



Decorated merge trees for persistent topology

Justin Curry¹ · Haibin Hang² · Washington Mio³ · Tom Needham³ · Osman Berat Okutan³

Received: 28 July 2021 / Revised: 17 January 2022 / Accepted: 9 February 2022
© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2022

Abstract

This paper introduces decorated merge trees (DMTs) as a novel invariant for persistent spaces. DMTs combine both π_0 and H_n information into a single data structure that distinguishes filtrations that merge trees and persistent homology cannot distinguish alone. Three variants on DMTs, which emphasize category theory, representation theory and persistence barcodes, respectively, offer different advantages in terms of theory and computation. Two notions of distance—an interleaving distance and bottleneck distance—for DMTs are defined and a hierarchy of stability results that both refine and generalize existing stability results is proved here. To overcome some of the computational complexity inherent in these distances, we provide a novel use of Gromov-Wasserstein couplings to compute optimal merge tree alignments for a combinatorial version of our interleaving distance which can be tractably estimated. We introduce computational frameworks for generating, visualizing and comparing decorated merge trees derived from synthetic and real data. Example applications include comparison of point clouds, interpretation of persistent homology of sliding window embeddings of time series, visualization of topological features in segmented brain tumor images and topology-driven graph alignment.

Keywords Topological data analysis · Persistent homology · Merge trees · Interleaving distance

Mathematics Subject Classification 55N31

✉ Tom Needham
tneedham@fsu.edu

¹ SUNY Albany, Albany, USA

² University of Delaware, Newark, USA

³ Florida State University, Tallahassee, USA

Contents

1	Introduction	1
2	Decorated merge trees three different ways	2
2.1	Preliminaries	2
2.2	The categorical decorated merge tree	2
2.3	The concrete decorated merge tree	2
2.4	The barcode decorated merge tree	2
3	Continuity and stability of decorated merge trees	3
3.1	Metrics on persistence modules	3
3.2	Continuity of the barcode transform	3
3.3	Interleavings of merge trees	3
3.4	A decorated bottleneck distance	3
3.5	Hierarchy of distances	3
4	Representations of tree posets and lift decorations	4
4.1	Indecomposables with totally ordered support	4
4.2	Injectivity of the barcode transform	4
4.3	Lift decorated merge trees	4
5	Computing interleaving distances	5
5.1	Labeled distance	5
5.2	Approximating merge tree interleaving via Gromov-Wasserstein distance	5
5.3	Matchings of DMTs via fused Gromov-Wasserstein matchings	5
6	Algorithmic details and examples	6
6.1	Merge tree interleaving distances	6
6.2	Decorated merge trees for point clouds	6
6.3	Sliding window embeddings	6
6.4	Clustering point clouds	6
6.5	Networks	6
6.6	Network matching	6
7	Discussion	7
A	Interval topology	A
B	Comparison and existence of merge trees	B
C	Technical proofs	C
C.1	Proof of Proposition 4	C.1
C.2	Proof of Proposition 8	C.2
	References	

1 Introduction

In this paper we introduce a new set of tools for Topological Data Analysis (TDA) called Decorated Merge Trees (DMTs). Not only do these new tools have a rich underlying theory that spans category theory and metric geometry, they also provide topological signatures for datasets such as point clouds, time series, grayscale images and networks which are more informative and interpretable than standard persistent homology barcodes. Figure 1 illustrates the main construction of the paper with a simple example. In this figure, two point clouds with different coarse topological structure are depicted. Their traditional TDA signatures—degree-0 and degree-1 Vietoris-Rips persistence diagrams—do not distinguish these point clouds. Our DMT construction illustrates the multiscale topology of each point cloud by overlaying a merge tree (a popular data structure in scientific visualization, used to summarize multiscale connectivity—see surveys Heine et al. 2016; Yan et al. 2021) with a degree-1 persistent homology barcode. This depicts not only the multiscale homological (H_1) data of each




Fig. 1 Decorated Merge Trees. The left column shows two point clouds. Their degree-0 and degree-1 Vietoris-Rips persistence diagrams in the middle column are essentially the same, despite clear topological differences in the point clouds. The decorated merge trees (DMTs) in the third column clearly distinguish the datasets topologically by fusing degree-0 and degree-1 information to track the topological *location* of the degree-1 features. Each red bar corresponds to a degree-1 persistent feature and its placement in the merge tree indicates the connected components over which the feature persists

point cloud, but also the (topological) location of each degree-1 feature in the dataset. This paper formalizes the DMT construction from several perspectives and extends classical lines of inquiry in the TDA literature—metric stability, decomposability and practical computational aspects—to this novel setting.

Although the construction of the decorated merge tree presented above is intuitive, it turns out that there are multiple ways of tracing births and deaths of homological features along an evolving set of connected components. To this end we provide in Sect. 2 three different definitions of a decorated merge tree:

1. The **categorical decorated merge tree** relies on the definition of a category of parameterized vector spaces $\mathbf{pVect}$. This definition fits squarely within the framework of generalized persistence modules (Bubenik et al. 2015; Bubenik and Scott 2014) as it is defined in terms of a functor from $(\mathbb{R}, \leq)$ to $\mathbf{pVect}$. This definition allows us to leverage existing results to define an interleaving distance and establish its stability.
2. The **concrete decorated merge tree** takes the perspective that the underlying merge tree, along with its poset structure $(\mathcal{M}_F, \preceq)$, should define the domain of a functor to the category of vector spaces $\mathbf{Vect}$, where the homology of each component at each time is recorded. This definition is equivalent to the categorical one, but is more intuitive and suggests computational approaches.

3. The **barcode decorated merge tree** takes a perspective similar to other TDA techniques (Landi 2018; Lesnick and Wright 2015; Turner et al. 2014) that reduce the study of complicated persistent spaces to ensembles of 1-dimensional persistence modules. The barcode DMT associates to each point in the merge tree the barcode gotten by restricting the filtration to the line that starts at that point and stretches to infinity.

In Sect. 3, we introduce an interleaving distance between decorated merge trees and establish several bounds that set decorated merge trees apart from existing TDA methods. In particular, let $f, g : X \rightarrow \mathbb{R}$ be functions whose sublevel sets are locally connected. For every homological degree n , one obtains categorical decorated merge trees $\tilde{F}_n$ and $\tilde{G}_n$, as well as classical merge trees and persistence modules, associated to the sublevel set filtrations of f and g . Theorems 2 and 3 allow one to extract the following statement:

Theorem The interleaving distance between $\tilde{F}_n$ and $\tilde{G}_n$ is stable with respect to L^∞ -distance between f and g , and is more sensitive than (i.e., lower bounded by) both the interleaving distance between the associated merge trees and the interleaving distance between the associated persistence modules.

After these stability results are proven, we focus on barcode decorated merge trees. We show in Theorem 1 that a barcode decorated merge tree can be understood as a Lipschitz map from the merge tree to the space of barcodes. Barcode decorated merge trees are amenable to a theory of matchings (Definition 19) and thus a new decorated bottleneck distance (Definition 20). This offers a tractable and approximable metric for comparing these enriched invariants. Theorem 2 shows that this matching distance is stable with respect to the interleaving distance between categorical decorated merge trees.

In Sect. 4, representation-theoretic aspects of DMTs take center stage. In general, one cannot hope for simple indecomposables such as the barcode decompositions that appear in standard persistent homology—see Example 5 for an illustration. We say that a DMT is *real interval decomposable* if it decomposes as a direct sum of DMTs with totally ordered support. Theorem 4 provides a condition which is equivalent to real interval decomposability. The class of real interval decomposable DMTs is of particular interest. As Theorem 5 shows, on this class the map taking a DMT to a barcode DMT is injective, thus providing a positive solution to a *topological inverse problem* (Oudot and Solomon 2020). We also give methods for generating real interval decomposable DMTs directly from data, with theoretical guarantees of their correctness (Proposition 8).

Section 5 shifts the focus to computational aspects of interleaving distance between decorated merge trees. Building on work of Gasparovic et al. (2019), we reformulate computation of the matching distance between barcode decorated merge trees as the search for an alignment between nodes of the trees which is optimal with respect to a certain cost function (Proposition 9). This reformulation allows us to introduce a method for estimating the metric via a continuous relaxation which can be solved within the Gromov-Wasserstein framework from optimal transport theory (Mémoli 2011). Our algorithm is novel even when estimating the interleaving distance between (undecorated) merge trees, but has close connections to other recent advances in the

literature (Li et al. 2021; Mémoli et al. 2019, 2021). The computational focus is continued in Sect. 6, where algorithms for computing and visualizing decorated merge trees from synthetic and real datasets are described. The Python code used to produce the figures and experiments for the paper are publicly available under an open source license at <https://github.com/trneedham/Decorated-Merge-Trees>.

The main paper concludes with a discussion of future directions of research in Sect. 7. In particular, we note that the DMT concept has natural generalizations such as Reeb or MAPPER graphs decorated with zig-zag modules (Carlsson and De Silva 2010), correspondence modules (Hang and Mio 2020) or Leray (co)sheaves (Curry 2015). These constructions will be the subject of future work as they require more substantive theoretical and algorithmic developments. By contrast, DMTs fit naturally into a pre-existing body of literature and a fleshed out code base. Regardless, decorated merge trees and decorated Reeb graphs are just a small part of a broader research program to construct enriched TDA invariants that are more informative than classical barcodes.

The paper includes three appendices. We draw readers' attention to Appendices A and B, which contain results on merge tree topology which are theoretically fundamental but fall outside of the narrative of the main body of the paper. Appendix C contains proofs of some technical results.

2 Decorated merge trees three different ways

In this paper we investigate topological signatures which go beyond standard persistent homology of filtered topological spaces. To this end, we formally define three notions of decorated merge trees—*categorical*, *concrete* and *barcode-decorated*—which were described informally in the introduction. Before doing so, we review some preliminary definitions.

2.1 Preliminaries

We now introduce the general class of objects to which our decorated merge tree constructions will be applied.

Definition 1 A **persistent space** is a functor $F: (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$, where $(\mathbb{R}, \leq)$ is considered as a poset category. Concretely, a persistent space associates to each $s \in \mathbb{R}$ a topological space $F(s)$ and to each ordered pair $s \leq t$ a continuous map $F(s \leq t): F(s) \rightarrow F(t)$. These maps collectively satisfy the usual composition rules of a functor. When each map $F(s \leq t)$ is an injection, we will call such a persistent space a **filtration**.

We now describe several situations where persistent spaces arise, focusing on the case of filtrations.

Example 1 Given a continuous function $f : X \rightarrow \mathbb{R}$, one can consider the **sublevel-set filtration**, which is a persistent space where

$$F(s) := f^{-1}(-\infty, s] = \{x \in X \mid f(x) \leq s\}.$$

By equipping each sublevel-set with the subspace topology, the inclusion maps $F(s) \subseteq F(t)$ are continuous and define the maps $F(s \leq t)$ for the persistent space.

The following gives a flexible class of instances of sublevel-set filtrations.

Example 2 Given a subset Z of a metric space X , we define the **offset function** $f_Z : X \rightarrow \mathbb{R}$ to be

$$f_Z(x) := \inf_{z \in Z} d(x, z).$$

The **offset filtration** F_Z , defined by

$$F_Z(s) := \{x \in X \mid d(x, Z) := \inf_{z \in Z} d(x, z) \leq s\},$$

is the sublevel-set filtration of f_Z .

Sublevel-set filtrations can be extended to certain non-continuous functions on simplicial complexes which arise frequently in topological data analysis.

Example 3 Let X be a simplicial complex and $f : X \rightarrow \mathbb{R}$ a function which is constant on each simplex. If f is **monotone**, that is whenever σ is a face of τ then $f(\sigma) \leq f(\tau)$, then the sublevel-set filtration $F(s) = |f^{-1}(-\infty, s]|$ defines a persistent space, where $|\cdot|$ denotes geometric realization.

We now introduce some basic concepts of Topological Data Analysis (TDA). These are invariants built to study persistent spaces. We assume that the reader is familiar with the fundamentals of TDA and mainly use the definitions here to set terminology and notation—see the survey (Carlsson 2014) for more background. For the remainder of this paper we let π_0 denote the connected components functor $\pi_0 : \mathbf{Top} \rightarrow \mathbf{Set}$ and let H_n denote some choice of homology theory with coefficients in a field $\mathbb{k}$.

Definition 2 A functor $S : (\mathbb{R}, \leq) \rightarrow \mathbf{Set}$ is called a **persistent set** (Carlsson and Mémoli 2013; Curry 2018). If $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ is a persistent space, then the **persistent set of components** is the composition of $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ with the connected components functor $\pi_0 : \mathbf{Top} \rightarrow \mathbf{Set}$, i.e.

$$\pi_0 \circ F : (\mathbb{R}, \leq) \rightarrow \mathbf{Set} \quad \text{where} \quad s \rightsquigarrow \pi_0(F(s)).$$

Associated to any persistent set $S : (\mathbb{R}, \leq) \rightarrow \mathbf{Set}$ is its **display poset**, which is simply the disjoint union of all the sets that appear in S , i.e.

$$\mathcal{S} := \bigsqcup_{t \in \mathbb{R}} S(t) := \bigcup S(t) \times \{t\}$$




Fig. 2 A persistent space is shown on the left, given by the offset filtration F_Z (Example 2), where $Z \subset \mathbb{R}$ consists of two points at distance $2L$. The associated persistent set is visualized through its generalized merge tree (a schematic Hasse diagram of the poset is shown here). The classical merge tree (visualized as the Reeb graph of the epigraph of the filtration function) is shown on the right

The poset structure on $\mathcal{S}$ is defined by declaring

$$(x, s) \preceq (y, t) \quad \text{if and only if} \quad S(s \leq t)(x) = y.$$

The **generalized merge tree** of F is the display poset $(\mathcal{M}_F, \preceq)$ associated to the persistent set $\pi_0 \circ F$.

We use the term “generalized” to distinguish it from the classical merge tree.

Definition 3 Let X be a topological space and $f : X \rightarrow \mathbb{R}$ a continuous function. The **epigraph** of f is the set $E_f := \{(x, r) \mid f(x) \leq r\}$. The **classical merge tree** $\mathcal{M}_f$ is the Reeb graph of the **projection function** $\pi_f : E_f \rightarrow \mathbb{R} : (x, r) \mapsto r$. That is, the merge tree is the quotient space $\mathcal{M}_f := E_f / \sim$, where $\sim$ is the equivalence relation $p = (x, r) \sim (x', r') = p'$ if and only if $r = r'$ and p and p' lie in the same connected component of the level set $\pi_f^{-1}(r)$. Let $\tilde{\pi}_f : \mathcal{M}_f \rightarrow \mathbb{R}$ denote the projection map induced by π_f .

A simple example showing a persistent space together with its persistent set and merge tree is shown in Fig. 2.

Remark 1 The differences between the classical and generalized merge tree have gone largely unmentioned on in the literature and are often treated as interchangeable. Observe that $\mathcal{M}_F$ is defined for any persistent space F and comes endowed with a poset structure. On the other hand, $\mathcal{M}_f$ is defined for a continuous function f on a space X and comes equipped with a quotient topology. This difference only creates problems when one notices that there are two notions of interleaving, which might not agree always. In Appendix A we prove that for functions on compact spaces with finitely many critical points, this difference can be ignored.

Beyond using π_0 to discriminate persistent spaces, we can use homology.

Definition 4 A **persistence module** is a functor from the poset category $(\mathbb{R}, \leq)$ into $\text{Vect}_{\mathbb{k}}$, the category of vector spaces over the field $\mathbb{k}$. A persistence module is **pointwise finite dimensional** if its image lies in $\text{vect}_{\mathbb{k}}$, the category of finite dimensional vector spaces over $\mathbb{k}$.

Let $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ be a persistent space. For any non-negative integer $n \geq 0$ the n^{th} **persistent homology module** $F_n : (\mathbb{R}, \leq) \rightarrow \mathbf{Vect}_{\mathbb{k}}$ is the persistence module

$$F_n := H_n \circ F : (\mathbb{R}, \leq) \rightarrow \mathbf{Vect}_{\mathbb{k}} \quad \text{with} \quad s \rightsquigarrow H_n(F(s) \mathbb{J} \mathbb{k}).$$

We will frequently drop the subscript $\mathbb{k}$, with the understanding that a field of coefficients has been fixed.

One of the central results in the theory of TDA is the theorem of Crawley-Boevey (2015, Theorem 1.1) which states that pointwise finite-dimensional persistence modules always decompose into direct sums of simple indecomposables. We introduce the relevant terminology and notation below.

Definition 5 Let $I \subset \mathbb{R}$ be an interval. The **interval module associated to** I is the persistence module $\mathbb{k}_I : (\mathbb{R}, \leq) \rightarrow \mathbf{vect}_{\mathbb{k}}$ with $\mathbb{k}_I(s) = \mathbb{k}$ if and only if $s \in I$ and otherwise $\mathbb{k}_I(s)$ is the zero vector space. We define $\mathbb{k}_I(s \leq t) = \text{id}_{\mathbb{k}}$ if and only if $s, t \in I$ and otherwise $\mathbb{k}_I(s \leq t)$ is the zero map.

Crawley-Boevey's theorem says that any pointwise finite-dimensional persistence module is isomorphic to a direct sum of interval modules, and that this representation is unique up to permuting factors. A **barcode** $B = \{(I, m_I)\}$ is a multiset of intervals in the real line, i.e. $I \subseteq \mathbb{R}$ is an interval and $m_I \in \mathbb{N}$ indicates its multiplicity. It follows from the discussion above that any pointwise finite dimensional $\mathbb{R}$ -module F has a uniquely associated barcode $B(F)$. Let **Barcodes** denote the set of all barcodes.

The existence of barcode representations of persistence modules allows for various methods of visualization and analysis of the multiscale topology of a filtered space. A barcode can be represented visually by drawing the collection of intervals in the plane (see the righthand column of Fig. 3) or as a multiset of points in the plane called a **persistence diagram** (see the righthand column of Fig. 1); here, the endpoints of each interval are plotted as an ordered pair.




Fig. 3 A motivating example for the need for decorated merge trees. Two subsets X and Y of $\mathbb{R}^2$ along with two times in their offset filtration are shown. Degree-0 and degree-1 persistent homology fails to distinguish them as the number of components and the number of holes are the same across all stages in the filtration. This is witnessed by their identical persistent homology barcodes, shown to the right

We now consider the simple example illustrated in Fig. 3 to see these concepts in action and to motivate the definitions introduced below. The spaces $X, Y \subset \mathbb{R}^2$ give rise to persistent spaces via their respective offset filtrations. Despite the fact that X and Y are topologically distinct, the degree-0 and degree-1 persistent homology barcodes extracted from these persistent spaces are the same. This brings us to the goal of defining richer topological signatures (decorated merge trees), which are able to track *interactions* of topological features in a persistent space.

2.2 The categorical decorated merge tree

The first definition of a decorated merge tree will be given as a purely category-theoretic construction. This perspective allows for streamlined proofs of stability theorems below, but has the downside that the connection to merge trees may not be as transparent. This is remedied with alternative constructions in the following subsections.

The key observation for defining the categorical decorated merge tree is that homology decomposes into a direct sum of the homology of each connected component. This is expressed in the well known fact that

$$\text{if } X \cong \bigsqcup X_i := \bigcup X_i \times \{i\} \quad \text{then} \quad H_n(X) \cong \bigoplus_i H_n(X_i).$$


Moreover, if $f : X \rightarrow Y$ is a continuous map of spaces then we can view f as a map between disjoint unions that send each factor in the domain to a unique factor in the range. In other words, the continuous map

$$f = \sqcup f_i : \bigsqcup_{i \in \pi_0(X)} X_i \rightarrow \bigsqcup_{j \in \pi_0(Y)} Y_j$$

can be *parameterized* by the underlying map of sets $\pi_0(f) : \pi_0(X) \rightarrow \pi_0(Y)$. This indicates that we can parameterize maps inside of a persistent space along its associated persistent set of connected components. This requires some further restrictions on properties of the topological spaces involved such as local connectedness. First we isolate an important categorical construction.

Definition 6 Let $\mathcal{C}$ be a category. The **category of discretely parameterized objects** in $\mathcal{C}$, written $\text{p}\mathcal{C}$, has for objects functors $I : S \rightarrow \mathcal{C}$ where S is a set viewed as a discrete category, i.e. the only morphisms in S are identity morphisms. The functor I amounts to a choice of object of $\mathcal{C}$ for each $s \in S$. We will refer to such a functor as an **S -parameterized object**. A morphism from an S -parameterized object $I : S \rightarrow \mathcal{C}$ to a T -parameterized object $J : T \rightarrow \mathcal{C}$ consists of a map of sets $m : S \rightarrow T$ and a natural transformation from the functor I to the pullback of J along m , i.e. a morphism is a natural transformation $\alpha : I \Rightarrow m^*J$ where $m^*J := J \circ m$.

We note that if $\mathcal{C}$ has coproducts, then $\text{p}\mathcal{C}$ participates in the following diagram of categories and functors:



The functor dom sends any S -parameterized object $I : S \rightarrow \text{C}$ to the underlying parameterizing set S . The functor cop sends the diagram $I : S \rightarrow \text{C}$ to its colimit, which is the coproduct in this case. Before exploiting the above diagram further, we state the result that was used implicitly at the outset of this subsection.

Lemma 1 *Let Top^{c} denote the category of connected and locally connected topological spaces. Let Top^{lc} denote the category of locally connected spaces. The coproduct functor induces an equivalence between these categories:*

$$\text{cop} : \text{pTop}^{\text{c}} \rightarrow \text{Top}^{\text{lc}} \quad \text{where} \quad I : S \rightarrow \text{Top}^{\text{c}} \rightsquigarrow \bigsqcup_{s \in S} I(s).$$

Proof To complete the proof, it suffices to show that $\text{cop} : \text{pTop}^{\text{c}} \rightarrow \text{Top}^{\text{lc}}$ is full, faithful and essentially surjective (Riehl 2017, Thm. 1.5.9). The essentially surjective property is true by virtue of the fact that every locally connected space is naturally homeomorphic to the coproduct (i.e., disjoint union) of its components. *Full* and *faithful* mean that if $I : S \rightarrow \text{Top}^{\text{c}}$ and $J : T \rightarrow \text{Top}^{\text{c}}$ are two parameterized connected spaces, then the map

$$\text{Hom}_{\text{pTop}^{\text{c}}}(I, J) \rightarrow \text{Hom}_{\text{Top}^{\text{lc}}}\left(\bigsqcup_{s \in S} I(s), \bigsqcup_{t \in T} J(t)\right)$$

is surjective and injective, respectively. To show surjectivity (fullness), we have to show that every continuous map

$$f : \bigsqcup_{s \in S} I(s) \rightarrow \bigsqcup_{t \in T} J(t)$$

is realized by some morphism (m, α) in pTop^{c} . Here connectivity of each $I(s)$ is an essential part of the hypothesis because it allows us to associate to each $s \in S$ a unique $t \in T$ so that $f(I(s)) \subseteq J(t)$. This specifies the map of sets $m : S \rightarrow T$. The restriction of the continuous map f to each $I(s)$ specifies the components of a natural transformation $\alpha : I \Rightarrow m^*J$. Injectivity is not difficult to see because if (m, α) and (n, β) are two morphisms that induce the same map between the disjoint unions, then set-theoretically they are equal as well. Recalling the set-theoretic definition of the disjoint union, this means that

$$\sqcup \alpha_s = \sqcup \ell_s : \bigcup_{s \in S} I(s) \times \{s\} \rightarrow \bigcup_{t \in T} J(t) \times \{t\}$$

and in particular that $m = n$ and $m^*\alpha = n^*\beta$.

One consequence of Lemma 1 is that we can define another functor that serves as sort of “inverse” to cop , up to natural isomorphism.

Definition 7 The **parameterized by components** functor

$$\text{pbc} : \text{Top}^{\text{lc}} \rightarrow \text{pTop}^{\text{c}}$$

takes each locally connected space X to the object $I : \pi_0(X) \rightarrow \text{Top}^{\text{c}}$ which takes the label $i \in \pi_0(X)$ for an equivalence class to the underlying subset of X carved out by this equivalence class, equipped with the subspace topology. A map of spaces $f : X \rightarrow Y$ is taken to the morphism (m, α) where $m = \pi_0(f)$ is the map recording which connected component of X maps to which component of Y and α is the natural transformation that records the restriction of f to each component.

By virtue of (Riehl 2017, Def. 1.5.4), an alternative proof to Lemma 1 is that

$$\text{pbc} \circ \text{cop} \cong \text{id}_{\text{pTop}^{\text{lc}}} \quad \text{and} \quad \text{cop} \circ \text{pbc} \cong \text{id}_{\text{Top}^{\text{lc}}}.$$

We leverage the above identities to provide our first refinement of persistent spaces into functors from $(\mathbb{R}, \leq) \rightarrow \text{pTop}^{\text{c}}$. This is the heart of the definition of a categorical decorated merge tree.

Lemma 2 Any persistent space $F : (\mathbb{R}, \leq) \rightarrow \text{Top}^{\text{lc}}$ has an associated **persistently parameterized space**

$$\tilde{F} := \text{pbc} \circ F : (\mathbb{R}, \leq) \rightarrow \text{pTop}^{\text{c}}.$$

The functor $\tilde{F}$ fits into the following diagram, which commutes up to natural isomorphism.

$$\begin{array}{ccc} & (\mathbb{R}, \leq) & \\ \pi_0 \circ \tilde{F} \swarrow & \downarrow \tilde{F} & \searrow F \\ & \text{pTop}^{\text{c}} & \\ \swarrow \text{dom} & \cong & \searrow \text{cop} \\ \text{Set} & \xleftarrow{\pi_0} & \text{Top}^{\text{lc}} \end{array}$$

Proof The natural isomorphism $\text{cop} \circ \text{pbc} \cong \text{id}_{\text{Top}^{\text{lc}}}$ from the remark above can be restricted to the image of F to yield

$$\text{cop} \circ \text{pbc} \circ F \cong F \quad \Leftrightarrow \quad \text{cop} \circ \tilde{F} \cong F.$$

Explicitly this means that for every $s \in \mathbb{R}$ the spaces $\text{cop} \circ \tilde{F}(s)$ and $F(s)$ are homeomorphic. Since homeomorphic spaces have isomorphic sets of components, we know that the persistent sets $\pi_0 \circ F$ and $\text{dom} \circ \tilde{F}$ are naturally isomorphic as well.




Fig. 4 Associated to the offset filtration of the two subsets X and Y of $\mathbb{R}^2$ from Fig. 3 are two categorical decorated merge trees $\tilde{F}_1$ and $\tilde{G}_1$, as defined in Definition 8. This figure shows the mechanics of checking if the parameterized vector spaces at offset 0, $\tilde{F}_1(0) = I : S \rightarrow \mathbf{Vect}$ and $\tilde{G}_1(0) = J : T \rightarrow \mathbf{Vect}$, are isomorphic. They are not, which proves that our categorical decorated merge tree can distinguish these spaces. See Example 4 for more details

We are now able to define the categorical decorated merge tree of a persistent (locally connected) space.

Definition 8 Let $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}^{\text{lc}}$ be a persistent space where every space is locally connected and let $\tilde{F} : (\mathbb{R}, \leq) \rightarrow \mathbf{pTop}^c$ denote the persistently parameterized space from Lemma 2. The **categorical decorated merge tree in degree n** is the functor

$$\tilde{F}_n := H_n \circ \tilde{F} : (\mathbb{R}, \leq) \rightarrow \mathbf{pVect} \quad \text{where} \quad s \rightsquigarrow I(s) : \pi_0(F(s)) \rightarrow \mathbf{Vect}.$$

The functor $I(s) : \pi_0(F(s)) \rightarrow \mathbf{Vect}$ is a parameterized vector space (in the sense of Definition 6) that sends each component index $i \in \pi_0(F(s))$ to the homology vector space of that component, i.e. $H_n(F(s)_i)$.

Example 4 (Our Motivating Example, Reconsidered) In Fig. 3, we considered the offset filtrations F and G associated to two different subsets of the plane X and Y , respectively. Following Definition 8 we can associate two categorical decorated merge trees in degree 1, $\tilde{F}_1$ and $\tilde{G}_1$, to X and Y . To verify that $\tilde{F}_1 \not\cong \tilde{G}_1$ it suffices to show that their values at filtration value 0, which we denote by $I : \{x, y\} \rightarrow \mathbf{Vect}$ and $J : \{a, b\} \rightarrow \mathbf{Vect}$, are not isomorphic in the category $\mathbf{pVect}$.

Two parameterized vector spaces $I : S \rightarrow \mathbf{Vect}$ and $J : T \rightarrow \mathbf{Vect}$ are **isomorphic** if there are set maps $m : S \rightarrow T$ and $n : T \rightarrow S$ and natural transformations $\alpha : I \Rightarrow m^*J$ and $\beta : J \Rightarrow n^*I$ satisfying

$$m^*\beta \circ \alpha = \text{id}_I \quad \text{and} \quad n^*\alpha \circ \beta = \text{id}_J$$

in particular,

$$n \circ m = \text{id}_S \quad \text{and} \quad m \circ n = \text{id}_T.$$

By considering the parameterized vector spaces at 0 in our example, $I : \{x, y\} \rightarrow \mathbf{Vect}$ and $J : \{a, b\} \rightarrow \mathbf{Vect}$, where $I(x) = \mathbb{k}^2$, $I(y) = 0$, $J(a) = \mathbb{k}$ and $J(b) = \mathbb{k}$, we can easily show that no isomorphism is possible because any bijection between $S = \{x, y\}$ and $Y = \{a, b\}$ will force a linear transformation of the form

$$\mathbb{k}^2 \rightarrow \mathbb{k} \rightarrow \mathbb{k}^2,$$

which can never be an isomorphism. This argument is summarized in Fig. 4.

2.3 The concrete decorated merge tree

The categorical notion of a decorated merge tree boils down to the following sequence of assignments: to each real number $s \in \mathbb{R}$ a set $I(s)$ is assigned and then to each element $i \in I(s)$ a (homology) vector space is assigned. This process is reminiscent of specifying an element of $\text{Hom}(A, \text{Hom}(B, C))$, which amounts to assigning to each element of A a map from B to C . The reader then might find it useful to consider the adjunction between products and exponentials gotten by currying:

$$\text{Hom}(A, \text{Hom}(B, C)) \cong \text{Hom}(A \times B, C).$$

In this section we work with, in essence, the right hand side of this isomorphism, where $A \times B$ is replaced with the generalized merge tree $\mathcal{M}_F$ and C is replaced with the category of vector spaces. The analog of currying in this section is *concretization*, the namesake of the *concrete decorated merge tree* defined below.

Definition 9 (The Concrete Decorated Merge Tree) Let $F : (\mathbb{R}, \leq) \rightarrow \text{Top}^{\text{lc}}$ be a persistent space with all $F(s)$ locally connected and let $(\mathcal{M}_F, \preceq)$ denote its generalized merge tree. The **concrete decorated merge tree in degree n** is the functor

$$\mathcal{F}_n : (\mathcal{M}_F, \preceq) \rightarrow \text{Vect}_{\mathbb{k}} \quad \text{where} \quad (i, s) \rightsquigarrow H_n(F(s)_i \mathbb{J} \mathbb{k})$$

that records the n^{th} homology of the i^{th} component of $F(s)$.

Remark 2 The definition of a concrete decorated merge tree can be abstracted in a way that does not refer to a persistent space F at all. That is, we can define a concrete decorated merge tree more generally to be a functor $\mathcal{F} : (\mathcal{M}_F, \preceq) \rightarrow \text{Vect}$ on a generalized merge tree (considered as a poset) associated to some persistent set. When we wish to displace emphasis on the originating persistent space F , we will use the term **tree module** to refer to $\mathcal{F}$.

Remark 3 A categorical merge tree can also be considered abstractly as a functor $\tilde{F} : (\mathbb{R}, \leq) \rightarrow \text{pVect}$. Moreover, these abstract constructions are equivalent in the sense that one can always be constructed from the other. Indeed, given a functor $\tilde{F} : (\mathbb{R}, \leq) \rightarrow \text{pVect}$, we define a persistent set $S : (\mathbb{R}, \leq) \rightarrow \text{Set}$ via $S(s) = \text{dom}(\tilde{F}(s))$. Let $\mathcal{M}$ denote the associated display poset and define $\mathcal{F} : (\mathcal{M}, \preceq) \rightarrow \text{Vect}$ via $\mathcal{F}(i, s) = \tilde{F}(s)(i)$. Likewise, from a functor $\mathcal{F} : (\mathcal{M}, \preceq) \rightarrow \text{Vect}$, where $\mathcal{M}$ is a generalized merge tree associated to a persistent set $S : (\mathbb{R}, \leq) \rightarrow \text{Vect}$, we define $\tilde{F} : (\mathbb{R}, \leq) \rightarrow \text{pVect}$ by setting $\text{dom}(\tilde{F}(s)) = S(s)$ and $\tilde{F}(s)(i) = \mathcal{F}(i, s)$.

2.4 The barcode decorated merge tree

To a persistent space $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}^{\text{lc}}$ we have already shown how to associate two (equivalent) devices to record homology in degree n as it varies across components and filtration values $s \in \mathbb{R}$. Unfortunately, unlike ordinary persistent homology modules, neither of these devices have simple summaries such as barcodes or persistence diagrams. This is due to the fact that the underlying poset $(\mathcal{M}_F, \preceq)$ is not totally-ordered. However, if one considers the restriction of a tree module $\mathcal{F}_n$ to the principal up set at a point $p = (i, s) \in \mathcal{M}_F$, then we *do* obtain a module indexed by a totally ordered set and can call this the “barcode at p .” This motivates the following definitions.

Definition 10 A **barcode decorated merge tree** is a map from a generalized merge tree to the set of barcodes,

$$\mathcal{B} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{Barcodes}.$$

We say that a barcode decorated merge tree is **determined by restriction** if whenever $(i, s) \preceq p \preceq q \preceq (j, t) \in \mathcal{M}_F$, we have that

$$\mathcal{B}(q) = \mathcal{B}(p) \cap [t, \infty).$$

If the generalized merge tree **has leaves**, meaning that every maximal chain in $(\mathcal{M}_F, \preceq)$ has a minimal element, then we call a barcode decorated merge tree that is determined by restriction a **leaf-decorated merge tree**.

Definition 11 Suppose $(\mathcal{M}_F, \preceq)$ is the generalized merge tree associated to the persistent set $\pi_0 \circ F$. Given a tree module $\mathcal{F} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{Vect}$ and a point $p = (i, s) \in \mathcal{M}_F$, we define the **restriction of $\mathcal{F}$ to the principal up set $\mathcal{U}_p$** to be the $\mathbb{R}$ -module

$$\mathcal{F}|_{\mathcal{U}_p} : (\mathbb{R}, \leq) \rightarrow \mathbf{Vect} \quad \text{where, for } s \leq t, \quad \mathcal{F}|_{\mathcal{U}_p}(t) = \mathcal{F}(\pi_0 \circ F(s \leq t)(i), t).$$

For $r < s$ we defined $\mathcal{F}|_{\mathcal{U}_p}(r) = 0$.

Proposition 1 Assume that the generalized merge tree $(\mathcal{M}_F, \preceq)$ has leaves. To any pointwise finite dimensional tree module $\mathcal{F} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{vect}$, we have a leaf-decorated merge tree

$$\mathcal{BF} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{Barcodes} \quad \text{where} \quad \mathcal{BF}(p) = BC(\mathcal{F}|_{\mathcal{U}_p}).$$

Here $\mathcal{U}_p := \{q \in \mathcal{M}_F \mid p \preceq q\}$ is the principal up set at p . Note that since $(\mathcal{M}_F, \preceq)$ has leaves, it suffices to compute this barcode for every leaf node v .

Definition 12 (Barcode Transform) The map $\mathcal{BF}$ whose existence is implied by Proposition 1 is referred to as the **barcode transform of $\mathcal{F}$** .

Proof of Proposition 1 Since the tree module is already pointwise finite dimensional, the restriction at each principal up set $\mathcal{U}_p$ will be a pointwise finite dimensional $\mathbb{R}$ -module. Crawley-Boevey’s Theorem (Crawley-Boevey 2015) then implies that this

restricted tree module has a barcode. Obviously the barcode decoration is determined by restriction because for any pair of comparable points $p \preceq q$ the restriction at $\mathcal{U}_q$ can be obtained by restricting the module at $\mathcal{U}_p$.

We have the following immediate corollary.

Corollary 1 *If $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ is a persistent space whose associated generalized merge tree $\mathcal{M}_F$ has leaves and whose associated concrete decorated merge tree in degree $n \geq 0$ $\mathcal{F}_n : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{vect}$ is pointwise finite dimensional, then F has an associated leaf-decorated merge tree in degree n , $\mathcal{BF}_n$.*

Remark 4 The barcode transform associates an (indexed) ensemble of barcodes to a filtered space through a certain “slicing” operation (i.e., slicing along upsets from leaves). This operation is analogous to several other recent methods in the TDA literature; we provide a few examples here. The *persistent homology transform* (Curry et al. 2018; Turner et al. 2014) associates to an embedded simplicial complex a barcode for each direction, obtained by using projection onto the direction axis as a filtration function. A *fibred barcode* (Hang and Mio 2020; Lesnick and Wright 2015) is a collection of barcodes associated to a multiparameter persistence module by slicing the parameter space by affine lines. The *barcode embedding* (Dey et al. 2015; Oudot and Solomon 2017) associates a barcode to each point in a metric graph by computing extended persistence with respect to a radial filtration centered at that point.

Leaf decorated merge trees are the most computationally tractable of the three variants of DMTs introduced in this section. Algorithms for analyzing DMTs are described in Sects. 5 and 6.

3 Continuity and stability of decorated merge trees

In this section we focus on metric theoretic aspects of decorated merge trees—namely, continuity of the barcode transform $\mathcal{BF} : \mathcal{M}_F \rightarrow \mathbf{Barcodes}$ of a decorated merge tree and stability of various pseudometrics on the space of decorated merge trees. To properly state our results, we first briefly review standard metrics in the TDA literature.

3.1 Metrics on persistence modules

Let $B, B' \in \mathbf{Barcodes}$ be barcodes.

Definition 13 A **matching** of barcodes B and B' is a bijection ξ between subsets $\text{dom}(\xi) \subset B$ and $\text{ran}(\xi) \subset B'$ —such a ξ is also commonly referred to as a **partial bijection** between B and B' . The **cost** of a matching ξ is

$$\max \left\{ \max_{I \in \text{dom}(\xi)} \|I - \xi(I)\|_\infty, \max_{I \in B \setminus \text{dom}(\xi)} \|I\|_\Delta, \max_{I' \in B' \setminus \text{ran}(\xi)} \|I'\|_\Delta \right\},$$

where, for I and I' with endpoints $b \leq d$ and $b' \leq d'$, respectively,

$$\|I - I'\|_{\infty} := \max\{|b - b'|, |d - d'|\} \quad \text{and} \quad \|I\|_{\Delta} := \frac{d - b}{2}.$$

If a matching ξ has cost less than or equal to ϵ , we refer to ξ as an ϵ -**matching**. The **bottleneck distance** between B and B' is

$$d_B(B, B') := \inf\{\epsilon \geq 0 \mid \text{there exists an } \epsilon\text{-matching of } B \text{ and } B'\}.$$

A matching realizing the bottleneck distance will be called an **optimal matching**.

The bottleneck distance is an especially natural metric due to its connection to a category-theoretic metric on persistence modules called *interleaving distance*. In fact, interleaving distance can be defined for more general functor categories (Bubenik et al. 2015; Bubenik and Scott 2014). Generalizing interleaving distances to increasingly abstract classes of functors is an active field of research (e.g., de Silva et al. 2018; Stefanou 2018), we define the interleaving distance below at the level of generality which will be most useful to us. In what follows, $\mathcal{C}$ denotes a fixed but arbitrary category.

Definition 14 (Interleaving Distance) Consider $(\mathbb{R}, \leq)$ as a poset category. For $\epsilon \in [0, \infty)$, we define the poset map $\sigma^{\epsilon} : \mathbb{R} \rightarrow \mathbb{R}$ via $\sigma^{\epsilon} : s \mapsto s + \epsilon$. Now let $F : (\mathbb{R}, \leq) \rightarrow \mathcal{C}$ be a functor, which is an object in the functor category $\text{Fun}(\mathbb{R}, \mathcal{C})$. We define the ϵ -**shift** of F , written $F^{\epsilon} : (\mathbb{R}, \leq) \rightarrow \mathcal{C}$, as $F^{\epsilon} := F \circ \sigma^{\epsilon}$. This shift is functorial, i.e. to every morphism $\phi : F \Rightarrow G$ in $\text{Fun}(\mathbb{R}, \mathcal{C})$ we have another morphism $\phi^{\epsilon} : F^{\epsilon} \Rightarrow G^{\epsilon}$, which defines the ϵ -**translation functor** $(\bullet)^{\epsilon} : \text{Fun}(\mathbb{R}, \mathcal{C}) \rightarrow \text{Fun}(\mathbb{R}, \mathcal{C})$. Since each object F has a naturally associated **internal ϵ -shift**,

$$\eta_F^{\epsilon} : F \Rightarrow F^{\epsilon} \quad \text{where} \quad \eta_F^{\epsilon}(s) : F(s) \rightarrow F(\sigma^{\epsilon}(s)) \quad \text{is} \quad F(s \leq s + \epsilon),$$

there is a natural transformation from the identity functor $\text{id}_{\mathcal{C}_{\mathbb{R}}}$ on $\text{Fun}(\mathbb{R}, \mathcal{C})$ to $(\bullet)^{\epsilon}$.

Two objects $F, G : (\mathbb{R}, \leq) \rightarrow \mathcal{C}$ are ϵ -**interleaved** if there are morphisms $\phi : F \Rightarrow G^{\epsilon}$ and $\psi : G \Rightarrow F^{\epsilon}$ such that

$$\psi^{\epsilon} \circ \phi = \eta_F^{2\epsilon} \quad \text{and} \quad \phi^{\epsilon} \circ \psi = \eta_G^{2\epsilon}.$$

We define the **interleaving distance between F and G** as

$$d_I(F, G) = \inf\{\epsilon \geq 0 \mid F \text{ and } G \text{ are } \epsilon\text{-interleaved}\}.$$

When considering persistence modules $F, G : (\mathbb{R}, \leq) \rightarrow \text{Vect}$, the above definition reduces to the classical interleaving distance of persistence modules as introduced in the landmark paper (Chazal et al. 2009). Building on this work, it was shown in Lesnick (2015) that the map taking a persistence module to its barcode is an isometry with respect to the interleaving and bottleneck distances.

3.2 Continuity of the barcode transform

We now establish the continuity of the barcode transform $\mathcal{BF} : \mathcal{M}_F \rightarrow \text{Barcodes}$ associated to a concrete decorated merge tree $\mathcal{F}$. Continuity is measured with respect to a whole family of ℓ^p -type metrics on the underlying merge tree $\mathcal{M}_F$ for $p \in [1, \infty]$. These metrics measure distance between pairs of points via their least common ancestor. The metric on **Barcodes** is the bottleneck distance (Definition 13).

Definition 15 Let $(\mathcal{M}_F, \preceq)$ be a generalized merge tree where $\pi_F : \mathcal{M}_F \rightarrow \mathbb{R}$ is the projection function. The **merge height** of points $u, v \in \mathcal{M}_F$ is

$$\text{merge}_F(u, v) := \inf\{\pi_F(w) \mid u, v \preceq w\}.$$

We define the **least common ancestor** of u and v to be the least common upper bound of u and v ,

$$\text{LCA}_F(u, v) := \text{argmin}\{\pi_F(w) \mid u, v \preceq w\},$$

when it exists. In this case,

$$\text{merge}_F(u, v) = \pi_F(\text{LCA}_F(u, v)).$$

We will drop the subscript F when convenient. Finally, if every pair of points in a merge tree have a least common ancestor, then we say the merge tree is **connected**.

Remark 5 Note that for a generalized merge tree the least common ancestor need not exist. For example if $S : (\mathbb{R}, \leq) \rightarrow \text{Set}$ is the persistent set that assigns $S(t) = \{x, y\}$ for $t \leq 0$ and $S(t) = \{z\}$ for $t > 0$, then there will not be a least upper bound for x, y . The assumption that S is constructible (Definition 24) will, however, guarantee the existence of least common ancestors.

Definition 16 If the generalized merge tree $\mathcal{M}_F$ is connected, then the ℓ^p **metric on** $(\mathcal{M}_F, \preceq)$ for $1 \leq p \leq \infty$ associates to every pair of points $u, v \in \mathcal{M}_F$ the value

$$d_{\mathcal{M}_F}^p(u, v) = \|(\text{merge}(u, v) - \pi(u), \text{merge}(u, v) - \pi(v))\|_p.$$

Using general estimates for ℓ^p norms on $\mathbb{R}^2$, one sees that the ℓ^p metrics are bi-Lipschitz equivalent and therefore induce the same topology. Moreover, we have the following characterization when $\mathcal{M}_F$ comes from a sublevel set filtration.

Proposition 2 Let $f : X \rightarrow \mathbb{R}$ be a continuous map such that the associated merge tree $\mathcal{M}_f$ has finitely many leaves. Then the topology induced by each ℓ^p metric coincides with the quotient space topology.

We defer the proof of the proposition to Appendix A, since it relies on other technical results about merge tree topologies.

Theorem 1 (Continuity of Barcode DMTs) *Let $(\mathcal{M}_F, \preceq)$ be a connected generalized merge tree associated to a persistent set $\pi_0 \circ F$. Endow $\mathcal{M}_F$ with the extended metric $d_{\mathcal{M}_F}^p$ and Barcodes with the bottleneck distance d_B . For any pointwise finite dimensional tree module $\mathcal{F} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{vect}$, the associated barcode decorated merge tree $\mathcal{BF} : \mathcal{M}_F \rightarrow \mathbf{Barcodes}$ is $2^{1-1/p}$ -Lipschitz for $p \in [1, \infty)$ and 2-Lipschitz for $p = \infty$.*

Proof Let $p = (i, s)$ and $q = (j, t)$ be elements of $\mathcal{M}_F$ and suppose that $\text{merge}_F(p, q) = r_0$, so that $d_{\mathcal{M}_F}^1(p, q) = (r_0 - s) + (r_0 - t)$. For arbitrary $\delta > 0$, choose $z = (k, r) \in \mathcal{M}_F$ such that z is a common ancestor of p and q and $r < r_0 + \delta$. The barcode $\mathcal{BF}(z)$ is obtained by truncating each bar of $\mathcal{BF}(p)$ at height r (see Definition 34 in the Appendix for a precise definition of *truncation*), and a simple consequence is that $d_B(\mathcal{BF}(z), \mathcal{BF}(p)) \leq r - s$; likewise, $d_B(\mathcal{BF}(z), \mathcal{BF}(q)) \leq r - t$. By the triangle inequality,

$$d_B(\mathcal{BF}(p), \mathcal{BF}(q)) \leq (r - s) + (r - t) < d_{\mathcal{M}_F}^1(p, q) + 2\delta.$$

Since this holds for arbitrary $\delta > 0$, this proves that $\mathcal{BF}$ is 1-Lipschitz with respect to $d_{\mathcal{M}_F}^1$. The remaining cases follow by the general bounds $\|\cdot\|_1 \leq 2^{1-1/p} \|\cdot\|_p$, for $p \in [1, \infty)$ and $\|\cdot\|_1 \leq 2 \|\cdot\|_\infty$ in $\mathbb{R}^2$.

3.3 Interleavings of merge trees

Implicit in Definition 14 is a way of comparing generalized merge trees. If F and G are persistent spaces, then the interleaving distance between $\pi_0 \circ F$ and $\pi_0 \circ G$, which are both objects of $\mathbf{Fun}(\mathbb{R}, \mathbf{Set})$, is covered by that definition. However, the original definition of the merge tree interleaving distance given in Morozov et al. (2013) used the classical merge tree construction (Definition 3) and assumed continuity of the interleaving maps. In this section we review the construction of Morozov et al. (2013) and prove that under suitable hypotheses this classical interleaving distance is the same as the interleaving distance implied by Definition 14, which we have isolated as Definition 18.

Definition 17 (Classical Merge Tree Interleaving Distance Morozov et al. 2013) Let $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ be continuous functions and let $\mathcal{M}_f$ and $\mathcal{M}_g$ be their associated merge trees, as defined in Definition 3. For $\epsilon \geq 0$ we define an ϵ -**map** to be a continuous (with respect to quotient space topologies) map $\alpha : \mathcal{M}_f \rightarrow \mathcal{M}_g$ such that

$$\tilde{\pi}_g \circ \alpha([x], t) = t + \epsilon$$

for all points $([x], t) \in \mathcal{M}_f$. Two ϵ -maps $\alpha : \mathcal{M}_f \rightarrow \mathcal{M}_g$ and $\beta : \mathcal{M}_g \rightarrow \mathcal{M}_f$ are said to be ϵ -**compatible** if $\beta \circ \alpha = \eta_f^{2\epsilon}$ and $\alpha \circ \beta = \eta_g^{2\epsilon}$. The **Morozov-Bekatajev-Weber (MBW) interleaving distance between merge trees $\mathcal{M}_f$ and $\mathcal{M}_g$** is

$$\begin{aligned} \theta_1^{\text{MBW}}(\mathcal{M}_f, \mathcal{M}_g) \\ := \inf\{\epsilon \geq 0 \mid \exists \epsilon\text{-compatible maps } \alpha : \mathcal{M}_f \rightarrow \mathcal{M}_g \text{ and } \beta : \mathcal{M}_g \rightarrow \mathcal{M}_f\}. \end{aligned}$$

Given two functions $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$, we note that their associated sublevel set filtrations determine persistent spaces $F, G : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$. By post-composing these functors with $\pi_0 : \mathbf{Top} \rightarrow \mathbf{Set}$, we get two persistent sets, which can be compared with Definition 14 when $C = \mathbf{Set}$. We isolate this special case now.

Definition 18 (Modern Merge Tree Interleaving Distance Bubenik et al. 2015) The **interleaving distance between merge trees** $\mathcal{M}_f$ and $\mathcal{M}_g$ associated to $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ is defined to be the interleaving distance between the persistent sets $\pi_0 \circ F$ and $\pi_0 \circ G$ where $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ is defined by $F(t) := f^{-1}(-\infty, t]$ and $G : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ is defined by $G(t) := g^{-1}(-\infty, t]$. We introduce the special notation

$$\theta_I(\mathcal{M}_f, \mathcal{M}_g) := d_I(\pi_0 \circ F, \pi_0 \circ G).$$

We now compare these two definitions of interleaving distance. As a preliminary observation, note that an ϵ -map $\alpha : \mathcal{M}_f \rightarrow \mathcal{M}_g$ carries components of $f^{-1}(-\infty, t]$ to $g^{-1}(-\infty, t + \epsilon]$. This is exactly the expression that α defines a natural transformation $\alpha : \pi_0 \circ F \Rightarrow \pi_0 \circ G^\epsilon$.

Lemma 3 If $\alpha : \mathcal{M}_f \rightarrow \mathcal{M}_g$ and $\beta : \mathcal{M}_g \rightarrow \mathcal{M}_f$ are ϵ -maps that are ϵ -compatible, then they induce an ϵ -interleaving between the persistent sets $\pi_0 \circ F$ and $\pi_0 \circ G$ of the sublevel set filtrations of $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$.

Proof The discussion above has already established that an ϵ -map α is equivalent to specifying a natural transformation $\alpha : \pi_0 \circ F \Rightarrow \pi_0 \circ G^\epsilon$ and that β is likewise tantamount to a natural transformation $\beta : \pi_0 \circ G \Rightarrow \pi_0 \circ F^\epsilon$ without referring to continuity. By comparing with Definition 14 it is now obvious that setting α to be ϕ and β to be ψ and enforcing the ϵ -compatibility condition then implies the interleaving condition defined there.

We note that because an ϵ -interleaving of persistent sets does not necessarily guarantee continuity of the maps between the merge trees $\mathcal{M}_f$ and $\mathcal{M}_g$, Lemma 3 establishes the bound:

$$\theta_I(\mathcal{M}_f, \mathcal{M}_g) \leq \theta_I^{\text{MBW}}(\mathcal{M}_f, \mathcal{M}_g).$$

In the next proposition we provide hypotheses under which these two interleaving distances are the same. This fills a gap in the literature that is not usually remarked upon and helps bridge the gap between (Morozov et al. 2013) and (Bubenik et al. 2015).

Proposition 3 Suppose $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ are continuous maps defined on compact spaces X and Y so that their (classical) merge trees $\mathcal{M}_f$ and $\mathcal{M}_g$ have finitely many leaves. Let F and G denote the sublevel set filtrations of f and g . Every ϵ -interleaving of the persistent sets $\pi_0 \circ F$ and $\pi_0 \circ G$ defines a pair of ϵ -compatible maps between the merge trees $\mathcal{M}_f$ and $\mathcal{M}_g$. Consequently,

$$\theta_I^{\text{MBW}}(\mathcal{M}_f, \mathcal{M}_g) = \theta_I(\mathcal{M}_f, \mathcal{M}_g) = d_I(\pi_0 \circ F, \pi_0 \circ G).$$

Proof Suppose $\phi : \pi_0 \circ F \Rightarrow \pi_0 \circ G^\epsilon$ and $\psi : \pi_0 \circ G \Rightarrow \pi_0 \circ F^\epsilon$ specify an ϵ -interleaving. Now consider the display posets of $\pi_0 \circ F$ and $\pi_0 \circ G$, written $\mathcal{M}_F$ and $\mathcal{M}_G$ to distinguish them from the merge trees $\mathcal{M}_f$ and $\mathcal{M}_g$. It is obvious that $\mathcal{M}_F$ and $\mathcal{M}_f$ are identical as posets, similarly for $\mathcal{M}_G$ and $\mathcal{M}_g$. Following the discussion in Appendix A we can equip $\mathcal{M}_F$ and $\mathcal{M}_G$ with the *interval topology* (see Definition 32). By Lemma 5 the principal up set of any point is closed in this topology. Since $\mathcal{M}_F$ and $\mathcal{M}_G$ have finitely many leaves, consider the closed cover of each by the principal up sets at each of the leaf nodes. We first show that the map $\phi : \mathcal{M}_F \rightarrow \mathcal{M}_G$ is continuous when restricted to the up set $\mathcal{U}_v$ of any leaf node $v \in \mathcal{M}_F$ because ϕ carries $\mathcal{U}_v$ bijectively onto $\mathcal{U}_{\phi(v)}$.

Bijectivity follows from the fact that ϕ is a natural transformation of functors. Indeed, consider a triple $v < p < q$ in the display poset $\mathcal{M}_F$ for F where $v \in F(r)$, $p \in F(s)$ and $q \in F(t)$. Necessarily, $r < s < t$ are strictly ordered in $\mathbb{R}$ and $\phi(v) \in G(r + \epsilon)$, $\phi(p) \in G(s + \epsilon)$, $\phi(q) \in G(t + \epsilon)$. Consequently, the images are all distinct in $\mathcal{M}_G$, because they are elements of sets with different real indices, and $\phi(v) < \phi(p) < \phi(q)$. Now if V is open in $\mathcal{M}_G$ and $p \in \phi^{-1}(V) \cap \mathcal{U}_v$, then we know there exists, for any $q > p$, a $w \in \mathcal{M}_G$ with $\phi(p) < w < \phi(q)$ and where the interval $[\phi(p), w] \subseteq V$. But ϕ is a bijection onto the image of $\mathcal{U}_v$, which contains $[\phi(p), w]$, so there exists a $w' \in \phi^{-1}(w)$ with $[p, w'] \subseteq \phi^{-1}(V)$.

Since a map is continuous if it's continuous when restricted to elements of a finite closed cover, we conclude that $\phi : \mathcal{M}_F \rightarrow \mathcal{M}_G$ is continuous with respect to the interval topology. A completely symmetric argument proves that $\psi : \mathcal{M}_G \rightarrow \mathcal{M}_F$ is continuous with respect to the interval topology as well. By Proposition 10 we can conclude that the interval topology and the quotient topology are the same, thus ϕ and ψ can be viewed as ϵ -maps (with the continuity assumption) that are ϵ -compatible.

3.4 A decorated bottleneck distance

In this section we leverage the barcode decorated merge tree perspective to define a second distance that uses an interleaving of merge trees along with an ϵ -matching of the barcodes over each of the merge trees, as determined by their decorations. Although finding optimal interleavings between merge trees is NP-Hard (Agarwal et al. 2018), we provide in Sect. 5 an algorithm for estimating them.

Definition 19 Assume that $\mathcal{M}_F$ and $\mathcal{M}_G$ are generalized merge trees associated to the persistent sets $\pi_0 \circ F$ and $\pi_0 \circ G$. Given two barcode decorated merge trees

$$\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes} \quad \text{and} \quad \mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$$

we define an (ϵ, δ) -**matching of $\mathcal{B}_F$ and $\mathcal{B}_G$** to consist of

- an ϵ -interleaving of the underlying generalized merge trees i.e. natural transformations $\phi : \pi_0 \circ F \Rightarrow \pi_0 \circ G^\epsilon$ and $\psi : \pi_0 \circ G \Rightarrow \pi_0 \circ F^\epsilon$ that satisfy $\psi^\epsilon \circ \phi = \eta_F^{2\epsilon}$ and $\phi^\epsilon \circ \psi = \eta_G^{2\epsilon}$. We will abuse notation and write $\phi : \mathcal{M}_F \rightarrow \mathcal{M}_G$ and $\psi : \mathcal{M}_G \rightarrow \mathcal{M}_F$ to make clear that interleaving of persistent sets always provide

- ϵ -compatible maps in the sense of Morozov et al. (2013) (without the continuity assumption) between the associated generalized merge trees;
- a δ -matching of the barcodes $\mathcal{B}_F(p)$ and $\mathcal{B}_G(\phi(p))$ for every $p \in \mathcal{M}_F$ and a δ -matching of the barcodes $\mathcal{B}_G(q)$ and $\mathcal{B}_F(\psi(q))$ for every $q \in \mathcal{M}_G$.

Definition 20 Let $p \in [1, \infty]$. The **decorated bottleneck p -distance** between two barcode decorated merge trees $\mathcal{B}_F$ and $\mathcal{B}_G$ is defined as

$$d_{B,p}(\mathcal{B}_F, \mathcal{B}_G) := \inf\{\|(\epsilon, \delta)\|_p \mid \exists (\epsilon, \delta)\text{-matching of } \mathcal{B}_F \text{ and } \mathcal{B}_G\}.$$

When $p = \infty$, we refer to the metric simply as the **decorated bottleneck distance** and write $d_B := d_{B,\infty}$.

Remark 6 We use the notation d_B for both the bottleneck distance between barcodes and for the decorated bottleneck ∞ -distance to emphasize the connection between the metrics. The meaning of d_B should always be clear from context.

Remark 7 We define the general family of decorated bottleneck p -distances in anticipation that they will be useful in data analysis tasks down the line. In this paper, we develop the theory of the $p = \infty$ metric d_B and will almost exclusively refer to the decorated bottleneck distance in what follows.

The ($p = \infty$) decorated bottleneck distance can be formulated more simply than the general p version. An (ϵ, δ) -matching of barcode decorated merge trees where $\epsilon = \delta$ will be referred to as an **ϵ -matching**. The proof of the following proposition is elementary but technical. We relegate it to Appendix C.1.

Proposition 4 Let $\mathcal{B}_F$ and $\mathcal{B}_G$ be barcode decorated merge trees which are determined by their leaves. The decorated bottleneck distance can be expressed as

$$d_B(\mathcal{B}_F, \mathcal{B}_G) = \inf\{\epsilon \geq 0 \mid \exists \epsilon\text{-matching of } \mathcal{B}_F \text{ and } \mathcal{B}_G\}.$$

The last result of this subsection says that the “determined by leaves” assumption allows us to check one of the matching conditions only at leaves. Its proof uses technical tools from the proof of Proposition 4, so we delay it to Appendix C.1.

Proposition 5 Let $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ and $\mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$ be two leaf-decorated merge trees. Let ϕ and ψ be an ϵ -interleaving of the underlying generalized merge trees $\mathcal{M}_F$ and $\mathcal{M}_G$ where for each leaf $v \in \mathcal{M}_F$ and each leaf $w \in \mathcal{M}_G$ we have a δ -matching between $\mathcal{B}_F(v)$ and $\mathcal{B}_G(\phi(v))$ and a δ -matching between $\mathcal{B}_G(w)$ and $\mathcal{B}_F(\psi(w))$. Then there is an (ϵ, δ) -matching between the entire barcode decorated merge trees $\mathcal{B}_F$ and $\mathcal{B}_G$.

3.5 Hierarchy of distances

In this section we prove the main two results of this paper.

Theorem 2 (*The Hierarchy of Stability Results for DMTs*) Let $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ be continuous functions whose sublevel sets are locally connected. Let $\mathcal{F}$ and $\mathcal{G}$ be the sublevel set filtrations, viewed as persistent spaces. The categorical decorated merge trees for homological degree n are denoted $\tilde{\mathcal{F}}_n$ and $\tilde{\mathcal{G}}_n$ and the associated barcode decorated merge trees are $\mathcal{BF}_n$ and $\mathcal{BG}_n$. Our various distances satisfy

$$\theta_I(\mathcal{M}_F, \mathcal{M}_G) \leq d_B(\mathcal{BF}_n, \mathcal{BG}_n) \leq d_I(\tilde{\mathcal{F}}_n, \tilde{\mathcal{G}}_n) \leq \delta_I(X_f, Y_g).$$

Moreover, if $X = Y$ then we have that

$$\delta_I(X_f, Y_g) \leq \|f - g\|_\infty.$$

The distance δ_I is a special case of the *persistent homotopy type distance* introduced in Frosini et al. (2019). We recall the definition below.

Definition 21 An $\mathbb{R}$ -space is a topological space X endowed with a continuous function $f : X \rightarrow \mathbb{R}$. An ϵ -interleaving of $\mathbb{R}$ -spaces $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ is a pair of continuous maps $\Phi : X \rightarrow Y$ and $\Psi : Y \rightarrow X$ along with homotopies $H_X : X \times [0, 1] \rightarrow X$ and $H_Y : Y \times [0, 1] \rightarrow Y$ connecting the identity maps id_X and id_Y with $\Psi \circ \Phi$ and $\Phi \circ \Psi$, respectively. We require further that the following four properties hold for Φ , Ψ , H_X and H_Y :

1. $\Phi(X_{\leq s}) \subseteq Y_{\leq s+\epsilon}$ for all $s \in \mathbb{R}$
2. $\Psi(Y_{\leq s}) \subseteq X_{\leq s+\epsilon}$ for all $s \in \mathbb{R}$
3. $f \circ H_X(x, t) \leq f(x) + 2\epsilon$ for all $x \in X$ and $t \in [0, 1]$
4. $g \circ H_Y(y, t) \leq g(y) + 2\epsilon$ for all $y \in Y$ and $t \in [0, 1]$

The **persistent homotopy type distance** between $\mathbb{R}$ -spaces $X_f := f : X \rightarrow \mathbb{R}$ and $Y_g := g : Y \rightarrow \mathbb{R}$ is defined as

$$\delta_I(X_f, Y_g) := \inf \{ \epsilon \mid X_f \text{ and } Y_g \text{ are } \epsilon\text{-interleaved} \}.$$

If no interleaving exists, we set $\delta_I(X_f, Y_g) = \infty$.

This metric was used in Hang et al. (2019) to develop persistent homology of $\mathbb{R}$ -spaces with different, but homotopic domains. A similar metric on $\mathbb{R}$ -spaces, called the **homotopy interleaving distance**, was defined in Blumberg and Lesnick (2017), but the exact connection between homotopy type distance and homotopy interleaving distance remains to be studied.

Proof of Theorem 2 The leftmost inequality $\theta_I(\mathcal{M}_F, \mathcal{M}_G) \leq d_B(\mathcal{BF}_n, \mathcal{BG}_n)$ follows from the definition of the decorated bottleneck distance, since the data of an interleaving of merge trees is part of the definition. Moreover it is obvious that $d_I(\tilde{\mathcal{F}}_n, \tilde{\mathcal{G}}_n) \geq \theta_I(\mathcal{M}_F, \mathcal{M}_G)$ since one can always apply the functor dom to any ϵ -interleaving $\tilde{\phi} : \tilde{\mathcal{F}}_n \Rightarrow \tilde{\mathcal{G}}_n^\epsilon$ and $\tilde{\psi} : \tilde{\mathcal{G}}_n \Rightarrow \tilde{\mathcal{F}}_n^\epsilon$ of categorical decorated merge trees to obtain an ϵ -interleaving ϕ and ψ of the underlying persistent sets or generalized merge trees; this is the Lipschitz stability result of (Bubenik et al. 2015, Thm. 2.3.6) applied to the functor dom . Moreover, if we take a point $p \in \mathcal{M}_F$ then $\tilde{\phi} : \tilde{\mathcal{F}}_n \Rightarrow \tilde{\mathcal{G}}_n^\epsilon$ specifies, by restriction, a morphism between the restriction of $\tilde{\mathcal{F}}_n$ to the upset at p and

the restriction of $\tilde{G}_n^\epsilon$ to the upset at $\phi(p)$. Similarly, since $\psi(\phi(p))$ must equal the 2ϵ translate of p in $\mathcal{M}_F$, we can conclude that the module $\tilde{F}_n|_{U_p}$ and $\phi^*\tilde{G}_n^\epsilon|_{U_p}$ are ϵ -interleaved, as $\mathbb{R}$ -modules. By the Algebraic Stability Theorem (Bauer and Lesnick 2020; Chazal et al. 2009), which states that an ϵ -interleaving of $\mathbb{R}$ -modules induces an ϵ -matching of barcodes, we have an ϵ -matching between the barcodes $\mathcal{BF}_n(p)$ and $\mathcal{BG}_n(\phi(p))$ for all p . A completely symmetric argument works where we consider $q \in \mathcal{M}_G$ and the natural transformation $\tilde{\psi} : \tilde{G}_n \Rightarrow \tilde{F}_n^\epsilon$. This proves that

$$d_B(\mathcal{BF}_n, \mathcal{BG}_n) \leq d_I(\tilde{F}_n, \tilde{G}_n).$$

To prove the rightmost inequality, first observe that if $\delta_I(X_f, Y_f) = \infty$ then we are done, so assume not. Given $\epsilon > \delta_I(X_f, Y_g)$, there exists an ϵ -interleaving of the $\mathbb{R}$ -spaces X_f and Y_g . In particular, for each $t \in \mathbb{R}$ we have that the map Φ restricts to a continuous map between the sublevel sets $\Phi : F(t) := X_{\leq t} \rightarrow Y_{\leq t+\epsilon} =: G(t+\epsilon)$. Since Φ is continuous and defined globally on X , it specifies a natural transformation of the persistent spaces $\phi : F \Rightarrow G^\epsilon$. Moreover, since the continuous image of a (path) connected set is connected, the natural transformation ϕ defines a natural transformation between the associated persistent spaces that are parameterized by their (path) components $\tilde{\phi} : \tilde{F} \Rightarrow \tilde{G}^\epsilon$. Interchanging the roles of X and Y in the above discussion implies that $\Psi : Y \rightarrow X$ defines a natural transformation $\tilde{\psi} : \tilde{G} \Rightarrow \tilde{F}^\epsilon$.

It should be noted that $\tilde{\phi}$ and $\tilde{\psi}$ *do not* define an interleaving of the functors $\tilde{F}$ and $\tilde{G}$ because that would require that for each $t \in \mathbb{R}$ the composition $\psi(t+\epsilon) \circ \phi(t)$ equals the inclusion map $X_{\leq t} \subseteq X_{\leq t+2\epsilon}$. This is not true, but condition (3) of Definition 21 does require that $\Psi \circ \Phi$ restrict to a map that is homotopic to the inclusion map $X_{\leq t} \subseteq X_{\leq t+2\epsilon}$ for all t . Since homotopic maps induce the same map on homology, we can conclude that the natural transformations $\tilde{\phi} : \tilde{F} \Rightarrow \tilde{G}^\epsilon$ and $\tilde{\psi} : \tilde{G} \Rightarrow \tilde{F}^\epsilon$ define an ϵ -interleaving of the persistent component homology modules $\tilde{F}_n$ and $\tilde{G}_n$.

In the case that $X = Y$, the desired inequality is (Frosini et al. 2019, Proposition 2.11). This completes the proof.

Although the above result establishes that categorical DMTs are more sensitive than undecorated merge trees, the relationship to higher dimensional persistent homology is not described above. The next result establishes that categorical DMTs are, in fact, more sensitive than ordinary persistent homology.

Theorem 3 *For locally connected persistent spaces $F, G : (\mathbb{R}, \leq) \rightarrow \text{Top}^{\text{lc}}$, the categorical decorated merge tree in homological degree n is more sensitive than the interleaving distance of the persistent homology modules in degree n , i.e*

$$d_I(H_n \circ F, H_n \circ G) \leq d_I(\tilde{F}_n, \tilde{G}_n).$$

Proof We have

$$\text{cop} \circ \tilde{F}_n(s) := \bigoplus_{i \in \pi^{-1}(s)} H_n(F(s)_i) \cong H_n \left(\bigsqcup_{i \in \pi^{-1}(s)} F(s)_i \right) \cong H_n \circ F(s),$$

and similarly for G .




Fig. 5 The Barcode Decorated Merge Tree comes from restricting a tree module to its leaf nodes, one at a time, and then calculating the barcode associated to the restriction of this tree module to the principal up set at each leaf node. In this figure two non-isomorphic tree modules are shown to have identical barcodes, when “viewed” from each of its leaf nodes. This example proves that the association of tree modules to their associated barcode decorations is not injective

Note that the last isomorphism on the right uses the assumption that each $F(s)$ is locally connected. Since naturally isomorphic modules have the same interleaving distance, the result now follows from the Lipschitz stability result of (Bubenik et al. 2015, Thm. 2.3.6) applied to the functor cop , i.e. applying cop is a 1-Lipschitz map from $\text{Fun}(\mathbb{R}, \text{pVect})$ to $\text{Fun}(\mathbb{R}, \text{Vect})$.

Remark 8 (No Contradiction to Universality) By reconsidering the example from Fig. 3, we know that the above inequality is not an equality, in general. Moreover, when taken together Theorems 3 and 2 may appear to contradict Lesnick’s Universality result (Lesnick 2015, Thm. 5.5), which asserts that the interleaving distance on persistent homology modules is the most sensitive distance that is bounded above by the L^∞ distance. However, there is no contradiction because universality theorems of the type stated and conjectured in Lesnick (2015) only apply to functors from $\mathbb{R}$ (or $\mathbb{R}^d$) to Vect and not pVect .

Remark 9 (No One-Way Relationship of Bottleneck Distances) In the case where locally connected persistent spaces F and G give rise to pointwise finite dimensional persistent homology modules, we know that $d_I(H_n \circ F, H_n \circ G)$ coincides with the bottleneck distance by virtue of the isometry theorem (Lesnick 2015, Thm. 3.4). Unfortunately, the decorated bottleneck distance $d_B(\mathcal{BF}_n, \mathcal{BG}_n)$ can be greater than or less than the bottleneck distance between $H_n \circ F$ and $H_n \circ G$. If one considers the motivating example in Fig. 3, then the decorated bottleneck distance there is obviously greater than the bottleneck distance between $H_1 \circ F_X$ and $H_1 \circ F_Y$. For this example, the bottleneck distance is 0, but the decorated bottleneck distance is $R/2$, corresponding to half the radius of the circles in Fig. 3. On the other hand, Fig. 5 gives an example of two non-isomorphic tree modules (concrete DMTs) with identical barcode decorations, so

the decorated bottleneck distance for $\mathcal{F}$ and $\mathcal{G}$ in those examples is 0. However, the underlying persistent homology modules, which arise via the pushforward to $\mathbb{R}$ are separated by the length of the edge from the merge point to w in the bottleneck distance.

4 Representations of tree posets and lift decorations

This section treats questions about the decomposability of decorated merge trees. We will see that the situation is much more subtle than in the classical persistence setting, due to the fact that merge tree posets are not totally ordered.

4.1 Indecomposables with totally ordered support

A crucial feature of pointwise finite dimensional (PFD) persistence modules is that they decompose as a direct sum of interval modules; these decompositions are used ubiquitously in visualization and analysis tasks in TDA. Although a PFD concrete decorated merge tree $\mathcal{F} : \mathcal{M}_F \rightarrow \mathbf{vect}$ also admits a (Remak) decomposition into a direct sum of indecomposables $\mathcal{F} \approx \bigoplus_i \mathcal{F}_i$, where the direct sum is taken pointwise (Botnan and Crawley-Boevey 2020, Theorem 1.1), the indecomposables need not have any special structure. However, we will consider a construction where the DMT does decompose into modules that are supported on totally ordered subsets of the merge tree, which will give us a structure theorem analogous to the usual one in persistence.

Definition 22 A concrete decorated merge tree $\mathcal{F} : \mathcal{M}_F \rightarrow \mathbf{Vect}_k$ is **real interval decomposable** if it can be expressed as a direct sum $\mathcal{F} = \bigoplus_i \mathcal{F}_i$, where each $\mathcal{F}_i : \mathcal{M}_F \rightarrow \mathbf{Vect}_k$ is indecomposable and has totally ordered support.

The next example shows that an arbitrary DMT does not have to be real interval decomposable.

Example 5 Consider the space and map shown in Fig. 6. The concrete decorated merge tree $\mathcal{F}_1 : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{vect}$ reduces to the study of the four vector spaces and the maps between them, indicated in red. It is a fact that this tree module is equivalent to one of the twelve indecomposable representations of the Dynkin Diagram D_4 . This gives a natural, Morse-theoretic example of a function whose associated tree module is not real-interval decomposable. This can be seen more intuitively by observing that the DMT fails to have the *untwisted* condition defined below: the image V_ℓ of the linear map from the left branch to the root (at the top) and the image V_r of the map from the right branch to the root do not satisfy $V_\ell \cap V_r = 0$.

The aim of this subsection is to establish a sufficient condition for a decorated merge tree to be real interval decomposable.

Definition 23 A collection of vector subspaces $V_1, \dots, V_m$ of a fixed vector space V is **independent** if for all i $V_i \cap (\sum_{j \neq i} V_j) = 0$. Under this hypothesis we have that

$$V_1 + \dots + V_m = V_1 \oplus \dots \oplus V_m.$$




Fig. 6 Consider a cylinder with height function induced by the implied embedding. The concrete decorated merge tree in degree-1 is an indecomposable tree module that is two dimensional at the merge point and for values up to, but not including, the maximum. This is equivalent to an indecomposable tree module, which is equivalent to one of the 12 indecomposable representations of the Dynkin Diagram D_4 . This shows that not every tree module is real interval decomposable

Let $(\mathcal{M}_F, \preccurlyeq)$ be a generalized merge tree and let $\mathcal{F} : (\mathcal{M}_F, \preccurlyeq) \rightarrow \mathbf{Vect}$ be a tree module. We say that $\mathcal{F}$ is **untwisted** if whenever we have a set of incomparable elements $p_1, p_2, \dots, p_n \in \mathcal{M}_F$ and any upper bound $p \in \mathcal{M}_F$ of these, the collection of vector subspaces

$$\mathrm{Im} \mathcal{F}(p_1 \preccurlyeq p), \mathrm{Im} \mathcal{F}(p_2 \preccurlyeq p) \cdots, \mathrm{Im} \mathcal{F}(p_n \preccurlyeq p)$$

is independent.

Remark 10 It should be noted that this notion of untwisted is completely unrelated to the notion used in Kim and Memoli (2018), which is used to describe Reeb graphs which admit sections when restricted to certain intervals.

To state our sufficient condition for real interval decomposability, we impose a mild geometric constraint on our generalized merge trees (Definition 2).

Definition 24 (cf. Defn. 2.2 of Patel 2018) A persistent set $S : (\mathbb{R}, \leq) \rightarrow \mathbf{set}$ is **constructible** if there exists a collection of times $\tau = \{t_0 < t_1 < \dots < t_n\}$ such that

- $S(t) = \emptyset$ for $t < t_0$,
- $S(t) = \{\star\}$ for $t > t_n$, and
- $S(t \leq s)$ is a bijection for every pair $t \leq s \subset [t_i, t_{i+1})$.

Assuming these conditions, the associated display poset $(S, \preccurlyeq)$ is said to be **tame**.

It is known that a constructible persistent set has an associated display poset that can be topologized as a finite cell complex, making it into a bona fide tree (Curry 2018; Stefanou 2018). There are a few other consequences of this tameness condition that are worth noting.

Proposition 6 Suppose $(\mathcal{M}_F, \preccurlyeq)$ is a tame generalized merge tree, i.e. the defining persistent set $\pi_0 \circ F$ is constructible. The following consequences of this assumption are obvious:

- Any maximal chain $C \subseteq \mathcal{M}_F$ has a minimum, which we call a **leaf node**. The set of leaf nodes of $\mathcal{M}_F$ is finite.
- The set of leaves has a (unique) least upper bound $p_\infty := (\star, t_n)$ (using the notation of Definition 24), which we call the **root node**.
- Any PFD tree module $\mathcal{F} : (\mathcal{M}_F, \preccurlyeq) \rightarrow \mathbf{vect}$ pushes forward along the natural projection map $\pi : \mathcal{M}_F \rightarrow \mathbb{R}$ to a PFD $\mathbb{R}$ -module $\pi_* \mathcal{F} : (\mathbb{R}, \leq) \rightarrow \mathbf{vect}$, which is defined pointwise as

$$\pi_* \mathcal{F}(s) := \bigoplus_{p \in \pi^{-1}(s)} \mathcal{F}(p).$$

Remark 11 When the need arises, we use notation $\pi_F : \mathcal{M}_F \rightarrow \mathbb{R}$ for the projection map. The extra notation will be necessary when comparing merge trees $\mathcal{M}_F$ and $\mathcal{M}_G$ later in the paper.

Remark 12 In Appendix B, we give an alternative construction of a merge tree which is guaranteed to produce a tree, in the metric sense, without any tameness assumption.

Theorem 4 Let $(\mathcal{M}_F, \preccurlyeq)$ be a tame generalized merge tree. A pointwise finite dimensional tree module $\mathcal{F} : (\mathcal{M}_F, \preccurlyeq) \rightarrow \mathbf{vect}$ is real interval decomposable if and only if it is untwisted.

Proof We proceed by induction on the number of leaf nodes n of a generalized merge tree $(\mathcal{M}_F, \preccurlyeq)$. If $n = 1$, then $\mathcal{M}_F$ is totally ordered and any tree module is automatically untwisted. The result then follows from the usual decomposition theorem for $\mathbb{R}$ -indexed persistence modules. If the result is true for any PFD tree module on a generalized merge tree with ℓ leaves where $1 \leq \ell \leq m$, then we now show the result is true for tree modules on trees with $m + 1$ leaf nodes.

To begin, we assume that $(\mathcal{M}_F, \preccurlyeq)$ has $m + 1$ leaf nodes and a root node p_∞ . Write $\mathcal{M}_F$ as the union of two subsets $\mathcal{M}_1$ and $\mathcal{M}_2$ where each subset has at least one leaf node and $\mathcal{M}_1 \cap \mathcal{M}_2 = \mathcal{U}_{p_\infty}$ —the intersection is the principal up set at the root node, which is totally ordered. Associated to any PFD tree module $\mathcal{F} : (\mathcal{M}_F, \preccurlyeq) \rightarrow \mathbf{vect}$ and such a decomposition $\mathcal{M}_F = \mathcal{M}_1 \cup \mathcal{M}_2$ are two natural submodules $\mathcal{F}_1$ and $\mathcal{F}_2$. We define $\mathcal{F}_1$ via three cases:

1. If $p \in \mathcal{M}_1$ and $p < p_\infty$, then $\mathcal{F}_1(p) := \mathcal{F}(p)$.
2. If $p \geq p_\infty$, then $\mathcal{F}_1(p) := \sum_{q \in \mathcal{M}_1 | q < p_\infty} \text{Im } \mathcal{F}(q < p)$.
3. If $p \in \mathcal{M}_2 \setminus \mathcal{M}_1$, then $\mathcal{F}_1(p) = 0$.

$\mathcal{F}_2$ is defined exactly in the same manner, simply by interchanging 1 and 2 in the above definition. By the induction hypothesis both $\mathcal{F}_1$ and $\mathcal{F}_2$ are untwisted and hence real interval decomposable, i.e. $\mathcal{F}_1 \cong \bigoplus_{\alpha \in I_1} \mathcal{F}_\alpha$ and $\mathcal{F}_2 \cong \bigoplus_{\beta \in I_2} \mathcal{F}_\beta$ and each of $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ are supported on totally ordered subsets of $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively.

By Proposition 6, $\pi_*\mathcal{F}$, $\pi_*\mathcal{F}_1$ and $\pi_*\mathcal{F}_2$ are each PFD $\mathbb{R}$ -modules. Moreover, since $\pi_* : \text{Fun}(\mathcal{M}_F, \text{vect}) \rightarrow \text{Fun}(\mathbb{R}, \text{vect})$ is an additive functor, we know that

$$\pi_*\mathcal{F}_1 \cong \pi_* \oplus_{\alpha} \mathcal{F}_{\alpha} \cong \oplus \pi_*\mathcal{F}_{\alpha},$$

which agrees, up to permuting factors with the Remak decomposition of $\pi_*\mathcal{F}_1$ guaranteed by the usual decomposition theorem in persistence. The symmetric statement is true for $\pi_*\mathcal{F}_2$ as well. This implies that we can collect the interval modules that appear in the Remak decomposition of $\pi_*\mathcal{F}$ into three terms:

$$\pi_*\mathcal{F} \cong F_1 \oplus F_2 \oplus F_3 \quad \text{where} \quad F_1 \cong \pi_*\mathcal{F}_1 \quad \text{and} \quad F_2 \cong \pi_*\mathcal{F}_2.$$

We claim that F_3 is the direct sum of real interval modules that are born at or after $\pi(p_{\infty})$. In particular, the support of F_3 is contained in $[\pi(p_{\infty}), \infty)$. This is clear because for any $t < \pi(p_{\infty})$ the set of pre-images $\pi^{-1}(t)$ can be partitioned into points that are in $\mathcal{M}_1$ or $\mathcal{M}_2$ exclusively, because $\pi^{-1}(t) \cap \mathcal{M}_1 \cap \mathcal{M}_2 = \emptyset$. This implies that

$$\text{for } t < \pi(p_{\infty}) \quad \pi_*\mathcal{F}(t) \cong \pi_*\mathcal{F}_1(t) \oplus \pi_*\mathcal{F}_2(t).$$

We now use this observation to define $\mathcal{F}_3 := \pi^*F_3$.

The proof of this direction concludes by showing that $\mathcal{F} \cong \mathcal{F}_1 \oplus \mathcal{F}_2 \oplus \mathcal{F}_3$. This decomposition obviously holds for the $\mathcal{M}_F$ -module when restricted to the set of points strictly below p_{∞} . Since $\mathcal{F}_i$ for $i = 1, 2, 3$ are naturally submodules of $\mathcal{F}$, it suffices to show that whenever $p \geq p_{\infty}$ and $v_i \in \mathcal{F}_i(p)$ are chosen such that $v_1 + v_2 + v_3 = 0$, then each $v_i = 0$. Of course if $v_1 + v_2 + v_3 = 0$, then $v_3 \in (\mathcal{F}_1(p) \oplus \mathcal{F}_2(p)) \cap \mathcal{F}_3(p) = 0$, so $v_3 = 0$. Moreover v_1 and v_2 are the images of w_1 and w_2 under the maps $\mathcal{F}_1(q_1 < p)$ and $\mathcal{F}_2(q_2 < p)$, so the untwisted condition implies that v_1 and v_2 are independent, hence 0 as well.

A real interval module satisfies the untwisted condition and this property is inherited by direct sums. This proves the converse statement and completes the proof of the theorem.

4.2 Injectivity of the barcode transform

The class of real interval decomposable tree modules is nice for various reasons, but one reason of interest to us is that the *barcode transform* is injective when restricted to real interval decomposable modules.

Theorem 5 *Suppose that $\mathcal{F} : (\mathcal{M}, \preceq) \rightarrow \text{vect}$ and $\mathcal{F}' : (\mathcal{M}', \preceq') \rightarrow \text{vect}$ are two real interval decomposable tree modules with finitely many intervals in their direct sum decompositions. If their associated barcode decorated merge trees $\mathcal{BF} : \mathcal{M} \rightarrow \text{Barcodes}$ and $\mathcal{BF}' : \mathcal{M}' \rightarrow \text{Barcodes}$ are isomorphic then $\mathcal{F} \cong \mathcal{F}'$.*

To prove the theorem, we introduce some notation. Let $\mathcal{J} \subset \mathcal{M}_F$ be a totally ordered subset. We use $\mathbb{k}_{\mathcal{J}} : (\mathcal{M}_F, \preceq) \rightarrow \text{vect}_{\mathbb{k}}$ to denote the functor that assigns the ground field $\mathbb{k}$ to points $p \in \mathcal{J}$ and where any ordered pair $p \preceq q$ in $\mathcal{J}$ is assigned the identity

map. If a decorated merge tree $\mathcal{F} : \mathcal{M}_F \rightarrow \mathbf{Vect}$ is real interval decomposable, then it can be expressed as

$$\mathcal{F} \approx \bigoplus_{\mathcal{R}(\mathcal{F})} \mathbb{k}_{\mathcal{J}}, \quad (1)$$

where $\mathcal{R}(\mathcal{F}) := \{(\mathcal{J}, m_{\mathcal{J}})\}$ indicates the collection of intervals $\mathcal{J}$ that appear in the Remak decomposition of $\mathcal{F}$ along with their multiplicity; the sum in (1) is taken with multiplicity.

The following proposition is obvious.

Proposition 7 *Suppose that the tree module $\mathcal{F} : (\mathcal{M}_F, \preceq) \rightarrow \mathbf{vect}$ is real interval decomposable. Then the associated barcode decorated merge tree $\mathcal{BF} : \mathcal{M}_F \rightarrow \mathbf{Barcodes}$ is determined at each point by the formula*

$$\mathcal{BF}(\mathfrak{p}) := \{(\mathcal{U}_{\mathfrak{p}} \cap \mathcal{J}, m_{\mathcal{J}}) \mid (\mathcal{J}, m_{\mathcal{J}}) \in \mathcal{R}(\mathcal{F})\}$$

where $\mathcal{U}_{\mathfrak{p}}$ is the principal up set at $\mathfrak{p}$ in $\mathcal{M}_F$. Moreover, when $\mathcal{M}_F$ is tame, the associated barcode decorated merge tree is completely described as the disjoint union of the barcodes when viewed from each leaf node $\mathfrak{v}$, i.e.

$$\mathcal{BF} = \bigsqcup_{\mathfrak{v} \in \mathbf{L}(\mathcal{M})} \mathcal{BF}(\mathfrak{v}),$$

where $\mathcal{BF}(\mathfrak{v})$ is determined via intersection with the principal up set $\mathcal{U}_{\mathfrak{v}}$ as indicated above. We note that every bar (I, m_I) in the barcode viewed from $\mathfrak{v}$ can be expressed as a sum

$$(I, m_I, \mathfrak{v}) = \sum_{\mathcal{J} \in \mathcal{R}(\mathcal{F})} \mathcal{U}_{\mathfrak{v}} \cap (\mathcal{J}, m_{\mathcal{J}}) \quad \text{where} \quad \mathcal{J} \cap \mathcal{U}_{\mathfrak{v}} = I.$$

Informally, we may write this as $m_I = m_{\mathcal{J}_1} + \cdots + m_{\mathcal{J}_k}$

Proof of Theorem 5 Two barcode decorated merge trees are isomorphic if there exists an underlying isomorphism of their generalized merge trees $\phi : \mathcal{M} \rightarrow \mathcal{M}'$ and $\psi : \mathcal{M}' \rightarrow \mathcal{M}$ so that $\mathcal{BF}(\mathfrak{x}) = \mathcal{BF}'(\phi(\mathfrak{x}))$ and $\mathcal{BF}'(\mathfrak{x}') = \mathcal{BF}(\psi(\mathfrak{x}'))$ for all $\mathfrak{x} \in \mathcal{M}$ and $\mathfrak{x}' \in \mathcal{M}'$. The argument proceeds by proving that $\mathcal{F} \cong \mathcal{F}' \circ \phi$ and symmetrically $\mathcal{F}' \cong \mathcal{F} \circ \psi$. Both arguments would be completely symmetric over the two generalized merge trees, so it suffices to consider the case where $\mathcal{M} = \mathcal{M}'$ and where we can assume that ϕ and ψ are identity maps. We now proceed to showing that if $\mathcal{BF} = \mathcal{BF}'$ then $\mathcal{F} \cong \mathcal{F}'$ as $\mathcal{M}$ -modules.

The argument proceeds by showing that every real interval $(\mathcal{J}, m_{\mathcal{J}})$ that appears in the Remak decomposition for $\mathcal{F}$ must appear in the Remak decomposition for $\mathcal{F}'$ with no intervals left over. This will establish the isomorphism of $\mathcal{F}$ and $\mathcal{F}'$. We start by considering an *oldest* interval $(I, m_I, \mathfrak{v})$ in $\mathcal{B}[\mathcal{F}]$, this means that if we consider the projection of I to $\mathbb{R}$ via the natural map $\pi : \mathcal{M} \rightarrow \mathbb{R}$ then $\inf\{\pi(I)\} \leq \inf\{\pi(\mathcal{J})\}$ for any

$(J, m_J, w) \in \mathcal{BF}$; such an oldest interval exists due to our finiteness assumption. By Proposition 7 we know that any (I, m_I, v) has a unique (up to permutation) expression

$$m_I = m_{J_1} + \cdots + m_{J_k}$$

for $J_j \in \mathcal{R}(\mathcal{F})$ with $J_j \cap U_v = I$. We claim that if (I, m_I, v) is an oldest interval then there exists an $(J_j, m_{J_j}) \in \mathcal{R}(\mathcal{F})$ with $(I, m_I, v) = (J_j, m_{J_j})$. Indeed, if there does not exist a unique contributor, then this contradicts the assumption that (I, m_I, v) was oldest. The interval J_j is the oldest contributing factor in the sum for m_I displayed above. Since the barcode decoration associated to $\mathcal{F}'$ is equal to the barcode decoration associated to $\mathcal{F}$, then (I, m_I, v) must also be an oldest interval and there must be the same oldest contributor $J_j \in \mathcal{R}(\mathcal{F}')$. This proves that $\mathcal{F}' \cong \mathcal{G}' \oplus \mathbb{k}_{J_j}^{m_{J_j}}$ and of course $\mathcal{F} \cong \mathcal{G} \oplus \mathbb{k}_{J_j}^{m_{J_j}}$ and $\mathcal{BG}' = \mathcal{BG}$. The argument repeats with the barcode decorated merge tree $\mathcal{B}[\mathcal{G}]$ until it is empty.

4.3 Lift decorated merge trees

The process of extracting a decorated merge tree from a dataset is an algorithmic challenge that we need to address. When a persistent space $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ is a filtration of a simplicial complex, existing software can be used to extract a merge tree $\mathcal{M}_F$ and a degree- k persistent homology barcode B . One would hope to use this easily accessible data to reverse engineer the underlying decorated merge tree $\mathcal{F} : \mathcal{M}_F \rightarrow \mathbf{vect}$, for which the persistent homology barcode is the pushforward; i.e., $B = \pi_* \mathcal{F}$, where $\pi : \mathcal{M}_F \rightarrow \mathbb{R}$ is the projection function. Indeed, algorithms presented in Sect. 6 are based on such a reverse engineering process and can be used to produce decorated merge trees from point cloud and filtered graph data. In this subsection, we provide the theoretical underpinning of these algorithms.

Let $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ be a filtration of a simplicial complex with a tame merge tree $\mathcal{M}_F$. The standard algorithm for extracting the barcode B records a **birth simplex** σ for each interval $I = [b, d]$ in the barcode (Edelsbrunner et al. 2000). This is to say that σ is a simplex with filtration value b whose birth creates a representative homology class for I . The **birth point** of I is the unique point $p \in \mathcal{M}_F$ with $\pi(p) = b$ and where p is an ancestor of all the vertices of σ . Let $q \in \mathcal{M}_F$ be the unique ancestor of p such that $\pi(q) = d$.

Definition 25 With the notation above, define the **lift** of I to be the decorated merge tree $\mathcal{J} : \mathcal{M}_F \rightarrow \mathbf{vect}$ defined by

$$\mathcal{J}(r) := \begin{cases} k & p \preceq r \preceq q \\ 0 & \text{otherwise.} \end{cases}$$

The **lift** of B is the decorated merge tree

$$\widehat{\mathcal{F}} := \bigoplus_{I \in B} \mathcal{J}.$$

A decorated merge tree is called a **lift decorated merge tree** if it is obtained as the lift of some barcode.

Based on the above discussion, it is straightforward to determine the lift decorated merge tree directly from the filtered simplicial complex. However, the lift DMT is not guaranteed to be isomorphic to the *true* DMT $\mathcal{F} : \mathcal{M} \rightarrow \mathbf{vect}$; indeed, we have $\widehat{\mathcal{F}} \approx \mathcal{F}$ if and only if $\mathcal{F}$ is real interval decomposable, in the sense of Definition 22. The task which we aim to address for the rest of this subsection is to determine a condition on the filtration F which

- can be verified directly from the merge tree $\mathcal{M}_F$ and degree- k barcode B , and
- implies that $\widehat{\mathcal{F}} \approx \mathcal{F}$.

Our proposed condition is the following.

Definition 26 Let $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ be a filtered simplicial complex with merge tree $\mathcal{M}_F$ and degree- k barcode B . We say that F has the H_k -**disjointness property** if for any pair of bars $I = [b, d)$ and $I' = [b', d')$ in B with incomparable birth points p and p' , one of their death times is less than their merge height (Definition 15), i.e.

$$\min\{d, d'\} < \text{merge}(p, p').$$

Intuitively, this means that the lifts of I and I' have disjoint support.

Proposition 8 *If a filtered simplicial complex $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ has the H_k -disjointness property, then the lift decorated merge tree and the concrete decorated merge tree are isomorphic, i.e. $\widehat{\mathcal{F}} \approx \mathcal{F}$.*

This result says that H_k -disjointness is sufficient to guarantee that the lifting procedure produces a valid decorated merge tree. Our proof is straightforward but technical, so we delay it to Appendix C.2.

5 Computing interleaving distances

The goal of this section is to develop a practically feasible algorithm for approximating the matching distance between barcode decorated merge trees (Definition 10). This relies on the notion of *labeling* a DMT.

5.1 Labeled distance

Throughout this section, we consider generalized merge trees $\mathcal{M}_F$ arising from persistent spaces $F : (\mathbb{R}, \leq) \rightarrow \mathbf{Top}$ such that the persistent set $\pi_0 \circ F$ is constructible and $\mathcal{M}_F$ is tame, as in Proposition 6. The finite set of leaves of $\mathcal{M}_F$ is denoted $L(\mathcal{M}_F)$. The canonical projection function $\mathcal{M}_F \rightarrow \mathbb{R}$ is denoted π_F .

Remark 13 The results of this section do not require that decorated merge trees arise from persistent spaces at all; they could be derived formally treating tree modules as the fundamental objects.

We restrict our attention to the case of barcode decorated merge trees $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ which are determined by restriction and which have leaves, i.e. $\mathcal{B}_F$ is a leaf-decorated merge tree. Recall that we denote an element of Barcodes as a multiset $\{(I, m_I)\}_{I \in \mathcal{J}}$, where I is an interval and m_I is its multiplicity. We make the (realistic, in practice) assumption that all barcodes are finite.

Inspired by a similar result for merge trees in Gasparovic et al. (2019), we will now explain how the bottleneck distance between leaf decorated merge trees can be expressed in terms of matrices obtained from labelings of merge trees.

Definition 27 A **labeling** of a merge tree $\mathcal{M}_F$ is a map

$$\lambda_F : [n] := \{1, \dots, n\} \rightarrow \mathcal{M}_F,$$

for some integer $n \geq 1$, which is surjective onto the set of leaves $L(\mathcal{M}_F)$; that is, $L(\mathcal{M}_F) \subseteq \text{Im}(\lambda_F)$.

To each labeling $\lambda_F : [n] \rightarrow \mathcal{M}_F$ one associates an $n \times n$ matrix Λ_F of merge heights (Definition 15), referred to as the *least common ancestor (LCA) matrix* and defined by

$$\Lambda_F(i, j) = \pi_F(\text{LCA}_F(\lambda_F(i), \lambda_F(j))) = \text{merge}_F(\lambda_F(i), \lambda_F(j)). \quad (2)$$

It is shown in (Gasparovic et al. 2019, Proposition 4.1) that interleaving distance between merge trees can be expressed in terms of the ℓ_∞ distance between these LCA matrices. To adapt this result to the decorated merge tree setting, we introduce a more involved objective function.

Definition 28 Let $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ and $\mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$ be leaf decorated merge trees. Define the **matching cost** of labelings $\lambda_F : [n] \rightarrow \mathcal{M}_F$ and $\lambda_G : [n] \rightarrow \mathcal{M}_G$ to be

$$\text{cost}(\lambda_F, \lambda_G) = \max \left\{ \|\Lambda_F - \Lambda_G\|_\infty, \max_i d_B(\mathcal{B}_F(\lambda_F(i)), \mathcal{B}_G(\lambda_G(i))) \right\}, \quad (3)$$

where $\|\cdot\|_\infty$ is the ℓ_∞ -norm on $n \times n$ matrices.

Proposition 9 The decorated bottleneck distance between leaf decorated merge trees $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ and $\mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$ is given by

$$d_B(\mathcal{B}_F, \mathcal{B}_G) = \min_{\lambda_F, \lambda_G} \text{cost}(\lambda_F, \lambda_G),$$

where the minimum is taken over labelings λ_F and λ_G with common domain $[n]$, where n is the sum of the number of leaves in $\mathcal{M}_F$ and the number of leaves in $\mathcal{M}_G$.

The proof adapts the proof of (Gasparovic et al. 2019, Proposition 4.1) to this new setting. We provide a sketch of the idea and omit technical, but ultimately routine, details.

Proof Sketch Throughout the proof sketch, fix leaf decorated merge trees $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ and $\mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$. Suppose that $\mathcal{M}_F$ has k leaves and $\mathcal{M}_G$ has ℓ leaves.

First assume that we have ϵ -compatible maps (ϕ, ψ) of $\mathcal{M}_F$ and $\mathcal{M}_G$; recall from Definition 19 that we slightly abuse notation and conflate the notion of natural transformation with that of ϵ -compatible maps. Also assume that

$$d_B(\mathcal{B}_F(u), \mathcal{B}_G(\phi(u)), d_B(\mathcal{B}_G(v), \mathcal{B}_F(\psi(v))) < \epsilon.$$

We construct labelings λ_F and λ_G on the common set of labels $[k + \ell]$ as follows.

1. Label each leaf of $\mathcal{M}_F$ with a unique element of $\{1, \dots, k\}$; this defines λ_F on $[k]$. Each label i is also assigned to $\mathcal{M}_G$ as $\lambda_G(i) = \psi(\lambda_F(i))$.
2. Similarly, for each $j \in \{k + 1, \dots, k + \ell\}$ define $\lambda_G(j)$ to be a unique leaf of $\mathcal{M}_G$ and define $\lambda_F(j) = \phi(\lambda_G(j))$.

One must then show that the labelings λ_F and λ_G defined above have the property that, for any $i, j \in [k + \ell]$,

$$|\pi_F(\text{LCA}_F(\lambda_F(i), \lambda_F(j))) - \pi_G(\text{LCA}_G(\lambda_G(i), \lambda_G(j)))| < \epsilon \quad (4)$$

and

$$d_B(\mathcal{B}_F(\lambda_F(i)), \mathcal{B}_G(\lambda_G(i))) < \epsilon. \quad (5)$$

Conversely, suppose that we have labelings λ_F and λ_G such that $\text{cost}(\lambda_F, \lambda_G) < \epsilon$. For each $i \in [k + \ell]$ set $u_i = \lambda_F(i)$ and $v_i = \lambda_G(i)$.

1. We first define the map $\phi : \mathcal{M}_F \rightarrow \mathcal{M}_G$. Let $x \in \mathcal{M}_F$ and choose an arbitrary labeled point $u_i \in \mathcal{M}_F$ with $u_i \leq x$ (always possible because all leaves of $\mathcal{M}_F$ are labeled). Define $\phi(x) \in \mathcal{M}_G$ to be the unique ancestor of v_i such that $\pi_G(\phi(x)) = \pi_F(x) + \epsilon$.
2. The map $\psi : \mathcal{M}_G \rightarrow \mathcal{M}_F$ is constructed similarly.

We first note that the maps ϕ and ψ are well defined. Indeed, if u_i and u_j are a pair of labeled points in $\mathcal{M}_F$ with $u_i, u_j \leq x$, then $\text{LCA}_F(u_i, u_j) \leq x$. This means that $\pi_G(\text{LCA}_G(v_i, v_j)) \leq \pi_F(x) + \epsilon$, by our assumption on the cost of the labelings λ_F and λ_G , so that v_i and v_j have the same ancestor at height $\pi_F(x) + \epsilon$. By the same reasoning, ψ is well defined. The proof is completed by showing that the maps (ϕ, ψ) define an ϵ -matching of barcode decorated merge trees.

5.2 Approximating merge tree interleaving via Gromov-Wasserstein distance

Computing interleaving distance between merge trees is known to be NP-Hard (Agarwal et al. 2018). A dynamic programming approach to computing interleaving distance is taken in Farahbakhsh Touli and Wang (2019), illustrating that the computation is at least fixed parameter tractable (roughly, precise upper bounds on interleaving distance are computable in time with polynomial growth in the number of the nodes, provided

one has a uniform upper bound on node degrees). An alternative and more efficient method for analyzing merge trees is to apply (Gasparovic et al. 2019, Proposition 4.1): given any labeling of merge trees, one obtains an upper bound on their interleaving distance, so one could rely on a good heuristic for estimating optimal labelings. Such a heuristic, based on replacing the problem with the simpler one of finding optimal bipartite graph matchings, is proposed in Yan et al. (2019), yielding a comprehensive framework for statistical analysis of merge tree ensembles.

In this section, we propose an alternative heuristic for estimating optimal labelings using the concept of Gromov-Wasserstein (GW) distances. Different formulations of GW distances were introduced around a decade ago by Sturm (2006, 2012), to study abstract convergence of sequences of metric measure spaces, and Mémoli (2007, 2011), with a view towards applications to object matching. GW distances have found recent popularity in the machine learning community (Alvarez-Melis et al. 2019; Chapel et al. 2020; Xu et al. 2019, 2020), largely due to the observation by Peyré et al. (2016) that they can be used to measure distances between general kernel matrices; this point of view was formalized mathematically in Chowdhury and Mémoli (2018). We now briefly present GW distances from a computational point of view—see the references above more thorough treatments.

Definition 29 (cf. Chowdhury and Mémoli 2018) A **(finite) measure network** is a pair $C = (C, \mu)$, where C is an $n \times n$ matrix and μ is a length- n probability vector (that is, $\sum \mu(i) = 1$ and $\mu(i) \geq 0$). A **coupling** between probability vectors $\mu_1 \in \mathbb{R}^n$ and $\mu_2 \in \mathbb{R}^m$ is a matrix $\nu \in \mathbb{R}^{n \times m}$ such that $\nu(i, j) \geq 0$ and

$$\sum_i \nu(i, j) = \mu_2(j), \quad \sum_j \nu(i, j) = \mu_1(i).$$

That is, ν is a joint probability distribution with marginals μ_1 and μ_2 . The set of couplings between fixed μ_1 and μ_2 is denoted $\mathcal{C}(\mu_1, \mu_2) \subset \mathbb{R}^{n \times m}$.

The **Gromov-Wasserstein p-distance** between finite measure networks (C_1, μ_1) and (C_2, μ_2) is defined by

$$d_{\text{GW},p}((C_1, \mu_1), (C_2, \mu_2)) = \min_{\nu \in \mathcal{C}(\mu_1, \mu_2)} J_p(\nu)^{1/p}, \quad (6)$$

where

$$J_p(\nu) = \sum_{i,j,k,\ell} (C_1(i, k) - C_2(j, \ell))^p \nu(i, j) \nu(k, \ell) \quad (7)$$

denotes the **Gromov-Wasserstein p-distortion functional**.

To handle computational aspects of merge tree interleaving, we introduce the following discrete representation of a merge tree.

Definition 30 (cf. Gasparovic et al. 2019) A **computational merge tree** is a (discrete) graph together with a height function on its nodes satisfying the following conditions:

1. the graph must be a tree—that is, there is a unique edge path between any two nodes;
2. exactly one of the leaf nodes (called the **root**) of the graph gets assigned height ∞ ;
3. adjacent nodes do not have equal function value; and
4. every non-root node has exactly one neighbor with higher function value.

A **computational decorated merge tree** is a computational merge tree together with an assignment of a barcode to each of its nodes.

Remark 14 In what follows, we use the notation $\mathcal{M}_F$ for a computational merge tree and $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ for a computational decorated merge tree. This is in agreement with the notation used for their non-computational counterparts earlier in the paper, but the distinction should be clear from context. We also generally drop the “computational” qualifier.

The GW framework is incorporated into a pipeline for estimating the interleaving distance between merge trees $\mathcal{M}_F$ and $\mathcal{M}_G$ as described by the following algorithm; this is augmented below in Sect. 5.3 to define an algorithm for estimating interleaving distance between decorated merge trees. The main idea is that the algorithm produces labelings of the merge trees, and the ℓ_∞ distance between the associated least common ancestor matrices upper bounds interleaving distance—of course, any pair of labelings produces such an upper bound, but the labelings constructed below are informed by a GW structural alignment. The algorithm is also summarized in Fig. 7.

Algorithm 1 *To estimate the interleaving distance between them merge trees $\mathcal{M}_F$ and $\mathcal{M}_G$:*

1. Probability distributions μ_F and μ_G , respectively, are chosen for the nodes of the trees. We use uniform distributions in our experiments
2. A cost matrix C_F (respectively, C_G) is associated to $\mathcal{M}_F$ (resp., $\mathcal{M}_G$) by choosing an arbitrary labeling of its nodes and taking the associated least common ancestor matrix. Then (C_F, μ_F) and (C_G, μ_G) define finite measure networks.




Fig. 7 Pipeline for estimating interleaving distance between merge trees. From left to right: we aim to compute the interleaving distance between merge trees shown in the left column. We sample the merge trees with a user-defined mesh (here the mesh is equal to 1). We assign arbitrary labels to the nodes in each tree independently; LCA matrices for each merge tree (with respect to these labels) are shown in the third column. The fourth column shows an optimal coupling between the sampled nodes of the merge tree. From the coupling, we estimate an optimal labeling of the merge trees (fifth column), and this easily yields an upper estimate on the interleaving distance between the merge trees

3. An optimal coupling γ of μ_F and μ_G is estimated by numerically solving the GW problem (6). Intuitively, the GW problem promotes large values of $\gamma(i, j)$ when node i of $\mathcal{M}_F$ and node j of $\mathcal{M}_G$ are structurally similar.
4. We use γ to estimate interleaving maps of $\mathcal{M}_F$ and $\mathcal{M}_G$: for each leaf u in $\mathcal{M}_F$, we locate the maximum entry of the row of γ corresponding to this leaf and define this to be $\phi(u) \in \mathcal{M}_G$. We likewise define a map ψ from the leaf set of $\mathcal{M}_G$ to $\mathcal{M}_F$ by examining columns of γ .
5. We use ϕ and ψ to construct labelings λ_F and λ_G with domain $[k + \ell]$ of the merge trees as in the algorithm in the proof of Proposition 9, where k is the number of leaves in $\mathcal{M}_F$ and ℓ is the number of leaves in $\mathcal{M}_G$.
6. From λ_F and λ_G , we construct matrices Λ_F and Λ_G and compute $\|\Lambda_F - \Lambda_G\|_\infty$, yielding a principled upper estimate of interleaving distance.

Remark 15 Computing GW distance is an instance of quadratic programming problem with a nonconvex objective function and is therefore NP-Hard to compute exactly (Mémoli 2011). However, since the space of couplings $\mathcal{C}(\mu_1, \mu_2)$ forms a convex polytope, it is possible to approximate GW distance via Frank-Wolfe-style projected gradient descent; our computations of GW distance will be handled by the Python Optimal Transport package (Flamary and Courty 2017). Gradient updates have computational complexity $O(n^3 \log(n))$ when $p = 2$ (Peyré et al. 2016), so we focus on this case for the sake of efficiency.

The quality of the gradient descent-based estimation in step (2) of the algorithm can be improved by upsampling the trees—i.e., adding degree-2 nodes at user-specified heights to better approximate the continuous nature of the true merge trees. We have found empirically that this procedure for estimating interleaving distance from above tends to give meaningful labelings of merge trees. This is illustrated by an experiment on synthetic data in Sect. 6.1.

Remark 16 Similar applications of Gromov-Wasserstein distance have recently appeared in the literature. In Li et al. (2021), a framework for summarizing sets of merge trees is developed using the Riemannian structure of the GW metric on measure networks developed in Chowdhury and Needham (2020). There, the authors use GW distance to align merge trees, but the measure network representation of the merge trees used there is different than the one proposed here and has a less clear connection to Gasparovic et al. (2019) and the computation of interleaving distance. The authors of Mémoli et al. (2021) study a variant of GW distance on the space of *ultrametric spaces* (metric spaces which satisfy a stronger version of the triangle inequality). Applications to merge tree interleaving are not directly discussed in Mémoli et al. (2021), but our matrix representations of merge trees are inherently utilizing an underlying ultrametric structure (or, using the terminology of Mémoli et al. (2021), ultradissimilarity structure), which is not the usual geodesic metric one considers when representing a merge tree as a metric tree. Similar ideas for comparing metric spaces go back to Smith et al. (2016) and Mémoli et al. (2019) (the latter describing specific connections to merge tree interleavings), but to our knowledge the specific algorithm and implementation described here for approximating merge tree interleaving is novel.

5.3 Matchings of DMTs via fused Gromov-Wasserstein matchings

The GW framework described above has thus far only been applied to estimate interleavings between (undecorated) merge trees. In order to incorporate barcodes into comparisons between *decorated* merge trees, we employ the more general Fused Gromov-Wasserstein (FGW) framework (Vayer et al. 2020). The FGW framework is used to compare measure networks with additional structure.

Definition 31 Let $Z = (Z, d_Z)$ be a metric space. A **Z-structured (finite) measure network** is a triple (C, μ, B) consisting of a finite measure network (C, μ) , where $C \in \mathbb{R}^{n \times n}$, together with an n -tuple B of points in Z .

Let (C_1, μ_1, B_1) and (C_2, μ_2, B_2) be Z -structured measure networks. For $\alpha \in [0, 1]$, we define the **Fused Gromov-Wasserstein (FGW) distance** for parameter $\zeta \in [0, 1]$ as

$$d_{\text{FGW}, \zeta}((C_1, \mu_1, B_1), (C_2, \mu_2, B_2))^2 = \min_{\gamma \in \mathcal{C}(\mu, \nu)} ((1 - \zeta)I_2(\gamma) + \zeta J_2(\gamma)), \quad (8)$$

where J_2 is the GW p -loss (7) with $p = 2$ and

$$I_2(\gamma) = \sum_{i,j} d_Z(B_1(i), B_2(j))^2 \gamma(i, j)$$

is the standard 2-Wasserstein loss from classical optimal transport (see, e.g., Villani 2008).

The FGW framework is used to augment Algorithm 1 to estimate matching distance between DMTs $\mathcal{B}_F : \mathcal{M}_F \rightarrow \text{Barcodes}$ and $\mathcal{B}_G : \mathcal{M}_G \rightarrow \text{Barcodes}$. The difference is that the optimal coupling of step (2) is obtained by solving (8) for some choice of hyperparameter ζ , where the functional J_2 is the GW 2-distortion functional with respect to the measure networks $(\bar{\Lambda}_F, \mu_F)$ and $(\bar{\Lambda}_G, \mu_G)$, exactly as in the previous subsection. The I_2 term is Wasserstein 2-loss with respect to bottleneck distance between node barcodes. Examples of FGW-based estimation of DMT matching distance are provided in Sects. 6.4 and 6.6.

6 Algorithmic details and examples

In this section, we outline theoretical and practical aspects of generating and comparing decorated merge trees. We provide several computational examples coming from real and synthetic data. Implementations (in Python) of all of these experiments as well as source code are freely available in our GitHub repository.¹ The code uses standard Python data science packages (scikit-learn Pedregosa et al. 2011, scipy Virtanen et al. 2020, etc.), as well as more specialized packages for topological data

¹ <https://github.com/trneedham/Decorated-Merge-Trees>.




Fig. 8 Merge tree classification experiment. The merge trees shown on the left are samples from 6 different classes generated from the family of scalar functions (9). Pairwise distance matrices and leave-one-out nearest neighbor classification scores for ℓ_∞ and ℓ_2 interleaving distances are shown on the right

analysis (gudhi GUDHI 2021, giotto-tda Tauzin et al. 2021, ripser Bauer 2019 and scikit-tda Saul and Tralie 2019) and optimal transport (Python Optimal Transport, pot Flamary and Courty 2017).

6.1 Merge tree interleaving distances

To test the reliability of our method for estimating merge tree interleaving distances, we create a simple classification experiment (summarized in Fig. 8); note that this example is only meant to demonstrate the viability of our method for computing merge tree interleaving distance, and does not yet incorporate any merge tree decorations. We define a parametric model for generating merge trees by considering merge trees $\mathcal{M}_F$ given by sublevel-set filtrations F of functions $f : [0, 1] \rightarrow \mathbb{R}$ given by

$$f(t) = \sin(\rho_1 \pi t) + \cos(\rho_2 \pi t) + g(t), \quad (9)$$

where ρ_j are parameters and $g(t) \in [0, 1/2]$ is uniformly distributed random additive noise. We construct a dataset of 120 random merge trees by drawing 10 samples for each combination of parameter choices $\rho_1 \in \{1, 2, 4, 8\}$ and $\rho_2 \in \{1, 3, 5\}$. The random additive noise has the effect of generating many spurious local minima for the functions which yield extraneous leaves in the merge trees, making the matching process nontrivial. We compute a 120×120 pairwise distance matrix using our interleaving distance estimator with uniform node weights and height sampling mesh size equal to 0.5. The leave-one-out nearest neighbor classification rate for this distance matrix is 98.33%, indicating that our interleaving distance computation provides a meaningful comparison for the merge tree dataset.

Since the labelings in the GW-based interleaving computation are obtained via optimization of an ℓ_2 loss, we also give a small tweak to the interleaving distance formulation and compute the ℓ_2 -interleaving distance between merge trees $\mathcal{M}_F$ and $\mathcal{M}_G$, defined to be

$$\min_{\lambda_F, \lambda_G} \|\Lambda_F - \Lambda_G\|_2,$$

where the minimum is, as usual, over labelings of the merge trees and $\|\cdot\|_2$ is the standard ℓ_2 norm on matrices. The classification rate for this distance is 100%. Exploring theoretical properties of ℓ_p versions of merge tree interleaving distance will be a direction of future work.

6.2 Decorated merge trees for point clouds

Vietoris-Rips complexes of point clouds in metric spaces provide a ubiquitous source for persistence diagrams in topological data analysis. In this subsection we provide details on the construction of decorated merge trees and on the methods used to visualize them. These constructions fit into the general theory by considering the DMTs as being generated by a filtered simplicial complex, viewed as a persistent space as in Example 3. Throughout the rest of this section, we will use the term DMT to refer specifically to tame leaf decorated merge trees (Definitions 10 and 24).

Let X be a finite subset of a metric space (in our examples, the metric space is $\mathbb{R}^n$, but the procedure we describe here can be performed on any distance matrix). From the Vietoris-Rips complex of X , we are easily able to compute a (discrete representation of a) merge tree $\mathcal{M}_F$ (F denoting the Vietoris-Rips filtration for the complete simplex on X)—i.e., a single linkage hierarchical clustering dendrogram of X (Carlsson and Mémoli 2013)—via off-the-shelf functions available in standard packages (e.g., `scipy`).

It remains to determine barcodes for the leaves of $\mathcal{M}_F$. A straightforward algorithm produces a lift decorated merge tree, as was considered in Sect. 4.3. The inputs are the tree $\mathcal{M}_F$ and a degree- k Vietoris-Rips barcode, generated by existing software (`ripser` or `gudhi`). For convenience, we assume the generic condition that all pairwise distances in X are distinct, but this condition can be removed with a bit more work. The basic idea of the algorithm is that we can *decorate* the merge tree with the bars of the barcode by locating the unique (under our generic assumption on the distances) point in the tree where each bar in the barcode is born. Recall from Sect. 4.3 that the lift decorated merge tree is not necessarily isomorphic to the true decorated merge tree—this will be the case if and only if the true DMT is real interval decomposable (Definition 22). Fortunately, Proposition 8 gives a simple method for certifying correctness of the lift decorated merge tree. If a given merge tree does not meet this certification, the true leaf decorations of the merge tree for the dataset can be computed via a more laborious algorithm, where the idea is to construct a modification of the Vietoris-Rips complex for each leaf and to then assign the associated persistent homology barcode to that leaf. This is guaranteed to produce the correct leaf barcodes, but comes with higher computational cost and less intuitive visualizations.

To visualize the lift decorated merge trees, we draw the merge tree together with offset edges indicating homology bars. Examples of decorated merge trees generated via this algorithm are shown in Fig. 1. Observe that the point clouds in this example have very similar degree-0 and degree-1 persistent homology barcodes, but that they are easily distinguished by the interactions between their connectivity and degree-1 data; this difference is clearly captured in their more refined DMT topological signatures.




Fig. 9 Summarizing the topology in all degrees. The figure on the left shows a toy point cloud dataset consisting of points sampled from a torus, a sphere and a Fig. 8. The figure on the right shows the associated merge tree decorated with both degree-1 and degree-2 persistent homology data. This summarizes the topology of the dataset in all degrees simultaneously, essentially giving a complete description of its overall topology. As described in the text, the merge tree and persistence diagrams have been thresholded for visual clarity: the merge tree was truncated at height 0.25 and the plotted bars have persistence at least 1.4 in degree-0 and 0.6 in degree-1

We can also visualize merge trees decorated with bars from several higher homology degrees—see Fig. 9. Our visualizations employ some tricks to improve legibility:

1. on the DMT, we typically only plot bars from the barcode which have persistence longer than a (tunable, user-defined) threshold;
2. we truncate the merge tree itself at a relatively low (tunable, user-defined) height and then extend the truncated edges to zero.

Both of these options are used to simplify the visualizations by removing unstable, noisy topological features (transient higher-dimensional features and changes in connectivity at small radii, respectively). Moreover, the simplified DMTs are frequently used in matching distance computations—since this is an ℓ_∞ -type distance, it is stable under these small adjustments, but the simplifications provide a significant computational speedup.

Figure 10 shows an example of a pointcloud where the hypotheses of Proposition 8 (i.e. H_1 -disjointness) fail. In fact, the true DMT for this dataset is not real interval decomposable and the true leaf barcodes are determined by the more computationally taxing algorithm. In this case, a different technique for visualizing decorated merge trees must be used—see the caption to the figure for details. We found in practice that almost all of our other examples naturally satisfied the hypotheses of Proposition 8.

6.3 Sliding window embeddings

Established in Perea and Harer (2015), sliding window embeddings provide a method for applying the techniques of point cloud TDA to 1-dimensional signals. Given a function $f : [0, T] \rightarrow \mathbb{R}$, one defines for each pair of parameters $d \in \mathbb{N}$ and




Fig. 10 A point cloud with non-real interval decomposable merge tree. The figure on the left shows a dataset in $\mathbb{R}^2$. Its classical degree-0 and degree-1 persistence diagrams are shown second from the left. The DMT shown in the second from the right is the one produced by the lifting algorithm—observe that this does not satisfy the hypotheses of Proposition 8. In fact, the output of the more involved algorithm disagrees with the lifted version, meaning that the true DMT for the point cloud is not real interval decomposable. The true DMT is depicted on the right. In the non-real interval decomposable case, we must visualize the leaf barcode for each leaf independently. All leaves colored blue (respectively, red) have the same barcode, also depicted in blue (resp., red). For visual clarity, the merge tree is truncated at height 0.01 in the display and only bars with persistence at least 0.2 are plotted

$\tau \in (0, T/d)$ the **sliding window embedding** (also referred to as the **Takens embedding**, in reference to Takens’s theorem from dynamical system theory (Takens 1981)) as

$$\begin{aligned} \text{SW}_{d,\tau}(f) : [0, T - d\tau] &\rightarrow \mathbb{R}^{d+1} \\ t &\mapsto (f(t), f(t + \tau), f(t + 2\tau), \dots, f(t + d\tau)). \end{aligned}$$

The sliding window embedding assigns a $(d + 1)$ -dimensional Euclidean point cloud to each collection of finite samples of f . Methods of TDA can then be applied to this point cloud; this approach to signal analysis has found success in applications such as wheeze detection in audio recordings of breathing (Emrani et al. 2014) and action recognition from scalar measurements of joint movement (Venkataraman et al. 2016).

When performing this sort of analysis, the focus is typically on degree-1 homological features, which indicate periodicity in the signal. We hypothesize that there may also be interesting degree-0 features in the case that the signal includes a sudden shift. We illustrate this behavior in the examples shown in Fig. 11. Given a signal with an apparent shift, we:

1. construct a point cloud via a sliding window embedding; parameters d and τ are chosen automatically via statistical tools in the `giotto-tda` package;
2. subsample the resulting point cloud by density—this has the effect of accentuating disconnected clusters, which correspond to the pre- and post-shift regimes in the original signal;
3. create a decorated merge tree from the result.

The examples in Fig. 11 (on both synthetic and real time series data) show that the DMTs resulting from this process uncover interesting features of the signals.

From a real interval decomposable DMT, we obtain a simple method for locating the connected component of a dataset which contains a given higher degree homology cycle. For an interval I in the higher degree barcode of the dataset, the points belonging to the connected component of the homology cycle are those corresponding to




Fig. 11 Decorated Merge Trees from Takens embeddings. Each row shows a different experiment. Starting with a time series (left), we produce an embedded point cloud via a Takens embedding (parameters determined automatically). The figure second from the left shows a PCA projection of the embedded point cloud. From the embedded point cloud, we produce a decorated merge tree (second from right) after light preprocessing via density-based subsampling in order to remove outliers (see text). Using the method described in the text, we are able to locate samples in the original time series which correspond to a point in the connected component where any particular degree-1 bar is born (shown on the right, color-coded to the bars on the DMT). The top row shows a synthetic time series and the bottom row shows a real world time series (heart rate data from Reiss et al. 2019). The heart rate data shows component-specific periodic behavior at the beginning and at the end of the time series, with the middle part apparently filtered out as a ‘transitory phase’ when subsampling. The short (blue) bar belonging to the early part of the time series intuitively indicates high frequency/low amplitude periodicity, whereas the more prominent (green) bar belonging to the latter part of the signal represents the more visually apparent periodicity

descendent leaves of the birth point of I in the merge tree. Using our cycle component location procedure, we can determine portions of the signals which generate various degree-1 homology classes—these are indicated by color coding in Fig. 11.

6.4 Clustering point clouds

In this example, we demonstrate the ability of DMT matching distance to distinguish point clouds with subtle topological differences. Figure 12a shows samples from six classes of synthetically-generated point clouds consisting of noisy blobs and circles. Each class contains 3 examples, with examples within class differing only due to noise effects. All classes have similar degree-0 topological structure, classes in the top row have three main degree-1 features and classes in the bottom row have two main degree-1 features. The distribution of the cycles among the connected components differ from class-to-class, but are the same within class.

For each point cloud, we construct various topological descriptors: a degree-0 persistence diagram, a degree-1 persistence diagram and a decorated merge tree. Pairwise distance matrices are computed for all 18 samples with respect to various metrics: bottleneck distance on degree-0 features, bottleneck distance on degree-1 features, maximum of bottleneck distance between degree-0 and degree-1 features and an estimation of decorated merge tree interleaving distance computed via the Fused Gromov-Wasserstein algorithm of Sect. 5.3. The resulting matrices and corresponding MDS plots are shown in Fig. 12b.




Fig. 12 **a** Samples from six classes of simple pointclouds in $\mathbb{R}^2$. Each class contains 3 examples of pointclouds with similar topological patterns. **b** Pairwise distance matrices for various methods shown in the top row, with corresponding MDS plots shown in the bottom row

6.5 Networks

Beyond Vietoris-Rips complexes of point clouds, we can generate DMTs for more general filtered simplicial complexes. A class of such objects which is important from a data science perspective is the class of filtered networks. To make this precise, a **filtered network** is a graph $G = (V, E)$ together with a function $f : V \rightarrow \mathbb{R}$. This function is extended to each edge $\{v, w\} \in E$ by the rule

$$f(\{v, w\}) := \max\{f(v), f(w)\},$$

and this yields a sublevel-set filtration in the sense of Example 3. We found it useful to add 2-dimensional simplices to the graphs by computing the 2-skeleton of the Vietoris-Rips complex with respect to shortest path distance; this has the effect of allowing degree-1 bars to die at finite times which reflect the size of their representative cycles. We are also able to threshold out small bars corresponding to triangles in the networks when displaying DMTs. This method is admittedly ad hoc, but it results in more informative topological summaries which capture both the topology of the filtration and the distance structure of the graph.

Given a filtered network, we produce a DMT by first building the merge tree itself by iteratively adding nodes to the tree in order of function value. Degree-1 homological information can be added to the merge tree by lifting bars (guaranteed to produce the correct DMT if the disjointness condition (Definition 26) is satisfied) or by a more involved construction, as in the case of point clouds.

An important source of network data comes from images represented as node-weighted graphs. Figure 13 shows a decorated merge tree describing the topological features of real fMRI brain scan image data. This is accomplished by considering an image as a weighted grid graph; that is, each pixel of the image is a node, nodes are connected in a grid to their neighbors and weights are assigned according to grayscale value. This procedure is also summarized in the figure. While the particular example shown in the Fig. 13 is only an illustration of how to apply DMTs to summarize an image network, there are potential realistic use cases. For example, there is a growing body of work on the use of TDA techniques for regularization of deep learning-based image segmentation (Hu et al. 2019, 2020); using DMTs instead of barcodes as features




Fig. 13 Grayscale images as a filtered networks. The **top row** illustrates the pipeline for extracting a decorated merge tree from a grayscale image. Beginning with a toy image on the left, we convert it to a filtered network (a regular graph with nodes weighted by grayscale value). From this, we produce a decorated merge tree (right). The **bottom row** shows two examples of decorated merge trees extracted from MRI images of Glioblastoma Multiforme tumors from The Cancer Imaging Archive (Clark et al. 2013; Scarpace et al. 2016), with the segmentations coming from Crawford et al. (2020). Each example shows the original image, with an inset showing a more detailed and lightly preprocessed version of the region of interest

would yield more detailed topological regularization. A full-fledged implementation of this pipeline is saved as a future direction.

6.6 Network matching

Computation of the decorated merge tree matching distance provides rich information about correspondences between points in the merge trees, since it involves the computation of an optimal coupling between these points. This is especially informative in the setting of filtered networks, since each node in the merge tree produced via our algorithm corresponds to a node in the original network. The optimal coupling then provides a probabilistic matching between nodes of the filtered networks which captures the topology of the filtration function.

Estimation of node correspondences between graphs is a classical problem which has recently been a focus of Gromov-Wasserstein-based methods. The optimal coupling used to match two graphs depends on the how they are represented as measure networks in the GW matching problem (6); for example a graph can be represented by its geodesic distance matrix (Hendrikson et al. 2016), its adjacency matrix (Xu et al. 2019) or by a heat kernel matrix at a chosen scale (Chowdhury and Needham 2020). The coupling produced in the process of computing the decorated merge tree matching distance differs from existing approaches in that it matches based on purely topological features of the graphs. A simple example illustrating the qualitative differences in graph matchings is shown in Fig. 14. In this example, the node filtration for each graph




Fig. 14 Matching networks. The top row shows two networks, constructed to have similar topological motifs. The bottom row illustrates node matchings between the graph; each matching is illustrated by transferring the node colors from the source graph (the node colors here are arbitrary) to the target graph via a probabilistic coupling of its nodes—node j in the target graph receives the color of target node i when the (i, j) entry is the largest in the j th column of the coupling matrix. The first color transfer utilizes the coupling obtained during the computation of DMT matching distance via the Fused GW framework, while the latter two couplings come from the GW framework applied to adjacency matrices and geodesic distance matrices, respectively. Qualitatively, the matching produced using DMTs most strongly preserves topological features

is given by a diffusion Fréchet function, which measures local node density (Martínez et al. 2019). This simple example suggests the possibility that including topology in the matching process may have a regularizing effect, allowing the matching to capture large-scale graph similarity which may be lost during the optimization process when using more localized similarity criteria. Learning the best filtration function for a given task in the context of classical TDA is an active field of research (Hajij et al. 2020; Hofer et al. 2020); extending this line of research to learn filtration functions which capture the interplay between degree-0 and degree-1 features will be taken up in future work.

7 Discussion

In this paper, we defined several variants of the notion of a decorated merge tree. We introduced a metric on the space of DMTs as well as a metric on the simpler space of leaf decorated merge trees which is more amenable to computation. A stability result was demonstrated for these metrics, and stability was extended to compare these metrics to several other metrics which have appeared in the literature. Several use cases for the DMT framework were demonstrated via computational examples.

There are many directions for future research on DMTs and related ideas. On the computational side, we plan to refine our algorithms to efficiently handle DMT computations on real datasets. This challenge comes with some interesting theoretical

questions. For example, we are currently using off-the-shelf methods for approximating Gromov-Wasserstein matchings between merge trees; the hierarchical structure of merge trees should allow for specialized algorithms which utilize this structure to produce faster and more accurate estimates. To scale DMT-based analyses to handle large datasets in a modern machine learning setting, we also plan to explore principled vectorizations of these signatures.

There are also several directions for research on the theoretical side. Two natural ways to generalize the concept of a decorated merge tree are to vary the object being decorated (the merge tree) and to vary the type of decoration (persistence barcodes). In future work, we plan to extend these ideas to treat decorated Reeb graphs. We also plan to explore decorations by zig-zag persistence modules and extended persistence modules.

Furthermore, Remark 8 suggests that perhaps the interleaving distance on categorical decorated merge trees is universal (Lesnick 2015). One avenue for studying this result could leverage an observation of Gunnar Carlsson: that the cup product structure on cohomology determines π_0 information. Perhaps a universality result that preserves the richer algebraic structure of cup products is the right framework for framing this result. Finally, one might take up the question of how to define “higher” decorated merge trees—via cochains, for example (Mandell 2006)—that might offer a complete homotopical invariant of persistent spaces. Repeating much of the standard template of results in algebraic topology from the past hundred years for persistent spaces is an active and interesting line of research.

Acknowledgements JC would like to thank Rachel Levanger for discussions dating back to 2017 when the module-theoretic and lift-theoretic approaches to DMTs were first considered. JC would also like to thank Gabriel Bainbridge for teaching him about the category of parameterized objects during the summer of 2020. Gabe’s use of the parameterized category is set to appear in Bainbridge (2022). Finally, JC would like to acknowledge NSF Grant CCF-1850052 and NASA Contract 80GRC020C0016 for supporting his research. WM acknowledges partial support by NSF grant DMS-1722995. TN acknowledges support by NSF grant DMS-2107808 and would like to thank Facundo Mémoli for useful feedback on an earlier draft of the paper. HH would like to acknowledge NSF grant DMS-1854683.

Declarations

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

A Interval topology

Let X be a compact topological space and $f : X \rightarrow \mathbb{R}$ be a continuous function. Let E_f be the epigraph of $f : X \rightarrow \mathbb{R}$, $\phi : E_f \rightarrow \mathbb{R}$ be the projection map, and $\pi : E_f \rightarrow T_f X$ be the Reeb quotient map, which identifies two points if they are in the same connected component of a level set.

Define a partial order on $T_f X$ by letting $p \leq q$ if there exists x in X and r, s in $\mathbb{R}$ such that $f(x) \leq r \leq s$ and $p = \pi(x, r)$ and $q = \pi(x, s)$. The fact that this defines a poset follows from the following lemma:

Lemma 4 *Let (x, r) and (y, s) be points in E_f such that $r \leq s$. The following are equivalent:*

- i) $\pi(x, r) \leq \pi(y, s)$.
- ii) $\pi(x, s) = \pi(y, s)$.
- iii) $\pi(x, t) = \pi(y, t)$ for all $t > s$.

Proof i) $\implies$ ii) : Let z be a point in X such that $\pi(z, r) = \pi(x, r)$ and $\pi(z, s) = \pi(y, s)$. Let C be the connected component of $\phi^{-1}(r)$ containing (x, r) and D be the connected component in $\phi^{-1}(s)$ containing (y, s) . Let C' be the subset of E_f obtained by shifting up the second coordinates of points in C to s . As C contains (z, r) and D contains (z, s) , both C' and D contains (z, s) . Since C' is connected, it is contained in D . Hence $\pi(x, s) = \pi(y, s)$.

ii) $\implies$ iii) Note that $\pi(x, s) = \pi(y, s) \leq \pi(y, t)$. Now, the result follows from the argument in the previous part.

iii) $\implies$ i) Let (s_n) be a decreasing sequence converging to s . Let $C_n \times \{s_n\}$ be the connected component of $\phi^{-1}(s_n)$ containing both (x, s_n) and (y, s_n) . Note that (C_n) is a decreasing sequence of closed connected sets in X . Let C be the intersection of C_n 's. Note that $C \times \{s\}$ is contained in E_f and x, y are elements of C . Let us show that C is connected. Let V and W be open subsets in X such that C is contained in the disjoint union of U of V and W . Assume x is in V . Let U_n be the complement of C_n in X . The collection consisting of U and all U_n 's is an open cover of X , hence it has a finite subcover. Since (U_n) is an increasing family of open sets, this implies that for n large enough X is the union of U and U_n , hence C_n is contained in U . Since C_n is connected and has non-empty intersection with V for all n , for n large enough C_n is contained in V . Therefore C is contained in V . This shows C is connected, hence $\pi(x, s) = \pi(y, s)$. By definition, $\pi(x, r) \leq \pi(y, s)$.

Corollary 2

- i) $(T_f X, \leq)$ is a tree poset in the sense that the upper set of any point is a chain (i.e., well ordered).
- ii) Let C be a chain in $(T_f X, \leq)$. Then C has an infimum (i.e., the maximum of its lower bounds), and the closure of C in the quotient topology is contained in the upper set of its infimum.

Proof i) : Assume that $\pi(x, r) \leq \pi(y, s)$, $\pi(x, r) \leq \pi(z, t)$ and $s \leq t$. Then, by Lemma 4, $\pi(y, t) = \pi(x, t) = \pi(z, t)$, and therefore $\pi(y, s) \leq \pi(z, t)$.

ii) : let (x_n, r_n) be a sequence in C such that $\lim r_n = r := \inf\{s : (x, s) \in C\}$. The sequence (x_n) has a subsequence converging to a point x in X . Let (y_n, s_n) be a convergent sequence in E_f such that $\pi(y_n, s_n)$ is in C for all n . Let $\lim(y_n, s_n) = (y, s)$. We need to show that $\pi(x, r) \leq \pi(y, s)$. Let $t > s$. Without loss of generality, we can assume that $r_n, s_n \leq t$ for all n and (x_n) converges to x . Since C is a chain, by Lemma 4, $\pi(y_i, t) = \pi(x_j, t)$ for all i, j . Let D be the connected component of $\phi^{-1}(t)$ containing (y_n, t) and (x_n, t) for all n . Since D is closed, it contains (x, t) and (y, t) , hence $\pi(x, t) = \pi(y, t)$. Since $t > s$ was arbitrary, by Lemma 4, $\pi(x, r) \leq \pi(y, s)$. This shows $\pi(x, r)$ is a lower bound for the closure of C . Let $\pi(x', r')$ be another lower bound for C . For any $t > r$, there exists n such that $r_n \leq t$, so $\pi(x, t) = \pi(x_n, t) = \pi(x', t)$. Therefore, by Lemma 4, $\pi(x', r') \leq \pi(x, r)$.

We introduce a new topology on $T_f X$, which we call the *interval topology*, defined as follows.

Definition 32 Let $T_f X$ be the merge tree, viewed as a poset. A subset U of $T_f X$ is open in the **interval topology** if for each p in U and each q in $T_f X \setminus \{p\}$ comparable to p , there exists w in U strictly in between p and q such that the interval between p and w is contained in U .

Proposition 10 *If $(T_f X, \leq)$ has finitely many leaves (i.e., minimal elements in the poset), then the quotient topology coincides with the interval topology.*

Lemma 5 *The upper set of any point is closed in the interval topology.*

Proof Let C be the upper set of a point $\pi(z, u)$ in $T_f X$. Let U be its complement. Let us show that U is open in the interval topology. Let $\pi(x, r)$ be in U . Let $\pi(y, s)$ be comparable to $\pi(x, r)$. If $\pi(y, s) < \pi(x, r)$, then everything in between them is already in U as U is the complement of an upper set. Let us assume then $\pi(x, r) < \pi(y, s)$. By Lemma 4, there exists t such that $r < t < s$ and $\pi(x, t)$ is not in the upper set of $\pi(z, u)$. Then $\pi(x, t)$ is strictly in between $\pi(x, s)$ and $\pi(y, s)$ and the interval in between $\pi(x, s)$ and $\pi(x, t)$ is in U . $\square$

Lemma 6 *The height map $h : T_f X \rightarrow \mathbb{R}$, $\pi(x, r) \mapsto r$ is continuous with respect to the interval topology.*

Proof Let r be a real number and $\epsilon > 0$. We need to show that the preimage of $(r - \epsilon, r + \epsilon)$ is open in the interval topology. Assume that $\pi(x, s)$ is in $T_f X$ and $|s - r| < \epsilon$. Let $\pi(y, t)$ be a point comparable to $\pi(x, s)$. If $\pi(y, t) < \pi(x, s)$, then $\pi(x, s) = \pi(y, s)$. Pick $s' > t$ such that $r - \epsilon < s' < s$. Then $\pi(y, s')$ is strictly in between $\pi(y, t)$ and $\pi(x, s)$ and the image of the interval between $\pi(y, s')$ and $\pi(x, s)$ is $[s', s]$, which is contained in $(r - \epsilon, r + \epsilon)$. If $\pi(y, t) > \pi(x, s)$, then $\pi(x, t) = \pi(y, t)$. Pick $s' < t$ such that $s < s' < r + \epsilon$. Then $\pi(x, s')$ is strictly in between $\pi(x, s)$ and $\pi(y, t)$, and the image of the interval between $\pi(x, s)$ and $\pi(x, s')$ is $[s, s']$, which is contained in $(r - \epsilon, r + \epsilon)$.

Lemma 7 *For each x in X , the map $\psi_x : [f(x), \infty) \rightarrow T_f X$, $r \mapsto \pi(x, r)$ is continuous with respect to the interval topology.*

Proof Let U be an open set in the interval topology. Assume $\psi_x(r) = \pi(x, r)$ is in U . The result follows from the comparability of $\pi(x, r)$ with $\pi(x, f(x))$ and $\pi(x, r + 1)$.

Proof of Proposition 10 Let $C_1, \dots, C_n$ be the upper sets of leaves in $T_f X$. In both topologies, they give a finite closed cover by Corollary 2 and Lemma 5. This implies that a set U is open in these topologies if and only if its intersection with each C_i 's is relatively open. Hence, it is enough to show that the relative topologies on the C_i 's coincide for the quotient and the interval topologies. Both the height map $\pi(x, r) \mapsto r$ of Lemma 6 and ψ_x of Lemma 7 are continuous with respect to both topologies. If we let $\pi(x, f(x))$ denote the leaf of C_i , and restrict h to C_i , then these maps becomes inverses of each other. This completes the proof.

We now return to the Proof of Proposition 2 from the main text, which says that the topology induced by the ℓ^p -metrics on a merge tree agrees with the quotient space topology.

Proof of Proposition 2 Using the logic from the Proof of Proposition 10, it suffices to show that the relative topologies on the C_i 's with respect to the interval and metric topologies coincide. Indeed, the intersection of an ℓ^p -ball with each upper set C_j of a leaf of the merge tree is relatively open with respect to the interval topology. This is clear, since this intersection is homeomorphic to an open interval via the height map. Similarly, relatively open sets in the interval topology are relatively open in the metric topology. $\square$

B Comparison and existence of merge trees

In this paper we used two primary perspectives on merge trees: the classical merge tree (Definition 3), which defines the merge tree as the Reeb graph of the epigraph, and the generalized merge tree (Definition 2), which is defined as the display poset of the persistent set. Each approach has its own advantages, but in order to ensure that the Reeb graph of a function is actually a graph requires specifying conditions on $f : X \rightarrow \mathbb{R}$ such as the Morse condition or piecewise linearity (De Silva et al. 2016, p. 857). In this section, we describe a third way of constructing the merge tree following the metric tree construction in Mémoli and Okutan (2018) and show that when X is locally path connected, the classical construction coincides with the metric construction described here.

For this section we assume that X is a compact path-connected topological space and $f : X \rightarrow \mathbb{R}$ is a continuous function. We are using the definition of metric tree given in (Bowditch 1991, Section 3.4), which is also called real tree or $\mathbb{R}$ -tree.

Remark 17 As the quotient of the epigraph of f , the merge tree is the wedge of the image of the graph of f and the half real line. Let us denote the image of the graph of $f : X \rightarrow \mathbb{R}$ inside the merge tree by $M_f X$.

Definition 33 (Metric Merge Tree) Let X be a compact path connected topological space and $f : X \rightarrow \mathbb{R}$ be a continuous function. Let us define $m_f : X \times X \rightarrow \mathbb{R}$ by

$$m_f(x, y) := \inf\{\max f \circ \gamma \mid \gamma : [a, b] \rightarrow X, \gamma(a) = x, \gamma(b) = y\}.$$

Define $t_f : X \times X \rightarrow \mathbb{R}$ by

$$t_f(x, y) = 2m_f(x, y) - f(x) - f(y).$$

It is easy to see that t_f is a pseudo-metric (symmetric, non-negative and satisfies triangle inequality). Let us denote the associated metric space by $(\mathcal{T}_f X, t_f)$ and call it *metric merge tree*. The function f is still well defined on $\mathcal{T}_f X$.

Proposition 11 Let X be a compact topological space and $f : X \rightarrow \mathbb{R}$ be a continuous function. Then $(\mathcal{T}_f X, t_f)$ is a metric tree.

Proof It is enough to show that t_f is a path metric and it has hyperbolicity zero (see Bowditch 1991, Proposition 3.4.2). To show that t_f is a path metric, it is enough to show that for each $x, y \in X$ and $\epsilon > 0$, there exists $z \in X$ such that

$$t_f(x, z) \leq \frac{t_f(x, y)}{2} + \epsilon, \quad t_f(y, z) \leq \frac{t_f(x, y)}{2} + \epsilon$$

(see BuragoBurago et al. 2001, Theorem 2.4.16). Take $x, y \in X$ and $\epsilon > 0$. Let $\gamma : [0, 1] \rightarrow X$ be a continuous curve from x to y such that

$$m_f(x, y) \geq \max f \circ \gamma \epsilon.$$

Let $M := \max f \circ \gamma$. Define

$$t_0 := \min\{t : f(\gamma(t)) = M\}, \quad t_1 := \max\{t : f(\gamma(t)) = M\}.$$

Let $d_0 := M - f(x)$ and $d_1 := M - f(y)$. Without loss of generality $d_0 \geq d_1$. Let $r := (d_0 + d_1)/2$ and

$$s := \min\{t : f \circ \gamma(t) = f(x) + r\}.$$

Let $z = \gamma(s)$. Now, we have:

$$\begin{aligned} t_f(x, y) &\geq 2(M - \epsilon) - f(x) - f(y) = d_0 + d_1 - 2\epsilon \\ t_f(x, z) &= r = \frac{d_0 + d_1}{2} \\ t_f(y, z) &\leq (M - f(x) - r) + d_1 = \frac{d_0 + d_1}{2}. \end{aligned}$$

Now let us show that t_f has 0-hyperbolicity. Let p be a point where f takes its maximum. The Gromov product $g_p : X \times X \rightarrow \mathbb{R}$ with respect to the pseudo-metric t_f is defined by:

$$g_p(x, y) = \frac{1}{2}(t_f(p, x) + t_f(p, y) - t_f(x, y)).$$

The metric t_f has zero hyperbolicity if

$$g_p(x, z) \geq \min(g_p(x, y), g_p(y, z))$$

for all $x, y, z \in X$ (see Gromov 1987, Corollary 1.1.B).

Claim 1: $g_p(x, y) = f(p) - m_f(x, y)$ for all $x, y \in X$:

By maximality of p , $m_f(p, x) = f(p)$. Hence,

$$t_f(p, x) = 2m_f(p, x) - f(p) - f(x) = f(p) - f(x).$$

So we have

$$\begin{aligned} g_p(x, y) &= \frac{1}{2}(t_f(p, x) + t_f(p, y) - t_f(x, y)) \\ &= \frac{1}{2}(f(p) - f(x) + f(p) - f(y) - 2m_f(x, y) + f(x) + f(y)) \\ &= f(p) - m_f(x, y) \end{aligned}$$

Claim 2: $g_p(x, z) \geq \min(g_p(x, y), g_p(y, z))$ for all x, y, z in X .

By its definition, $m_f(x, z) \leq \max(m_f(x, y), m_f(y, z))$. Claim 2 follows from this and Claim 1.

Therefore the hyperbolicity of $\mathcal{T}_f X$ is 0. $\square$

Theorem 6 *If X is locally path connected, then $M_f X$ is the underlying topological space of $\mathcal{T}_f X$.*

Proof It is enough to show that $t_f(x, y) = 0$ if and only if $(x, f(x))$ and $(y, f(y))$ are identified under the Reeb quotient map π .

Assume $t_f(x, y) = 0$. Then, $m_f(x, y) = f(x) = f(y)$. So, for each $t > f(x) = f(y)$, there is a path in the sublevel set $X^{\leq t}$ connecting x, y . Therefore, $\pi(x, t) = \pi(y, t)$, and by Lemma 4, $\pi(x, r) = \pi(y, r)$.

Assume $\pi(x, f(x)) = \pi(y, f(y))$. Let $r = f(x)$ and $t > r$. Since x and y are in the same connected component of the (closed) sublevel set $X^{\leq r}$ and X is locally path connected, x, y are in the same path component of the (open) sublevel set $X^{< t}$. This implies that $m_f(x, y) < t$. Since $t > r$ arbitrary, $m_f(x, y) = r = f(x) = f(y)$, so $t_f(x, y) = 0$. $\square$

C Technical proofs

This section contains proofs from the main body of the paper which are technical, but essentially straightforward. These results are focused on barcodes and bottleneck distances. To keep the exposition clean, we deal with barcodes satisfying the simplifying assumptions:

- the barcode is a *set* (all multiplicities of intervals are equal to one);
- the barcode is finite;
- each interval in the barcode is a half open interval of the form $[b, d)$, where $0 \leq b \leq d \leq \infty$.

C.1 Proof of Proposition 4

To prove the proposition, we will introduce the notion of truncating a barcode.

Definition 34 Let B be a barcode and let $I = [b, d) \in B$. The **truncation** of I at height h , denoted $\text{trunc}_h(I)$ is

- equal to I if $h \leq b$;
- the interval $[h, d]$ if $b \leq h \leq d$;
- the empty interval if $h > d$.

The **truncation of B at height h** , $\text{trunc}_h(B)$, is the barcode obtained by truncating all of the intervals of B at height h .

Lemma 8 *Let B and B' be barcodes whose intervals $[b, d]$ all satisfy $b \geq H$ for some constant $H \geq 0$. Suppose that there exists a δ -matching between B and B' . For $\epsilon \leq \delta$, there is a δ -matching between B and $\text{trunc}_{H+\epsilon}(B')$.*

Proof Let ξ be a δ -matching of B and B' . Define a matching $\widehat{\xi}$ of B and $\text{trunc}_{H+\epsilon}(B')$ by setting

$$\begin{aligned}\text{dom}(\widehat{\xi}) &:= \{I \in \text{dom}(\xi) \mid \text{trunc}_{H+\epsilon}(\xi(I)) \neq \emptyset\} \\ \text{ran}(\widehat{\xi}) &:= \{I' \in \text{ran}(\xi) \mid \text{trunc}_{H+\epsilon}(I') \neq \emptyset\}, \text{ and} \\ \widehat{\xi}(I) &:= \text{trunc}_{H+\epsilon}(\xi(I)).\end{aligned}$$

Let $I = [b, d] \in \text{dom}(\widehat{\xi})$ with $I' = [b', d'] = \xi(I)$. Then $b' \leq H + \epsilon \leq d'$ and $\text{trunc}_{H+\epsilon}(I') = [H + \epsilon, d']$. We have $|d - d'| \leq \delta$, by the assumption that ξ was a δ -matching. Moreover,

$$\delta \geq b - b' \geq b - (H + \epsilon)$$

and

$$\delta \geq \epsilon \geq \epsilon + (H - b) = (H + \epsilon) - b$$

imply that $|b - (H + \epsilon)| \leq \delta$, so that the cost of matching I with $\text{trunc}_{H+\epsilon}(I')$ is less than δ . On the other hand, if $I \notin \text{dom}(\widehat{\xi})$ then either $I \notin \text{dom}(\xi)$, in which case we are done, or $\text{trunc}_{H+\epsilon}(\xi(I)) = \emptyset$. In the latter case, let $I = [b, d]$ and $I' = [b', d'] = \xi(I)$. We have

$$d - b \leq d' + \delta - b \leq d' + \delta - H \leq H + \epsilon + \delta - H \leq 2\delta,$$

hence $\| [b, d] \|_{\Delta} \leq \delta$. Similar arguments handle intervals $I' \notin \text{ran}(\widehat{\xi})$. □

Proof of Proposition 4 Let

$$\widehat{d}_B(\mathcal{B}_F, \mathcal{B}_G) := \inf\{\epsilon \geq 0 \mid \exists \epsilon\text{-matching of } \mathcal{B}_F \text{ and } \mathcal{B}_G\},$$

so that our goal is to show $d_B = \widehat{d}_B$. Clearly $d_B \leq \widehat{d}_B$, since the former infimizes over a larger set of matchings than the latter. To see the reverse inequality, let Φ, Ψ define an (ϵ, δ) -matching. If $\epsilon \geq \delta$, then each of the δ -matchings between $\mathcal{B}_F(p)$ and $\mathcal{B}_G(\Phi(p))$ and between $\mathcal{B}_G(q)$ and $\mathcal{B}_F(\Psi(q))$ is, in particular, an ϵ -matching. This implies the existence of an ϵ -matching between $\mathcal{B}_F$ and $\mathcal{B}_G$. Finally, suppose that $\delta \geq \epsilon$. Let $\widehat{\Phi}$ and $\widehat{\Psi}$ denote the δ -interleaving maps obtained by composing Φ and Ψ

with flows in their respective target merge trees. We claim that there exist δ -matchings between all pairs $\mathcal{B}_F(p)$ and $\mathcal{B}_G(\widehat{\Phi}(p))$ and between $\mathcal{B}_G(q)$ and $\mathcal{B}_F(\widehat{\Psi}(q))$. Indeed, setting $H = \pi_F(p)$, with $\pi_F : \mathcal{M}_F \rightarrow \mathbb{R}$ denoting projection, we have that

$$\mathcal{B}_G(\widehat{\Phi}(p)) = \text{trunc}_{H+\delta-\epsilon}(\mathcal{B}_G(\Phi(p))).$$

Lemma 8 therefore implies that there exists a δ -matching of $\mathcal{B}_F(p)$ and $\mathcal{B}_G(\widehat{\Phi}(p))$. The existence of a δ -matching of $\mathcal{B}_G(q)$ and $\mathcal{B}_F(\widehat{\Psi}(q))$ follows for the same reason. $\square$

We now have the tools at hand to prove Proposition 5.

Proof of Proposition 5 Let $p \in \mathcal{M}_F$ be an arbitrary point. Choose a leaf v of $\mathcal{M}_F$ with $v \preceq p$ and a leaf $w \preceq \Phi(p)$. Let $\pi_F : \mathcal{M}_F \rightarrow \mathbb{R}$ be the projection map. Then the barcodes at these points can be expressed as

$$\begin{aligned}\mathcal{B}_F(p) &= \text{trunc}_{\pi_F(p)}(\mathcal{B}_F(v)), \text{ and} \\ \mathcal{B}_G(\Phi(p)) &= \text{trunc}_{\pi_F(p)+\epsilon}(\mathcal{B}_G(w)).\end{aligned}$$

It is straightforward to check that a δ -matching of $\mathcal{B}_F(p)$ and $\mathcal{B}_F(\Phi(p))$ induces a δ -matching of the truncated barcodes. A similar argument applies to pairs $q \in \mathcal{M}_G$ and $\Psi(q) \in \mathcal{M}_F$. $\square$

C.2 Proof of Proposition 8

Let $B = \{I_j = [b_j, d_j]\}_{j=1}^N$ be the degree- k barcode for F . To ease exposition, assume that the births in the barcode are distinct—this assumption is easily removed at the cost of necessitating more involved notation. Suppose that the indices have been chosen so that $b_1 < b_2 < \dots < b_N$. For each bar I_j , we choose a representative cycle c_j which generates the persistent homology class represented by I_j ; in particular, choose c_j so that all of its simplices are contained in the connected component of the birth simplex for I_j . Let $[c_j]$ denote the homology class of c_j . Let $p_j \in \mathcal{M}_F$ denote the birth point of I_j .

We can use these representative cycles to build bases for various homology vector spaces which agree with the barcode decomposition. In the following, let $\iota_{b,d}$ denote the inclusion map $F(b) \hookrightarrow F(d)$ for each $b < d$. By an abuse of notation, we also let $\iota_{b,d}$ denote the induced map on homology $H_k(F(b)) \rightarrow H_k(F(d))$. We now construct our bases:

1. $\{v_1 := [c_1]\}$ is a basis for $H_k(F(b_1))$,
2. $\{[c_2], \iota_{b_1,b_2}(v_1)\}$ is a spanning set for $H_k(F(b_2))$ and is linearly independent if

$$\iota_{b_1,b_2}(v_1) \neq 0.$$

Moreover, we can choose $\lambda_1^{(2)} \in \mathbb{k}$ so that

$$v_2 := [c_2] + \lambda_1^{(2)} \iota_{b_1,b_2}(v_1)$$

has the property that

$$\iota_{b_2,d}(v_2) \notin \text{span}\{\iota_{b_1,d}(v_1)\}$$

for any $d < d_2$. In particular, $v_2 \notin \ker(\iota_{b_2,d})$ for $d < d_2$.

3. Suppose that we have defined $v_1, \dots, v_{j-1}$. Then

$$\{[c_j], \iota_{b_{j-1},b_j}(v_{j-1}), \dots, \iota_{b_1,b_j}(v_1)\}$$

is a spanning set for $H_k(F(b_j))$ which is linearly independent once any zero vectors have been removed. We define

$$v_j := [c_j] + \lambda_{j-1}^{(j)} \iota_{b_{j-1},b_j}(v_{j-1}) + \dots + \lambda_1^{(j)} \iota_{b_1,b_j}(v_1),$$

where coefficients $\lambda_i^{(j)}$ are chosen so that

$$\iota_{b_j,d}(v_j) \notin \text{span}\{\iota_{b_{j-1},d}(v_{j-1}), \dots, \iota_{b_1,d}(v_1)\} \quad (10)$$

for all $d < d_j$. In particular, $v_j \notin \ker(\iota_{b_j,d})$ for $d < d_j$.

The lifespan of the vector v_j defined via this process represents the persistent homology bar I_j .

To prove that $\widehat{\mathcal{F}} \approx \mathcal{F}$, we will show that $\mathcal{F}$ is real interval decomposable. To do so, it suffices to show that each of these v_j can be chosen to be a linear combination of classes with cycle representatives that lie in the same connected component as the birth simplex of I_j . This will be achieved via induction.

For the base case of the induction argument, we have $v_1 = [c_1]$, and the claim follows immediately. Suppose that $v_1, \dots, v_{j-1}$ have the desired property and consider v_j . We sort the $v_1, \dots, v_{j-1}$ into a collection $v_{r_1}, \dots, v_{r_m}$ of vectors corresponding to bars in B which are born in the same component as c_j (i.e., bars whose birth points satisfy $p_{r_1}, \dots, p_{r_m} \preccurlyeq p_j$) and a collection $v_{s_1}, \dots, v_{s_n}$ of vectors which do not have this property, so that

$$v_j = [c_j] + \lambda_{r_1}^{(j)} \iota_{b_{r_1},b_j}(v_{r_1}) + \dots + \lambda_{r_m}^{(j)} \iota_{b_{r_m},b_j}(v_{r_m}) + \lambda_{s_1}^{(j)} \iota_{b_{s_1},b_j}(v_{s_1}) + \dots + \lambda_{s_n}^{(j)} \iota_{b_{s_n},b_j}(v_{s_n}).$$

We claim that taking $\lambda_{s_i}^{(j)} = 0$ for all $i = 1, \dots, n$ results in a valid vector v_j ; i.e., that this yields a v_j satisfying (10). Indeed, let $d < d_j$ and suppose that

$$\begin{aligned} \iota_{b_j,d} \left([c_j] + \lambda_{r_1}^{(j)} \iota_{b_{r_1},b_j}(v_{r_1}) + \dots + \lambda_{r_m}^{(j)} \iota_{b_{r_m},b_j}(v_{r_m}) \right) \\ = \mu_1 \iota_{b_j,d}(v_{s_1}) + \dots + \mu_m \iota_{b_j,d}(v_{s_m}) \end{aligned} \quad (11)$$

for some coefficients μ_i . We first note that there must be some $v_{s_i} \notin \ker(\iota_{s_i,d})$ —otherwise the right-hand side of (11) is zero, hence left-hand side is as well and

we arrive at the contradiction that $\iota_{b_j, d}(v_j) = 0$. We therefore have a well defined number

$$d' := \min\{\text{merge}(p_j, p_{s_i}) \mid \iota_{b_{s_i}, d}(v_{s_i}) \neq 0\}.$$

By the H_k -disjointness assumption, either $d_j < d'$ or $d_{s_i} < d'$ for all s_i . The former case implies $d < d'$. By the induction hypothesis, there is a cycle representation of the left-hand side of (11) whose vertices all belong to the same connected component as those of $\iota_{b_j, d}(c_j)$ and a cycle representation of the right-hand side whose vertices belong to a different connected component than c_j . Moreover, these cycle representatives are homologous, and we have arrived at a contradiction. On the other hand, if $d_{s_i} < d'$ for all s_i then it must be that $d < d_{s_i}$ for some s_i (since there exists some $v_{s_i} \notin \ker(\iota_{b_{s_i}, d})$) and a similar argument yields a contradiction. This completes the proof.

References

- Agarwal, P.K., Fox, K., Nath, A., Sidiropoulos, A., Wang, Y.: Computing the Gromov-Hausdorff distance for metric trees. *ACM Trans. Algorithms (TALG)* **14**(2), 1–20 (2018)
- Alvarez-Melis, D., Jegelka, S., Jaakkola, T.S.: Towards optimal transport with global invariances. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1870–1879. PMLR (2019)
- Bainbridge, G.: Directional cellular sheaves for multicast network routing. In Preparation (2022)
- Bauer, U.: Ripser: efficient computation of vietoris-rips persistence barcodes (2019). Preprint
- Bauer, U., Lesnick, M.: Persistence diagrams as diagrams: a categorification of the stability theorem. In: *Topological Data Analysis*, pp. 67–96. Springer (2020)
- Blumberg, A.J., Lesnick, M.: Universality of the homotopy interleaving distance. *arXiv preprint arXiv:1705.01690* (2017)
- Botnan, M., Crawley-Boevey, W.: Decomposition of persistence modules. *Proc. Am. Math. Soc.* **148**(11), 4581–4596 (2020)
- Bowditch, B.H.: Notes on Gromov’s hyperbolicity criterion for path-metric spaces. In: *Group Theory From a Geometrical Viewpoint* (1991)
- Bubenik, P., De Silva, V., Scott, J.: Metrics for generalized persistence modules. *Found. Comput. Math.* **15**(6), 1501–1531 (2015)
- Bubenik, P., Scott, J.A.: Categorification of persistent homology. *Discret. Comput. Geom.* **51**(3), 600–627 (2014)
- Burago, D., Burago, I.D., Burago, Y., Ivanov, S., Ivanov, S.V., Ivanov, S.A.: *A course in metric geometry*, vol. 33. American Mathematical Society (2001)
- Carlsson, G.: Topological pattern recognition for point cloud data. *Acta. Numer.* **23**, 289–368 (2014)
- Carlsson, G., De Silva, V.: Zigzag persistence. *Found. Comput. Math.* **10**(4), 367–405 (2010)
- Carlsson, G., Mémoli, F.: Classifying clustering schemes. *Found. Comput. Math.* **13**(2), 221–252 (2013)
- Chapel, L., Alaya, M.Z., Gasso, G.: Partial Gromov-Wasserstein with applications on positive-unlabeled learning. *arXiv preprint arXiv:2002.08276* (2020)
- Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J., Oudot, S.Y.: Proximity of persistence modules and their diagrams. In: *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, pp. 237–246 (2009)
- Chowdhury, S., Mémoli, F.: The Gromov-Wasserstein distance between networks and stable network invariants. *CoRR abs/1808.04337* (2018). *arXiv:1808.04337*
- Chowdhury, S., Needham, T.: Generalized spectral clustering via Gromov-Wasserstein learning. *arXiv preprint arXiv:2006.04163* (2020)
- Chowdhury, S., Needham, T.: Gromov-wasserstein averaging in a riemannian framework. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 842–843 (2020)

- Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., et al.: The cancer imaging archive (tcia): maintaining and operating a public information repository. *J. Digit. Imaging* **26**(6), 1045–1057 (2013)
- Crawford, L., Monod, A., Chen, A.X., Mukherjee, S., Rabadán, R.: Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *J. Am. Stat. Assoc.* **115**(531), 1139–1150 (2020)
- Crawley-Boevey, W.: Decomposition of pointwise finite-dimensional persistence modules. *J. Algebr. Appl.* **14**(05), 1550066 (2015)
- Curry, J.: The fiber of the persistence map for functions on the interval. *J. Appl. Comput. Topol.* **2**(3–4), 301–321 (2018)
- Curry, J., Mukherjee, S., Turner, K.: How many directions determine a shape and other sufficiency results for two topological transforms. arXiv preprint [arXiv:1805.09782](https://arxiv.org/abs/1805.09782) (2018)
- Curry, J.M.: Topological data analysis and cosheaves. *Jpn. J. Ind. Appl. Math.* **32**(2), 333–371 (2015)
- De Silva, V., Munch, E., Patel, A.: Categorified reeb graphs. *Discret. Comput. Geom.* **55**(4), 854–906 (2016)
- de Silva, V., Munch, E., Stefanou, A.: Theory of interleavings on categories with a flow. *Theory Appl. Categories* **33**(21), 583–607 (2018)
- Dey, T.K., Shi, D., Wang, Y.: Comparing graphs via persistence distortion. In: 31st International Symposium on Computational Geometry (SoCG 2015). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2015)
- Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. In: Proceedings 41st Annual Symposium on Foundations of Computer Science, pp. 454–463. IEEE (2000)
- Emrani, S., Gentimis, T., Krim, H.: Persistent homology of delay embeddings and its application to wheeze detection. *IEEE Signal Process. Lett.* **21**(4), 459–463 (2014)
- Farahbakhsh Toul, E., Wang, Y.: FPT-algorithms for computing Gromov-Hausdorff and interleaving distances between trees. In: 27th Annual European Symposium on Algorithms (ESA 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019)
- Flamary, R., Courty, N.: Pot python optimal transport library (2017). <https://pythonot.github.io/>
- Frosini, P., Landi, C., Mémoli, F.: The persistent homotopy type distance. *Homol. Homot. Appl.* **21**(2), 231–259 (2019)
- Gasparovic, E., Munch, E., Oudot, S., Turner, K., Wang, B., Wang, Y.: Intrinsic interleaving distance for merge trees. arXiv preprint [arXiv:1908.00063](https://arxiv.org/abs/1908.00063) (2019)
- Gromov, M.: Hyperbolic groups. In: Essays in group theory, pp. 75–263. Springer (1987)
- Hajj, M., Munch, E., Rosen, P.: Graph similarity using pagerank and persistent homology. arXiv preprint [arXiv:2002.05158](https://arxiv.org/abs/2002.05158) (2020)
- Hang, H., Mémoli, F., Mio, W.: A topological study of functional data and fréchet functions of metric measure spaces. *J. Appl. Comput. Topol.* **3**(4), 359–380 (2019)
- Hang, H., Mio, W.: Correspondence modules and persistence sheaves: a unifying framework for one-parameter persistent homology. arXiv preprint [arXiv:2006.08557](https://arxiv.org/abs/2006.08557) (2020)
- Heine, C., Lette, H., Hlawitschka, M., Iuricich, F., De Floriani, L., Scheuermann, G., Hagen, H., Garth, C.: A survey of topology-based methods in visualization. In: Computer Graphics Forum, vol. 35, pp. 643–667. Wiley Online Library (2016)
- Hendrikson, R., et al.: Using Gromov-Wasserstein distance to explore sets of networks. University of Tartu, Master Thesis **2** (2016)
- Hofer, C., Graf, F., Rieck, B., Niethammer, M., Kwitt, R.: Graph filtration learning. In: International Conference on Machine Learning, pp. 4314–4323. PMLR (2020)
- Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. In: Wallach, H., Larochelle, H., Beygelzimer, A., Alché, F., Fox, E., Garnett, R. (eds.) Advances in neural information processing systems, vol. 32. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/file/2d95666e2649fcfc6e3af75e09f5adb9-Paper.pdf> (2019)
- Hu, X., Wang, Y., Fuxin, L., Samaras, D., Chen, C.: Topology-aware segmentation using discrete morse theory. In: International Conference on Learning Representations (2020)
- Kim, W., Memoli, F.: Generalized persistence diagrams for persistence modules over posets. arXiv preprint [arXiv:1810.11517](https://arxiv.org/abs/1810.11517) (2018)
- Landi, C.: The rank invariant stability via interleavings. In: Chambers, E.W., Terese, F.B., Lori, Z. (eds.) Research in computational topology, pp. 1–10. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-89593-2_1 (2018)
- Lesnick, M.: The theory of the interleaving distance on multidimensional persistence modules. *Found. Comput. Math.* **15**(3), 613–650 (2015)

- Lesnick, M., Wright, M.: Interactive visualization of 2-d persistence modules. arXiv preprint [arXiv:1512.00180](https://arxiv.org/abs/1512.00180) (2015)
- Li, M., Palande, S., Wang, B.: Sketching merge trees. arXiv e-prints [arXiv:2101](https://arxiv.org/abs/2101) (2021)
- Mandell, M.A.: Cohains and homotopy type. *Publications Mathématiques de l'IHÉS* **103**, 213–246 (2006). <https://doi.org/10.1007/s10240-006-0037-6>. http://www.numdam.org/item/PMIHES_2006__103__213_0/
- Martínez, D.H.D., Lee, C.H., Kim, P.T., Mio, W.: Probing the geometry of data with diffusion fréchet functions. *Appl. Comput. Harmon. Anal.* **47**(3), 935–947 (2019)
- Mémoli, F.: On the use of Gromov-Hausdorff distances for shape comparison. *The Eurographics Association* (2007)
- Mémoli, F.: Gromov-Wasserstein distances and the metric approach to object matching. *Found. Comput. Math.* **11**(4), 417–487 (2011)
- Mémoli, F., Munk, A., Wan, Z., Weitkamp, C.: The ultrametric Gromov-Wasserstein distance. arXiv preprint [arXiv:2101.05756](https://arxiv.org/abs/2101.05756) (2021)
- Mémoli, F., Okutan, O.B.: Metric graph approximations of geodesic spaces. arXiv preprint [arXiv:1809.05566](https://arxiv.org/abs/1809.05566) (2018)
- Mémoli, F., Smith, Z., Wan, Z.: Gromov-Hausdorff distances on p-metric spaces and ultrametric spaces. arXiv preprint [arXiv:1912.00564](https://arxiv.org/abs/1912.00564) (2019)
- Morozov, D., Beketayev, K., Weber, G.: Interleaving distance between merge trees. *Discret. Comput. Geom.* **49**, 22–45 (2013)
- Oudot, S., Solomon, E.: Barcode embeddings for metric graphs. arXiv preprint [arXiv:1712.03630](https://arxiv.org/abs/1712.03630) (2017)
- Oudot, S., Solomon, E.: Inverse problems in topological persistence. In: *Topological Data Analysis*, pp. 405–433. Springer (2020)
- Patel, A.: Generalized persistence diagrams. *J. Appl. Comput. Topol.* **1**(3), 397–419 (2018)
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
- Perea, J.A., Harer, J.: Sliding windows and persistence: An application of topological methods to signal analysis. *Found. Comput. Math.* **15**(3), 799–838 (2015)
- Peyré, G., Cuturi, M., Solomon, J.: Gromov-wasserstein averaging of kernel and distance matrices. In: *International Conference on Machine Learning*, pp. 2664–2672 (2016)
- Reiss, A., Indlekofer, I., Schmidt, P., Van Laerhoven, K.: Deep ppg: large-scale heart rate estimation with convolutional neural networks. *Sensors* **19**(14), 3079 (2019)
- Riehl, E.: *Category theory in context*. Courier Dover Publications, New York (2017)
- Saul, N., Tralie, C.: Scikit-tda: Topological data analysis for python (2019). <https://doi.org/10.5281/zenodo.2533369>
- Scarpace, L., Mikkelsen, L., Cha, T., Rao, S., Tekchandani, S., Gutman, S., Pierce, D.: Radiology data from the cancer genome atlas glioblastoma multiforme [tcga-gbm] collection. *The Cancer Imaging Archive* **11**(4), 1 (2016)
- Smith, Z., Chowdhury, S., Mémoli, F.: Hierarchical representations of network data with optimal distortion bounds. In: *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp. 1834–1838. IEEE (2016)
- Stefanou, A.: *Dynamics on categories and applications*. Ph.D. Thesis (2018)
- Sturm, K.T.: On the geometry of metric measure spaces. *Acta Math.* **196**(1), 65–131 (2006)
- Sturm, K.T.: The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces. arXiv preprint [arXiv:1208.0434](https://arxiv.org/abs/1208.0434) (2012)
- Takens, F.: Detecting strange attractors in turbulence. In: *Dynamical Systems and Turbulence*, Warwick 1980, pp. 366–381. Springer (1981)
- Tauzin, G., Lupo, U., Tunstall, L., Pérez, J.B., Caorsi, M., Medina-Mardones, A., Dassatti, A., Hess, K.: giotto-tda: A topological data analysis toolkit for machine learning and data exploration. *J. Mach. Learn. Res.* **22**, 39–1 (2021)
- The GUDHI Project: GUDHI User and Reference Manual, 3.4.1 edn. GUDHI Editorial Board (2021). <https://gudhi.inria.fr/doc/3.4.1/>
- Turner, K., Mukherjee, S., Boyer, D.M.: Persistent homology transform for modeling shapes and surfaces. *Inf. Inf. A J. IMA* **3**(4), 310–344 (2014)
- Vayer, T., Chapel, L., Flamary, R., Tavenard, R., Courty, N.: Fused Gromov-Wasserstein distance for structured objects. *Algorithms* **13**(9), 212 (2020)

- Venkataraman, V., Ramamurthy, K.N., Turaga, P.: Persistent homology of attractors for action recognition. In: 2016 IEEE International Conference on Image Processing (ICIP), pp. 4150–4154. IEEE (2016)
- Villani, C.: *Optimal Transport: Old and New*, vol. 338. Springer, Berlin (2008)
- Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020). doi: <https://doi.org/10.1038/s41592-019-0686-2>
- Xu, H., Luo, D., Carin, L.: Scalable Gromov-Wasserstein learning for graph partitioning and matching. In: *Advances in Neural Information Processing Systems*, pp. 3052–3062 (2019)
- Xu, H., Luo, D., Henao, R., Shah, S., Carin, L.: Learning autoencoders with relational regularization. *arXiv preprint arXiv:2002.02913* (2020)
- Yan, L., Masood, T.B., Sridharamurthy, R., Rasheed, F., Natarajan, V., Hotz, I., Wang, B.: Scalar field comparison with topological descriptors: properties and applications for scientific visualization. *Comput. Graph. Forum* **40**(3), 599–633 (2021)
- Yan, L., Wang, Y., Munch, E., Gasparovic, E., Wang, B.: A structural average of labeled merge trees for uncertainty visualization. *IEEE Trans. Visual Comput. Graphics* **26**(1), 832–842 (2019)

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