text{BC}_k(Z^\bullet)$ is $\delta(\tau) - \beta(\tau)$. Further, for a filtered topological space with N filtration levels, when $\delta(\tau) = N + 1$, it is common to write instead $\delta(\tau) = \infty$ and say that τ has *infinite* length or lifetime. We will require this convention in Sect. 4.1 when we discuss a filtration for which the final level is not a contractible space.

While the barcode $\text{BC}_k(Z^\bullet)$ is unique, the interval decomposition is not. For example, let Z^\bullet denote the filtered simplicial complex in Figure 1. In dimension 1, the corresponding barcode module $\mathbb{I}_{\text{BC}_1(Z^\bullet)}$ and the persistence module $P\tilde{H}_1(Z^\bullet)$ are

$$\mathbb{F} \xrightarrow{\begin{bmatrix} 1 \\ 0 \end{bmatrix}} \mathbb{F}^2 \xrightarrow{\begin{bmatrix} 0 & 1 \end{bmatrix}} \mathbb{F}.$$

³ That is, set isomorphism on $\text{BC}_k(Z^\bullet)$ along with precomposition by the inverse of that isomorphism for β and δ .

Fig. 1 A filtered finite simplicial complex



The following illustrates two possible interval decompositions \mathcal{B} and \mathcal{C} .

$$\begin{array}{ccccc}
 \mathbb{I}_{\text{BC}_k(Z^\bullet)} : \mathbb{F} & \longrightarrow & \mathbb{F}^2 & \longrightarrow & \mathbb{F} \\
 \downarrow \mathcal{B} & & \downarrow \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} & & \downarrow \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \\
 P\tilde{H}_k(Z^\bullet) : \mathbb{F} & \longrightarrow & \mathbb{F}^2 & \longrightarrow & \mathbb{F}
 \end{array}
 \qquad
 \begin{array}{ccccc}
 \mathbb{I}_{\text{BC}_k(Z^\bullet)} : \mathbb{F} & \longrightarrow & \mathbb{F}^2 & \longrightarrow & \mathbb{F} \\
 \downarrow \mathcal{C} & & \downarrow \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} & & \downarrow \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \\
 P\tilde{H}_k(Z^\bullet) : \mathbb{F} & \longrightarrow & \mathbb{F}^2 & \longrightarrow & \mathbb{F}
 \end{array}$$

Bars and barcodes are the central actors in our story. Throughout this paper, we will carefully think about how bars correspond to homology classes. In particular, it will be important to think about how the vector space structure on homology is reflected by the bars in a barcode. To that end, we require the following terminology regarding persistence modules.

Definition 2.1.5 Let V^\bullet be a persistence module with structure maps $\{\phi^\ell\}$. A *persistent subspace* (respectively, *persistent affine subspace*) of V^\bullet is given by a choice of $\beta < \delta \in \mathbb{Z}$ and $W^\bullet = \bigoplus_{\ell=-\infty}^{\infty} W^\ell$, where

1. $W^\ell \subseteq V^\ell$ is a subspace (respectively, affine subspace),
2. $W^\ell = 0$ if $\ell < \beta$ or $\ell \geq \delta$, and
3. $\phi^\ell(W^\ell) = W^{\ell+1}$ if $\ell \neq \beta - 1, \delta - 1$ ⁴.

That is, a persistent (affine) subspace of V^\bullet is a choice of (affine) subspaces of each V^ℓ that are consistent with the structure maps and supported on a bounded interval⁵. If $\dim(W^\ell) = d$ for some fixed d whenever $\beta \leq \ell < \delta$, we say W^\bullet is a *persistent d -dimensional subspace*.

Observe that given a persistence module $P\tilde{H}_k(Z^\bullet)$, a bar $\tau \in \text{BC}_k(Z^\bullet)$, and an interval decomposition $\mathcal{B} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$, the image of the corresponding interval module $I[\beta(\tau), \delta(\tau)]$ under \mathcal{B} is a persistent 1-dimensional subspace of $P\tilde{H}_k(Z^\bullet)$. Thus, another way to describe an interval decomposition is as a choice of identification of bars in the barcode with persistent 1-dimensional subspaces in the persistence module. Before we move on, it will be useful to take a closer look at what this choice entails.

⁴ The third condition distinguishes a persistent subspace from a submodule (Carlsson and de Silva 2010).

⁵ A persistent subspace is a submodule of the persistence module V^\bullet .

2.2 The relationship between bars and homology classes

Theorem 2.1.4 describes $P\tilde{H}_k(Z^\bullet)$ as a finitely generated $\mathbb{F}[x]$ -module by providing a minimal set of generators for the module and their annihilators. While the barcode which enumerates these pairs of births and deaths is unique, the explicit identification of these generators as elements of the module is usually quite the opposite. Indeed, this is precisely a choice of basis for a finite-dimensional vector space that is compatible with the structure maps of the persistence module. Since we need cycle representatives, we must check that our constructions are invariant under this choice. Thus, we need analogues of the usual change of basis formalism from linear algebra.

Definition 2.2.1 Let Z^\bullet be a filtered topological space, and let $BC_k(Z^\bullet)$ and $\mathcal{B} : \mathbb{I}_{BC_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be the barcode and interval decomposition for $P\tilde{H}_k(Z^\bullet)$, made up of the following commutative diagram of vector spaces and linear isomorphisms

$$\begin{CD}
 \mathbb{I}_{BC_k(Z^\bullet)} @>>> I^1 @>>> \dots @>>> I^\ell @>>> \dots @>>> I^N \\
 @VV\mathcal{B}V @VV\mathcal{B}^1V @. @VV\mathcal{B}^\ell V @. @VV\mathcal{B}^N V \\
 P\tilde{H}_k(Z^\bullet) @>>> \tilde{H}_k(Z^1) @>>> \dots @>>> \tilde{H}_k(Z^\ell) @>>> \dots @>>> \tilde{H}_k(Z^N)
 \end{CD} \tag{1}$$

where I^ℓ is a direct sum of copies of \mathbb{F} indexed over $BC_k(Z^\ell)$,

$$I^\ell = \bigoplus_{BC_k(Z^\ell)} \mathbb{F}.$$

For each $\ell \in \mathbb{Z}$, let $\{e_\tau^\ell : \tau \in BC_k(Z^\ell)\}$ be the standard basis of I^ℓ , where e_τ^ℓ is the standard basis vector corresponding to the τ -summand. We call e_τ^ℓ the *basis vector for τ at ℓ* .

Conceptually, we do not require an ordering on the bars and adding one requires another layer of notation that we would like to avoid. However, when we are working with matrix computations, it will be convenient to set a linear order $\tau_1 < \tau_2 < \dots < \tau_{|BC_k(Z^\ell)|}$ on the bars alive at parameter ℓ . In this case, $e_{\tau_r}^\ell$ is the vector whose r^{th} component is 1 and all other components are zero.

We can now set some terminology for the correspondence between bars in $BC_k(Z^\bullet)$ and classes in $P\tilde{H}_k(Z^\bullet)$.

Definition 2.2.2 Let Z^\bullet be a filtered topological space, and let $BC_k(Z^\bullet)$ and $\mathcal{B} : \mathbb{I}_{BC_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be the barcode and interval decomposition of $P\tilde{H}_k(Z^\bullet)$. Let $\ell \in \mathbb{Z}$ and $\tau \in BC_k(Z^\ell)$. Take e_τ^ℓ to be the basis vector for τ at ℓ . Write $[\tau^{\mathcal{B},\ell}] = [\mathcal{B}^\ell(e_\tau^\ell)] \in \tilde{H}_k(Z^\ell)$ for the homology class that \mathcal{B} -corresponds to τ at parameter ℓ and say that any cycle representative of $[\tau^{\mathcal{B},\ell}]$ is a \mathcal{B} -representative of the bar τ at parameter ℓ .

Definition 2.2.2 finds the homology class that corresponds to a given bar. Conversely, we can also find bar representations of a homology class. To that end, we first define the notion of a weighted bar.

Definition 2.2.3 A *weighted bar* is an element in $\mathbb{F} \times \text{BC}_k(Z^\bullet)$.

We denote a weighted bar (c, τ) by $c\tau$ for brevity. A bar representation is a collection of weighted bars as defined in the following.

Definition 2.2.4 Let Z^\bullet be a filtered topological space, and let $\text{BC}_k(Z^\bullet)$ and $\mathcal{B} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be the barcode and interval decomposition of $P\tilde{H}_k(Z^\bullet)$. Given any homology class $[z] \in \tilde{H}_k(Z^\ell)$, we say that the collection

$$S_{[z]}^{\mathcal{B}} = \{c_1\tau_1, \dots, c_K\tau_K \mid \tau_1, \dots, \tau_K \in \text{BC}_k(Z^\ell), c_1, \dots, c_K \in \mathbb{F} \setminus \{0\}, [z] = \sum_{i=1}^K c_i[\tau_i^{\mathcal{B}, \ell}]\}$$

is the $(\mathcal{B}$ -)bar representation of $[z]$ (at ℓ).

We can visually represent the persistent 1-dimensional subspace generated by a \mathcal{B} -bar representation $S_{[z]}^{\mathcal{B}} = \{c_1\tau_1, \dots, c_K\tau_K\}$ by highlighting the right half-bars of $\tau_1, \dots, \tau_K \in \text{BC}_k(Z^\bullet)$ starting from parameter ℓ (see Figs. 9 and 10 in Sect. 6 for such visual depictions). When we work with \mathbb{F}_2 coefficients where bar representations are in correspondence with various subsets of the bars, this visual representation provides complete information.

The correspondence between bars and homology classes in Definitions 2.2.2 and 2.2.4 depend on the specific interval decomposition \mathcal{B} . To avoid choice-dependent constructions, we consider all homology classes corresponding to τ under different interval decompositions of $P\tilde{H}_k(Z^\bullet)$. As noted, interval decompositions are vector space isomorphisms that are compatible with the structure maps. Thus, any interval decomposition \mathcal{C} of $P\tilde{H}_k(Z^\bullet)$ can be obtained from a given interval decomposition \mathcal{B} by the composition $\mathcal{C} = \mathcal{B} \circ \mathcal{L}$ for some automorphism $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$ of the barcode module, given by the following commutative diagram of vector spaces and linear isomorphisms.

$$\begin{CD} \mathbb{I}_{\text{BC}_k(Z^\bullet)} @>I^1>> \dots @>I^\ell>> \dots @>I^N>> I^N \\ @V\mathcal{L}VV @VV\mathcal{L}^1V @VV\mathcal{L}^\ell V @VV\mathcal{L}^N V \\ \mathbb{I}_{\text{BC}_k(Z^\bullet)} @>I^1>> \dots @>I^\ell>> \dots @>I^N>> I^N \end{CD} \tag{2}$$

We will require the following two technical lemmas⁶ to work with the set of all such automorphisms in Sect. 4.2. As their statements involve matrices, we fix linear orderings on each $\text{BC}_k(Z^\ell)$. Let L^ℓ denote the corresponding matrix for the linear automorphisms \mathcal{L}^ℓ , and let $L_{r,c}^\ell$ indicate the matrix entry at row r and column c .

Lemma 2.2.5 Given a filtered topological space Z^\bullet , let $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$ be an endomorphism of its barcode module. Let $\ell \in \mathbb{Z}$, and let $\tau_1, \dots, \tau_{|\text{BC}_k(Z^\ell)|}$ be some

⁶ Per Remark 1, a generalized version of these results can be found in Jacquard et al. (2022).

ordering of the bars in $\text{BC}_k(Z^\ell)$. If $L_{r,c}^\ell \neq 0$ for $r \neq c$, then the bars τ_r and τ_c must satisfy the relations

$$\beta(\tau_r) \leq \beta(\tau_c) < \delta(\tau_r) \leq \delta(\tau_c). \tag{3}$$

Proof Proof can be found in Proposition 16 of Bubenik and Vergili (2018). □

Lemma 2.2.6 *Given a filtered topological space Z^\bullet , assume that if $\tau \neq \tau' \in \text{BC}_k(Z^\bullet)$ then either $\beta(\tau) \neq \beta(\tau')$ or $\delta(\tau) \neq \delta(\tau')$. Let $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$ be an automorphism of the barcode module. Then $L_{j,j}^\ell \neq 0$ for every $1 \leq j \leq |\text{BC}_k(Z^\ell)|$ and $\ell \in \mathbb{Z}$.*

Proof Deferred to Appendix A.1. □

In Sect. 4.2, we will enumerate the matrices of isomorphisms \mathcal{L}^ℓ that extend to a barcode module automorphism \mathcal{L} . The following lemma will be useful.

Lemma 2.2.7 *Let $\text{BC}_k(Z^\bullet)$ be the barcode of some persistence module $P\tilde{H}_k(Z^\bullet)$. Assume that $\beta(\tau) \neq \beta(\tau')$ or $\delta(\tau) \neq \delta(\tau')$ for all $\tau \neq \tau' \in \text{BC}_k(Z^\bullet)$. Given $\ell \in \mathbb{Z}$, let $\tau_1, \dots, \tau_{|\text{BC}_k(Z^\ell)|}$ be the bars in $\text{BC}_k(Z^\ell)$. For $1 \leq t \leq |\text{BC}_k(Z^\ell)|$, let $L \in \mathbb{F}^{|\text{BC}_k(Z^\ell)| \times t}$ be any matrix whose entries satisfy*

$$L_{r,r} \neq 0 \text{ and } L_{r,c} = 0 \text{ if } \beta(\tau_r) > \beta(\tau_c) \text{ or } \delta(\tau_r) > \delta(\tau_c) \text{ for all } r, c. \tag{4}$$

Then, the columns of L are linearly independent.

Proof Deferred to Appendix 1. □

To apply the above lemmas, we will need to assume that all bars have unique birth-death pairs. To ensure that the enumeration of \mathcal{L}^ℓ is finite, we assume that we are working over a finite field. If we further restrict our attention to homology with \mathbb{F}_2 -coefficients, as is common in applied topology, our persistent affine subspaces will correspond to the power set of some subset of the bars in a barcode. This is appealing both for intuition building and from the perspective of studying data visually using these tools.

Therefore, for the remainder of the paper we will assume that all fields are finite and all bars have unique birth-death pairs. In addition, our computational examples will use \mathbb{F}_2 -coefficients. In all cases, we will suppress the field from our notation unless we need to refer to it specifically. We will occasionally recall these assumptions when they are pertinent for our discussion.

2.3 Birth and death parameters of homology classes.

The barcode of a persistence module has been the primary method for characterizing its structure in practice, so it is common to discuss the birth and death time of bars. The same terminology is often applied to persistent homology classes only informally to avoid dealing directly with the interval decomposition. However, to work with these

classes algorithmically we will need to make this terminology precise. Thus, we will use the correspondence between bars and homology classes in Definition 2.2.2 to state and describe how to compute the birth and death parameters of a class $[w] \in \tilde{H}_k(Z^\ell)$.

Definition 2.3.1 Let $P\tilde{H}_k(Z^\bullet)$ be a persistence module with structure maps $\{\phi^\ell\}$. Given a class $[w] \in \tilde{H}_k(Z^\ell)$, let W^\bullet be the persistent subspace of $P\tilde{H}_k(Z^\bullet)$ satisfying the following:

1. $W^\ell = \text{span}_{\mathbb{F}}\{[w]\}$
2. $W^i = \text{span}_{\mathbb{F}}\{[x] \in \tilde{H}_k(Z^i) \mid [w] = \phi^{\ell-1} \circ \dots \circ \phi^i([x])\}$ for $i < \ell$
3. $W^i = \phi^{i-1}(W^{i-1})$ for $i > \ell$.

The *birth* and *death* parameters of $[w]$ are given by $\beta([w]) = \min\{i \leq \ell \mid W^i \neq 0\}$ and $\delta([w]) = \min\{i > \ell \mid W^i = 0\}$.

That is, W^\bullet is the maximal persistent subspace of $P\tilde{H}_k(Z^\bullet)$ containing $[w]$. To compute the birth and death parameters, we use the following lemma.

Lemma 2.3.2 Given a persistence module $P\tilde{H}_k(Z^\bullet)$, let $\mathcal{B} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be any interval decomposition. Given $[w] \in \tilde{H}_k(Z^\ell)$, express $[w]$ using the bars alive at ℓ as $[w] = \sum_{\tau \in \text{BC}_k(Z^\ell)} f_\tau [\tau^{\mathcal{B}, \ell}]$ for some coefficients $f_\tau \in \mathbb{F}$. The birth and death times of $[w]$ are given by $\max\{\beta(\tau) \mid \tau \in \text{BC}_k(Z^\ell) \text{ and } f_\tau \neq 0\}$ and $\max\{\delta(\tau) \mid \tau \in \text{BC}_k(Z^\ell) \text{ and } f_\tau \neq 0\}$, respectively.

Note that even though we fix an interval decomposition \mathcal{B} for the computation, the birth and death times of $[w]$ is independent of the choice of \mathcal{B} . The proof of this lemma is not particularly enlightening to read, so we leave it as an exercise to the interested reader.

2.4 Clique and witness complexes

In this paper we restrict our attention to two special classes of filtered topological spaces: filtered *clique complexes* and filtered *witness complexes*. The former is the standard tool for constructing simplicial complexes from data represented by a collection of pairwise dissimilarity measures among elements of a population. The ubiquitous example is the filtered Vietoris-Rips complex of a point cloud in a metric space. The latter is less commonly used but naturally encodes common cross-population dissimilarity measures in a fashion that is compatible with the usual applications of the clique complex.

Definition 2.4.1 Let A be an $(n \times n)$ matrix with zero diagonal and N unique non-negative real off-diagonal entries $\alpha_1 < \dots < \alpha_N$. The (filtered) *clique complex* for A is the filtered simplicial complex $X(A)^\bullet$ where $X(A)^\ell$ has vertices $[n] = \{1, \dots, n\}$ and faces $\{\sigma \subseteq [n] : \max_{\{r,c\} \in \sigma} A_{r,c} \leq \alpha_\ell\}$.

Our standard context for constructing a clique complex will involve a population $P = \{p_1, \dots, p_n\}$ endowed with a pairwise dissimilarity matrix M_P . In this case, we will write X_P^\bullet for $X(M_P)^\bullet$.

Definition 2.4.2 (Dowker 1952; de Silva and Carlsson 2004; Chowdhury and Mémoli 2018) Let B be an $(n \times m)$ matrix with N unique non-negative real entries $\beta_1 < \dots < \beta_N$. The (filtered) witness complex of B is the filtered simplicial complex $W(B)^\bullet$ where $W(B)^\ell$ has vertices $[n] = \{1, \dots, n\}$ and faces $\{\{r \in [n] : B_{r,c} \leq \beta_\ell\} : c \in [m]\}$. We refer to the set $[n]$ as the set of landmarks and $[m]$ as the set of witnesses.

Note that some of the faces of $W(B)^\ell$ may be duplicates or subfaces of other faces. When working with two populations $P = \{p_1, \dots, p_n\}$ and $Q = \{q_1, \dots, q_m\}$ equipped with a $(n \times m)$ cross-dissimilarity matrix $M_{P,Q}$, we will write $W_{P,Q}^\bullet$ for $W(M_{P,Q})^\bullet$.

The apparent asymmetry between rows and columns—the sets of landmarks and witnesses—in the construction may seem off-putting: there's no intrinsic reason to choose one of P or Q as the vertex set. Fortunately, this choice no longer matters once we pass to persistent homology.

Theorem 2.4.3 (Functorial Dowker Theorem (Dowker 1952; Chowdhury and Mémoli 2018), paraphrased) Let B be an $(n \times m)$ matrix with non-negative real entries. Then $P\hat{H}_k(W(B)^\bullet) \cong P\hat{H}_k(W(B^T)^\bullet)$ as persistence modules for every dimension k . In particular $\text{BC}_k(W(B)^\bullet) \cong \text{BC}_k(W(B^T)^\bullet)$.

This result is remarkable: it states that the roles of landmark and witness are interchangeable from a topological perspective. Dowker's theorem was originally stated for unweighted complexes, and it has been extended to the persistent homology setting (Chowdhury and Mémoli 2018). Further, the isomorphism is explicit, so we can compute it. We will leverage this perspective shift as a foundation for comparing barcodes across populations ⁷.

Note for both X_P^\bullet and $W_{P,Q}^\bullet$, the corresponding persistence modules are necessarily finitely generated. Therefore, their homology groups are finite-dimensional vector spaces. Further, observe that both X_P^\bullet and $W_{P,Q}^\bullet$ are contractible for large enough filtration parameters. Since we are working with reduced homology, neither construction admits a bar τ with $\delta(\tau) = \infty$.

3 Understanding the problem

Now that we have the appropriate language, we pause to discuss more thoroughly the motivation for this work. Recall that we are interested in comparing two populations, $P = \{p_1, \dots, p_n\}$ and $Q = \{q_1, \dots, q_m\}$; for example, P and Q can represent point clouds sampled from some distributions on metric spaces or agents in some complex systems. We will take as our data three matrices: M_P and M_Q , symmetric non-negative dissimilarity matrices for P and Q , and $M_{P,Q}$, a non-negative $(n \times m)$ matrix whose (r, c) entry measures dissimilarity between $p_r \in P$ and $q_c \in Q$.

Suppose now that we compute the degree k persistent homology of the clique complexes X_P^\bullet and X_Q^\bullet and obtain barcodes $\text{BC}_k(X_P^\bullet)$ and $\text{BC}_k(X_Q^\bullet)$. Given a bar of interest $\tau \in \text{BC}_k(X_P^\bullet)$, corresponding to some 1-dimensional persistent subspace of $P\hat{H}_k(X_P^\bullet)$, how can we determine which bars in $\text{BC}_k(X_Q^\bullet)$ represent topological features similar to τ , as measured by $M_{P,Q}$?

⁷ Throughout this paper, we will use the term *Dowker's theorem* to refer to both the original theorem and the functorial theorem.

Perhaps the most straightforward option is to apply recently developed methods for comparing barcodes using the functoriality of homology, including induced matching (Bauer and Lesnick 2014, 2020), cycle registration (Reani and Bobrowski 2021), basis-independent partial matching (Gonzalez-Díaz and Soriano-Trigueros 2020), or quiver-representations of correspondences (Takeuchi 2021). In this section, we take P and Q as point clouds in a common metric space to explore why these methods may not provide the answers a practitioner might expect or desire.

3.1 The problem of comparing clique complexes on P and Q

First, let us briefly recall how the above-mentioned methods work. Given a map of persistence modules $f : U^\bullet \rightarrow V^\bullet$, induced matching (Bauer and Lesnick 2014, 2020; Bauer and Schmahl 2021) factorizes f as $U^\bullet \rightarrow (\text{im } f)^\bullet \hookrightarrow V^\bullet$ and defines a matching between $\text{BC}_k(U^\bullet)$ and $\text{BC}_k(V^\bullet)$ via canonical injections through $\text{BC}_k((\text{im } f)^\bullet)$. The canonical injections match bars only when their endpoints are aligned. Cycle registration (Reani and Bobrowski 2021) uses induced matching to find matchings between $\text{BC}_k(U^\bullet)$ and $\text{BC}_k(V^\bullet)$ given maps of the form $U^\bullet \rightarrow Y^\bullet \leftarrow V^\bullet$. Basis-independent partial matching (Gonzalez-Díaz and Soriano-Trigueros 2020), on the other hand, considers the map $f : U^\bullet \rightarrow V^\bullet$ as a commutative diagram and reports the dimension of the appropriate subspace that corresponds to given bars in $\text{BC}_k(U^\bullet)$ and $\text{BC}_k(V^\bullet)$. Lastly, if X and Y are topological spaces with a continuous map $g : X \rightarrow Y$ and P is a finite sample of X , then one can use quiver-representations of correspondences (Takeuchi 2021) to approximate $g_* : H_k(X) \rightarrow H_k(Y)$ from the sampled data $g|_P$. We can adapt this construction to the context of persistent homology by considering maps of persistence modules $PH_k(X^\bullet) \leftarrow PH_k(G^\bullet) \rightarrow PH_k(Y^\bullet)$, where G^\bullet represents a filtered simplicial complex associated to the graph of the sampled map $g|_P$.

When we have access to explicit maps on persistence modules or maps on finite samples of topological spaces, these methods provide the best possible comparison of topological features; they recapitulate the usual functoriality on homology in the context of barcodes. However, we often build our persistence modules using samples from unknown distributions on unknown spaces, so the required maps of persistence modules of the form $U^\bullet \rightarrow V^\bullet$ or $U^\bullet \rightarrow Y^\bullet \leftarrow V^\bullet$ are not known. Indeed, if the point clouds P and Q represent time series data on two different systems, there may not be an explicit map between P and Q at all.

In simple cases, we might still attempt to impute the necessary maps from the data. If $P = \{p_1, \dots, p_n\}$ and $Q = \{q_1, \dots, q_m\}$ are sampled from a common metric space, then the metric provides enough information to naïvely construct a diagram of persistence modules as

$$P\tilde{H}_k(X_P^\bullet) \rightarrow P\tilde{H}_k(X_{P \cup Q}^\bullet) \leftarrow P\tilde{H}_k(X_Q^\bullet). \tag{5}$$

However, attempts to apply the above-listed methods to this diagram suffer from two major issues.

First, the lifetimes of the topological features in P and Q must be very closely aligned in the parameter range. Consider two probability distributions on an annulus

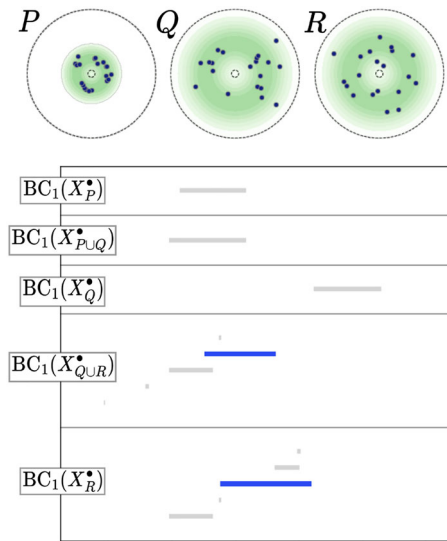


Fig. 2 Complexities arise in applying induced matching with related point clouds. (top) Point clouds P , Q , R sampled on an annulus. P is sampled from the distribution in green concentrated around the inner boundary, while Q and R are sampled from identical, centrally concentrated distributions. (bottom) Barcodes in dimension 1 for the Vietoris-Rips complexes of the point clouds P , $P \cup Q$, Q , $Q \cup R$, and R . Because the unique bar in $BC_1(X_Q^\bullet)$ is misaligned with the others, induced matching through the unions will fail to identify the feature in Q with non-trivial features in either P or R . Blue bars indicate the only possible induced match between $BC_1(X_{Q \cup R}^\bullet)$ and $BC_1(X_R^\bullet)$

illustrated in green on the top of Fig. 2. Let P and Q be samples from the two probability distributions. Because they surround the hole in the middle of the annulus, we intuit that the two samples share a common topological feature. However, as the bottom panel in as Fig. 2 illustrates, the bars in $BC_1(X_{P \cup Q}^\bullet)$ and $BC_1(X_Q^\bullet)$ are not aligned because the scale of the two point clouds is different. As a result, existing methods will fail to identify the longest bar in $BC_1(X_P^\bullet)$ with the longest bar in $BC_1(X_Q^\bullet)$, even though intuition suggests they represent the same feature. When working with measurements taken from multiple sources, different calibrations of instruments or small changes in experimental conditions could easily result in these kinds of shifts in scale.

The second issue with the construction in Eq. 5 is that we must have a dense sample of points to ensure that persistent features have similar scale. Consider the populations Q and R in Fig. 2 sampled from the same distribution. Each individual point cloud has a barcode that suggests a strong 1-dimensional feature. However, when we apply induced matching, only the highlighted blue bars in $BC_1(X_{Q \cup R}^\bullet)$ and $BC_1(X_R^\bullet)$ can possibly match; the unique bar in $BC_1(X_Q^\bullet)$ does not have any matching due to an apparent scale difference. Note, however, that P and R can be matched using existing methods; thus, the existence or non-existence of a matching is unreliable even in this simple setting.

3.2 The problem of comparing clique complexes and witness complex

Some of the earliest papers in applied topology suggest an alternative to the union for investigating the relationship between two point clouds in a common metric space: the witness complex $W_{P,Q}^\bullet$ (de Silva and Carlsson 2004). Empirically, witness complexes have been observed to mitigate many (though not all) of the issues with scale and density we describe above. However, we no longer have an inclusion map and cannot directly apply induced matching to compare persistent homology.⁸

Simple examples show that features in $BC_k(W_{P,Q}^\bullet)$ need not correspond in any meaningful way to those in $BC_k(X_P^\bullet)$. Taking P and Q to be dense samples from complementary semi circles on the unit circle, the barcode $BC_1(W_{P,Q}^\bullet)$ will have a single nontrivial bar, while $BC_1(X_P^\bullet)$ and $BC_1(X_Q^\bullet)$ will both be empty. On the other hand, taking P to be concentrated along the equator of a sphere and at its poles, and Q only at the poles, $BC_1(W_{P,Q}^\bullet)$ will be empty while $BC_1(X_P^\bullet)$ will have a persistent bar. To get traction, we will need a method for comparing homology classes between these complexes.

In Sect. 4, we propose a solution to this problem we call the *persistent extension* method. Persistent extension uses a zig-zag through the intersection of two simplicial complexes on the same vertex set to identify classes which contain identical representatives at the chain level for some range of parameters, thus identifying topological features one could reasonably consider to be “the same.” In Sect. 5, we apply this method and Dowker’s theorem to develop the method of *analogous bars*, which provides our desired method for comparing topological features of X_P^\bullet and X_Q^\bullet using only the data in the cross-dissimilarity measure $M_{P,Q}$.

4 Persistent extension

In this section, we develop the persistent extension method for comparing persistent homology classes between two filtered simplicial complexes supported on the same vertex set P . In Sect. 4.1, we introduce the class-to-class extension method, which, given a homology class in one simplicial complex, enumerates classes in the other that share its cycle representatives. In Sect. 4.2, we extend this algorithm to a bar-to-bars extension method, which, given a bar in one barcode, enumerates representations in the other barcode. In Sect. 4.3 we discuss considerations for reducing the computational overhead of applying the persistent extension method.

4.1 class-to-class extension

Suppose Z^\bullet and Y^\bullet are filtered simplicial complexes on vertices P . Fix some parameter ψ and take $[\tau] \in \tilde{H}_k(Z^\psi)$. The goal of the class-to-class extension is to find all parameters ℓ and classes $[y] \in \tilde{H}_k(Y^\ell)$ for which $[\tau]$ and $[y]$ share cycle representa-

⁸ One can construct a map $P\tilde{H}_k(W_{P,Q}^\bullet) \rightarrow P\tilde{H}_k(X_P^\bullet)$ by doubling the parameter. However, doubling the scale at which we consider a complex can be very destructive to homology.

tives. We begin by setting some terminology with which to describe the steps in our algorithm.

To simplify this section, we assume that the filtered simplicial complex Y^\bullet has the property that Y^N is the complete simplex on its vertices, as is the case for both clique and witness complexes as defined in Sect. 2.4. It is straightforward but tedious to modify the following for settings where this is not the case.

Definition 4.1.1 Let Z^\bullet and Y^\bullet be filtered simplicial complexes on common vertex set P and fix a parameter ψ of Z^\bullet . Given a parameter ℓ , let χ_ℓ and Υ_ℓ each denote the following induced maps

$$\tilde{H}_k(Z^\psi) \xleftarrow{\chi_\ell} \tilde{H}_k(Z^\psi \cap Y^\ell) \xrightarrow{\Upsilon_\ell} \tilde{H}_k(Y^\ell).$$

Given a homology class $[\tau] \in \tilde{H}_k(Z^\psi)$, define the *set of restrictions (of $[\tau]$) at ℓ* as

$$R_\ell = \{[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell) \mid [\tau] = \chi_\ell[w]\},$$

and define the *set of class extensions (of $[\tau]$) at ℓ* as

$$E_\ell = \{[y] \in \tilde{H}_k(Y^\ell) \mid [y] = \Upsilon_\ell[w] \text{ for } [w] \in R_\ell\}.$$

We refer to any $[y] \in E_\ell$ for which $[y] = \Upsilon_\ell[w]$ as a *class extension (via $[w]$)*. We refer to $E = \bigcup_\ell E_\ell$ as the *set of class extensions (of $[\tau]$)*.

For a fixed parameter ℓ , the set R_ℓ can be expressed as the solution to a chain-level linear equation, and E_ℓ is then given by the solutions to a set of linear equations indexed by R_ℓ . To naively compute E , we would have to iterate this process across every parameter ℓ , which would be infeasible for even moderately sized data. However, the following two lemmas characterize how restrictions and class extensions at different parameters are related, allowing us to limit our computations to a much smaller set of linear systems.

Given parameters $\ell \leq \ell'$, let

$$\begin{aligned} \eta_{\ell'}^\ell &: \tilde{H}_k(Z^\psi \cap Y^\ell) \rightarrow \tilde{H}_k(Z^\psi \cap Y^{\ell'}), \\ \lambda_{\ell'}^\ell &: \tilde{H}_k(Y^\ell) \rightarrow \tilde{H}_k(Y^{\ell'}) \end{aligned}$$

denote the structure maps for the persistence modules, induced by inclusion.

Lemma 4.1.2 *Let Z^\bullet, Y^\bullet be filtered simplicial complexes on a common vertex set P . Given $[\tau] \in \tilde{H}_k(Z^\psi)$, fix parameter ℓ and let $[y] \in \tilde{H}_k(Y^\ell)$ be a class extension of $[\tau]$ at ℓ via some $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$. If $\ell' > \ell$, then $\lambda_{\ell'}^\ell[y]$ is a class extension of $[\tau]$ at ℓ' via $\eta_{\ell'}^\ell[w]$.*

That is, any class extension at ℓ passes to a class extension at any later parameter ℓ' via the structure maps.

Lemma 4.1.3 *Let Z^\bullet, Y^\bullet be filtered simplicial complexes on a common vertex set P . Given $[\tau] \in \tilde{H}_k(Z^\psi)$, let $[y'] \in \tilde{H}_k(Y^{\ell'})$ be a class extension of $[\tau]$ at ℓ' via some $[w'] \in \tilde{H}_k(Z^\psi \cap Y^{\ell'})$. Assume that there exists a parameter $\ell < \ell'$ and $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$ such that $\eta_{\ell'}^\ell[w] = [w']$. Then, there exists a class extension $[y] \in \tilde{H}_k(Y^\ell)$ of $[\tau]$ at ℓ via $[w]$ satisfying $[y'] = \lambda_{\ell'}^\ell[y]$.*

The proofs of both lemmas follows directly from commutativity of the relevant diagrams of structure maps and morphisms induced by inclusions. Details are left to the interested reader.

Lemma 4.1.3 states that if a restriction $[w']$ is in the image of the structure map $\eta_{\ell'}^\ell$, then any class extension at parameter ℓ' via $[w']$ is in the image of the structure map $\lambda_{\ell'}^\ell$, so this extension is induced by an extension at ℓ .

In particular, suppose $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$ is a restriction at ℓ . Let $\hat{\beta}[w] \leq \ell$ denote the birth parameter of $[w]$ in $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$. Then, there exists some $[\hat{w}] \in \tilde{H}_k(Z^\psi \cap Y^{\hat{\beta}[w]})$ satisfying $[w] = \eta_{\ell}^{\hat{\beta}[w]}[\hat{w}]$. Let $[\hat{y}]$ be a class extension of $[\tau]$ at $\hat{\beta}[w]$ via $[\hat{w}]$. Together, Lemmas 4.1.2 and 4.1.3 tell us that any class extension $[y]$ of $[\tau]$ at ℓ via $[w]$ must satisfy $[y] = \lambda_{\ell}^{\hat{\beta}[w]}[\hat{y}]$. That is, in order to find class extensions of $[\tau]$ via $[w]$ at all parameters, it suffices to find the class extensions via $[\hat{w}]$ at $\hat{\beta}[w]$.

We leverage this observation to enumerate all possible class extensions while considering only a substantially restricted set of parameters. To do so, we utilize the auxiliary filtration

$$Z^\psi \cap Y^\bullet : Z^\psi \cap Y^1 \rightarrow Z^\psi \cap Y^2 \rightarrow \dots \rightarrow Z^\psi \cap Y^N = Z^\psi \tag{6}$$

along with its barcode $(BC_k(Z^\psi \cap Y^\bullet), \hat{\beta}, \hat{\delta})$. The benefit of using this filtration is that if $\hat{\delta}(\rho) = \infty$ for $\rho \in BC_k(Z^\psi \cap Y^\bullet)$, then ρ corresponds to a homology class in $\tilde{H}_k(Z^\psi)$ under any choice of interval decomposition of $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$.

We now refer the reader to Algorithm 1, where we describe the class-to-class extension method in full. In step (1), we compute the auxiliary filtration in Eq. 6. In step (2), we find the collection of parameters that correspond to the birth times of all restrictions of $[\tau]$. In step (3), we find all restrictions at a given birth parameter and find the corresponding class extensions.

There are several facets of Algorithm 1 that require justification. We need to show that the collections p_Y and $\mathfrak{R}_\ell^{\mathcal{F}}$ are independent of the choice of interval decomposition \mathcal{F} of $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$, that the affine subspace $\mathfrak{R}_\ell^{\mathcal{F}}$ coincides with the set R_ℓ of restrictions⁹ as defined in Definition 4.1.1, and that the output \mathfrak{E} of the algorithm provides the complete set E of class extensions of $[\tau]$ from Definition 4.1.1, as stated in the following Theorem.

Theorem 4.1.4 *Let Z^\bullet, Y^\bullet be filtered simplicial complexes on a common vertex set P . Fix a parameter ψ and homology class $[\tau] \in \tilde{H}_k(Z^\psi)$. The output of Algorithm 1 suffices to recover all class extensions of $[\tau]$. That is, let E be the set of all class extensions of $[\tau]$ per Definition 4.1.1. Given parameters $\ell < \ell'$, let $\lambda_{\ell'}^\ell : \tilde{H}_k(Y^\ell) \rightarrow$*

⁹ In particular, note that we avoid solving a linear system to extract the set of solutions.

Algorithm 1 Extension method (class-to-class extensions)**Input:**

- filtered simplicial complexes Z^\bullet, Y^\bullet on vertex set P ,
- a parameter ψ , and
- a homology class $[\tau] \in \tilde{H}_k(Z^\psi)$

Output:

- a collection of parameters p_Y , and
- a collection \mathfrak{E} of class extensions of $[\tau]$ to $P\tilde{H}_k(Y^\bullet)$.

Steps:

1. Compute the persistent homology $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$, the barcode $(\text{BC}_k(Z^\psi \cap Y^\bullet), \widehat{\beta}, \widehat{\delta})$, and an interval decomposition $\mathcal{F} : \mathbb{I}_{\text{BC}_k(Z^\psi \cap Y^\bullet)} \rightarrow P\tilde{H}_k(Z^\psi \cap Y^\bullet)$ of the auxiliary filtration in Equation (6),
2. Find a sufficient collection of parameters p_Y at which to look for class extensions of $[\tau]$.
 - (a) Let $S_{[\tau]}^{\mathcal{F}} = \{f_1^* \rho_1, \dots, f_m^* \rho_m\}$ be an \mathcal{F} -bar representation of $[\tau]$ at N , where N is the final filtration parameter of Y^\bullet .
 - (b) Let $\text{BARS}_{\tau}^{\mathcal{F}} = \{\rho_1, \dots, \rho_m\}$ be the set of bars in this \mathcal{F} -bar representation.
 - (c) Let $\ell_0 = \max_{\rho \in \text{BARS}_{\tau}^{\mathcal{F}}} \{\widehat{\beta}(\rho)\}$.
 - (d) Let $\text{BARS}_{>\ell_0} = \{\mu \in \text{BC}_k(Z^\psi \cap Y^\bullet) \mid \ell_0 < \widehat{\beta}(\mu) < \widehat{\delta}(\mu) < \infty\}$,
 - (e) Let $p_Y = \{\ell_0\} \cup \{\widehat{\beta}(\mu) \mid \mu \in \text{BARS}_{>\ell_0}\}$.
3. Find a complete set of class extensions of $[\tau]$ to $P\tilde{H}_k(Y^\bullet)$.
 - (a) Let $\text{BARS}_{\text{short}} = \{\rho \in \text{BC}_k(Z^\psi \cap Y^\bullet) \mid \ell_0 < \widehat{\delta}(\rho) < \infty\}$.
 - (b) For each $\ell \in p_Y$,
 - (i) Let $\text{BARS}_{\text{short}}^{\ell} = \{\rho \in \text{BARS}_{\text{short}} \mid \widehat{\beta}(\rho) \leq \ell < \widehat{\delta}(\rho) < \infty\}$.
 - (ii) Let $V_{\text{short}}^{\ell} = \{[\rho^{\mathcal{F}, \ell}] \mid \rho \in \text{BARS}_{\text{short}}^{\ell}\}$.
 - (iii) Let $\mathfrak{R}_{\ell}^{\mathcal{F}} = \sum_{i=1}^m f_i^*[\rho_i^{\mathcal{F}, \ell}] + \text{span}_{\mathbb{F}} V_{\text{short}}^{\ell}$.
 - (iv) Let $\mathfrak{E}_{\ell} = \{\Upsilon_{\ell}([w]) \mid [w] \in \mathfrak{R}_{\ell}^{\mathcal{F}}\}$ where $\Upsilon_{\ell} : \tilde{H}_k(Z^\psi \cap Y^{\ell}) \rightarrow \tilde{H}_k(Y^{\ell})$ is the map induced by inclusion.
4. Return p_Y and $\mathfrak{E} = \bigcup_{\ell \in p_Y} \mathfrak{E}_{\ell}$.

$\tilde{H}_k(Y^{\ell'})$ be the map induced by inclusion. Let

$$\mathfrak{E}^* = \{[y] \in \tilde{H}_k(Y^{\ell'}) \mid \text{there exists } \ell \in p_Y, [y_{\ell}] \in \mathfrak{E}_{\ell} \text{ such that } [y] = \lambda_{\ell'}^{\ell}([y_{\ell}])\},$$

where p_Y and $\mathfrak{E} = \bigcup_{\ell \in p_Y} \mathfrak{E}_{\ell}$ are the outputs of Algorithm 1. Then $E = \mathfrak{E}^*$.

The proof is technical, so we defer it to Appendix B where we break it down into a sequence of smaller theorems, and we instead proceed with an example that illustrates the method.

Example 1 We visualize the key steps of Algorithm 1 using a set of points P sampled from a flat torus and a set of points Q sampled near a essential circle on that torus, as depicted in Fig. 3. We assume that homology is computed with \mathbb{F}_2 coefficients and that all bars have unique birth-death pairs.

Let $Z^\bullet = W_{P,Q}^\bullet$, the Witness filtration with P as landmarks and Q as witnesses computed from the cross-dissimilarity matrix $M_{P,Q}$, and let $Y^\bullet = X_P^\bullet$, computed from the pairwise dissimilarity matrix M_P . Let $[\tau] \in \tilde{H}_1(W_{P,Q}^\psi)$ be the homology class illustrated in Fig. 3.

We now execute Algorithm 1. The right panel in Fig. 3 illustrates the barcode $BC_1(W_{P,Q}^\psi \cap X_P^\bullet)$, and we will refer to it throughout. In step (1), we fix an interval decomposition $\mathcal{F} : \mathbb{I}_{BC_1(W_{P,Q}^\psi \cap X_P^\bullet)} \rightarrow P\tilde{H}_1(W_{P,Q}^\psi \cap X_P^\bullet)$. In step (2), we gather information about this barcode. For our choice of \mathcal{F} , the unique \mathcal{F} -bar representation of $[\tau]$ at N is $S_{[\tau]}^\mathcal{F} = \{\rho_1\}$, where ρ_1 is the green bar. The parameter $\ell_0 = \hat{\beta}(\rho_1)$ at which this bar is born is indicated by the vertical dotted line. Note that ℓ_0 is the minimum parameter at which there exists a restriction of $[\tau]$. The set $BARS_{>\ell_0}$ consists of the blue bars above the green bar, each being born and dying within the interval (ℓ_0, ∞) . Finally, take our list of parameters p_Y to be the set of birth times of the green and blue bars. The collection p_Y corresponds to the birth parameters of restrictions of $[\tau]$ in $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$.

For step (3), we take $BARS_{short}$ to be the set of all bars that die after the green bar is born. In the first pass through the loop in step (3-b), taking $\ell_0 \in p_Y$ as our parameter, the purple bars in Fig. 3 indicate the intervals $BARS_{short}^{\ell_0} = \{\rho_2, \dots, \rho_{15}\}$ selected in step (3-b-i). The restrictions at ℓ_0 in step (3-b-iii) thus constitute the affine subspace

$$\mathfrak{R}_{\ell_0} = \{[\rho_1^{\mathcal{F},\ell_0}] + c_2[\rho_2^{\mathcal{F},\ell_0}] + \dots + c_{15}[\rho_{15}^{\mathcal{F},\ell_0}] \mid c_2, \dots, c_{15} \in \mathbb{F}_2\}.$$

The inset panel in Fig. 3 illustrates cycle representatives of $[\rho_1^{\mathcal{F},\ell_0}]$ in green and $[\rho_2^{\mathcal{F},\ell_0}]$ in purple. Here, we report the default cycle representatives returned by software Eirene (Henselman and Ghrist 2016). While the algorithm is independent of the cycle representatives, we included the illustrations to show that elements of \mathfrak{R}_{ℓ_0} are homology classes that are similar to $[\tau]$. Note that both $[\rho_1^{\mathcal{F},\ell_0}]$ and $[\rho_1^{\mathcal{F},\ell_0}] + [\rho_2^{\mathcal{F},\ell_0}]$ describe homology classes whose representatives are similar to that of $[\tau]$; the latter is simply a deformation of the former to pass through a local cycle with short lifetime. Indeed, all elements of \mathfrak{R}_{ℓ_0} are such deformations. Finally, in step (3-b-iv), we push these into $\tilde{H}_k(Y^{\ell_0})$ as elements in \mathfrak{E}_{ℓ_0} , which is the collection of class extensions at ℓ_0 .

Step 3(b) repeats the above process for all $\ell \in p_Y$. For a fixed ℓ , the set \mathfrak{R}_ℓ corresponds to the set of restrictions at ℓ , and \mathfrak{E}_ℓ corresponds to the set of class extensions at ℓ .

4.2 Bar-to-bars extensions

It is common in applications to refer to bars, not homology classes, as the features of a data set. Thus, for practitioners, it may be of use to understand relations between bars rather than homology classes. The bar-to-bars extension method enumerates ways in which a selected bar $\tau \in BC_k(Z^\bullet)$ coincides at the chain level with elements of $P\tilde{H}_k(Y^\bullet)$ and represents the corresponding homology classes in terms of $BC_k(Y^\bullet)$. To do so, we first find all homology class representations of τ in $P\tilde{H}_k(Z^\bullet)$. For each homology class representation, we find all class extensions, and, for each class extension, we find the corresponding bar representation in $BC_k(Y^\bullet)$. We discuss the notion of bar extensions before presenting the algorithm.

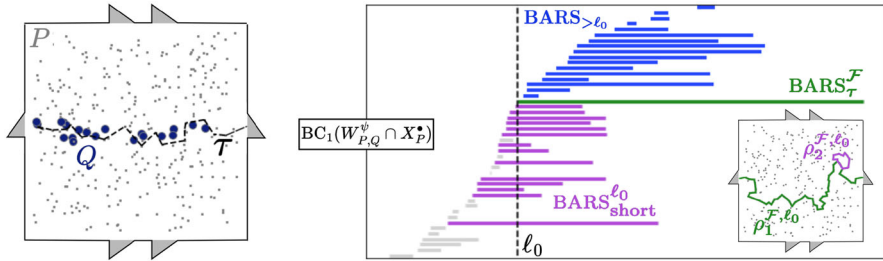


Fig. 3 Elements of Algorithm 1 in Example 1. (left) Grey points P sampled from a flat torus and blue points Q sampled near an essential circle of that torus, along with a representative τ of a class $[\tau] \in \tilde{H}_1(W_{P,Q}^\psi)$ that represents the essential circle. (right) The barcode of the auxiliary filtration, $BC_1(W_{P,Q}^\psi \cap X_P^\bullet)$. For a fixed interval decomposition \mathcal{F} of $BC_1(W_{P,Q}^\psi \cap X_P^\bullet)$, the long green bar corresponds to the only element $\rho_1 \in BARS_\tau^{\mathcal{F}}$, which is born at parameter ℓ_0 . Short purple bars with lifetimes intersecting ℓ_0 make up $BARS_{short}^{\ell_0}$, and finite blue bars born after ℓ_0 constitute the set $BARS_{>\ell_0}$. (inset) Representatives of $[\rho_1^{\mathcal{F},\ell_0}]$ in green and $[\rho_2^{\mathcal{F},\ell_0}]$ in purple corresponding to a bar in $BARS_{short}^{\ell_0}$

The first step is to find all homology classes that correspond to a given bar $\tau \in BC_k(Z^\bullet)$. If we fix an interval decomposition \mathcal{B} of $P\tilde{H}_k(Z^\bullet)$ and some parameter ψ for Z^\bullet , then we automatically have the homology class $[\tau^{\mathcal{B},\psi}]$ that \mathcal{B} -corresponds to τ .

Which parameter ψ should we choose? In the absence of further information, the safest choice is to select ψ which maximizes the size of our homology class of interest¹⁰ but avoiding introducing enough equivalences to trivialize our feature of interest. Thus, we consider the last parameter at which the bar τ is alive.

Definition 4.2.1 Let $\mathcal{B} : \mathbb{I}_{BC_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be an interval decomposition. The terminal class of a bar τ in $BC_k(Z^\bullet)$ (under \mathcal{B}) is the class $[\tau_*^{\mathcal{B}}] = [\tau^{\mathcal{B},\delta(\tau)-1}] \in \tilde{H}_k(Z^{\delta(\tau)-1})$. When \mathcal{B} is unambiguous, we will suppress it and write $[\tau_*]$.

Recall that a bar $\tau \in BC_k(Z^\bullet)$ corresponds to a 1-dimensional persistent subspace V^\bullet of $P\tilde{H}_k(Z^\bullet)$ under \mathcal{B} .

Remark 2 Without a preferred interval decomposition \mathcal{B} , it will be necessary to consider all possible choices of \mathcal{B} to obtain a comprehensive list of possible terminal classes of τ . In Sect. 4.2.3, we will show that the terminal class is independent of the interval decomposition \mathcal{B} if homology is computed with \mathbb{F}_2 coefficients and if the bars of $BC_k(Z^\bullet)$ have unique death parameters. Thus, for readers interested in that very common case, it is reasonable to ignore the dependence on \mathcal{B} in what follows.

Once we fix the terminal class $[\tau_*^{\mathcal{B}}]$ of τ , one can compute the class extensions of $[\tau_*^{\mathcal{B}}]$. In the context of comparing barcodes, we need to describe these class extensions in terms of the bars in $BC_k(Y^\bullet)$. Given an interval decomposition $\mathcal{D} : \mathbb{I}_{BC_k(Y^\bullet)} \rightarrow P\tilde{H}_k(Y^\bullet)$ and a class extension $[y] \in \tilde{H}_k(Y^\ell)$ of $[\tau_*^{\mathcal{B}}]$ at ℓ , recall from Definition

¹⁰ In specific applications where we understand the complexes involved more completely, it may be sensible to select a different parameter at which to consider classes corresponding to τ .

2.2.4 that $S_{[y]}^{\mathcal{D}}$ denotes the \mathcal{D} -bar representation of $[y]$. Recall that $S_{[y]}^{\mathcal{D}}$ is a collection of weighted bars. Using this terminology, we can describe class extensions and bar extensions of a bar $\tau \in BC_k(Z^\bullet)$.

Definition 4.2.2 Let Z^\bullet, Y^\bullet be filtered simplicial complexes on a common vertex set P . Let $\tau \in BC_k(Z^\bullet)$. Given an interval decomposition \mathcal{B} of $P\tilde{H}_k(Z^\bullet)$, let $E_{\mathcal{B}}$ denote the set of class extensions of $[\tau_{*}^{\mathcal{B}}]$. Let B and D each denote the set of all possible interval decompositions of $P\tilde{H}_k(Z^\bullet)$ and $P\tilde{H}_k(Y^\bullet)$. Define the *class extensions* of τ as

$$E(\tau, Y^\bullet) = \bigcup_{\mathcal{B} \in B} E_{\mathcal{B}},$$

and define the *bar extensions* of τ as

$$S(\tau, Y^\bullet) = \{S_{[y]}^{\mathcal{D}} \mid [y] \in E(\tau, Y^\bullet), \mathcal{D} \in D\}.$$

When given a homology class $[\tau] \in \tilde{H}_k(Z^\psi)$, we use the corresponding terminology *bar extensions* of $[\tau]$ to refer to

$$S([\tau], Y^\bullet) = \{S_{[y]}^{\mathcal{D}} \mid [y] \in E, \mathcal{D} \in D\},$$

where E denotes the set of class extensions of $[\tau]$ from Definition 4.1.1.

Now that we defined the bar extensions of τ , we can discuss the overall structure of our bar-to-bars extension algorithm. Given $\tau \in BC_k(Z^\bullet)$ we first find all terminal classes of τ by enumerating all possible interval decompositions of $P\tilde{H}_k(Z^\bullet)$. We then use the resulting homology classes as input for Algorithm 1 and generate all class extensions of τ . Finally, for each class extension, we enumerate all interval decompositions of $P\tilde{H}_k(Y^\bullet)$ to find the bar representations of the class extensions.

In Sects. 4.2.1 and 4.2.2, we discuss details of performing this enumeration efficiently and summarize the process in Algorithm 2. Finally, in Sect. 4.2.3, we show that a much simpler process (Algorithm 3) suffices when we compute persistent homology with \mathbb{F}_2 coefficients and when our bars have unique death parameters. We first discuss the computational details of enumerating the interval decompositions.

4.2.1 Finding all terminal classes of a bar

Given a bar $\tau \in BC_k(Z^\bullet)$, we must find an explicit parameter ψ and homology class $[\tau] \in \tilde{H}_k(Z^\psi)$ to be used as input to the class-to-class extension method (Algorithm 1). As discussed in Definition 4.2.1, we fix $\psi = \delta(\tau) - 1$ and the terminal class $[\tau_{*}^{\mathcal{B}}] \in \tilde{H}_k(Z^{\delta(\tau)-1})$ under some fixed interval decomposition $\mathcal{B} : \mathbb{I}_{BC_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$. To avoid imposing an arbitrary choice of \mathcal{B} , it is necessary to find all terminal classes $[\tau_{*}^{\mathcal{C}}]$ under different interval decompositions \mathcal{C} of $P\tilde{H}_k(Z^\bullet)$ and apply Algorithm 1 to these terminal classes. Note that the iteration over all interval decompositions of $P\tilde{H}_k(Z^\bullet)$ allows us to find all 1-dimensional persistent subspaces of $P\tilde{H}_k(Z^\bullet)$ that

can be represented by τ . We then find all homology classes representing such persistent subspaces.

Recall from Sect. 2.2 that any interval decomposition \mathcal{C} of $P\tilde{H}_k(Z^\bullet)$ can be obtained from \mathcal{B} as $\mathcal{C} = \mathcal{B} \circ \mathcal{L}$ for some automorphism \mathcal{L} of the barcode module $\mathbb{I}_{\text{BC}_k(Z^\bullet)}$. The terminal class under \mathcal{C} can be computed by $[\tau_*^{\mathcal{C}}] = [\mathcal{C}^\psi(\mathbf{e}_\tau^\psi)] = [\mathcal{B}^\psi \circ \mathcal{L}^\psi(\mathbf{e}_\tau^\psi)] = [\mathcal{B}^\psi \cdot \mathcal{L}^\psi \cdot \mathbf{e}_\tau^\psi]$, where $\psi = \delta(\tau) - 1$, \mathbf{e}_τ^ψ is the basis vector for τ at ψ , and \mathcal{B}^ψ and \mathcal{L}^ψ are the matrix representations of the maps \mathcal{B}^ψ and \mathcal{L}^ψ . Thus, to find all possible terminal class $[\tau_*^{\mathcal{C}}]$, it suffices to find all linear isomorphisms $\mathcal{L}^\psi : (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\psi \rightarrow (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\psi$ resulting from the restriction of the automorphism $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$. We use the following lemma for this purpose.¹¹

Lemma 4.2.3 *Given persistence module $P\tilde{H}_k(Z^\bullet)$, let $\text{BC}_k(Z^\psi) = \{\tau_1, \dots, \tau_m\}$ be some ordering of bars alive at ψ . Let L_Z be the collection of $m \times m$ matrices*

$$L_Z = \{L \in \mathbb{F}^{m \times m} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\tau_r) > \beta(\tau_c) \text{ or } \delta(\tau_r) > \delta(\tau_c) \text{ for all } r, c\}.$$

The collection L_Z is precisely the collection of matrices representing the map $\mathcal{L}^\psi : (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\psi \rightarrow (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\psi$ under all possible choices of barcode module automorphism $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$.

Proof Deferred to Appendix C. □

Since L_Z contains all of the pertinent linear isomorphisms \mathcal{L}^ψ , we can enumerate the terminal classes of interest via $T = \{[B^\psi \cdot L \cdot \mathbf{e}_\tau^\psi] \in \tilde{H}_k(Z^\psi) \mid L \in L_Z\}$, where B^ψ denotes the matrix representation of \mathcal{B}^ψ . Step (2) of Algorithm 2 performs this process.¹²

Once we find the collection T , we apply the class-to-class extension method (Algorithm 1) to each terminal class in T , which results in a collection of class extensions $[y] \in \tilde{H}_k(Y^\ell)$ at various parameters ℓ . The final step we need to discuss is finding all bar representations of class extensions.

4.2.2 Finding all bar-representations of a class extension.

Let $[y] \in \tilde{H}_k(Y^\ell)$ be a class extension output from Algorithm 1. Fix some interval decomposition \mathcal{D} of $P\tilde{H}_k(Y^\bullet)$, and let $S_{[y]}^{\mathcal{D}}$ denote the \mathcal{D} -bar representation of $[y]$. Our goal is to find all \mathcal{G} -bar representation of $[y]$ for all possible interval decompositions \mathcal{G} of $P\tilde{H}_k(Y^\bullet)$.

Any interval decomposition \mathcal{G} of $P\tilde{H}_k(Y^\bullet)$ can be obtained by $\mathcal{G} = \mathcal{D} \circ \mathcal{L}^{-1}$, where \mathcal{L} is an automorphism of $\mathbb{I}_{\text{BC}_k(Y^\bullet)}$. Given \mathcal{L}^ℓ , we can find the \mathcal{G} -bar representation of $[y]$ from $S_{[y]}^{\mathcal{D}}$ as follows. Let D^ℓ, G^ℓ , and L^ℓ each denote the matrix representing the linear isomorphism $\mathcal{D}^\ell, \mathcal{G}^\ell$, and \mathcal{L}^ℓ . Let $S_{[y]}^{\mathcal{D}} = \{d_1\gamma_{j_1}, \dots, d_t\gamma_{j_t}\}$. That is, $d_1[\gamma_{j_1}^{\mathcal{D},\ell}] + \dots + d_t[\gamma_{j_t}^{\mathcal{D},\ell}] = [y]$ in $\tilde{H}_k(Y^\ell)$ for nonzero d_1, \dots, d_t , i.e.,

$$D^\ell(d_1\mathbf{e}_{j_1}^\ell + \dots + d_t\mathbf{e}_{j_t}^\ell) = [y] \text{ in } \tilde{H}_k(Y^\ell) \text{ for nonzero } d_1, \dots, d_t,$$

¹¹ Per Remark 1, a generalization of this result was established independently in Jacquard et al. (2022).

¹² Since we are working with finite fields, it is possible to enumerate L_Z .

Algorithm 2 Extension method (bar-to-bars extension)

Input:

- filtered simplicial complexes Z^\bullet, Y^\bullet on vertex set P ,
- a bar $\tau \in BC_k(Z^\bullet)$

Output:

- a collection $E(\tau, Y^\bullet)$ of all class extensions of τ to $P\tilde{H}_k(Y^\bullet)$, and
- a collection $S(\tau, Y^\bullet)$ of all bar extensions¹³ of τ to $BC_k(Y^\bullet)$.

Steps:

1. Fix interval decompositions \mathcal{B} of $P\tilde{H}_k(Z^\bullet)$ and \mathcal{D} of $P\tilde{H}_k(Y^\bullet)$ ¹⁴.
2. Fix an ordering on $BC_k(Z^{\delta(\tau)-1}) = \{\tau_1, \dots, \tau_m\}$.
3. Enumerate the collection of matrices

$$L_Z = \{L \in \mathbb{F}^{m \times m} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\tau_r) > \beta(\tau_c) \text{ or } \delta(\tau_r) > \delta(\tau_c) \text{ for all } r, c\}.$$

4. Enumerate the collection of possible terminal classes

$$T = \{[B^{\delta(\tau)-1} \cdot L \cdot \mathbf{e}_\tau^{\delta(\tau)-1}] \in \tilde{H}_k(Z^{\delta(\tau)-1}) \mid L \in L_Z\},$$

where $\mathbf{e}_\tau^{\delta(\tau)-1}$ is the basis vector for τ at parameter $\delta(\tau) - 1$ (Definition 2.2.1) and $B^{\delta(\tau)-1}$ denotes the matrix representation of $\mathcal{B}^{\delta(\tau)-1}$.

5. For each $[\tau] \in T$,
 - (a) Let $p_{Y, [\tau]}$ and $\mathfrak{C}_{[\tau]}$ be the output of class-to-class extension (Algorithm 1) with inputs $Z^\bullet, Y^\bullet, \delta(\tau) - 1$ and $[\tau]$.
 - (b) For each $\ell \in p_{Y, [\tau]}$,
 - (i) Fix an ordering on $BC_k(Y^\ell) = \{\gamma_1, \dots, \gamma_n\}$.
 - (ii) Enumerate the collection of matrices

$$L_Y = \{L \in \mathbb{F}^{n \times n} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\gamma_r) > \beta(\gamma_c) \text{ or } \delta(\gamma_r) > \delta(\gamma_c) \text{ for all } r, c\}$$

- (iii) For each $[y] \in (\mathfrak{C}_{[\tau]})_\ell$, compute the \mathcal{D} -bar representation $S_{[y]}^{\mathcal{D}} = \{d_1\gamma_{j_1}, \dots, d_t\gamma_{j_t}\}$.
- (iv) For each $L \in L_Y$, compute

$$S_{[y]}^{\mathcal{D} \circ L^{-1}} = \{c_1\gamma_{i_1}, \dots, c_s\gamma_{i_s} \mid L(d_1\mathbf{e}_{j_1}^\ell + \dots + d_t\mathbf{e}_{j_t}^\ell) = c_1\mathbf{e}_{i_1}^\ell + \dots + c_s\mathbf{e}_{i_s}^\ell\},$$

where \mathbf{e}_j^ℓ is the basis vector for γ_j at parameter ℓ (Definition 2.2.1).

6. Return $E(\tau, Y^\bullet) = \bigcup_{[\tau] \in T} \mathfrak{C}_{[\tau]}$ and $S(\tau, Y^\bullet) = \{S_{[y]}^{\mathcal{D} \circ L^{-1}} \mid [\tau] \in T, \ell \in p_{Y, [\tau]}, [y] \in (\mathfrak{C}_{[\tau]})_\ell, L \in L_Y\}$.

where \mathbf{e}_j^ℓ is the basis vector for γ_j at parameter ℓ . Note that $S_{[y]}^{\mathcal{G}} = \{c_1\gamma_{i_1}, \dots, c_s\gamma_{i_s}\}$ if $G^\ell(c_1\mathbf{e}_{i_1}^\ell + \dots + c_s\mathbf{e}_{i_s}^\ell) = [y]$ for nonzero c_1, \dots, c_s . Since $L^\ell = (G^\ell)^{-1} \circ D^\ell$, this means that $L^\ell(d_1\mathbf{e}_{j_1}^\ell + \dots + d_t\mathbf{e}_{j_t}^\ell) = c_1\mathbf{e}_{i_1}^\ell + \dots + c_s\mathbf{e}_{i_s}^\ell$. Thus, given the matrix L^ℓ and $S_{[y]}^{\mathcal{D}}$, we can directly compute $S_{[y]}^{\mathcal{G}}$.

¹³ Recall from Definition 4.2.2 that bar extensions is a collection of sets.

¹⁴ The output of the algorithm does not depend on \mathcal{B} and \mathcal{D} . We use \mathcal{B} and \mathcal{D} as the initial points for enumerating all possible interval decompositions of $P\tilde{H}_k(Z^\bullet)$ and $P\tilde{H}_k(Y^\bullet)$.

Algorithm 3 Extension method (bar-to-bars extension, \mathbb{F}_2 -coeffs and unique death times)

Input and Output:

Same as bar-to-bars extension (Algorithm 2).

Steps:

1. Fix interval decompositions \mathcal{B} of $P\tilde{H}_k(Z^\bullet)$ and \mathcal{D} of $P\tilde{H}_k(Y^\bullet)$.
2. Perform step (5) of bar-to-bars extension (Algorithm 2) with $T = \{[\tau_*^{\mathcal{B}}]\}$.
3. Return $E(\tau, Y^\bullet) = \mathfrak{C}_{[\tau_*^{\mathcal{B}}]}$ and $S(\tau, Y^\bullet) = \{S_{[y]}^{\mathcal{D} \circ L^{-1}} \mid \ell \in p_Y, [y] \in (\mathfrak{C}_{[\tau_*^{\mathcal{B}}]})_\ell, L \in L_Y\}$.

As in the previous case, we can enumerate the matrix L^ℓ . Let $n = \dim \tilde{H}_k(Y^\ell)$ and fix an order of the bars $\text{BC}_k(Y^\ell) = \{\gamma_1, \dots, \gamma_n\}$. Let L_Y :

$$L_Y = \{L \in \mathbb{F}^{n \times n} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\gamma_r) > \beta(\gamma_c) \text{ or } \delta(\gamma_r) > \delta(\gamma_c) \text{ for all } r, c\}. \tag{7}$$

Lemma 4.2.3 showed that L_Y finds all \mathcal{L}^ℓ that are restrictions of a persistence module automorphism \mathcal{L} to the parameter ℓ . For each $L \in L_Y$, we can compute the bar representation of $[y]$ under a new interval decomposition as

$$S_{[y]}^{\mathcal{D} \circ L^{-1}} = \{c_1 \gamma_{i_1}, \dots, c_s \gamma_{i_s} \mid L(d_1 \mathbf{e}_{j_1}^\ell + \dots + d_t \mathbf{e}_{j_t}^\ell) = c_1 \mathbf{e}_{i_1}^\ell + \dots + c_s \mathbf{e}_{i_s}^\ell \text{ for } c_1, \dots, c_s \neq 0\}.$$

We return the collection of all bar representations of $[y]$

$$S_{[y]} = \{S_{[y]}^{\mathcal{D} \circ L^{-1}} \mid L \in L_Y\}.$$

Step (5-b) of Algorithm 2 performs this process.

Lemma 4.2.3 guarantees that the enumerations in Sects. 4.2.1 and 4.2.2 find all terminal classes and bar-representations. Combining the two, we obtain Algorithm 2, which provides a complete enumeration of bar extensions of $\tau \in \text{BC}_k(Z^\bullet)$ to $\text{BC}_k(Y^\bullet)$. In step (1), we fix initial interval decompositions $P\tilde{H}_k(Z^\bullet)$ and $P\tilde{H}_k(Y^\bullet)$. In steps (2) to (4), we compute all terminal classes of bar τ . In step (5), we compute the class extensions of each terminal class and find their bar representations in $\text{BC}_k(Y^\bullet)$.

4.2.3 The case of \mathbb{F}_2 -coefficients and bars with unique death parameters

Algorithm 2 can be substantially simplified when we make two assumptions that are often satisfied when working with real data: that homology is computed with \mathbb{F}_2 coefficients, and that all bars in $\text{BC}_k(Z^\bullet)$ have unique death parameters. Under such assumptions, the following lemma says we can omit the loop over interval decompositions of $P\tilde{H}_k(Z^\bullet)$.

Lemma 4.2.4 *Assume that $\mathbb{F} = \mathbb{F}_2$ and that all bars of $\text{BC}_k(Z^\bullet)$ have unique death parameters. Let $\mathcal{B}, \mathcal{C} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be two different interval decompositions.*

Given a bar $\tau \in BC_k(Z^\bullet)$, let $[\tau_*^B]$ and $[\tau_*^C]$ denote the terminal class for τ under B and C respectively. Then, $[\tau_*^B] = [\tau_*^C]$ in $H_k(Z^{\delta(\tau)-1})$.

Proof Deferred to Appendix D. □

Algorithm 3 describes the modified bar-to-bars extension method under these assumptions.

Example 2 Recall the point clouds P and Q on the torus from Example 1 and Fig. 3, and again let $Z^\bullet = W_{P,Q}^\bullet$, and $Y^\bullet = X_P^\bullet$. Let τ be the highlighted bar in $BC_1(W_{P,Q}^\bullet)$, as shown in Fig. 4. We will apply Algorithm 3 to determine how the selected bar τ can be related to bars in $BC_1(X_P^\bullet)$.

In step (1), we fix the interval decompositions of $P\tilde{H}_k(Z^\bullet)$ and $P\tilde{H}_k(Y^\bullet)$. In step (2), we find the terminal class $[\tau_*]$ of τ ; the cycle representative depicted in the left panel of Fig. 3 is a representative of this $[\tau_*]$. We apply Algorithm 1 to this terminal class and find all class extensions. Finally, we enumerate possible bar representations of the resulting class extensions in $BC_1(Y^\bullet)$. The bottom pane of Fig. 4 illustrates two possible bar-representations of some class extension $[y]$. The union of the green and purple bars illustrates the \mathcal{D} -bar representation $S_{[y]}^{\mathcal{D}}$ under some fixed interval decomposition \mathcal{D} . The collection of green bars represent an alternative bar-representation $S_{[y]}^{\mathcal{D} \circ L^{-1}}$ for some $L \in L_Y$.

The output $S(\tau, Y^\bullet)$ is large, as it is a collection of $S_{[y]}^{\mathcal{D} \circ L^{-1}}$ for all possible combinations of $\ell \in p_Y$, $[y] \in (\mathfrak{E}_{[\tau_*]})_\ell$, and $L \in L_Y$. In this example, p_Y is a set with 18 parameters. Fig. 4 illustrates an example \mathcal{D} -bar representative $S_{[y]}^{\mathcal{D}}$ for $\ell = \min p_Y$, a randomly selected $[y] \in (\mathfrak{E}_{[\tau_*]})_\ell$, and an identity matrix L . For the fixed ℓ and L , there are 2^{14} possibilities for $S_{[y]}^{\mathcal{D} \circ L^{-1}}$. However, every possible $S_{[y]}^{\mathcal{D} \circ L^{-1}}$ is a sum of the green bars and a short bar in $BC_1(X_P^\bullet)$, much like the presentation in Figure 4. There are 14 short bars in $BC_1(X_P^\bullet)$ that one can highlight as the purple bar. Any subsets of the 14 short bars, when combined with the green bars, are valid bar representatives. In all 2^{14} different possibilities, the longest green bar is always present. The observation that every bar representative $S_{[y]}^{\mathcal{D}}$ can be written as a combination of the green bars and a selection of short bars motivated the development of component-wise extension discussed in the following section.

4.3 Component-wise extension

The process of computing the \mathcal{D} -bar representation $S_{[y]}^{\mathcal{D}}$ can be computationally intensive. In step (5-b-iii) of Algorithm 2, we find the bar extension $S_{[y]}^{\mathcal{D}}$ for every $[y] \in (\mathfrak{E}_{[\tau]})_\ell$. From steps (3-b-iii) - (3-b-iv) of Algorithm 1, we know that any $[y] \in (\mathfrak{E}_{[\tau]})_\ell$ has the form

$$[y] = \Upsilon_\ell \left(\sum_{i=1}^m f_i^* [\rho_i^{\mathcal{F}, \ell}] \right) + \sum_{[\rho_j^{\mathcal{F}, \ell}] \in V_{\text{short}}^\ell} \Upsilon_\ell \left(f_j [\rho_j^{\mathcal{F}, \ell}] \right) \tag{8}$$




Fig. 4 Input and output of the bar-to-bars extension in Example 2. (top) Barcode $BC_1(W_{P,Q}^\bullet)$ for the witness complex of point clouds P and Q from Fig. 3. We apply the bar-to-bars extension method to the blue long bar. (bottom) Barcode for the clique complex of the point cloud P from Fig. 3. Green and purple bars together are the \mathcal{D} -bar representative $S_{[y]}^{\mathcal{D}}$ produced by Algorithm 3. Omitting the purple bar produces the alternative bar representation $S_{[y]}^{\mathcal{D} \circ L^{-1}}$ for one choice of $L \in L_Y$

for some $f_j \in \mathbb{F}$. Note that $f_1^*, \dots, f_m^* \in \mathbb{F}$ are fixed from $S_{[\tau]}^{\mathcal{F}} = \{f_1^* \rho_1, \dots, f_m^* \rho_m\}$ in step (2-a) of Algorithm 1. Thus, the number of class extensions $[y] \in (\mathfrak{C}_{[\tau]})_\ell$ we must consider grows exponentially as $|\mathbb{F}|^{|\text{BARS}_{\text{short}}^\ell|}$.

Fortunately, we can leverage the vector space structure on homology to mitigate the resulting explosion in computation time. Suppose that $[y_1], [y_2] \in \tilde{H}_k(Y^\ell)$ and define $S_{[y_1]}^{\mathcal{D}} \oplus S_{[y_2]}^{\mathcal{D}} = S_{[y_1+y_2]}^{\mathcal{D}}$. Given $c \in \mathbb{F}$, define $c \cdot S_{[y]}^{\mathcal{D}} = S_{c[y]}^{\mathcal{D}}$ for any $[y] \in \tilde{H}_k(Y^\ell)$. From Definition 2.2.4, we see that \oplus and \cdot defines a \mathbb{F} -vector space structure on the set $\{S_{[y]}^{\mathcal{D}} \mid [y] \in \tilde{H}_k(Y^\ell)\}$, and that the collection of \mathcal{D} -bar representations of $(\mathfrak{C}_{[\tau]})_\ell$ can be regarded as an affine subspace of this space.

Therefore, instead of computing bar representations for every $[y] \in (\mathfrak{C}_{[\tau]})_\ell$, we can compute bar representations $S_{\Upsilon_\ell(\sum_{i=1}^m f_i^* [\rho_i^{\mathcal{F}, \ell}])}^{\mathcal{D}}$ and $S_{\Upsilon_\ell[\rho^{\mathcal{F}, \ell}]}^{\mathcal{D}}$ for each $[\rho^{\mathcal{F}, \ell}] \in V_{\text{short}}^\ell$. We'll refer to

$$B = \Upsilon_\ell \left(\sum_{i=1}^m f_i^* [\rho_i^{\mathcal{F}, \ell}] \right) \text{ and } O = \left\{ \Upsilon_\ell([\rho^{\mathcal{F}, \ell}]) \mid [\rho^{\mathcal{F}, \ell}] \in V_{\text{short}}^\ell \right\}$$

as the *baseline* and *offset class extensions*, respectively, with $(\mathfrak{C}_{[\tau]})_\ell = B \oplus \text{span}_{\mathbb{F}} O$. We then take

$$B^{\mathcal{D}} = \left\{ S_{\Upsilon_\ell(\sum_{i=1}^m f_i^* [\rho_i^{\mathcal{F}, \ell}])}^{\mathcal{D}} \right\} \text{ and } O^{\mathcal{D}} = \left\{ S_{\Upsilon_\ell([\rho^{\mathcal{F}, \ell}])}^{\mathcal{D}} \mid [\rho^{\mathcal{F}, \ell}] \in V_{\text{short}}^\ell \right\}$$

to be the collections of *baseline* and *offset bar extensions*. Note that $\{S_{[y]}^{\mathcal{D}} \mid [y] \in (\mathfrak{C}_{[\tau]})_\ell\} = \text{span}_{\mathbb{F}} B^{\mathcal{D}} \oplus \text{span}_{\mathbb{F}} O^{\mathcal{D}}$. Perhaps unsurprisingly, we have found it simpler

and faster to visualize the sets $B^{\mathcal{D}}$ and $O^{\mathcal{D}}$ for understanding extensions in examples. Note that the individual members of the baseline and offset class extensions depend on the interval decomposition \mathcal{F} . However, when the baseline and offset class extensions are considered collectively, the collection of class extensions is independent of \mathcal{F} according to Theorem B.2.2. Similarly, the collection of baseline and offset bar extensions is independent of \mathcal{F} .

Remark 3 When the number of bars in the barcode is large, the number of linear systems we must solve to carry out the extension method can be prohibitive, even when we consider only these basis elements. Therefore, it may be useful to omit classes of bars based on length statistics for a null model or domain knowledge. We have found this to be a useful strategy in practice, but leave a detailed discussion for future work.

5 Analogous bars

With the persistent extension method developed in Sect. 4 in hand, we can move on to addressing the motivating question for this paper. Let $Q = \{q_1, \dots, q_m\}$ and $P = \{p_1, \dots, p_n\}$ be two populations equipped with dissimilarity matrices M_Q and M_P , as well as a cross-dissimilarity matrix $M_{Q,P}$. Let X_Q^\bullet and X_P^\bullet each denote the Vietoris-Rips complexes on Q and P . Given a bar $\tau \in \text{BC}_k(X_Q^\bullet)$, our goal is to find all persistent 1-dimensional subspaces of $P\tilde{H}_k(X_P^\bullet)$, represented as collections of bars in $\text{BC}_k(X_P^\bullet)$, which could correspond to features similar to τ under the information provided by the cross-dissimilarity measure.

As discussed in Sect. 3, we will encode the information in $M_{Q,P}$ as a witness complex and utilize the isomorphism in Dowker’s theorem to identify $\text{BC}_k(W_{Q,P}^\bullet)$ and $\text{BC}_k(W_{P,Q}^\bullet)$. We then apply the persistent extension method to associate bars in $\text{BC}_k(X_Q^\bullet)$ and $\text{BC}_k(X_P^\bullet)$ to those in the witness barcodes. There are two perspectives that we can take.

1. First, a *feature-centric* approach, where the goal is to study how a particular bar $\tau \in \text{BC}_k(X_Q^\bullet)$ is represented in $\text{BC}_k(X_P^\bullet)$. To streamline this approach, we give two truncated variations of Algorithm 2. Algorithm 4, the bar-to-class extension method, returns only the collection of class extensions, thereby eliminating the need to iterate over the interval decompositions of $P\tilde{H}_k(Y^\bullet)$. Algorithm 5, the class-to-bar extension method, takes as input an explicit homology class $[\tau] \in \tilde{H}_k(Z^\psi)$, which eliminates the iteration over interval decompositions of $P\tilde{H}_k(Z^\bullet)$. Given a bar $\tau \in \text{BC}_k(X_Q^\bullet)$, we first apply Algorithm 4 to τ , producing a family $E(\tau, W_{Q,P}^\bullet)$ of class extensions in $P\tilde{H}_k(W_{Q,P}^\bullet)$. We find the corresponding collection of homology classes $E(\tau, W_{P,Q}^\bullet)$ in $P\tilde{H}_k(W_{P,Q}^\bullet)$ by applying Dowker’s Theorem to elements of $E(\tau, W_{Q,P}^\bullet)$. Finally, to each class extension $[\sigma] \in E(\tau, W_{P,Q}^\bullet)$, we apply Algorithm 5 to find a family of bar extensions $S(\tau, X_P^\bullet) = \bigcup_{[\sigma] \in E(\tau, W_{P,Q}^\bullet)} S([\sigma], X_P^\bullet)$ in $\text{BC}_k(X_P^\bullet)$. Any bar extension in $S([\sigma], X_P^\bullet) \subseteq S(\tau, X_P^\bullet)$ is *analogous to τ through $[\sigma]$* . Algorithm 6 provides a complete account of this approach.

2. Second, a *similarity-centric* approach, where the goal is to understand what relationship a particular bar $\tau \in \text{BC}_k(W_{Q,P}^\bullet)$ indicates between bars in $\text{BC}_k(X_Q^\bullet)$ and $\text{BC}_k(X_P^\bullet)$. In this case, we apply Dowker's Theorem to find the corresponding bar $\tau' \in \text{BC}_k(W_{P,Q}^\bullet)$, and then we apply the persistent extension method to τ and τ' respectively. We end up with a collections of bar extensions $S(\tau, X_Q^\bullet)$ of τ to $\text{BC}_k(X_Q^\bullet)$ and $S(\tau', X_P^\bullet)$ of τ' to $\text{BC}_k(X_P^\bullet)$. Any pair of extensions in $S(\tau, X_Q^\bullet) \times S(\tau', X_P^\bullet)$ can be viewed as *analogous through* τ . This approach is described in Algorithm 7.

Algorithm 4 Extension method (bar-to-class extension)

(H) **Input:**

- filtered simplicial complexes Z^\bullet and Y^\bullet on vertex set P ,
- a bar $\tau \in \text{BC}_k(Z^\bullet)$,

Output:

- a family $E(\tau, Y^\bullet)$ of collections of class extensions of τ to $P\tilde{H}_k(Y^\bullet)$.

Steps:

1. Fix an interval decomposition \mathcal{B} of $P\tilde{H}_k(Z^\bullet)$ ¹⁵.
 2. Perform steps (2) to (5-a) of bar-to-bars extension (Algorithm 2). Let T denote the collection of possible terminal classes computed in Algorithm 2 step (4). For each $[\tau] \in T$, let $\mathcal{E}_{[\tau]}$ denote the class extensions of $[\tau]$ computed in Algorithm 2 step (5-a).
 3. Return $E(\tau, Y^\bullet) = \bigcup_{[\tau] \in T} \mathcal{E}_{[\tau]}$
-

Both methods use Dowker's Theorem to compare $W_{Q,P}^\bullet$ and $W_{P,Q}^\bullet$ and use the extension method to compare the Witness complexes and Vietoris-Rips complexes. Because the involved algorithms enumerate different collections of homology classes, in general, we expect that the relationship between these two approaches may be quite intricate. Nonetheless, we believe it is useful to study how they relate in a simple example. In the following section, we provide an empirical comparison of the two perspectives on a slightly modified version of our earlier torus example which demonstrates that, in simple settings, the outputs of the two methods can coincide in a reasonable way.

¹⁵ The output of the algorithm does not depend on \mathcal{B} . We use \mathcal{B} as an initial point to enumerate all possible interval decompositions of $P\tilde{H}_k(Z^\bullet)$ in step (2).

Algorithm 5 Extension method (class-to-bar extension)

Input:

- filtered simplicial complexes Z^\bullet and Y^\bullet on vertex set P ,
- a parameter ψ , and
- a homology class $[\tau] \in \tilde{H}_k(Z^\psi)$.

Output:

- a collection $S([\tau], Y^\bullet)$ of bar extensions of $[\tau]$ to $\text{BC}_k(Y^\bullet)$.

Steps:

1. Let p_Y and $\mathfrak{E} = \bigcup_{\ell \in p_Y} \mathfrak{E}_\ell$ be the outputs of class-to-class extension (Algorithm 1).
2. Fix interval decomposition \mathcal{D} of $P\tilde{H}_k(Y^\bullet)$ ¹⁶.
3. For each $\ell \in p_Y$ ¹⁷,
 - (a) Fix an ordering on $\text{BC}_k(Y^\ell) = \{\gamma_1, \dots, \gamma_n\}$.
 - (b) Enumerate the collection of matrices

$$L_Y = \{L \in \mathbb{F}^{n \times n} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\gamma_r) > \beta(\gamma_c) \text{ or } \delta(\gamma_r) > \delta(\gamma_c) \text{ for all } r, c\}$$

- (c) For each $[y] \in \mathfrak{E}_\ell$, compute the \mathcal{D} -bar representation $S_{[y]}^{\mathcal{D}} = \{d_1\gamma_{j_1}, \dots, d_t\gamma_{j_t}\}$.
- (d) For each $L \in L_Y$, compute

$$S_{[y]}^{\mathcal{D} \circ L^{-1}} = \{c_1\gamma_{i_1}, \dots, c_s\gamma_{i_s} \mid L(d_1\mathbf{e}_{j_1}^\ell + \dots + d_t\mathbf{e}_{j_t}^\ell) = c_1\mathbf{e}_{i_1}^\ell + \dots + c_s\mathbf{e}_{i_s}^\ell\},$$

where \mathbf{e}_j^ℓ is the basis vector for γ_j at parameter ℓ (Definition 2.2.1).

4. Return $S([\tau], Y^\bullet) = \{S_{[y]}^{\mathcal{D} \circ L^{-1}} \mid \ell \in p_Y, [y] \in \mathfrak{E}_\ell, L \in L_Y\}$.

Example 3 Take P to be the collection of points sampled uniformly from a flat torus considered in Examples 1 and 2, and let Q be a set of points sampled along a deformation of one of its essential circles, as illustrated in Fig. 5. We have chosen Q to take a substantial detour from the usual “taut” essential circle so that both the Vietoris-Rips complex X_Q^\bullet and the witness complex $W_{Q,P}^\bullet$ will, for some range of parameters, contain homology classes representing both of the essential circles in the torus, but that the primary feature will be the vertical circle. The left panel in Fig. 5 shows the three relevant barcodes $\text{BC}_1(X_Q^\bullet)$, $\text{BC}_1(W_{Q,P}^\bullet) = \text{BC}_1(W_{P,Q}^\bullet)$, and $\text{BC}_1(X_P^\bullet)$. Note that all bars have unique death times.

Feature-centric analogous bars: We will apply the feature-centric analogous bars method (Algorithm 6) to investigate the representation of the longest bar $\tau \in \text{BC}_1(X_Q^\bullet)$ in $\text{BC}_1(X_P^\bullet)$. Step (1) is to apply the bar-to-class extension (Algorithm 4) to find all class extensions $E(\tau, W_{Q,P}^\bullet)$. Since we are working with \mathbb{F}_2 coefficients, the terminal class $[\tau_*] \in \tilde{H}_1(X^{\delta(\tau)-1})$ is independent of the interval decomposition of $P\tilde{H}_1(X_Q^\bullet)$ by Lemma 4.2.4. Algorithm 4 simplifies since T , the collection of possible terminal

¹⁶ The output of the algorithm does not depend on \mathcal{D} . We use \mathcal{D} as an initial point to enumerate all possible interval decompositions of $P\tilde{H}_k(Y^\bullet)$ in step (3).

¹⁷ Step (3) of Algorithm 5 is step (5-b) of the bar-to-bars extension (Algorithm 2).

Algorithm 6 Analogous bars method (feature-centric)

Input:

- filtered clique complexes X_Q^\bullet and X_P^\bullet ,
- cross-dissimilarity witness complexes $W_{Q,P}^\bullet$ and $W_{P,Q}^\bullet$, and
- a bar $\tau \in \text{BC}_k(X_Q^\bullet)$.

Output:

- a family $S(\tau, X_P^\bullet)$ of collections of weighted bars (Definition 2.2.3) in $\text{BC}_k(X_P^\bullet)$ related to τ

Steps:

1. Apply bar-to-class extension (Algorithm 4) with inputs $Z^\bullet = X_Q^\bullet$, $Y^\bullet = W_{Q,P}^\bullet$, and τ . Denote by $E(\tau, W_{Q,P}^\bullet)$ the resulting collection of class extensions of τ in $P\tilde{H}_k(W_{Q,P}^\bullet)$.
2. Apply Dowker’s Theorem to each homology class in $E(\tau, W_{Q,P}^\bullet)$ to find the collection

$$E(\tau, W_{P,Q}^\bullet) = \{[x_P] \in W_{P,Q}^\epsilon \mid [x_P] \text{ is Dowker dual to } [x_Q] \in W_{Q,P}^\epsilon, \text{ where } [x_Q] \in E(\tau, W_{Q,P}^\bullet)\}.$$

3. For each $[\sigma] \in E(\tau, W_{P,Q}^\bullet)$,
 - (a) Find the death time $\delta([\sigma])$.
 - (b) Apply class-to-bar extension (Algorithm 5) with $Z^\bullet = W_{P,Q}^\bullet$, $Y^\bullet = X_P^\bullet$, $\psi = \delta([\sigma]) - 1$, and $[\sigma] \in \tilde{H}_k(W_{P,Q}^\psi)$ as input to find a family of bar extensions $S([\sigma], X_P^\bullet)$ in $\text{BC}_k(X_P^\bullet)$.
4. Return $S(\tau, X_P^\bullet) = \bigcup_{[\sigma] \in E(\tau, W_{P,Q}^\bullet)} S([\sigma], X_P^\bullet)$.

Algorithm 7 Analogous bars method (similarity-centric)

Input:

- filtered clique complexes X_Q^\bullet and X_P^\bullet ,
- cross-dissimilarity witness complexes $W_{Q,P}^\bullet$ and $W_{P,Q}^\bullet$, and
- a bar $\tau \in \text{BC}_k(W_{Q,P}^\bullet)$.

Output:

- collections $S(\tau, X_Q^\bullet)$ and $S(\tau', X_P^\bullet)$ of weighted bars (Definition 2.2.3) in $\text{BC}_k(X_Q^\bullet)$ and $\text{BC}_k(X_P^\bullet)$ respectively, for which pairs $(\sigma_Q, \sigma_P) \in S(\tau, X_Q^\bullet) \times S(\tau', X_P^\bullet)$ are related through τ

Steps:

1. Apply Dowker’s Theorem to τ to find the corresponding bar $\tau' \in \text{BC}_k(W_{P,Q}^\bullet)$.
2. Apply bar-to-bars extension (Algorithm 2) with $Z^\bullet = W_{Q,P}^\bullet$, $Y^\bullet = X_Q^\bullet$, and τ as input and output a family $S(\tau, X_Q^\bullet)$ of bar extensions of τ in $\text{BC}_k(X_Q^\bullet)$.
3. Apply bar-to-bars extension (Algorithm 2) with $Z^\bullet = W_{P,Q}^\bullet$, $Y^\bullet = X_P^\bullet$, and τ' as input and output a family $S(\tau', X_P^\bullet)$ of bar extensions of τ' in $\text{BC}_k(X_P^\bullet)$.
4. Return $S(\tau, X_Q^\bullet)$ and $S(\tau', X_P^\bullet)$.

classes, contains only one element $[\tau_*]$. The left-middle panel in the right column in Fig. 5 illustrates an example cycle in the class extension $[x_Q] \in E(\tau, W_{Q,P}^\bullet)$. The class $[x_Q]$ is the baseline class extension.

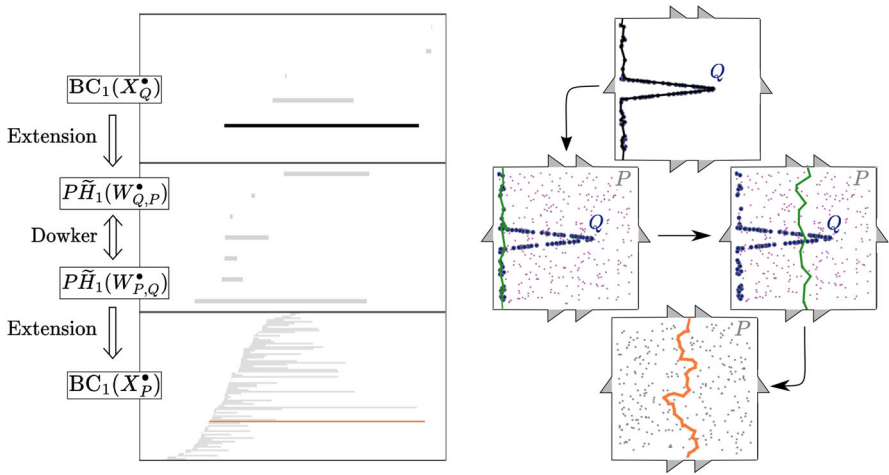


Fig. 5 Illustration of the feature-centric analogous bars method. Let P be a point cloud sampled from a torus from Example 1, and let Q be a deformation of an essential circle that witnesses both essential circles of the torus. (left column) Beginning with the black bar in $BC_1(X_Q^\bullet)$, apply persistent extension to obtain class extensions in $W_{Q,P}^\bullet$. Using the isomorphism in Dowker’s theorem, find the corresponding class extension in $W_{P,Q}^\bullet$. Apply persistent extension again to obtain the orange bar in $BC_1(X_P^\bullet)$. The highlighted bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous bars. (right column) Cycle representatives illustrating the classes obtained in this example. (top) Chain in X_Q^\bullet representing the long bar in $BC_1(X_Q^\bullet)$. (middle-left) Cycle representative in $W_{Q,P}^\bullet$ obtained by persistent extension of the bar in the top row is transformed into (middle-right) a cycle in $W_{P,Q}^\bullet$, supported on P , via Dowker’s theorem. (bottom) Applying persistent extension to this cycle, we obtain a corresponding cycle in X_P^\bullet

In step (2) of Algorithm 6, we find the family of homology classes that are dual to $E(\tau, W_{Q,P}^\bullet)$ via Dowker’s Theorem. The right-middle panel in the right column of Fig. 5 illustrates the dual $[x_P] \in E(\tau, W_{P,Q}^\bullet)$ to the previously illustrated $[x_Q]$.

In step (3) of Algorithm 6, for each class extension $[\sigma] \in E(\tau, W_{P,Q}^\bullet)$, we apply the class-to-bar extension method (Algorithm 5) to find the bar extensions $S([\sigma], X_P^\bullet)$ in $BC_1(X_P^\bullet)$. The baseline bar extension is illustrated on $BC_1(X_P^\bullet)$ in the left column of Fig. 5. The highlighted bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous through $[\sigma]$. The bottom panel in the right column of Fig. 5 illustrates the corresponding cycle representative of the baseline bar extension.

Note that Fig. 5 illustrates only one example output of the feature-centric analogous bars method. Specifically, it illustrates the output obtained by computing the baseline persistence extensions in steps (1) and (3-b) of Algorithm 6. Other analogous bars in $BC_1(X_P^\bullet)$ are combinations of the highlighted orange bar and short bars in $BC_1(X_P^\bullet)$. *Similarity-centric analogous bars:* We now illustrate the similarity-centric analogous bars method on the same dataset. The presence of a long bar in $BC_1(W_{P,Q}^\bullet)$ in Fig. 6 suggests that there is a significant feature present in both P and Q . Thus, we select the long interval $\tau \in BC_1(W_{Q,P}^\bullet)$ to which to apply the similarity-centric analogous bars method (Algorithm 7).

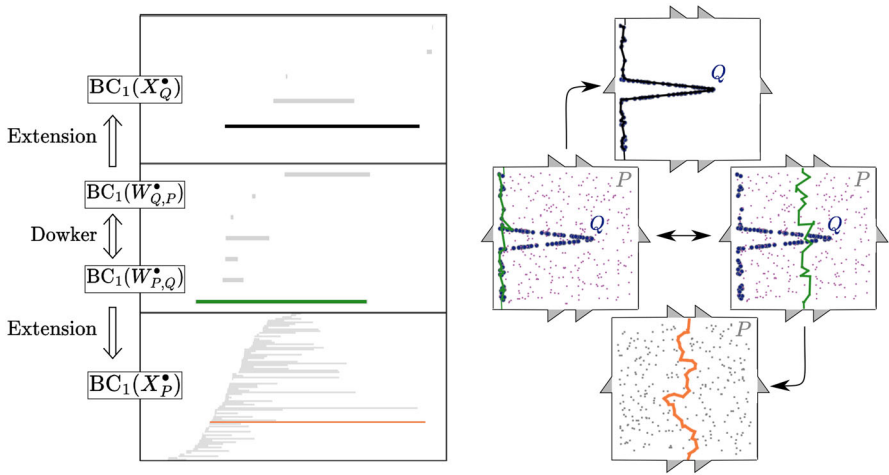


Fig. 6 Illustration of the similarity-centric analogous bars method. Let P be a point cloud sampled from a torus from in Examples 1, and let Q be a deformation of an essential circle that witnesses both essential circles of the torus. (left column) Beginning with the green bar in $BC_1(W_{P,Q}^\bullet) = BC_1(W_{Q,P}^\bullet)$, we apply persistent extension twice to obtain baseline black bar extension in $BC_1(X_Q^\bullet)$ and orange bar extension in $BC_1(X_P^\bullet)$. The highlighted bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous bars. (right column) Cycle representatives illustrating the classes obtained in this example. (middle) Cycle in $W_{Q,P}^\bullet$ representing the selected bar in the left column and a corresponding representative in $W_{P,Q}^\bullet$ obtained via Dowker's theorem. (top, bottom) Cycles in X_Q^\bullet and X_P^\bullet respectively representing the corresponding bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$

We apply Dowker's Theorem to identify the corresponding bar $\tau' \in BC_1(W_{P,Q}^\bullet)$, though morally these are identical. Computationally, we need to change representations to proceed. We then apply the bar-to-bar extension method (Algorithm 2) to find bar extensions of τ in $BC_1(X_Q^\bullet)$ and τ' in $BC_1(X_P^\bullet)$. The left column of Fig. 6 illustrates the selected bar $\tau \in BC_1(W_{Q,P}^\bullet)$ and the baseline bar extensions in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ respectively. The highlighted bars of $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous through τ . The cycle representatives show that these analogous bars, indeed, trace qualitatively similar homology classes.

Note that Fig. 6 illustrates only one example output of the similarity-centric analogous bars method. Specifically, it illustrates the output obtained by computing the baseline persistence extensions in steps (2) and (3) of Algorithm 7.

Comparing Figs. 5 and 6, we see that at the level of barcodes on simple examples, the feature- and similarity-centric methods produce comparable and reasonable results.

6 Applications

Finally, we demonstrate the use of the methods we have developed on some simple data sets: we apply the analogous bars method to identify corresponding features from two samples of the same space and to determine whether topological features are

retained under clustering; then, we use persistent extension on its own to investigate how topological features are transformed by dimensionality reduction.

6.1 Two samples from the same distribution

One of the most common tasks in topological data analysis is the recovery of topological features from a finite sample of a probability distribution on a metric space. If we have multiple samples, for example, from different measurements of a system, understanding how the features in those samples are related would provide utility in understanding what these features mean in terms of the system being studied. This is one of the simplest settings in which we can apply the analogous bars method, using the block sub-matrix of the distance matrix that measures distance between the two samples as a measure of cross-dissimilarity.

In Fig. 7, we illustrate this approach using two point clouds, P and Q , sampled uniformly and independently from a flat torus. We consider a choice of feature $\tau \in \text{BC}_1(X_Q^\bullet)$, and locate the corresponding feature in $\text{BC}_1(X_P^\bullet)$ using the feature-centric analogous bars method with τ as input. The highlighted bars of $\text{BC}_1(X_Q^\bullet)$ and $\text{BC}_1(X_P^\bullet)$ are analogous bars that represent the same essential circle.

Note that Fig. 7 illustrates only one example output of the feature-centric analogous bars method. The particular output was obtained by computing the baseline persistence extensions in steps (1) and (3-b) of Algorithm 6.

6.2 Clustering

The analogous bars method can similarly be applied to determine if topological features of a point cloud Q are preserved by clustering. In Fig. 8, we demonstrate such a comparison using the distance matrix between points Q and cluster centroids P as our cross-dissimilarity matrix. Because the witness barcode $\text{BC}_1(W_{P,Q}^\bullet)$ has a long bar, we expect that the two point clouds share a feature. Applying the similarity-centric analogous bars method, we find the analogous bars in $\text{BC}_1(X_Q^\bullet)$ and $\text{BC}_1(X_P^\bullet)$ and conclude that the major topological feature of Q is, in fact, retained by P .

6.3 Dimensionality reduction

Most dimensionality reduction techniques retain the identities of the projected points. Thus, we can apply the persistent extension method to the Vietoris-Rips complexes of the original point cloud and the projected point cloud to study whether, and how, the resulting lower-dimensional point cloud retains topological features of interest. Here, we provide two toy examples using the extension method in the context of dimensionality reduction via principal component analysis.

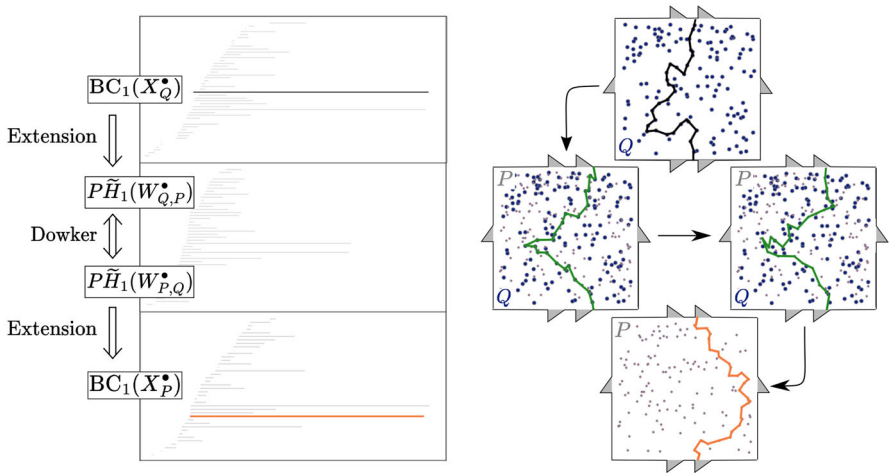


Fig. 7 Using the feature-centric analogous bars method to identify features from independent samples of a torus. Grey points P and blue points Q are independently sampled from the uniform distribution on the torus. (left column) Dimension-one barcodes for X_Q^\bullet , $W_{Q,P}^\bullet$ and X_P^\bullet . The black bar in $\tau \in BC_1(X_Q^\bullet)$ is the input to the method, and the orange bar in $BC_1(X_P^\bullet)$ is the baseline bar extension. The highlighted bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous bars. No bars in $BC_1(W_{P,Q}^\bullet)$ are highlighted as Algorithm 6 bypasses translation through this barcode. (right column) Cycle representatives at each stage of the feature-centric analogous bars method. (top) Cycle representative of τ supported on Q . (middle-left) Representatives of a homology class $[x_Q] \in E(\tau, W_{Q,P}^\bullet)$ produced by step (1) of Algorithm 6 and (middle-right) its Dowker dual $[x_P] \in E(\tau, W_{P,Q}^\bullet)$. (bottom) Representative of $\tau' \in S(\tau, X_P^\bullet)$ produced by step (3) of Algorithm 6 applied to $[x_P]$ in (middle-right)

6.3.1 Trefoil knot

First, we consider a collection of points $Q \subseteq \mathbb{R}^3$ sampled from the trefoil knot depicted in Fig. 9 and take $\pi(Q) \subseteq \mathbb{R}^2$ to be its 2-dimensional projection via PCA. We apply the persistent extension method to these two point clouds to identify how the dominant 1-dimensional homological feature of the trefoil knot is represented in its two-dimensional projection.

The 1-dimensional barcodes $BC_1(X_Q^\bullet)$ and $BC_1(X_{\pi(Q)}^\bullet)$ are depicted in the middle panes of Fig. 9. We apply the persistent extension method to the long bar in $BC_1(X_Q^\bullet)$. The vertical line in the bottom pane indicates the ℓ_0 parameter, and the highlighted bars in $BC_1(X_{\pi(Q)}^\bullet)$ correspond to one of the bar extensions; in fact, the top bar of $BC_1(X_{\pi(Q)}^\bullet)$ corresponds to the baseline bar extension, and the three bars whose interiors intersect the ℓ_0 -line independently make up the offset bar extensions and so can be included in any combination. The pictured extension indicates that the S^1 described by the 3-dimensional trefoil knot can be represented by a combination of the four features detected in $BC_1(X_{\pi(Q)}^\bullet)$. Examining the cycle representatives of the highlighted bars of $BC_1(X_{\pi(Q)}^\bullet)$ depicted on the right of Fig. 9 illustrates why this is a reasonable answer: the homology class of the baseline bar encloses the region “internal” to the knot projection.

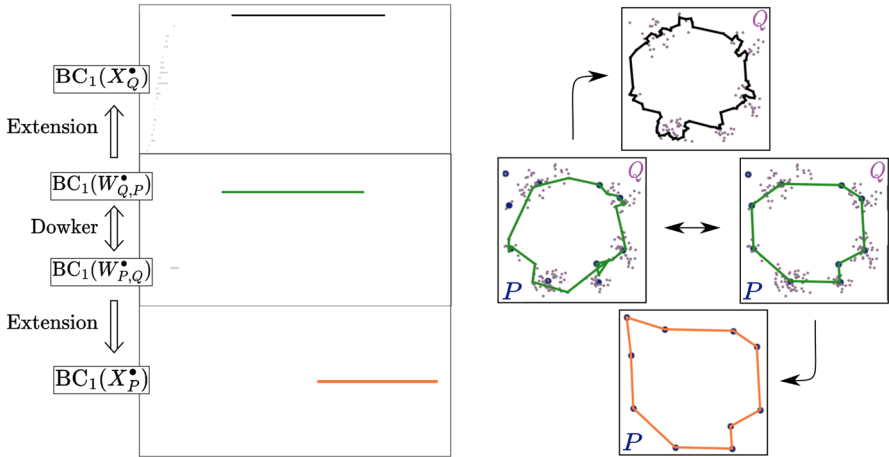


Fig. 8 Application of similarity-centric analogous bars to detect topological features preserved by clustering. Grey points Q sampled from a multimodal distribution in the plane, and blue centroids P induced by clustering. (left column) Barcodes for X_Q^\bullet , $W_{Q,P}^\bullet$ and X_P^\bullet . The green bar $\tau \in BC_1(W_{Q,P}^\bullet)$ is the input to the method, and the black bar in $BC_1(X_Q^\bullet)$ and orange bar in $BC_1(X_P^\bullet)$ are the baseline bar extensions. These highlighted bars in $BC_1(X_Q^\bullet)$ and $BC_1(X_P^\bullet)$ are analogous bars. (right column) Cycle representatives for the input, its Dowker dual, and the outputs. (middle-left) Representative of the cycle in $W_{Q,P}^\bullet$ corresponding to the selected bar and (middle-right) its Dowker dual in $W_{P,Q}^\bullet$. (top) Cycle representatives in the original data Q and (bottom) in the centroids of the clusters P confirm that this topological feature of Q is preserved by clustering




Fig. 9 Application of persistent extension method to planar projection of the trefoil knot via PCA. (left) Blue points Q sampled from a standard trefoil knot in three dimensions. (middle) One-dimensional barcodes for the sample Q and its projection $\pi(Q)$ to the first two loading vectors via PCA. Selecting the long orange bar $\tau \in BC_1(X_Q^\bullet)$ and applying bar-to-bar persistent extension (Algorithm 3), we find that the top black bar in $BC_1(X_{\pi(Q)}^\bullet)$ is our baseline bar extension and the remainder are individually offsets. (right) Cycle representatives illustrating the four bars appearing in the set of possible extensions

6.3.2 Cylinder

Sometimes, applying dimensionality reduction to a point cloud flattens the object too much to easily visualize what happens to the topological features. Let Q be a set of points sampled from a cylinder in \mathbb{R}^3 , as shown in the left panel of Fig. 10, and let $\pi(Q)$ be the result of projecting the points onto their first two loading vectors. As before, we consider X_Q^\bullet and $X_{\pi(Q)}^\bullet$, with 1-dimensional barcodes as shown in the




Fig. 10 Application of persistent extension method to planar projection of a sampled cylinder via PCA. (left) Blue points Q sampled from a cylinder in three dimensions. (middle) One-dimensional barcodes for the sample Q and its projection $\pi(Q)$ to the first two loading vectors via PCA. Selecting the long bar $\tau \in BC_1(X_Q^\bullet)$ and applying bar-to-bar persistent extension (Algorithm 3), we find that the green bars in $BC_1(X_{\pi(Q)}^\bullet)$ form our baseline bar extension. (right) Cycle representative for the baseline class extension. Green represents the cycle representative of the sum of the homology classes corresponding to the two green bars. Cycle representative illustrates how the homology class in the cylinder is flattened

middle pane of Fig. 10. Selecting the highlighted long bar in $BC_1(X_Q^\bullet)$ and applying the persistent extension method produces the highlighted bars in $BC_1(X_{\pi(Q)}^\bullet)$ as our baseline bar extensions. The right panel in Fig. 10 shows the cycle representative of the baseline class extension corresponding to the baseline bar extension in $BC_1(X_{\pi(Q)}^\bullet)$. Other choices of extension and representatives will recover how paths “around” the cylinder project onto this plane.

7 Conclusion

In this paper, we addressed the problem of comparing topological features computed from a pair of finite systems equipped with internal and cross-dissimilarity measures. We focused on the case where a triple of matrices provides a summary of the observations about such systems, a setting that commonly arises in applications. To leverage this information we developed two methods:

1. the *persistent extension* method, a zig-zag computation which provides a means for testing potential relationships between two complexes built on the same vertex set; and,
2. the *analogous bars* method, which leverages the Dowker complex of a cross-dissimilarity matrix to bridge between complexes supported on two distinct vertex sets.

Our examples in this paper focused on point cloud data, a common setting for topological data analysis. However, there are a variety of settings where these tools are likely to be of use. In particular, there is no explicit requirement that we use clique complexes; we focused on this example because it is the most commonly studied and easily implemented with existing software packages.

For some time there has been discussion of using zig-zag based analyses for time-varying simplicial complexes. The persistent extension method provides a concrete approach for studying how features evolve within a single system. It can also be applied to study how different filtrations affect the topology of a complex. Indeed,

the dimensionality reduction examples in Sect. 6.3 can be cast in this light. We can also envision using this approach to investigate two-dimensional persistence modules, studying how features in given one-dimensional submodules manifest in others.

The analogous bars method provides a tool for studying simultaneously observed systems and determining whether their structure and dynamics capture common features. It is worth noting that while many potential applications will involve an underlying function, the Dowker complex was originally conceived as a topological method for studying relations (Dowker 1952), and so our method does not require such an assumption to provide sensible answers. For example, it could be used to study a pair of systems driven by a common topologically interesting input signal and detect signatures of that signal which are preserved in each. In forthcoming work, the authors apply these methods for precisely this kind of analysis.

While we have laid out the theoretical and computational foundation for using these tools, we have not attempted to address a large number of immediate questions that arose while we were writing this paper. We close with a selection of these that we find compelling.

1. Can we systematically choose the parameter ψ used in Algorithm 1 using features of the data? That is, given a homology class $[\tau]$ and its death time $\delta([\tau])$, what is the optimal parameter ψ ? In this paper, we suggested $\psi = \delta([\tau]) - 1$ as a safe choice. Empirically, we have found that when the homology class $[\tau]$ is represented by multiple bars, for example, as $S_{[\tau]}^{\mathcal{B}} = \{\tau_1, \dots, \tau_n\}$, then the class extension that one intuitively expects may arise only when we consider a parameter ψ that is smaller than $\delta([\tau])$. This appears to be the result of homology classes becoming too large as the parameter increases so that the intuitive class extension becomes homologous to unexpected homology classes.
2. How do we characterize the quality of the output of these methods? For example, suppose we consider Vietoris-Rips complexes for samples from some distribution on underlying metric spaces X and Y , where our cross-dissimilarity measure is given by $d(f(x), y) + \eta$ for a noise distribution η and some underlying map $f : X \rightarrow Y$. Can we analytically or computationally give a measure of confidence that a pair of analogous bars is correct?
3. If we consider a sequence of Vietoris-Rips complexes built from increasingly large samples from a distribution on a metric space, do the sequence of analogous bars obtained from these methods stabilize in some manner, as we might hope they would?
4. Under what conditions do the feature-centric and similarity-centric analogous bars methods produce comparable results? As the feature-centric method is more intuitive and the similarity-centric method can be much more computationally efficient, understanding how they compare would be useful.
5. How stable are the analogous bars under deformation of the original complexes and the cross-dissimilarity matrix? We expect the standard stability theorems of persistent homology (Chazal et al. 2009; Cohen-Steiner et al. 2005) and of Dowker persistence diagrams (Chowdhury and Mémoli 2018) will provide a first approximation to this answer, but can we say more?

6. There is a form of inverse problem suggested by this work: can we use information about semantics of topological features, such as knowledge that we have two samples from the same distribution, to infer information about relationships between the underlying points when no explicit pointwise cross-dissimilarity measure is available? That is, what constraints does the existence of an analogous class put on cross-dissimilarity measures?
7. Can we generalize the persistent extension and analogous bars method to barcodes with non-unique birth-death pairs? Recent work by Jacquard et al. (2022) characterizes interval decompositions of a barcode, including barcodes with repeated bars. Incorporating the characterizations will generalize the methods. The challenge lies in navigating the large search space of interval decompositions.
8. Recent work by Gregory Henselman-Petrusek and the second author Ghrist and Henselman-Petrusek (2021) develops a framework for persistence over more general (non-field) coefficient systems. Can we leverage this to generalize the algorithms developed in the current paper to, for example, the setting of lattices?

Acknowledgements The authors would like to thank Gregory Henselman-Petrusek for a great many enlightening conversations. CG is supported by NSF-1854683 and AFOSR FA9550-21-1-0266. HY is supported by NSF-1854683 and ONR N00014-16-1-2010. RG is supported by ONR N00014-16-1-2010 and NSF-1934960.

Declarations

Declarations On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Proofs of Lemmas 2.2.6 and 2.2.7

A.1 Proof of Lemma 2.2.6

Lemma 2.2.6 *Given a filtered topological space Z^\bullet , assume that if $\tau \neq \tau' \in \text{BC}_k(Z^\bullet)$ then either $\beta(\tau) \neq \beta(\tau')$ or $\delta(\tau) \neq \delta(\tau')$. Let $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$ be an automorphism of the barcode module. Then $L_{j,j}^\ell \neq 0$ for every $1 \leq j \leq |\text{BC}_k(Z^\ell)|$ and $\ell \in \mathbb{Z}$.*

Proof Fix $\ell \in \mathbb{Z}$ and $1 \leq j \leq |\text{BC}_k(Z^\ell)|$. Order the bars of $\text{BC}_k(Z^\ell)$ that are alive at ℓ by decreasing birth parameters. If multiple bars share the same birth parameter, then order them by decreasing death parameters. Let $\tau_1, \dots, \tau_{|\text{BC}_k(Z^\ell)|}$ be the resulting order. Without loss of generality, assume that each bar corresponds to the the basis

vectors $\mathbf{e}_1^\ell, \dots, \mathbf{e}_{|\text{BC}_k(Z^\ell)|}^\ell$ of $I^\ell = (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\ell$. For every $r < c$, either $\beta(\tau_r) > \beta(\tau_c)$ or $\delta(\tau_r) > \delta(\tau_c)$. Thus, by Lemma 2.2.5, $L_{r,c}^\ell = 0$ for every $r < c$.

Proceed by induction. From Lemma 2.2.5, we know that $L_{1,c}^\ell = 0$ for all $c > 1$. Since L^ℓ is a linear isomorphism, it cannot have a row of zeros. Thus, $L_{1,1}^\ell \neq 0$.

For the inductive step, assume that $L_{j-1,j-1}^\ell \neq 0$. We will show that $L_{j,j}^\ell \neq 0$ by assuming the contrary. Assume that $L_{j,j}^\ell = 0$. Let $\mathbf{r}_1, \dots, \mathbf{r}_j$ each denote the first j row vectors of L^ℓ . Let \mathbf{r}_i^* denote the vector in \mathbb{F}^{j-1} that consists of the first $j - 1$ components of \mathbf{r}_i for $i = 1, \dots, j$. Since $\mathbf{r}_1^*, \dots, \mathbf{r}_j^*$ are j vectors in \mathbb{F}^{j-1} , the collection is linearly dependent. By construction, each \mathbf{r}_i can be obtained from \mathbf{r}_i^* by appending zeros. Thus, the collection $\mathbf{r}_1, \dots, \mathbf{r}_j$ is also linearly dependent. This contradicts the fact that L^ℓ is a linear isomorphism. Thus, it must be the case that $L_{j,j}^\ell \neq 0$. □

A.2 Proof of Lemma 2.2.7

Lemma 2.2.7 *Let $\text{BC}_k(Z^\bullet)$ be the barcode of some persistence module $P\tilde{H}_k(Z^\bullet)$. Assume that $\beta(\tau) \neq \beta(\tau')$ or $\delta(\tau) \neq \delta(\tau')$ for all $\tau \neq \tau' \in \text{BC}_k(Z^\bullet)$. Given $\ell \in \mathbb{Z}$, let $\tau_1, \dots, \tau_{|\text{BC}_k(Z^\ell)|}$ be the bars in $\text{BC}_k(Z^\ell)$. For $1 \leq t \leq |\text{BC}_k(Z^\ell)|$, let $L \in \mathbb{F}^{|\text{BC}_k(Z^\ell)| \times t}$ be any matrix whose entries satisfy*

$$L_{r,r} \neq 0 \text{ and } L_{r,c} = 0 \text{ if } \beta(\tau_r) > \beta(\tau_c) \text{ or } \delta(\tau_r) > \delta(\tau_c) \text{ for all } r, c. \tag{4}$$

Then, the columns of L are linearly independent.

Proof Without loss of generality, assume that the bars of $\text{BC}_k(Z^\ell)$ are ordered by decreasing birth parameter, and if multiple bars share the same birth parameters, assume that the bars are ordered by decreasing death parameters. So if $r < c$, then $\beta(\tau_r) > \beta(\tau_c)$ or $\delta(\tau_r) > \delta(\tau_c)$. Fix $1 \leq t \leq |\text{BC}_k(Z^\ell)|$. Let L be any matrix satisfying Eq. 4. By construction, $L_{r,c} = 0$ for all $r < c$. Let $\{L_{*,1}, \dots, L_{*,t}\}$ denote the column vectors of L . Assume that

$$c_1 L_{*,1} + \dots + c_t L_{*,t} = \mathbf{0} \tag{A1}$$

Considering the first component of Eq. A1, we know that

$$c_1 L_{1,1} + \dots + c_t L_{1,t} = 0$$

By construction, we know that $L_{1,1} \neq 0$, and $L_{1,2} = \dots = L_{1,t} = 0$. So $c_1 = 0$.

Proceed by induction. Assume $c_1 = \dots = c_{j-1} = 0$. The j^{th} component of Eq. A1 is

$$c_j L_{j,j} + \dots + c_t L_{j,t} = 0$$

Again, by construction, we know that $L_{j,j} \neq 0$, and $L_{j,j+1} = \dots = L_{j,t} = 0$. So $c_j = 0$. By induction, $c_1 = \dots = c_t = 0$, and the column vectors of L are linearly independent. \square

Appendix B Proof of Theorem 4.1.4

To prove Theorem 4.1.4, we need to show the following facts:

- the collections p_Y and $\mathfrak{R}_\ell^{\mathcal{F}}$ are independent of the choice of interval decomposition \mathcal{F} ,
- the collection $\mathfrak{R}_\ell^{\mathcal{F}}$ coincides with R_ℓ , the set of restrictions as defined in Definition 4.1.1, and
- the output of Algorithm 1 suffices to find the set of class extensions E of $[\tau]$ as given in Definition 4.1.1.

We prove each of these facts in turn.

B.1 The collection of parameters p_Y constructed in Algorithm 1 is well-defined

Lemma B.1.1 *The collection p_Y is independent of the choice of interval decomposition \mathcal{F} .*

Proof We first show that ℓ_0 is well-defined. Let $\mathcal{F}, \mathcal{G} : \mathbb{I}_{\text{BC}_k(Z^\psi \cap Y^\bullet)} \rightarrow P\tilde{H}_k(Z^\psi \cap Y^\bullet)$ be interval decompositions consisting of the following diagram. Recall that $\tilde{H}_k(Z^\psi \cap Y^N) = \tilde{H}_k(Z^\psi)$.

$$\begin{array}{ccccccc}
 I^1 & \longrightarrow & \dots & \longrightarrow & I^\ell & \longrightarrow & \dots & \longrightarrow & I^N \\
 \mathcal{F} \downarrow \downarrow \mathcal{G}^1 & & & & \mathcal{F}^\ell \downarrow \downarrow \mathcal{G}^\ell & & & & \mathcal{F}^N \downarrow \downarrow \mathcal{G}^N \\
 \tilde{H}_k(Z^\psi \cap Y^1) & \longrightarrow & \dots & \longrightarrow & \tilde{H}_k(Z^\psi \cap Y^\ell) & \longrightarrow & \dots & \longrightarrow & \tilde{H}_k(Z^\psi)
 \end{array} \tag{B2}$$

Let M be the dimension of $\tilde{H}_k(Z^\psi)$, and let ρ_1, \dots, ρ_M denote some ordering of the bars of $\text{BC}_k(Z^\psi \cap Y^\bullet)$ at parameter N . That is, $\widehat{\delta}(\rho_1) = \dots = \widehat{\delta}(\rho_M) = \infty$. Recall our assumption that all bars have unique birth-death parameters. In particular, all ρ_1, \dots, ρ_M have unique birth parameters. Let $\{\mathbf{e}_1^N, \dots, \mathbf{e}_M^N\}$ denote the corresponding basis vectors of I^N . Then, the collections $F = \{\mathcal{F}^N(\mathbf{e}_1^N), \dots, \mathcal{F}^N(\mathbf{e}_M^N)\}$ and $G = \{\mathcal{G}^N(\mathbf{e}_1^N), \dots, \mathcal{G}^N(\mathbf{e}_M^N)\}$ each form a basis of $\tilde{H}_k(Z^\psi)$.

Without loss of generality, let $S_{[\tau]}^{\mathcal{F},N} = \{f_1\rho_1, \dots, f_m\rho_m\}$ and $S_{[\tau]}^{\mathcal{G},N} = \{g_{i_1}\rho_{i_1}, \dots, g_{i_t}\rho_{i_t}\}$ be the \mathcal{F} -bar and \mathcal{G} -bar representations of $[\tau]$ at N . That is,

$$\begin{aligned}
 [\tau] &= f_1[\rho_1^{\mathcal{F},N}] + \dots + f_m[\rho_m^{\mathcal{F},N}] \\
 [\tau] &= g_{i_1}[\rho_{i_1}^{\mathcal{G},N}] + \dots + g_{i_t}[\rho_{i_t}^{\mathcal{G},N}],
 \end{aligned}$$

and all coefficients are nonzero. It follows that

$$\text{BARS}_\tau^{\mathcal{F}} = \{\rho_1, \dots, \rho_m\} \text{ and } \text{BARS}_\tau^{\mathcal{G}} = \{\rho_{i_1}, \dots, \rho_{i_t}\}.$$

Let $[\tau]_F$ and $[\tau]_G$ each denote the coordinate vectors relative to the basis F and G . Given a vector \mathbf{v} , we'll use $(\mathbf{v})_r$ to refer to its r^{th} coordinate. Then, $([\tau]_F)_r$ is f_r for $1 \leq r \leq m$ and 0 otherwise. Similarly, $([\tau]_G)_r$ is g_r for $r \in \{i_1, \dots, i_t\}$ and 0 otherwise.

Let $\mathcal{L} = (\mathcal{G})^{-1} \circ \mathcal{F}$. Let F^N, G^N, L^N each denote the matrix representation of $\mathcal{F}^N, \mathcal{G}^N$, and \mathcal{L}^N . Consider the matrix $L^N = (G^N)^{-1} \circ F^N$. Note that $[\tau]_G = L^N[\tau]_F$. Then, $[\tau]_G$ is a linear combination of the first m column vectors of L^N . That is, for any r , we have

$$([\tau]_G)_r = f_1(L^N)_{r,1} + \dots + f_m(L^N)_{r,m}. \tag{B3}$$

We will now show that $\max_{\rho \in \text{BARS}_\tau^{\mathcal{F}}} \{\widehat{\beta}(\rho)\} = \max_{\rho \in \text{BARS}_\tau^{\mathcal{G}}} \{\widehat{\beta}(\rho)\}$. Without loss of generality, assume that $\widehat{\beta}(\rho_1) = \max_{\rho \in \text{BARS}_\tau^{\mathcal{F}}} \{\widehat{\beta}(\rho)\}$. That is,

$$\widehat{\beta}(\rho_j) \leq \widehat{\beta}(\rho_1) \text{ for all } 1 \leq j \leq m. \tag{B4}$$

We first show that $\rho_1 \in \text{BARS}_\tau^{\mathcal{G}}$ by showing $([\tau]_G)_1 \neq 0$. From Eq. B3, we know that

$$([\tau]_G)_1 = f_1(L^N)_{1,1} + \dots + f_m(L^N)_{1,m}.$$

By Eq. B4 and the fact that all bars have unique birth parameters, we have $\widehat{\beta}(\rho_j) < \widehat{\beta}(\rho_1)$ for all $2 \leq j \leq m$. From Lemmas 2.2.5 and 2.2.6, we know that $(L^N)_{1,1} \neq 0$ and $(L^N)_{1,j} = 0$ for all $2 \leq j \leq m$. Thus, $([\tau]_G)_1 = f_1(L^N)_{1,1} \neq 0$. Recall that $([\tau]_G)_r$ is g_r for $r \in \{i_1, \dots, i_t\}$ and 0 otherwise. Thus, $([\tau]_G)_1 \neq 0$ implies that $\rho_1 \in \text{BARS}_\tau^{\mathcal{G}}$.

We now show that $\widehat{\beta}(\rho_i) \leq \widehat{\beta}(\rho_1)$ for every $\rho_i \in \text{BARS}_\tau^{\mathcal{G}}$. Let $\rho_i \in \text{BARS}_\tau^{\mathcal{G}}$. If $\rho_i \in \text{BARS}_\tau^{\mathcal{F}}$ as well, then $\widehat{\beta}(\rho_i) \leq \widehat{\beta}(\rho_1)$ by assumption. If $\rho_i \notin \text{BARS}_\tau^{\mathcal{F}}$, recall from Eq. B3 that

$$([\tau]_G)_i = f_1(L^N)_{i,1} + \dots + f_m(L^N)_{i,m}$$

Since $([\tau]_G)_i \neq 0$, there must be some $1 \leq j \leq m$ such that $(L^N)_{i,j} \neq 0$. By Lemma 2.2.5, this implies that $\widehat{\beta}(\rho_i) \leq \widehat{\beta}(\rho_j)$, and by Eq. B4, we have $\widehat{\beta}(\rho_i) \leq \widehat{\beta}(\rho_j) \leq \widehat{\beta}(\rho_1)$.

So far, we showed that $\widehat{\beta}(\rho_i) \leq \widehat{\beta}(\rho_1)$ for every $\rho_i \in \text{BARS}_\tau^{\mathcal{G}}$, and that $\rho_1 \in \text{BARS}_\tau^{\mathcal{G}}$. Then, $\widehat{\beta}(\rho_1) = \max_{\rho \in \text{BARS}_\tau^{\mathcal{G}}} \{\widehat{\beta}(\rho)\}$, which coincides with $\max_{\rho \in \text{BARS}_\tau^{\mathcal{F}}} \{\widehat{\beta}(\rho)\}$. Thus, ℓ_0 is independent of the choice of the interval decomposition $\mathcal{F} : \mathbb{I}_{\text{BC}_k}(Z^\psi \cap Y^\bullet) \rightarrow P\widetilde{H}_k(Z^\psi \cap Y^\bullet)$.

The remaining parameters of p_Y are the birth parameters of bars in $\text{BARS}_{>\ell_0}$, which is independent of the interval decomposition. Thus, p_Y is independent of the interval decomposition $\mathcal{F} : \mathbb{I}_{\text{BC}_k(Z^\psi \cap Y^\bullet)} \rightarrow P\tilde{H}_k(Z^\psi \cap Y^\bullet)$. \square

Note that by construction, ℓ_0 is the smallest parameter at which a restriction of $[\tau]$ can be found.

Lemma B.1.2 *Given parameter ℓ , let $\chi_\ell : \tilde{H}_k(Z^\psi \cap Y^\ell) \rightarrow \tilde{H}_k(Z^\psi)$ be the map induced by inclusion. Given $[\tau] \in \tilde{H}_k(Z^\psi)$, if there exists parameter ℓ and $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$ such that $[\tau] = \chi_\ell[w]$, then $\ell \geq \ell_0$*

Proof Assume that $S_{[\tau]}^{\mathcal{F}} = \{f_1\rho_1, \dots, f_m\rho_m\}$ is the \mathcal{F} -bar representation of $[\tau]$ at N . That is,

$$[\tau] = f_1[\rho_1^{\mathcal{F},N}] + \dots + f_m[\rho_m^{\mathcal{F},N}] = f_1[\mathcal{F}^N(\mathbf{e}_{\rho_1}^N)] + \dots + f_m[\mathcal{F}^N(\mathbf{e}_{\rho_m}^N)] \tag{B5}$$

in $\tilde{H}_k(Z^\psi)$ for nonzero f_1, \dots, f_m .

We prove via contradiction. Assume that there exists some $\ell < \ell_0$ and $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$ such that $[\tau] = \chi_\ell([w])$. Consider the following portion of the diagram for the interval decomposition of $P\tilde{H}_k(Z^\psi \cap Y^\bullet)$.

$$\begin{array}{ccc} I^\ell & \xrightarrow{\iota} & I^N \\ \downarrow \mathcal{F}^\ell & & \downarrow \mathcal{F}^N \\ \tilde{H}_k(Z^\psi \cap Y^\ell) & \xrightarrow{\chi_\ell} & \tilde{H}_k(Z^\psi) \end{array}$$

Since \mathcal{F}^ℓ is an isomorphism, there exists $w' \in I^\ell$ such that $[w] = [\mathcal{F}^\ell(w')]$. By commutativity, $\mathcal{F}^N \circ \iota(w') = [\tau]$. By Eq. B5 and the fact that \mathcal{F}^N is an isomorphism, $\iota(w') = f_1\mathbf{e}_{\rho_1}^N + \dots + f_m\mathbf{e}_{\rho_m}^N$. In particular, this implies that bars ρ_1, \dots, ρ_m are present at parameter ℓ . That is, $\widehat{\beta}(\rho_1), \dots, \widehat{\beta}(\rho_m) \leq \ell$. Recall that $\ell_0 = \max\{\widehat{\beta}(\rho_1), \dots, \widehat{\beta}(\rho_m)\}$. This contradicts $\ell < \ell_0$. \square

B.2 $\mathfrak{R}_\ell^{\mathcal{F}}$ is well-defined and is the desired set of restrictions

We now show that $\mathfrak{R}_\ell^{\mathcal{F}}$ from Algorithm 1 step (3-b-iii) is independent of the choice of interval decomposition $\mathcal{F} : \mathbb{I}_{\text{BC}_k(Z^\psi \cap Y^\bullet)} \rightarrow P\tilde{H}_k(Z^\psi \cap Y^\bullet)$ and that $\mathfrak{R}_\ell^{\mathcal{F}}$ coincides with the set of restrictions R_ℓ defined in Definition 4.1.1.

We first set the notations. Let $Z^\psi \cap Y^\bullet, \mathcal{F}, \text{BARS}_\tau^{\mathcal{F}}, \text{BARS}_{\text{short}}$, and p_Y be as in Algorithm 1. Let

$$\text{BARS}_{\text{long}}^{\mathcal{F}} = \{\rho \in \text{BC}_k(Z^\psi \cap Y^\bullet) \mid \widehat{\delta}(\rho) = \infty \text{ and } \rho \notin \text{BARS}_\tau^{\mathcal{F}}\}.$$

At any parameter ℓ , we partition the bars of $\text{BC}_k(Z^\psi \cap Y^\bullet)$ at parameter ℓ into following three sets

$$\text{BARS}_\tau^{\ell, \mathcal{F}} = \{\rho \in \text{BARS}_\tau^{\mathcal{F}} \mid \widehat{\beta}(\rho) \leq \ell\}$$

$$\begin{aligned} \text{BARS}_{\text{long}}^{\ell, \mathcal{F}} &= \{\rho \in \text{BARS}_{\text{long}}^{\mathcal{F}} \mid \widehat{\beta}(\rho) \leq \ell\} \\ \text{BARS}_{\text{short}}^{\ell} &= \{\rho \in \text{BARS}_{\text{short}} \mid \widehat{\beta}(\rho) \leq \ell < \widehat{\delta}(\rho) < \infty\}. \end{aligned} \tag{B6}$$

Finally, we state a lemma that is frequently used.

Lemma B.2.1 *Let $\chi_\ell : \widetilde{H}_k(Z^\psi \cap Y^\ell) \rightarrow \widetilde{H}_k(Z^\psi)$ be the map induced by inclusion. Then,*

$$\begin{aligned} \chi_\ell([\rho^{\mathcal{F}, \ell}]) &= [\rho^{\mathcal{F}, N}] \text{ for } \rho \in \text{BARS}_\tau^{\ell, \mathcal{F}} \cup \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}, \\ \chi_\ell([\rho^{\mathcal{F}, \ell}]) &= 0 \text{ for } \rho \in \text{BARS}_{\text{short}}^{\ell} \end{aligned}$$

In particular, if $[w] \in \mathfrak{R}_\ell^{\mathcal{F}}$, then $\chi_\ell([w]) = [\tau]$.

Proof Follows from the commutativity of the following portion of the diagram for the interval decomposition of $P\widetilde{H}_k(Z^\psi \cap Y^\bullet)$.

$$\begin{array}{ccc} I^\ell & \xrightarrow{\iota} & I^N \\ \downarrow \mathcal{F}^\ell & & \downarrow \mathcal{F}^N \\ \widetilde{H}_k(Z^\psi \cap Y^\ell) & \xrightarrow{\chi_\ell} & \widetilde{H}_k(Z^\psi) \end{array}$$

□

Theorem B.2.2 *Let $\mathcal{F}, \mathcal{G} : \mathbb{I}_{\text{BC}_k(Z^\psi \cap Y^\bullet)} \rightarrow P\widetilde{H}_k(Z^\psi \cap Y^\bullet)$ be two different interval decompositions. Given $\ell \in p_Y$, let $\mathfrak{R}_\ell^{\mathcal{F}}$ and $\mathfrak{R}_\ell^{\mathcal{G}}$ each denote the set of restrictions in step (3-b-iii) of Algorithm 1. Then, $\mathfrak{R}_\ell^{\mathcal{F}} = \mathfrak{R}_\ell^{\mathcal{G}}$.*

Proof Without loss of generality, let $S_{[\tau]}^{\mathcal{F}} = \{f_1^* \rho_1, \dots, f_m^* \rho_m\}$ be the \mathcal{F} -bar representation of $[\tau]$ at N . That is,

$$[\tau] = f_1^*[\rho_1^{\mathcal{F}, N}] + \dots + f_m^*[\rho_m^{\mathcal{F}, N}] \text{ in } \widetilde{H}_k(Z^\psi) \text{ for nonzero } f_1^*, \dots, f_m^*. \tag{B7}$$

Recall the partition of $\text{BC}_k(Z^\psi \cap Y^\bullet)$ at parameter ℓ into sets in Eq. B6. The collection $F = \{[\rho^{\mathcal{F}, \ell}] \mid \rho \in \text{BARS}_\tau^{\ell, \mathcal{F}} \cup \text{BARS}_{\text{long}}^{\ell, \mathcal{F}} \cup \text{BARS}_{\text{short}}^{\ell}\}$ form a basis of $\widetilde{H}_k(Z^\psi \cap Y^\ell)$. To prove Theorem B.2.2, we will take $[w] \in \mathfrak{R}_\ell^{\mathcal{G}}$, express it using the basis F , and show that $[w] \in \mathfrak{R}_\ell^{\mathcal{F}}$.

Let $[w] \in \mathfrak{R}_\ell^{\mathcal{G}}$. When we express $[w]$ using basis F , it will have the form

$$[w] = \sum_{\rho_j \in \text{BARS}_\tau^{\ell, \mathcal{F}}} f_j^\tau [\rho_j^{\mathcal{F}, \ell}] + \sum_{\rho_j \in \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}} f_j^{\text{long}} [\rho_j^{\mathcal{F}, \ell}] + \sum_{\rho_j \in \text{BARS}_{\text{short}}^{\ell}} f_j^{\text{short}} [\rho_j^{\mathcal{F}, \ell}] \tag{B8}$$

for some coefficients in \mathbb{F} . From Lemma B.2.1, we know that $\chi_\ell([w]) = [\tau]$. From Eq. B8 and Lemma B.2.1,

$$\chi_\ell([w]) = \sum_{\rho_j \in \text{BARS}_\tau^{\ell, \mathcal{F}}} f_j^\tau [\rho_j^{\mathcal{F}, N}] + \sum_{\rho_j \in \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}} f_j^{\text{long}} [\rho_j^{\mathcal{F}, N}],$$

and this must equal $[\tau]$. From Eq. B7, we know that $f_j^\tau = f_j^*$ for all $\rho_j \in \text{BARS}_\tau^{\ell, \mathcal{F}}$, and $f_j^{\text{long}} = 0$ for all $\rho_j \in \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}$. We can re-write Eq. B8 as

$$[w] = \sum_{i=1}^m f_i^* [\rho_i^{\mathcal{F}, \ell}] + \sum_{\rho_j \in \text{BARS}_{\text{short}}^{\ell, \mathcal{F}}} f_j^{\text{short}} [\rho_j^{\mathcal{F}, \ell}],$$

for some coefficients f_j^{short} in \mathbb{F} . Thus, $[w] \in \mathfrak{R}_\ell^{\mathcal{F}}$.

Conversely, if $[w] \in \mathfrak{R}_\ell^{\mathcal{F}}$, the same argument shows that $[w] \in \mathfrak{R}_\ell^{\mathcal{G}}$ as well. Thus, $\mathfrak{R}_\ell^{\mathcal{G}} = \mathfrak{R}_\ell^{\mathcal{F}}$. □

Since the set of restrictions $\mathfrak{R}_\ell^{\mathcal{F}}$ is independent of the interval decomposition \mathcal{F} , we can omit the \mathcal{F} .

Theorem B.2.3 *Given parameter $\ell \in p_Y$, the set \mathfrak{R}_ℓ found in Algorithm 1 coincides with R_ℓ from Definition 4.1.1.*

Proof If $[w] \in \mathfrak{R}_\ell$, then by Lemma B.2.1, we know that $\chi_\ell([w]) = [\tau]$. So $[w] \in R_\ell$.

If $[w] \in R_\ell$, then $[w] \in \tilde{H}_k(Z^\psi \cap Y^\ell)$. Recall the partition of bars in Eq. B6. We can express $[w]$ using the basis $\{[\rho^{\mathcal{F}, \ell}] \mid \rho \in \text{BARS}_\tau^{\ell, \mathcal{F}} \cup \text{BARS}_{\text{long}}^{\ell, \mathcal{F}} \cup \text{BARS}_{\text{short}}^{\ell, \mathcal{F}}\}$ of $\tilde{H}_k(Z^\psi \cap Y^\ell)$ as

$$[w] = \sum_{\rho_j \in \text{BARS}_\tau^{\ell, \mathcal{F}}} f_j^\tau [\rho_j^{\mathcal{F}, \ell}] + \sum_{\rho_j \in \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}} f_j^{\text{long}} [\rho_j^{\mathcal{F}, \ell}] + \sum_{\rho_j \in \text{BARS}_{\text{short}}^{\ell, \mathcal{F}}} f_j^{\text{short}} [\rho_j^{\mathcal{F}, \ell}]$$

for some coefficients in \mathbb{F} .

Recall $S_{[\tau]}^{\mathcal{F}} = \{f_1^* \rho_1, \dots, f_m^* \rho_m\}$, the \mathcal{F} -bar representation of $[\tau]$ at N from step (2-a) of Algorithm 1. Since $[w] \in R_\ell$, we know $\chi_\ell([w]) = [\tau]$. By Lemma B.2.1, the coefficients f_j^τ must agree with f_j^* for all $\rho_j \in \text{BARS}_\tau^{\ell, \mathcal{F}}$, and $f_j^{\text{long}} = 0$ for all $\rho_j \in \text{BARS}_{\text{long}}^{\ell, \mathcal{F}}$. Thus, $[w] \in \mathfrak{R}_\ell$. □

B.3 The output of Algorithm 1 suffices to find all class extensions of $[\tau]$

We now prove the main theorem.

Theorem 4.1.4 *Let Z^\bullet, Y^\bullet be filtered simplicial complexes on a common vertex set P . Fix a parameter ψ and homology class $[\tau] \in \tilde{H}_k(Z^\psi)$. The output of Algorithm 1 suffices to recover all class extensions of $[\tau]$. That is, let E be the set of all class*

extensions of $[\tau]$ per Definition 4.1.1. Given parameters $\ell < \ell'$, let $\lambda_{\ell'}^{\ell} : \tilde{H}_k(Y^{\ell}) \rightarrow \tilde{H}_k(Y^{\ell'})$ be the map induced by inclusion. Let

$$\mathfrak{E}^* = \{[y] \in \tilde{H}_k(Y^{\ell'}) \mid \text{there exists } \ell \in p_Y, [y_{\ell}] \in \mathfrak{E}_{\ell} \text{ such that } [y] = \lambda_{\ell'}^{\ell}([y_{\ell}])\},$$

where p_Y and $\mathfrak{E} = \cup_{\ell \in p_Y} \mathfrak{E}_{\ell}$ are the outputs of Algorithm 1. Then $E = \mathfrak{E}^*$.

Proof We first show that if $[y] \in \mathfrak{E}^*$, then $[y] \in E$. Assume $[y] \in \mathfrak{E}^*$. So $[y] \in \tilde{H}_k(Y^{\ell'})$ for some parameter ℓ' , and there exists some $\ell \in p_Y$ and $[y_{\ell}] \in \mathfrak{E}_{\ell}$ such that $[y] = \lambda_{\ell'}^{\ell}([y_{\ell}])$. That is, $[y_{\ell}]$ is a class extension at ℓ via some $[w] \in \mathfrak{R}_{\ell} = R_{\ell}$. By Lemma 4.1.2, we know that $[y] = \lambda_{\ell'}^{\ell}([y_{\ell}])$ is a class extension via $\eta_{\ell'}^{\ell}([w])$ at ℓ' , where $\eta_{\ell'}^{\ell} : \tilde{H}_k(Z^{\psi} \cap Y^{\ell}) \rightarrow \tilde{H}_k(Z^{\psi} \cap Y^{\ell'})$ is the map induced by inclusion. Note that $\eta_{\ell'}^{\ell}([w]) \in R_{\ell'}$. Thus, $[y] \in E_{\ell'}$, and hence, $[y] \in E$.

Now, assume that $[y] \in E$. There exists some ℓ and $[w] \in R_{\ell}$ such that $[y] = \Upsilon_{\ell}([w])$, where $\Upsilon_{\ell} : \tilde{H}_k(Z^{\psi} \cap Y^{\ell}) \rightarrow \tilde{H}_k(Y^{\ell})$ is the map induced by inclusion. Assume that the birth parameter of $[w]$ in $H_k(Z^{\psi} \cap Y^{\bullet})$ is $\bar{\ell}$. Let $\eta_{\bar{\ell}}^{\ell} : \tilde{H}_k(Z^{\psi} \cap Y^{\bar{\ell}}) \rightarrow \tilde{H}_k(Z^{\psi} \cap Y^{\ell})$ be the map induced by inclusion. Then, $[w] = \eta_{\bar{\ell}}^{\ell}([\bar{w}])$ for some $[\bar{w}] \in \tilde{H}_k(Z^{\psi} \cap Y^{\bar{\ell}})$. By Lemma 4.1.3, there exists a class extension $[\bar{y}] \in \tilde{H}_k(Y^{\bar{\ell}})$ at $\bar{\ell}$ via $[\bar{w}]$, and $\lambda_{\bar{\ell}}^{\ell}([\bar{y}]) = [y]$, where $\lambda_{\bar{\ell}}^{\ell} : \tilde{H}_k(Y^{\bar{\ell}}) \rightarrow \tilde{H}_k(Y^{\ell})$ is the map induced by inclusion. One can check that $[\bar{w}]$ is a restriction of $[\tau]$. We thus know from Lemma B.1.2 that $\ell_0 \leq \bar{\ell}$. Furthermore, since a new homology class is born at $\bar{\ell}$, the parameter should coincide with the birth of a new bar of $\text{BC}_k(Z^{\psi} \cap Y^{\bullet})$. Thus, $\bar{\ell} \in p_Y$, $[\bar{w}] \in \mathfrak{R}_{\bar{\ell}}$, and $[\bar{y}] \in \mathfrak{E}_{\bar{\ell}}$. That is, we must have found $[\bar{y}]$ as a class extension in Algorithm 1. Thus, $[y] = \lambda_{\bar{\ell}}^{\ell}([\bar{y}]) \in \mathfrak{E}^*$. □

Appendix C Proof of Lemma 4.2.3

Lemma 4.2.3 Given persistence module $P\tilde{H}_k(Z^{\bullet})$, let $\text{BC}_k(Z^{\psi}) = \{\tau_1, \dots, \tau_m\}$ be some ordering of bars alive at ψ . Let L_Z be the collection of $m \times m$ matrices

$$L_Z = \{L \in \mathbb{F}^{m \times m} \mid L_{r,r} \neq 0, L_{r,c} = 0 \text{ if } \beta(\tau_r) > \beta(\tau_c) \text{ or } \delta(\tau_r) > \delta(\tau_c) \text{ for all } r, c\}.$$

The collection L_Z is precisely the collection of matrices representing the map $\mathcal{L}^{\psi} : (\mathbb{I}_{\text{BC}_k(Z^{\bullet})})^{\psi} \rightarrow (\mathbb{I}_{\text{BC}_k(Z^{\bullet})})^{\psi}$ under all possible choices of barcode module automorphism $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^{\bullet})} \rightarrow \mathbb{I}_{\text{BC}_k(Z^{\bullet})}$.

Proof Let M_Z be the collection of matrix representations of \mathcal{L}^{ψ} for every possible persistence module automorphism \mathcal{L} .

We first show that $L_Z \subseteq M_Z$. Let $L \in L_Z$. From Lemma 2.2.7, we know that L is invertible. Further, any such isomorphism L immediately extends to an automorphism of persistence modules $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^{\bullet})} \rightarrow \mathbb{I}_{\text{BC}_k(Z^{\bullet})}$ because $\mathbb{I}_{\text{BC}_k(Z^{\bullet})}$ is a direct sum of interval modules. Thus, $L_Z \subseteq M_Z$.

We now show that $M_Z \subseteq L_Z$. Let M be a matrix in M_Z . That is, M is the matrix representation of the linear isomorphism \mathcal{L}^{ψ} for some automorphism of persistence

modules $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$. It follows from Lemmas 2.2.5 and 2.2.6 that $M \in L_Z$. Thus, $M_Z \subseteq L_Z$. □

Appendix D Proof of Lemma 4.2.4

Lemma 4.2.4 *Assume that $\mathbb{F} = \mathbb{F}_2$ and that all bars of $\text{BC}_k(Z^\bullet)$ have unique death parameters. Let $\mathcal{B}, \mathcal{C} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow P\tilde{H}_k(Z^\bullet)$ be two different interval decompositions. Given a bar $\tau \in \text{BC}_k(Z^\bullet)$, let $[\tau_*^{\mathcal{B}}]$ and $[\tau_*^{\mathcal{C}}]$ denote the terminal class for τ under \mathcal{B} and \mathcal{C} respectively. Then, $[\tau_*^{\mathcal{B}}] = [\tau_*^{\mathcal{C}}]$ in $H_k(Z^{\delta(\tau)-1})$.*

Proof For notational convenience, let $\psi = \delta(\tau) - 1$, and let $\mathcal{L} : \mathbb{I}_{\text{BC}_k(Z^\bullet)} \rightarrow \mathbb{I}_{\text{BC}_k(Z^\bullet)}$ be the isomorphism $\mathcal{L} = \mathcal{C}^{-1} \circ \mathcal{B}$. Let $I^\psi = (\mathbb{I}_{\text{BC}_k(Z^\bullet)})^\psi$. Let B^ψ, C^ψ, L^ψ each be the matrix representation of linear isomorphisms $\mathcal{B}^\psi, \mathcal{C}^\psi : I^\psi \rightarrow \tilde{H}_k(Z^\psi)$ and $\mathcal{L}^\psi : I^\psi \rightarrow I^\psi$.

Let τ_1, \dots, τ_m denote the bars of $\text{BC}_k(Z^\bullet)$ alive at parameter ψ . Let B and C each denote the collection of homology classes that \mathcal{B} (or \mathcal{C})-correspond to the bars at parameter ψ :

$$B = \{[\tau_1^{\mathcal{B},\psi}], \dots, [\tau_m^{\mathcal{B},\psi}]\}, \text{ and } C = \{[\tau_1^{\mathcal{C},\psi}], \dots, [\tau_m^{\mathcal{C},\psi}]\}.$$

Note that both B and C are valid basis for $\tilde{H}_k(Z^\psi)$.

Without loss of generality, assume that $\tau_1 = \tau$. Then,

$$[\tau_*^{\mathcal{B}}] = [\tau_1^{\mathcal{B},\psi}], \quad [\tau_*^{\mathcal{C}}] = [\tau_1^{\mathcal{C},\psi}]. \tag{D9}$$

Now consider the change of basis matrix $L^\psi = (C^\psi)^{-1} \circ B^\psi$. Using the coordinates vectors with respect to the basis B and C , we know that

$$[\tau_*^{\mathcal{B}}]_C = L^\psi [\tau_*^{\mathcal{B}}]_B.$$

We will show that $[\tau_*^{\mathcal{B}}]_C = [\tau_*^{\mathcal{C}}]_C$. From Eq. D9, we know that $[\tau_*^{\mathcal{B}}]_B = [1, 0, \dots, 0]^T$, so $[\tau_*^{\mathcal{B}}]_C$ coincides with the 1st column vector of L^ψ . By Lemma 2.2.6 and the fact that $\mathbb{F} = \mathbb{F}_2$, we know that $(L^\psi)_{(1,1)} = 1$. Recall that bars τ_1, \dots, τ_m are alive at parameter $\psi = \delta(\tau) - 1$, so $\delta(\tau) \leq \delta(\tau_j)$ for all $j > 1$. Since the death parameters are unique, we have $\delta(\tau) < \delta(\tau_j)$ for all $j > 1$. By Lemma 2.2.5, we know that $(L^\psi)_{(j,1)} = 0$ for all $j \neq 1$. Thus, $[\tau_*^{\mathcal{B}}]_C = [1, 0, \dots, 0]^T$, which coincides with $[\tau_*^{\mathcal{C}}]_C$ by Eq. D9. Thus, $[\tau_*^{\mathcal{B}}]_C = [\tau_*^{\mathcal{C}}]_C$, and $[\tau_*^{\mathcal{B}}] = [\tau_*^{\mathcal{C}}]$ in $\tilde{H}_k(Z^\psi)$. □

References

Bauer, U., Lesnick, M.: Induced matchings of barcodes and the algebraic stability of persistence. In: Proceedings of the Thirtieth Annual Symposium on Computational Geometry. SOCG'14, pp. 355–364. Association for Computing Machinery, New York, NY, USA (2014). <https://doi.org/10.1145/2582112.2582168>

- Bauer, U., Lesnick, M.: Persistence diagrams as diagrams: A categorification of the stability theorem. In: *Topological Data Analysis*, 67–96 (2020). https://doi.org/10.1007/978-3-030-43408-3_3
- Bauer, U., Schmahl, M.: Lifespan functors and natural dualities in persistent homology. *arXiv:2012.12881* (2021) [math.AT]
- Bubenik, P., Vergili, T.: Topological spaces of persistence modules and their properties. *J. Appl. Comput. Topol.* **2**, 233–269 (2018)
- Carlsson, G.: Topology and data. *Bull. Am. Math. Soc.* **46**, 255–308 (2009). <https://doi.org/10.1090/S0273-0979-09-01249-X>
- Carlsson, G., de Silva, V.: Zigzag persistence. *Found. Comput. Math.* **10**, 367–405 (2010). <https://doi.org/10.1007/s10208-010-9066-0>
- Chazal, F., Cohen-Steiner, D., Guibas, L., Mémoli, F., Oudot, S.: Gromov-hausdorff stable signatures for shapes using persistence. *Comput. Graph. Forum* **28**, 1393–1403 (2009). <https://doi.org/10.1111/j.1467-8659.2009.01516.x>
- Chowdhury, S., Mémoli, F.: A functorial Dowker theorem and persistent homology of asymmetric networks. *J. Appl. Comput. Topol.* (2018). <https://doi.org/10.1007/s41468-018-0020-6>
- Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. *Discrete Comput. Geom.* **37**, 263–271 (2005). <https://doi.org/10.1007/s00454-006-1276-5>
- de Silva, V., Carlsson, G.: Topological estimation using witness complexes. In: *SPBG'04 Symposium on Point - Based Graphics 2004* (2004). <https://doi.org/10.2312/SPBG/SPBG04/157-166>
- Dowker, C.H.: Homology groups of relations. *Ann. Math.* **56**(1), 84–95 (1952). <https://doi.org/10.2307/1969768>
- Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pp. 454–463 (2000). <https://doi.org/10.1109/SFCS.2000.892133>
- Ghrist, R., Henselman-Petrusek, G.: Saecular persistence. *arXiv:2112.04927* (2021)
- Ghrist, R.: Barcodes: The persistent topology of data. *Bull. Am. Math. Soc.* (2008). <https://doi.org/10.1090/S0273-0979-07-01191-3>
- Gonzalez-Díaz, R., Soriano-Trigueros, M.: Basis-independent partial matchings induced by morphisms between persistence modules. *arXiv:2006.11100v1* (2020)
- Henselman, G., Ghrist, R.: Matroid Filtrations and Computational Persistent Homology. (2016) *arXiv:1606.00199* [math.AT]
- Jacquard, E., Nanda, V., Tillmann, U.: The space of barcode bases for persistence modules. *J. Appl. Comput. Topol.* (2022). <https://doi.org/10.1007/s41468-022-00094-6>
- Milnor, J.W.: Construction of universal bundles, ii. *Ann. Math.* **63**, 272 (1956)
- Perea, J.A.: Sparse circular coordinates via principal \mathbb{Z} -bundles. In: *Topological Data Analysis*, pp. 435–458 (2020)
- Reani, Y., Bobrowski, O.: Cycle registration in persistent homology with applications in topological bootstrap. *arXiv:2101.00698v1* (2021)
- Takeuchi, H.: The persistent homology of a sampled map: from a viewpoint of quiver representations. *J. Appl. Comput. Topol.* (2021). <https://doi.org/10.1007/s41468-021-00065-3>
- Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discret. Comput. Geom.* **33**, 249–274 (2005). <https://doi.org/10.1007/s00454-004-1146-y>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.