Topologically Attributed Graphs for Shape Discrimination

Justin Curry [* 1] **Washington Mio** [* 2] **Tom Needham** [* 2] **Osman Berat Okutan** [* 3] **Florian Russold** [* 4]

Abstract

In this paper we introduce a novel family of attributed graphs for the purpose of shape discrimination. Our graphs typically arise from variations on the Mapper graph construction, which is an approximation of the Reeb graph for point cloud data. Our attributions enrich these constructions with (persistent) homology in ways that are provably stable, thereby recording extra topological information that is typically lost in these graph constructions. We provide experiments which illustrate the use of these invariants for shape representation and classification. In particular, we obtain competitive shape classification results when using our topologically attributed graphs as inputs to a simple graph neural network classifier.

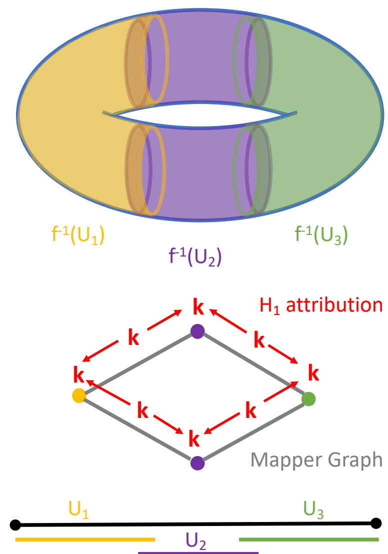

Figure 1. The Mapper graph construction applied to the torus admits a natural decoration with homology (with coefficients in a field k).

1. Introduction

Topological Data Analysis studies finite spaces by associating topological invariants to them that serve as intuitive structural summaries for unsupervised analysis or as nonlinear featurizations for downstream supervised learning applications. The most common such invariant is a persistence diagram, which, roughly, gives a concise representation of homological features that are apparent in the data at multiple scales. Graphical topological summaries form another important collection of tools for representing data; these include merge trees, Mapper graphs, and, their continuous counterparts, Reeb graphs.

In this paper, we combine Mapper graphs with persistence diagrams in order to define new, highly discriminative shape representations. The main ideas of these constructions are illustrated in Figures 1 and 2. Roughly, the Mapper graph gives a large-scale structural summary of connected components, while the persistence diagram attributions encode finer-scale topological structure.

The structure of the paper is as follows. In Sections 2 and 3, we introduce precise mathematical formalism for attributed graphs and their continuous analogues—decorated Reeb graphs—in the language of category theory. We then introduce novel constructions of topologically attributed graphs and prove their stability in Section 4. Sections 5 and 6 are devoted to computational considerations; in particular, how our topologically attributed graphs are constructed and compared in practice. We also provide a classification experiment, where we show that our constructions achieve competitive shape classification performance when they are fed as inputs into a simple graph neural network.

2. Categorically Attributed Graphs

We view a simple undirected graph $G = (V, E)$ as a category[1] **G**, with objects corresponding to elements of $V \cup E$

[*]Equal contribution [1]Department of Mathematics, SUNY Albany, Albany, NY, USA [2]Department of Mathematics, Florida State University, Tallahassee, FL, USA [3]Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany [4]Institute of Geometry, Graz University of Technology, Graz, Austria. Correspondence to: Tom Needham <tneedham@fsu.edu>.

Proceedings of the 2^{nd} Annual Workshop on Topology, Algebra, and Geometry in Machine Learning (TAG-ML) at the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. 2023.

[1]See (Riehl, 2017) for a good introduction to category theory.

and with a unique morphism $e \to v$ whenever a node v is incident to an edge e. This makes $\mathbf{G}$ equivalent to a poset $(G, \leq)$ where $e \leq v$.

Definition 2.1 (Attributed Graph)**.** An *attributed graph* is a functor $F\colon \mathbf{G} \to \mathbf{C}$ which assigns to each vertex v and each edge e of G objects $F(v)$ and $F(e)$ in $\mathbf{C}$, along with a morphism $F(e \leq v) : F(e) \to F(v)$ in $\mathbf{C}$, provided $e \leq v$.

Example 2.2. Suppose we are considering a social media platform, such as Facebook. Users correspond to nodes of a graph G and edges correspond to friendships. We can define an attribution valued in the category **Set** as follows: Let $F(v)$ to be the set of interests or pages that v follows and let $F(e)$ be the intersection of these interests or pages. The inclusion $F(e) \hookrightarrow F(v)$ makes this an attribution.

Attributed Graphs for Representing Shapes. In this paper, we are primarily interested in attributed graphs which capture aspects of the geometry and/or topology of a given space (or finite approximation thereof). As such, we will mostly work with attributions that come from homology[2], which is an attribution valued in **Vec**—the category of vector spaces and linear maps over a field k. A first example of the type of attributed graph we are interested in is as follows.

Example 2.3 (Decorated Mapper Graphs). Let X be a compact space and assume $f\colon X \to \mathbb{R}$ is a continuous map. Let $\mathcal{U} = (U_i)_{i \in I}$ be a cover of $\mathbb{R}$ with no (non-empty) triple intersections. We can pullback this cover along f to obtain $f^{-1}(\mathcal{U})$ as a cover of X, where each cover element $f^{-1}(U_i)$ is further refined into its connected components. The nerve of this cover defines the *Mapper graph* $\mathcal{M}_{\mathcal{U},f}$ of $X \xrightarrow{f} \mathbb{R}$ with respect to $\mathcal{U}$ (Singh et al., 2007). It has vertices V corresponding to components C of $f^{-1}(U_i)$ and edges $E = \{C \cap C' \mid C, C' \in f^{-1}(\mathcal{U}) \text{ and } C \cap C' \neq \emptyset\}$ corresponding to non-empty intersections of these components. The *decorated Mapper graph (DMG)* $F\colon \mathcal{M}_{\mathcal{U},f} \to \mathbf{Vec}$ augments the Mapper graph by assigning to each component $C \in V$ and $C \cap C' \in E$ the homology (with coefficients in a field k) of the corresponding components, i.e. $F(C) \coloneqq H_n(C)$ and $F(C \cap C') \coloneqq H_n(C \cap C')$. The inclusion $C \cap C' \subseteq C$ of components induces a map in homology $F(C \cap C' \leq C) \coloneqq H_n(C \cap C' \subseteq C)$.

An example of a DMG is shown in Figure 1.

Discrete Shape Representations and TDA. In our next example, we consider an extension of the DMG concept which applies to finite spaces. As a reminder, Topological Data Analysis (TDA) provides a tool for homology inference that replaces homology with *persistent homology*; we assume that the reader is familiar with the basic concepts of TDA, but review one key construction below.

Definition 2.4 (Rips Persistence)**.** Given a finite metric space (X, d_X) the *Vietoris-Rips complex at scale* r is the simplicial complex $VR(X, r)$ whose simplices consist of subsets $\sigma \subseteq X$ where $d_X(x, x') \leq 2r$ for all $x, x' \in \sigma$. Notice that if $r \leq s$, then there is an inclusion $VR(X, r) \subseteq VR(X, s)$. Passing to the geometric realization of these complexes and the induced continuous maps makes $VR(X) := VR(X, \bullet)$ into a functor from $(\mathbb{R}, \leq)$—the poset category of the reals—to **Top**—the category of topological spaces and maps, i.e. $VR(X)$ is an object in the functor category $\mathbf{Top}^{\mathbb{R}}$. Applying homology H_n then defines the Rips persistent homology $PH_n(X)$, which is an object in $\mathbf{Vec}^{\mathbb{R}}$. This latter object can then be faithfully encoded as a persistence diagram, which records births and deaths of homological features across scales.

[2]See (Hatcher, 2002) for a textbook treatment.

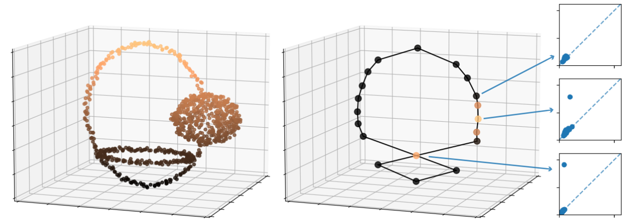

(a) (b) (c)

Figure 2. Persistent Decorated Mapper Graph. (a) A synthetic point cloud data set, nodes colored by the value of its filtration function, height along the z-axis. (b) A Mapper graph of the dataset. (c) The nodes of the Mapper graph are attributed with persistence diagrams; each node corresponds to a connected component of a level set of the dataset, and the (degree-1) persistent homology of this subset gives the attribution. Nodes of the Mapper graph are colored by total persistence (i.e., $\sum(d_{\mathrm{I}} - b_i)$, where the sum is over points (b_i, d_{I}) in the diagram).

Definition 2.5 (Persistent DMGs)**.** Given a finite metric space (X, d_X) and function $f : X \to \mathbb{R}$, we can construct the (discrete version of the) Mapper graph $\mathcal{M}_{\mathcal{U},f}$ in a manner similar to Example 2.3 by inferring components via a chosen clustering algorithm applied to $f^{-1}(U_i)$. The clusters then replace the components $C \in f^{-1}(U_i)$ in the construction above. This allows us to define a *persistent Decorated Mapper Graph* $F\colon \mathcal{M}_{\mathcal{U},f} \to \mathbf{Vec}^{\mathbb{R}}$ that assigns to each vertex and each edge—corresponding to a cluster and an intersection of clusters, respectively—the persistent homology of each, i.e. $F(C) \coloneqq PH_n(C)$ and $F(C \cap C' \leq C) \coloneqq PH_n(C \cap C' \subseteq C)$.

This structure is illustrated in Figure 2. In this figure, only the node attributions are included.

Persistent DMGs give intuitive and informative summaries of discrete shapes, and are the main object that we use in applications below (see Sections 5 and 6). The next part of the paper is concerned with establishing basic theory for these objects, focusing on stability.

3. Decorated Reeb Graphs

In this section, we introduce continuous versions of the Mapper graphs and attributions described above.

Reeb Graphs. In practice, choosing the cover and clustering schema for Mapper can be an art with sometimes hard to interpret and unstable behavior. These defects are then inherited by the decoration process. These issues have been mostly handled (Munch & Wang, 2016; Carriere & Oudot, 2018) by viewing Mapper graphs as discrete approximations of the Reeb graph (Reeb, 1946), which we review next.

Definition 3.1. A *Reeb graph* is a pair (R, f) consisting of a compact 1-dimensional geometric simplicial complex R and a piecewise linear map $f : R \to \mathbb{R}$. A metric d_f on R is defined by $d_f(x, x') = \inf_\gamma \max f \circ \gamma - \min f \circ \gamma$, where the infimum is over all paths from x to x'.

Example 3.2. Let X be a compact geometric simplicial complex and let $f : X \to \mathbb{R}$ be a continuous piecewise linear map. The Reeb graph associated to $f : X \to \mathbb{R}$ starts by defining R to be set of equivalence classes $X/\sim$, where $x \sim x'$ if x and x' lie in the same connected component of $f^{-1}(v)$. Since f is constant on equivalence classes, it factors to define a map $\hat{f} : R \to \mathbb{R}$ where $f = \hat{f} \circ q$ and q is the quotient map $q : X \to X/\sim$. The pair $(R, \hat{f})$ then defines a Reeb graph in the sense of Definition 3.1.

We now make geometric graphs the domain of attribution.

Definition 3.3. Let R be a compact 1D geometric complex and let $\mathcal{O}(R)$ be its poset category of open sets. A *continuous attribution* is a functor $F : \mathcal{O}(R) \to \mathbf{C}$.

Example 3.4 (Decorated Reeb Graph (DRG)). When (R, f) is a Reeb graph and $\mathbf{C} = \mathbf{Vec}$, we refer to a continuous attribution $F : \mathcal{O}(R) \to \mathbf{Vec}$ as a *decorated Reeb graph* or *DRG*. Specifically, let $(R, \hat{f})$ be a Reeb graph arising from the construction of Example 3.2. The *homology decorated Reeb graph* is the continuously attributed graph that assigns to each open set $U \subset R$ in the Reeb graph the homology $F(U) = H_n(q^{-1}(U))$.

Categorical Reeb Graphs. In order to prove that the Reeb graph is stable, (de Silva et al., 2016) used the following definition of a Reeb graph.

Definition 3.5. A *categorical Reeb graph* is a functor $\mathcal{R} : \mathcal{O}(\mathbb{R}) \to \mathbf{Set}$, where $\mathcal{O}(\mathbb{R})$ is the category of open subsets of $\mathbb{R}$ ordered by inclusion, that satisfies *constructibility*—there exists some finite collection of critical values $\tau = \{t_0, \dots, t_n\} \subset \mathbb{R}$ such that if $I \subseteq J$ are two intervals with equal intersection with τ, then the map $\mathcal{R}(I \subseteq J)$ is an isomorphism—and the *cosheaf axiom*—if $\mathcal{U} = \{U_i\}_{i \in I}$ is a cover of an open set $U \in \mathcal{O}(\mathbb{R})$ then the universal map from the colimit $\varinjlim \mathcal{R}(U_i) \to \mathcal{R}(U)$ is an isomorphism.

Example 3.6. Every Reeb graph (R, f) gives rise to a categorical Reeb graph $\mathcal{R}$ via $\mathcal{R}(U) := \pi_0(f^{-1}(U))$, where $\pi_0 : \mathbf{Top} \to \mathbf{Set}$ is the path components functor.

We now unify Definitions 3.3 and 3.5 to provide an alternative description of Example 3.4. This involves engineering a category that can track both components and homology vector spaces.

Definition 3.7. Let **PVec** denote the category of *discretely parameterized vector spaces*. Objects of **PVec** are functors $\sigma : S \to \mathbf{Vec}$, where S is a set regarded as a discrete category and a morphism from $\sigma : S \to \mathbf{Vec}$ to $\tau : T \to \mathbf{Vec}$ consists of a set map $\mu : S \to T$ and a natural transformation $\sigma \Rightarrow \tau \circ \mu$. This category has a functor $\mathbf{dom} : \mathbf{PVec} \to \mathbf{Set}$ that sends $\sigma : S \to \mathbf{Vec}$ to S.

Definition 3.8. A *categorical decorated Reeb graph* is a functor $\mathcal{F} : \mathcal{O}(\mathbb{R}) \to \mathbf{PVec}$ such that $\mathbf{dom} \circ \mathcal{F}$ satisfies the axioms of Definition 3.5.

Example 3.9. Let $F : \mathcal{O}(R) \to \mathbf{Vec}$ be a DRG. This gives rise to a categorical DRG $\mathcal{F} : \mathcal{O}(\mathbb{R}) \to \mathbf{PVec}$ where $\mathcal{F}(U)$ is the object of **PVec** that maps $\pi_0(\hat{f}^{-1}(U)) \to \mathbf{Vec}$ by taking a connected component of $\hat{f}^{-1}(U) \ni A \subset R$ to $F(A)$.

4. Persistent Decorations and Stability

One of the main contributions of TDA has been the observation that connected components and homology are stable only when considered as part of a family of topological spaces. We now review the concepts used to quantify this.

Metrics. Two of the most prominent distance metrics used in TDA are Gromov-Hausdorff distance and interleaving distance, which we now define.

Definition 4.1 (Gromov-Hausdorff Distance)**.** Let (X, d_X) and (Y, d_Y) be metric spaces. The *distortion* of a pair of (not necessarily continuous) maps $\Phi : X \to Y$ and $\Psi : Y \to X$ is the quantity dist(Φ, Ψ) defined by

$$\sup\{|d_X(x, x') - d_Y(y, y')| \mid (x, y), (x', y') \in C(\Phi, \Psi)\},$$

where

$$C(\Phi, \Psi) := \{(x, y) \in X \times Y \mid y = \Phi(x) \text{ or } x = \Psi(y)\}.$$

The *Gromov-Hausdorff distance* between X and Y is

$$d_{\mathrm{GH}}(X, Y) := \inf_{\Phi, \Psi} \frac{1}{2} \mathrm{dist}(\Phi, \Psi).$$

The stability results we are interested in are based on the interleaving construction of TDA.

Definition 4.2 (Interleaving Distance)**.** Let $\mathcal{P}$ be a poset, $\mathbf{C}$ a category and $\mathbf{C}^{\mathcal{P}}$ the functor category equipped with a notion of shifting/smoothing for any $\epsilon \geq 0$, i.e. $(\bullet)^\epsilon : \mathbf{C}^{\mathcal{P}} \to \mathbf{C}^{\mathcal{P}}$, is a functor that sends $F \mapsto F^\epsilon$ and this functor is equipped with a natural transformation $\eta^\epsilon : \mathrm{id}_{\mathbf{C}^{\mathcal{P}}} \Rightarrow (\bullet)^\epsilon$ that interacts in compatible ways[3]. We say that two objects $F, G \in \mathbf{C}^{\mathcal{P}}$ are *ϵ-interleaved* if there exist morphisms $\phi : F \to G^\epsilon$ and $\psi : G \to F^\epsilon$ such that $\eta_F^{2\epsilon} = \psi^\epsilon \circ \phi$ and $\eta_G^{2\epsilon} = \phi^\epsilon \circ \psi$. The *interleaving distance* between F and G is then defined as

$$d_{\mathrm{I}}(F, G) = \inf\{\epsilon \mid F \text{ and } G \text{ are } \epsilon\text{-interleaved}\}.$$

Example 4.3 (Rips Persistence). If we choose $\mathcal{P} = (\mathbb{R}, \leq)$ and $\mathbf{C} = \mathbf{Vec}$ in Definition 4.2 we obtain the usual interleaving distance for 1-parameter persistence modules. Rips persistent homology for a finite metric space X (see Definition 2.4), written $PH_n(X) \in \mathbf{Vec}^{\mathbb{R}}$, has a natural notion of shifting by defining $PH_n(X)^\epsilon(r) = H_n(VR(X)(r + \epsilon))$. The fact that $VR(X)$ is a functor provides a map from $VR(X, r) \to VR(X, r + \epsilon)$, which gives the data of the natural transformation η^ϵ. Thus the interleaving distance between Rips persistent homology functors is well-defined.

Interleavings between Rips persistent homology leads to a foundational stability result of TDA (Chazal et al., 2009; 2014): for finite metric spaces X and Y

$$d_{\mathrm{I}}(PH_n(X), PH_n(Y)) \leq d_{\mathrm{GH}}(X, Y).$$

The type of stability we're interested in is not only governed by the Gromov-Hausdorff distance between point clouds, but scalar functions on these. This is expressed in the following definition, which is equivalent to a metric used in (Chazal et al., 2009); see also (Bauer et al., 2014), where a similar metric is used in the context of Reeb graphs.

Definition 4.4. Let X and Y be metric spaces equipped with functions $f : X \to \mathbb{R}$ and $g : Y \to \mathbb{R}$, written X_f and Y_g, respectively. If $\Phi : X \to Y$ and $\Psi : Y \to X$ are maps, then the functional distortion of Φ and Ψ is

$$\mathrm{FunDist}\big(\Phi, \Psi\big) := \max \begin{cases} \frac{1}{2}\mathrm{dist}(\Phi, \Psi) \\ ||f - g \circ \Phi||_\infty \\ ||g - f \circ \Psi||_\infty. \end{cases}$$

The functional distortion distance is then

$$d_{\mathrm{FD}}(X_f, Y_g) := \inf_{\Phi, \Psi} \mathrm{FunDist}\big(\Phi, \Psi\big).$$

It is straightforward to show that d_{FD} is a pseudometric on the space of pairs (X, f).

Stability of Persistent Discrete DRGs. We now define a persistent discrete Decorated Reeb Graph construction, which refines the notion of a persistent DMG (Definition 2.5), and will be stable under perturbations of the functional distortion distance.

[3]This is called a flow structure on the category $\mathbf{C}^{\mathcal{P}}$, whose full details are explored in (Stefanou, 2018; De Silva et al., 2018).

Definition 4.5 (Persistent Discrete DRG)**.** Given a finite metric space endowed with a scalar-valued function X_f, we define $f_r^{-1}(U)$ to be the full subcomplex of $VR(X, r)$ on all vertices $x \in f^{-1}(U)$, for each open subset $U \subset \mathbb{R}$. We then define the *persistent (discretized) decorated Reeb graph* of X_f to be the following 2-parameter family of categorical DRGs (Definition 3.8):

$$DF : (\mathbb{R}^2, \leq) \to \mathbf{Fun}(\mathcal{O}(\mathbb{R}), \mathbf{PVec}) \qquad (r, s) \mapsto \mathcal{F}(r, s)$$

where

$$\mathcal{F}(r, s)(U) := \{A \in \pi_0(f_r^{-1}(U)) \to H_\bullet(VR(A, s))\}.$$

Remark 4.6. For a fixed $r \geq 0$, $\mathcal{F}(r, s)$ assigns to each connected component of $f_r^{-1}(U)$ the Vietoris-Rips persistent homology of that point cloud at scale s. Then $DF(r, \bullet)$ can be considered as a DRG $DF(r, \bullet) : \mathcal{O}(\mathbb{R}) \to \mathbf{PVec}^{\mathbb{R}}$, by setting $DF(r, \bullet)(U)(s) = DF(r, s)(U)$. If we also fix a cover $\mathcal{U}$, we recover the persistent DMG of Definition 2.5 by choosing clusters associated to $U \in \mathcal{U}$ to be given by the connected components of $f_r^{-1}(U)$.

In the above sense, the persistent discrete DRG refines the notion of a persistent DMG. This relaxation to a more continuous and categorical setting is crucial to our proof of the stability result below.

Theorem 4.7 (Stability of Persistent Discrete DRGs)**.** *Let X_f and Y_g be finite metric spaces endowed with scalar-valued functions. Let DF and DG be their respective persistent discrete DRGs (Definition 4.5). Then we have*

$$d_{\mathrm{I}}(DF, DG) \leq d_{\mathrm{FD}}(X_f, Y_g).$$

The ϵ-smoothing of DF is defined by

$$DF^\epsilon(r, s)(U) = DF(r + \epsilon, s + \epsilon)(U^\epsilon)$$

where $U^\epsilon := \{t \in \mathbb{R} \mid \exists v \in U \text{ s.t. } |t - v| < \epsilon\}$ is the ϵ-thickening of the open set $U \in \mathcal{O}(\mathbb{R})$. This leads to the notion of the interleaving distance d_{I} used in the theorem.

Proof Sketch. Suppose the functional distortion distance of Definition 4.4 between X_f and Y_g is less than δ. This means that for every $\epsilon > \delta$ there are maps $\Phi : X \to Y$ and $\Psi : Y \to X$ whose distortion is less than 2ϵ. Also, $||f - g \circ \Phi||_\infty \leq \epsilon$, which implies that $\forall U \in \mathcal{O}(\mathbb{R})$ we have

$$f^{-1}(U) \subseteq \Phi^{-1}(g^{-1}(U^\epsilon)).$$

This implies that if $\sigma \subseteq f^{-1}(U)$ is a subset with $d(x_i, x_j) \leq 2r$ for all pairs of points in σ, i.e. $\sigma \in VR(X, r) \cap f_r^{-1}(U)$, then $\Phi_{r,s}(\sigma) \in g_{r+\epsilon}^{-1}(U^\epsilon) \cap VR(Y, s + \epsilon)$. Moreover, this containment holds when restricted to a component $A \in \pi_0(f_r^{-1}(U))$. Symmetric reasoning using the condition $||g - f \circ \Psi||_\infty \leq \epsilon$ guarantees that $\forall U \in \mathcal{O}(\mathbb{R})$

$$g^{-1}(U) \subseteq \Psi^{-1}(f^{-1}(U^\epsilon))$$

and in particular $\Phi_{r,s}(\sigma) \in g_{r+\epsilon}^{-1}(U^\epsilon) \cap VR(Y, s+\epsilon)$ is carried to a simplex $\Psi_{r+\epsilon,s+\epsilon} \circ \Phi(\sigma) \in f_{r+2\epsilon}^{-1}(U^{2\epsilon}) \cap VR(X, s+2\epsilon)$ that is contiguous to σ inside $f_{r+2\epsilon}^{-1}(U^{2\epsilon}) \cap VR(X, s+2\epsilon)$, thus guaranteeing that the induced map on homology $VR(A, s) \to VR(A, s+\epsilon)$ for each component $A \in \pi_0(f_r^{-1}(U))$ is the same as $\Psi_{r+\epsilon,s+\epsilon} \circ \Phi_{r,s}$. This establishes half of the interleaving condition and the other half is argued *mutatis mutandi*. □

Stability of Barcode Transforms. We end this theoretical section with another stability result, which deals more directly with DRGs. While the constructions involved are somewhat more straightforward, we discuss their limitations in practice at the end of the section.

Definition 4.8 (Barcode Transform)**.** Let $F\colon \mathcal{O}(R) \to \mathbf{Vec}$ be the decorated Reeb graph associated to $f : X \to \mathbb{R}$, where each open set $U \subseteq R$ is assigned a finite-dimensional vector space. We define the *barcode transform* of F to be the map

$$\begin{aligned} BF\colon R &\to \mathbf{Vec}^{\mathbb{R}} \\ r \in R &\mapsto \Big(t \in \mathbb{R}_{\geq 0} \mapsto F\big(B_{d_f}(r,t)\big)\Big) \end{aligned}$$

Since every persistence module can be identified with a barcode, we can view the barcode transform as an assignment of a barcode to each point in the Reeb graph.

Using ϵ-smoothings of open sets, i.e. setting $\mathcal{P} = \mathcal{O}(\mathbb{R})$ and $F^\epsilon(U) := F(U^\epsilon)$ in Definition 4.2, we can define interleaving distances for categorical Reeb graphs and categorical decorated Reeb graphs. Moreover, the interleaving distance of categorical Reeb graphs gives rise to an interleaving distance of concrete Reeb graphs as defined in (de Silva et al., 2016). In the following we define the functional distortion distance for barcode transforms and show that it is controlled by the interleaving distance of the Reeb graphs and their corresponding categorical decorated Reeb graphs.

Definition 4.9. Let $F\colon \mathcal{O}(R) \to \mathbf{Vec}$ and $G\colon \mathcal{O}(S) \to \mathbf{Vec}$ be concrete decorated Reeb graphs over (R, f) and (S, g). We define the functional distortion distance of the corresponding barcode transforms by

$$d_{\mathrm{FD}}\big(BF, BG\big) := \inf_{\Phi,\Psi} \max \begin{cases} \mathrm{FunDist}(\Phi, \Psi) \\ \sup_{r\in R} d_{\mathrm{I}}\big(BF(r), BG \circ \Phi(r)\big) \\ \sup_{s\in S} d_{\mathrm{I}}\big(BF \circ \Psi(s), BG(s)\big) \end{cases}$$

where FunDist is taken w.r.t. f and g and the infimum is over all functions $\Phi\colon R \to S$ and $\Psi\colon S \to R$.

Theorem 4.10. *Let* $F\colon \mathcal{O}(R) \to \mathbf{Vec}$ *and* $G\colon \mathcal{O}(S) \to \mathbf{Vec}$ *be concrete decorated Reeb graphs over* (R, f) *and* (S, g) *and* $\mathcal{F}, \mathcal{G}\colon \mathcal{O}(\mathbb{R}) \to \mathbf{PVec}$ *the corresponding categorical decorated Reeb graphs, then*

$$d_{\mathrm{I}}\big(R, S\big) \leq d_{\mathrm{FD}}\big(BF, BG\big) \leq 6 d_{\mathrm{I}}\big(\mathcal{F}, \mathcal{G}\big).$$

Proof Sketch. We begin with the inequality on the left. As shown in (Bauer et al., 2015), $d_{\mathrm{I}}\big(R, S\big) \leq d_{\mathrm{FD}}\big(R_f, S_g\big)$ and, since the functional distortion distance on Reeb graphs corresponds to the first part of the functional distortion distance of barcode transforms (Definition 4.9), we obviously get $d_{\mathrm{FD}}\big(R_f, S_g\big) \leq d_{\mathrm{FD}}\big(BF, BG\big)$.

To demonstrate the inequality on the right, let $\mathbf{dom}\colon \mathbf{PVec} \to \mathbf{Set}$ be the functor that sends functors in $\mathbf{PVec}$ to its domain. We observe that $\mathbf{dom}\mathcal{F} = \mathcal{R}$ the categorical Reeb graph of (R, f). Hence, given an ϵ-interleaving between $\mathcal{F}$ and $\mathcal{G}$, applying $\mathbf{dom}$ yields an ϵ-interleaving between $\mathcal{R}$ and $\mathcal{S}$ and, furthermore, an ϵ-interleaving between (R, f) and (S, g). As shown in (Bauer et al., 2015), $d_{\mathrm{FD}}(R_f, S_g) \leq 3 d_{\mathrm{I}}(R, S)$. One can now check that the functions $\Phi\colon R \to S$ and $\Psi\colon S \to R$ constructed in the proof of this inequality satisfy $d_{\mathrm{I}}\big(BF(r), BG \circ \Phi(r)\big) \leq 6\epsilon$ and $d_{\mathrm{I}}\big(BF \circ \Psi(s), BG(s)\big) \leq 6\epsilon$ for all $r \in R$ and for all $s \in S$. □

The details of the last part of the proof are quite technical, and we provide full details in the Appendix.

Remark 4.11. We remark that this result is interesting from a theoretical perspective, but has some shortcomings in practice. In particular, the functional distortion distance used here is infinite if the ranks of $BF(r)$ and $BG(r)$ do not agree for all sufficiently large r.

5. Computation

We now describe constructions of attributed graphs from point cloud data. In the following, we generically refer to such attributed graphs as Decorated Reeb Graphs (DRGs).

Creating Reeb Graphs. Reeb graphs are most naturally defined for continuous metric spaces, so one needs to approximate a Reeb graph structure for discrete data. We provide a construction similar to the Mapper algorithm (Singh et al., 2007) for estimating Reeb graphs. We first fix a scale r for the Vietoris-Rips complex $VR(X, r)$. Choosing an appropriate value of r is treated as a hyperparameter tuning process; similar ideas for Reeb graph estimation go back at least to (Ge et al., 2011). For the shape datasets considered in this paper, we used a simple heuristic which took $r = m \cdot r_0$, where r_0 is the smallest scale at which the VR complex is connected and m is a small integer (we typically took $m = 2$ or 3). Next, we choose a resolution parameter n and uniformly subdivided the image of f into n bins, $U_1, \ldots, U_n$ (we treat this as a partition of the range, but one could instead thicken slightly and work with an open cover, similar to the usual Mapper construction). This is used to approximate the Reeb graph G of $(VR(X, r), f)$: each node v of G corresponds to a connected component $A_v \subset X$ of one of the sets $f^{-1}(U_i)$, and there is an edge between nodes v and w if $A_v \subset f^{-1}(U_i)$, $A_w \subset f^{-1}(U_{i+1})$ and A_v and A_w are connected by an edge in the 1-skeleton of $VR(X,r)$. An alternative approach would be to compute the exact Reeb graph for $VR(X,r)$ via algorithms of (Harvey et al., 2010) or (Parsa, 2012); these algorithms are very efficient, but we found that they did not scale well in the Vietoris-Rips setting due to blowup in the size of the simplicial complexes.

Creating DRGs - Local Approach. Let $G=(V,E)$ be an estimated Reeb graph for $VR(X,r)$. A simple approach to adding persistent homology decorations to G is as follows. In this *local approach*, degree-n persistent homology is computed for the subset of X corresponding to each node in G. The resulting data structure (G,D) consisting of a finite graph $G=(V,E)$ and an attribution function $D: V \to \mathbb{R} \times$ **Barcodes**, where **Barcodes** is the set of (persistent homology) barcodes. The attribution takes a node $v \in V$ to $D(v)=(\bar{f}(v), B(v))$, where $\bar{f}(v)=\frac{1}{|A_v|}\sum_{x\in A_v} f(x)$ and $B(v)$ is the persistent homology barcode of A_v.

This method for constructing a DRG can be seen as an approximation of a particular slice of the persistent discrete DRG structure $DF : (\mathbb{R}^2, \leq) \to \mathbf{Fun}(\mathcal{O}(\mathbb{R}), \mathbf{PVec})$ introduced in Definition 4.5, as we observed in Remark 4.6.

Creating DRGs - Barcode Transform Approach. The following is an alternative approach to adding persistent homology decorations to $G=(V,E)$, an estimated Reeb graph for $VR(X,r)$. For each $v \in V$, we define a filtration on $VR(X,r)$ by distance to the set A_v and compute the degree-n persistent homology of the resulting filtered simplicial complex. This once again results in a data structure of the form (G,D) with the attribution function $D: V \to \mathbb{R} \times$ **Barcodes** now recording average function value and persisent homology of the distance-to-A_v function. This method for constructing the DRG is a simplification of the true barcode transform for the decorated Reeb graph of $VR(X,r)$ (see Definition 4.8).

Remark 5.1. The local DRG algorithm easily scales to handle datasets with thousands of points and provides intuitive data summaries consisting of an approximation of the Reeb graph skeleton, attributed with persistence diagrams encoding local structural information (see Figures 2 and 3). However, we note that this representation is an aggressive simplification of the structure described in Definition 4.5, so that the theoretical stability result of Theorem 4.7 does not directly apply in this setting. On the other hand, the barcode transform approach to constructing DRGs gives an approximation of the true barcode transform of $VR(X,r)$, and is therefore much more closely tied to theory. This construction requires several computations of persistent homology on the full complex $VR(X,r)$ (endowed with different filtrations), so that it is not scalable to large datasets. As such, most of our computational experiments below will focus on the local DRG construction.

Comparing DRGs. Since the interleaving distance considered in Theorem 4.7 is not applicable to our data representation and is computationally intractable, we use the *Fused Gromov-Wasserstein (FGW) framework* (Vayer et al., 2020) for metric-based analysis of DRGs. Intuitively, the FGW distance, defined below, is a more easily approximable proxy for the functional distortion distance d_{FD}.

Let (G_1,D_1), (G_2,D_2) be DRGs constructed as described above (via either the local or barcode transform approaches). The *α-FGW distance* is defined by

$$d_{\mathrm{FGW},\alpha}((G_1,D_1),(G_2,D_2))^2 = \inf_{\pi\in\mathcal{C}(V_1,V_2)} \alpha\cdot L_{\mathrm{gr}}(\pi)+(1-\alpha)\cdot L_{\mathrm{bc}}(\pi) \tag{1}$$

where the set $\mathcal{C}(G_1,G_2)$ and the loss functions L_{gr} and L_{bc} are defined as follows. An element $\pi \in \mathcal{C}(G_1,G_2)$ is a matrix of size $|V_1|\times|V_2|$ satisfying $\sum_{v_1\in V_1}\pi(v_1,v_2)=\frac{1}{|V_2|}$ for all $v_2 \in V_2$ and $\sum_{v_2\in V_2}\pi(v_1,v_2) = \frac{1}{|V_1|}$ for all $v_1 \in V_1$—intuitively, this is the space of probabilistic couplings of the uniform measures on V_1 and V_2, respectively. The *graph loss* $L_{\mathrm{gr}}(\pi)$ is defined by

$$\sum (d_1(v_1,w_1)-d_2(v_2,w_2))^2\,\pi(v_1,v_2)\pi(w_1,w_2), \tag{2}$$

where the sum is over all $v_i, w_i \in V_i$ and $d_{\mathrm{I}}: V_i\times V_i \to \mathbb{R}$ is a choice of function representing graph structure—for example, we typically use the shortest path distance, where each edge (v_i,w_i) in E_i is weighted by $|\overline{f}_i(v_i)-\overline{f}_i(w_i)|$. Finally, the *barcode loss* $L_{\mathrm{bc}}(\pi)$ is defined by

$$\sum_{v_i\in V_i} d_b(D_1(v_1),D_2(v_2))^2\pi(v_1,v_2), \tag{3}$$

where d_b is the standard bottleneck distance between barcodes. The intuition for the distance is as follows: $\pi \in \mathcal{C}(V_1,V_2)$ is interpreted as a probabilistic registration of the nodes of G_1 and G_2, the loss L_{gr} measures how well the registration preserves the graph structure, the loss L_{bc} measures how well the registration preserves attributions, and the hyperparameter α balances contributions of graph structure and attributions; the optimization problem therefore searches for a probabilistic registration which incurs the least total distortion of these structures. Fixing a methodology for assigning a distance graph function $d: V\times V\to\mathbb{R}$ to a DRG (G,D), it follows from Theorem 1 of (Vayer et al., 2020) that $d_{\mathrm{FGW},\alpha}$ defines a pseudometric on the space of DRGs, for any choice of $\alpha\in[0,1]$.

The idea for FGW distance originates from the Gromov-Wasserstein (GW) distances introduced by Mémoli in (Mémoli, 2007); roughly, the GW 2-distance is obtained by setting $\alpha = 1$ in the FGW formula. The GW distances can be seen as L^p-relaxations of the Gromov-Hausdorff (GH) distance and are used for the comparison of metric measure spaces (mm-spaces). The FGW distance was introduced in (Vayer et al., 2020) to adapt GW distance to the setting of mm-spaces whose points come with feature attributions. In particular, our use of $d_{\mathrm{FGW},\alpha}$ can be seen as a relaxation of the functional distortion distance of Definition 4.9. Such a relaxation makes the discrete optimization problem arising in GH distances amenable to approximation by gradient descent, as the search space of the optimization becomes a compact, convex polytope $\mathcal{C}(V_1, V_2) \subset \mathbb{R}^{|V_1|\times|V_2|}$. For finite mm-spaces of size $O(n)$, a gradient descent iteration for approximation of GW distance has $O(n^3 \log(n))$ cost (Peyré et al., 2016). For the FGW distance in our setting, computation is complicated by the need to evaluate $|V_1| \times |V_2|$ bottleneck distances, each of which carries a cubic cost in the number of points in the diagrams. To ease this computational burden, we frequently replace the bottleneck distance computations with Euclidean distance between persistence image vectorizations of the diagrams (Adams et al., 2017). We remark that (Fused) Gromov-Wasserstein distances have been used successfully in several recent works to compare other topological invariants—see, e.g., (Li et al., 2021; Curry et al., 2022; Li et al., 2023).

Related Constructions. Our constructions have a similar flavor to other enriched topological invariants in the literature. Most similar are the Decorated Merge Trees (DMTs) introduced by the first four authors in (Curry et al., 2022). A DMT is a certain rooted tree attributed with homological information which can be extracted from a dataset endowed with a filter function—roughly, a DMT captures the topology of sublevel sets of the filtration, while the constructions in the present paper are concerned with level set topology. Our constructions also share features of Persistent Homology Transforms (PHTs) (Turner et al., 2014), which associate a collection of persistence diagrams to a shape in Euclidean space by using projections to various 1-dimensional subspaces as filter functions. Our constructions also yield families of persistence diagrams, but the families are parameterized by nodes of Mapper graphs, rather than by collections of lines.

6. Experiments

We now illustrate our computational pipeline with several experiments. Our source code is available at our GitHub repository[4].

[4]https://github.com/trneedham/Topologically-Attributed-Graphs

Example DRGs. In Figure 3, we provide a few examples of DRGs constructed via the methods described in Section 5.

The first row of Figure 3 shows a shape from ModelNet10 (Wu et al., 2015), a curated collection of CAD models of household objects. The CAD model has been converted to point cloud data by sampling. We show the DRG computed via the local approach, with respect to height along the z-axis; nodes of the DRG are colored by total persistence of their diagram attributes. We show the diagrams associated to two of the nodes, as well as the associated subsets of the original dataset.

The second row of the figure shows a humanoid figure from the SHREC14 dataset (Pickup et al., 2014); once again, the pointcloud data is obtained by sampling a triangulated surface. For this example, the filtration function is the p-eccentricity (with $p = 100$) (see (Mémoli, 2011), Definition 5.3) of the shortest path distance on the 1-skeleton of the underlying Vietoris-Rips complex. We show the associated DRG (via the local approach), some persistence diagram attributes, and the persistence image vectorizations of these diagrams.

Finally, the third row of Figure 3 shows the synthetic dataset from Figure 2, once again endowed with the height function. In this case, the DRG is computed via the Barcode Transform approach. Observe that the associated persistence diagrams capture the global topology of the shape—note that the death time of each point is ∞ and that one point in each diagram is of multiplicity 2. The difference between the diagrams is the birth times of the features; the DRG nodes are colored by average birth time in their diagrams.

Synthetic Shape Comparison. To explore the behavior of the Fused Gromov-Wasserstein (FGW) distance $d_{\mathrm{FGW},\alpha}$ defined in (1), we consider a synthetic point cloud dataset consisting of four classes: torus, solid torus, cylinder and solid cylinder. Each torus shape consists of 400 points sampled from a toroidal surface with minor radius 1 and major radius 6 and solid tori are generated similarly, but we take 1600 samples to get a comparable density. Cylinders consist of 400 points sampled from the surface of a cylinder of radius 1 and length $2 \cdot \pi \cdot (6 + 1)$ so that the surface area is comparable to the torus and solid cylinders are generated similarly with 1600 sample points. Each shape in the dataset has its point coordinates perturbed independently at random and the resulting point cloud is then randomly rigidly rotated. The full dataset consists of 20 samples of each shape class.

Each shape is converted to a DRG using the local approach, with filter functions given by 1st PCA coordinates. Node attributes are converted to persistence images, for computational efficiency. For each $\alpha \in \{0.0, 0.25, 0.5, 1.0\}$, we construct the shape-to-shape distance matrix with respect

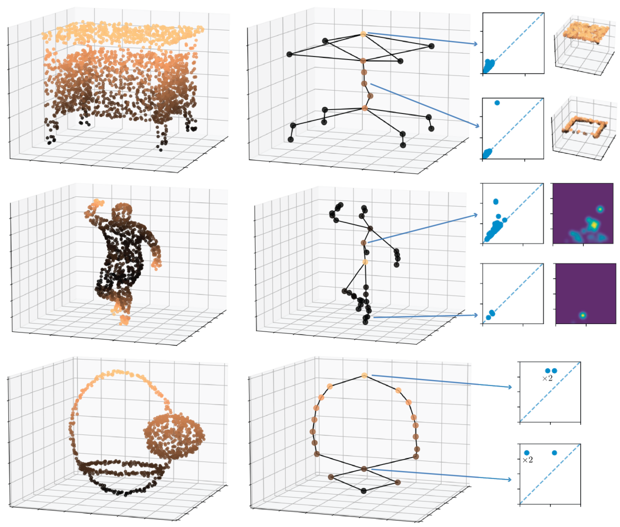

Figure 3. Examples of DRGs created with our methods. See Section 6 for a detailed description.

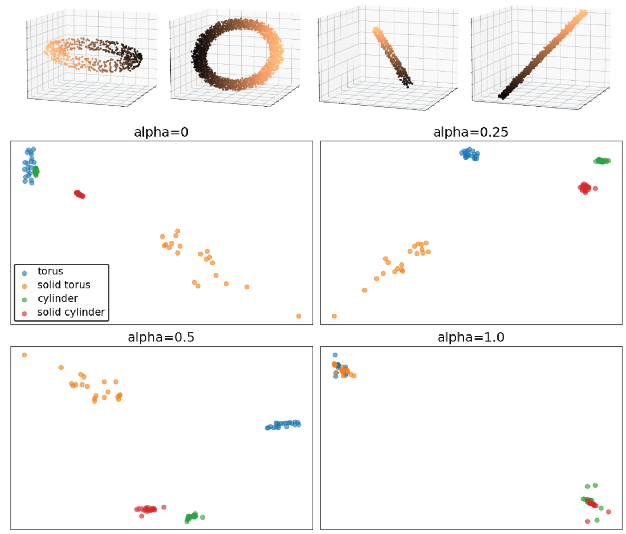

Figure 4. Synthetic shape dataset results. The top row shows a sample from each of the shape classes in the experiment; points are colored by their filter function values (first PCA coordinate). The remaining figures show Multidimensional Scaling embeddings of the dataset, coming from pairwise distance matrices with respect to $d_{\mathrm{FGW},\alpha}$ for various α-values.

to $d_{\mathrm{FGW},\alpha}$, where the d_{I}'s in (2) are shortest path distances in the Reeb graph and d_b in (3) is replaced by Euclidean distance between persistence images.

Multidimensional Scaling plots for the distance matrices are shown in 4. When $\alpha = 0$, the distance only sees the persistence image attributes, and the torus and cylinder classes are indistinguishable (the topology of each connected component of each of their level sets are very similar). For $\alpha \in \{0.25, 0.5\}$, both the global structure of the Reeb graph and the local persistence image structures are considered, and all classes are able to be distinguished. Finally, when $\alpha = 1$, only the Reeb graph structures are considered in the metric; here, the torus and solid torus classes are unable to be distinguished, and the same is true of the cylinder and solid cylinder classes. This experiment suggests that there is a fairly robust range of α-values where the metric meaningfully takes into account both Reeb graph and local persistence structures for distinguishing shapes.

We remark that we ran the same experiment using DRGs constructed via the Barcode Transform method; here the results were much less interpretable. This is due to the fact that the distance between Barcode Transform DRGs is strongly controlled by the homotopy types of the shapes (cf. Remark 4.11); it was difficult to tune parameters so that the underlying Vietoris-Rips complexes used in the construction consistently had the correct homotopy type.

Graph Neural Network-Based Classification. To more thoroughly test the capability of DRGs to distinguish shapes, we use DRGs as inputs to a Graph Neural Network (GNN) classifier. Our data comes from the ModelNet10 dataset (Wu et al., 2015) of CAD models of household objects. The data consists of 10 classes of objects, pre-partitioned into a training set (3991 objects) and a testing set (908 objects). We sample 1024 points form each object to form a dataset of point clouds. From each point cloud, we extract two DRGs: one using z-coordinates as the filtration function and the other using x-coordinates (see the first row of Figure 3). The nodes of the DRGs are attributed with vectorizations of the persistence diagrams—in this case, we used summary statistics of the persistence diagrams, following (Ali et al., 2022), Definition 2.1. We also attributed each node with the average Euclidean position of its associated pointcloud. These vector-attributed graphs were used as input to a simple GNN, consisting of four convolutional layers of width 256, implemented in PyTorch. All neural networks in this experiment were trained on a single CPU.

Table 1. ModelNet10 Classification Results.

PointNet	Reeb	Dgms	DRG_z	DRG_x	$\mathrm{DRG}_{\mathrm{xz}}$
89.43	77.64	63.55	84.69	85.24	87.11

Results of the classification experiment are reported in Table 1. Besides results for the DRGs with filter function given by z-coordinate (DRG_z) and x-coordinate (DRG_x), we also report the combined prediction from the two models (DRG_{xz}). The combined prediction was made by averaging the predictions of the DRG_z and DRG_x models—we also trained a GNN using a disjoint pair of DRGs (with x and z-coordinate filtrations) to represent each shape and got similar classification scores. To test the contributions of the diagram attributes, we also trained a network on graphs where nodes were only attributed with average Euclidean coordinates (Reeb). Likewise, we tested the contribution of the graph structure by converting each DRG into a complete graph on its nodes, each of which is attributed with persistence statistics and Euclidean coordinates (Dgms). We see that the combination of graph structure and topological attributes provide a large boost in classification accuracy, with the best accuracy obtained by the combination of x and z-filtrations.

To test against a baseline, we use the popular PointNet architecture for point cloud classification (Qi et al., 2017). The PointNet classification accuracy is essentially state-of-the-art for ModelNet10 classification (when using only point clouds as input, without additional structure from the CAD models), and we see that it has a slight edge over the DRG classification score. However, we note that the PointNet model[5] contains 3,463,763 parameters, compared to our 209,162 parameter GNN. Moreover, achieving this level of accuracy took $\sim$12 hours of training time for PointNet, while our GNN model contained an order of trained in $\sim$10 minutes; we preprocessed the data to extract Reeb graphs, which took $\sim$1.5 hours, bringing the total time for processing and training to around $\sim$3 hours (processing using both x and z filtrations). This suggests that the DRG representations have a rich structure with easily learnable features.

Table 2. ModelNet10 Subset Classification Results.

PointNet	DRG_z	DRG_x	$\mathrm{DRG}_{\mathrm{xz}}$
83.26	78.52	70.70	82.93

We tested the representational richness of DRGs further by retraining the DRG and PointNet models on only 10% of the ModelNet10 training data, then testing classification accuracy on the full training set. In this sparse training data setting, we see that all models still perform reasonably well, but point out that the gap between PointNet and DRG_{xz} has essentially vanished, even though the latter is less complex by an order of magnitude.

7. Discussion

In this paper, we introduced formalism for topologically attributed graphs and provided theoretical results on their stability. We also demonstrated the potential applicability of these ideas through proof-of-concept experiments. Future work will involve building a closer connection between the theory of topologically attributed Reeb graphs and their computational execution. Notably, our computational pipeline does not incorporate the more sheaf-theoretic or categorical features of decorated Reeb graphs, and integration of these aspects is an important goal. We also plan to continue to develop the computational pipeline toward more robust applications. One interesting direction will be to develop the pipeline to handle more general filtration functions, or to incorporate discovery of effective filter functions into a machine learning framework. We also intend to extend this framework to handle more general attributed simplicial complexes, to which newly developed tools of Topological Deep Learning (Hajij et al., 2023) will apply.

[5] A PyTorch implementation, following `https://github.com/nikitakaraevv/pointnet`.

References

Adams, H., Emerson, T., Kirby, M., Neville, R., Peterson, C., Shipman, P., Chepushtanova, S., Hanson, E., Motta, F., and Ziegelmeier, L. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18, 2017.

Ali, D., Asaad, A., Jimenez, M.-J., Nanda, V., Paluzo-Hidalgo, E., and Soriano-Trigueros, M. A survey of vectorization methods in topological data analysis. *arXiv preprint arXiv:2212.09703*, 2022.

Bauer, U., Ge, X., and Wang, Y. Measuring distance between Reeb graphs. In *Proceedings of the thirtieth annual symposium on Computational geometry*, pp. 464–473, 2014.

Bauer, U., Munch, E., and Wang, Y. Strong Equivalence of the Interleaving and Functional Distortion Metrics for Reeb Graphs. In Arge, L. and Pach, J. (eds.), *31st International Symposium on Computational Geometry (SoCG 2015)*, volume 34 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 461–475, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-939897-83-5. doi: 10.4230/LIPIcs.SOCG.2015.461. URL `http://drops.dagstuhl.de/opus/volltexte/2015/5146`.

Carriere, M. and Oudot, S. Structure and stability of the one-dimensional Mapper. *Foundations of Computational Mathematics*, 18:1333–1396, 2018.

Chazal, F., Cohen-Steiner, D., Guibas, L. J., Mémoli, F., and Oudot, S. Y. Gromov-Hausdorff stable signatures for shapes using persistence. In *Computer Graphics Forum*, volume 28, pp. 1393–1403. Wiley Online Library, 2009.

Chazal, F., De Silva, V., and Oudot, S. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1): 193–214, 2014.

Curry, J., Hang, H., Mio, W., Needham, T., and Okutan, O. B. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, 6(3): 371–428, 2022.

de Silva, V., Munch, E., and Patel, A. Categorified Reeb graphs. *Discrete & Computational Geometry*, 55(4):854–906, 2016. doi: 10.1007/s00454-016-9763-9. URL https://doi.org/10.1007/s00454-016-9763-9.

De Silva, V., Munch, E., and Stefanou, A. Theory of interleavings on categories with a flow. *Theory and Applications of Categories*, 33(21):583–607, 2018.

Ge, X., Safa, I., Belkin, M., and Wang, Y. Data skeletonization via Reeb graphs. *Advances in Neural Information Processing Systems*, 24, 2011.

Hajij, M., Zamzmi, G., Papamarkou, T., Miolane, N., Guzmán-Sáenz, A., Ramamurthy, K. N., Birdal, T., Dey, T. K., Mukherjee, S., Samaga, S. N., Livesay, N., Walters, R., Rosen, P., and Schaub, M. T. Topological deep learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2023.

Harvey, W., Wang, Y., and Wenger, R. A randomized O(m log m) time algorithm for computing Reeb graphs of arbitrary simplicial complexes. In *Proceedings of the twenty-sixth annual symposium on Computational geometry*, pp. 267–276, 2010.

Hatcher, A. *Algebraic Topology*. Cambridge University Press, 2002.

Li, M., Palande, S., Yan, L., and Wang, B. Sketching merge trees for scientific data visualization. *arXiv preprint arXiv:2101.03196*, 2021.

Li, M., Yan, X., Yan, L., Needham, T., and Wang, B. Flexible and probabilistic topology tracking with partial optimal transport. *arXiv preprint arXiv:2302.02895*, 2023.

Mémoli, F. On the use of Gromov-Hausdorff Distances for Shape Comparison. In Botsch, M., Pajarola, R., Chen, B., and Zwicker, M. (eds.), *Eurographics Symposium on Point-Based Graphics*. The Eurographics Association, 2007. ISBN 978-3-905673-51-7. doi: 10.2312/SPBG/SPBG07/081-090.

Mémoli, F. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.

Munch, E. and Wang, B. Convergence between categorical representations of Reeb space and Mapper. In *32nd International Symposium on Computational Geometry (SoCG 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.

Parsa, S. A deterministic O(m log m) time algorithm for the Reeb graph. In *Proceedings of the twenty-eighth annual symposium on Computational geometry*, pp. 269–276, 2012.

Peyré, G., Cuturi, M., and Solomon, J. Gromov-Wasserstein averaging of kernel and distance matrices. In *International conference on machine learning*, pp. 2664–2672. PMLR, 2016.

Pickup, D., Sun, X., Rosin, P. L., Martin, R. R., Cheng, Z., Lian, Z., Aono, M., Ben Hamza, A., Bronstein, A., Bronstein, M., Bu, S., Castellani, U., Cheng, S., Garro, V., Giachetti, A., Godil, A., Han, J., Johan, H., Lai, L., Li, B., Li, C., Li, H., Litman, R., Liu, X., Liu, Z., Lu, Y., Tatsuma, A., and Ye, J. SHREC'14 track: Shape retrieval of non-rigid 3d human models. In *Proceedings of the 7th Eurographics workshop on 3D Object Retrieval*, EG 3DOR'14. Eurographics Association, 2014.

Qi, C. R., Su, H., Mo, K., and Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660, 2017.

Reeb, G. Sur les points singuliers d'une forme de pfaff completement integrable ou d'une fonction numerique [on the singular points of a completely integrable pfaff form or of a numerical function]. *Comptes Rendus Acad. Sciences Paris*, 222:847–849, 1946.

Riehl, E. *Category theory in context*. Courier Dover Publications, 2017.

Singh, G., Mémoli, F., Carlsson, G. E., et al. Topological methods for the analysis of high dimensional data sets and 3d object recognition. *PBG@ Eurographics*, 2:091–100, 2007.

Stefanou, A. *Dynamics on categories and applications*. State University of New York at Albany, 2018.

Turner, K., Mukherjee, S., and Boyer, D. M. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4): 310–344, 2014.

Vayer, T., Chapel, L., Flamary, R., Tavenard, R., and Courty, N. Fused Gromov-Wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020.

Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. 3d Shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.

A. Proof of Theorem 4.10

Left inequality

As shown in (Bauer et al., 2015), $d_{\mathrm{I}}\big(R, S\big) \leq d_{\mathrm{FD}}\big(R_f, S_g\big)$ and, since the functional distortion distance on Reeb graphs corresponds to the first part of the functional distortion distance of barcode transforms (Definition 4.9), we obviously get $d_{\mathrm{FD}}\big(R_f, S_g\big) \leq d_{\mathrm{FD}}\big(BF, BG\big)$.

Right inequality

Let $S_\epsilon\colon \mathbf{Fun}(\mathcal{O}(\mathbb{R}), \mathbf{PVec}) \to \mathbf{Fun}(\mathcal{O}(\mathbb{R}), \mathbf{PVec})$ be the smoothing functor, defined by $S_\epsilon\mathcal{F}(U) := \mathcal{F}(U^\epsilon)$ and $\iota\colon \mathcal{F} \to S_\epsilon\mathcal{F}$ be defined by $\iota(U) := \mathcal{F}(U \subseteq U^\epsilon)$. In the following we denote an object of $\mathbf{PVec}$ by a tuple (I, D) representing a set I and a functor $D\colon I \to \mathbf{Vec}$. Suppose $\mathcal{F}$ and $\mathcal{G}$ are ϵ-interleaved, i.e. we have the following commutative diagram:

$$
\begin{tikzcd}
\mathcal{F} \arrow[r, "\iota"] \arrow[dr, "{(\alpha,\eta)}"'] & S_\epsilon\mathcal{F} \arrow[r, "\iota_\epsilon"] \arrow[dr, "{(\alpha_\epsilon,\eta_\epsilon)}"'] & S_{2\epsilon}\mathcal{F} \\
\mathcal{G} \arrow[r, "\iota"] \arrow[ur, "{(\beta,\rho)}"] & S_\epsilon\mathcal{G} \arrow[r, "\iota_\epsilon"] \arrow[ur, "{(\beta_\epsilon,\rho_\epsilon)}"] & S_{2\epsilon}\mathcal{G}
\end{tikzcd} \tag{4}
$$

where α denotes the morphisms between the parameterizing sets and η denotes the morphisms between the parameterized vector spaces. Let $\mathbf{dom}\colon \mathbf{PVec} \to \mathbf{Set}$ be a forgetful functor defined on an object $(I, D) \in \mathbf{PVec}$ by $\mathbf{dom}\big((I, D)\big) := I$ and on a morphism $(\alpha, \eta)\colon (I, D) \to (I', D')$ by $\mathbf{dom}\big((\alpha, \eta)\big) := \alpha$. If we postcompose $\mathcal{F}$ with $\mathbf{dom}$ we obtain $\mathbf{dom} \circ \mathcal{F}(U) = \mathbf{dom}(\mathcal{F}(U)) = \pi_0(f^{-1}(U))$ and $\mathbf{dom} \circ \mathcal{F}(U \subseteq V) = \pi_0(f^{-1}(U) \subseteq f^{-1}(V))$. Hence, we get $\mathbf{dom} \circ \mathcal{F} = \mathcal{R}$ the categorical Reeb graph corresponding to (R, f). Denote by $\mathbf{R}\colon \mathbf{concreteReebgraphs} \to \mathbf{categoricalReebgraphs}$ the functor that sends a concrete Reeb graph (R, f) to the corresponding categorical Reeb graph $\mathcal{R}$ (see (de Silva et al., 2016)). We now apply $\mathbf{dom}$ on Equation (4) and obtain the following commutative diagram of $\mathbf{Set}$-valued functors:

$$
\begin{tikzcd}
\mathbf{R}(R,f) \arrow[r, "\iota"] \arrow[dr, "\alpha"'] & S_\epsilon\mathbf{R}(R,f) \arrow[r, "\iota_\epsilon"] \arrow[dr, "\alpha_\epsilon"'] & S_{2\epsilon}\mathbf{R}(R,f) \\
\mathbf{R}(S,g) \arrow[r, "\iota"] \arrow[ur, "\beta"] & S_\epsilon\mathbf{R}(S,g) \arrow[r, "\iota_\epsilon"] \arrow[ur, "\beta_\epsilon"] & S_{2\epsilon}\mathbf{R}(S,g)
\end{tikzcd} \tag{5}
$$

By Proposition 4.29 in (de Silva et al., 2016), the smoothing of open sets S_ϵ is equivalent to the smoothing of the underlying geometric Reeb graphs. Let $T_\epsilon(R, f)$ be the ϵ-thickening of (R, f) defined by $T_\epsilon R := R \times [-\epsilon, \epsilon]$ and $\mathcal{U}_\epsilon(R, f)$ be the Reeb graph of $T_\epsilon(R, f)$ (the ϵ-smoothing of (R, f)). These spaces can be summarized by the following commutative diagram:

$$
\begin{tikzcd}
R \arrow[dr, "f"'] & T_\epsilon R \arrow[l, "p_1"'] \arrow[r, "q"] \arrow[d, "\hat{f}_\epsilon"] & \mathcal{U}_\epsilon R \arrow[dl, "f_\epsilon"] \\
& \mathbb{R} &
\end{tikzcd} \tag{6}
$$

where p_1 is the projection to the first factor and q is the quotient map to the Reeb space. The map p_1 induces a natural isomorphism $\mathbf{R}T_\epsilon \implies S_\epsilon\mathbf{R}$ such that

$$
\begin{aligned}
&\Big(\mathbf{R}T_\epsilon(R, f)(U) \to S_\epsilon\mathbf{R}(R, f)(U)\Big) \\
=&\Big(\pi_0(\hat{f}_\epsilon^{-1}(U)) \xrightarrow{\pi_0(p_1)} \pi_0(f^{-1}(U^\epsilon))\Big)
\end{aligned}
$$

; (de Silva et al., 2016) Theorem 4.2. Moreover, the map q induces a natural isomorphism $\mathbf{R}T_\epsilon \implies \mathbf{R}\mathcal{U}_\epsilon$ such that

$$
\begin{aligned}
&\Big(\mathbf{R}T_\epsilon(R, f)(U) \to \mathbf{R}\mathcal{U}_\epsilon(R, f)(U)\Big) \\
=&\Big(\pi_0(\hat{f}_\epsilon^{-1}(U)) \xrightarrow{\pi_0(q)} \pi_0(f_\epsilon^{-1}(U))\Big)
\end{aligned}
$$

; (de Silva et al., 2016) Theorem 3.15. Let h denote the composition of the following natural isomorphisms:

$$h\colon S_\epsilon\mathbf{R} \Longrightarrow \mathbf{R}T_\epsilon \implies \mathbf{R}\mathcal{U}_\epsilon$$

$$
\begin{tikzcd}
h(U)\colon \pi_0(f^{-1}(U^\epsilon)) \arrow[dr, "\pi_0(p_1)^{-1}"'] & & \pi_0(f_\epsilon^{-1}(U)) \\
& \pi_0(\hat{f}_\epsilon^{-1}(U)) \arrow[ur, "\pi_0(q)"'] &
\end{tikzcd} \tag{7}
$$

Applying h to Equation (5) yields

$$
\begin{tikzcd}
\mathbf{R}(R,f) \arrow[r, "h(\iota)"] \arrow[dr, "h(\alpha)"'] & \mathbf{R}\mathcal{U}_\epsilon(R,f) \arrow[r, "h(\iota_\epsilon)"] \arrow[dr, "h(\alpha_\epsilon)"'] & \mathbf{R}\mathcal{U}_{2\epsilon}(R,f) \\
\mathbf{R}(S,g) \arrow[r, "h(\iota)"] \arrow[ur, "h(\beta)"] & \mathbf{R}\mathcal{U}_\epsilon(S,g) \arrow[r, "h(\iota_\epsilon)"] \arrow[ur, "h(\beta_\epsilon)"] & \mathbf{R}\mathcal{U}_{2\epsilon}(S,g)
\end{tikzcd} \tag{8}
$$

By Theorem 3.20 in (de Silva et al., 2016), the functor $\mathbf{R}$ is one part of an equivalence between the categories of concrete Reeb graphs and categorical Reeb graphs. If we apply the inverse functor $\mathbf{R}^{-1}$ (the display locale functor) to Equation (8) we obtain the following ϵ-interleaving of Reeb graphs:

$$
\begin{tikzcd}
(R,f) \arrow[r, "\iota"] \arrow[dr, "\varphi"'] & \mathcal{U}_\epsilon(R,f) \arrow[r, "\iota_\epsilon"] \arrow[dr, "\varphi_\epsilon"'] & \mathcal{U}_{2\epsilon}(R,f) \\
(S,g) \arrow[r, "\iota"] \arrow[ur, "\psi"] & \mathcal{U}_\epsilon(S,g) \arrow[r, "\iota_\epsilon"] \arrow[ur, "\psi_\epsilon"] & \mathcal{U}_{2\epsilon}(S,g)
\end{tikzcd} \tag{9}
$$

Note that by the proof of Theorem 3.20 in (de Silva et al., 2016) and the following discussion the functors $\mathbf{R}$ and $\mathbf{R}^{-1}$ are actually inverse to each other, i.e. $\mathbf{R} \circ \mathbf{R}^{-1} = \text{id}$ and $\mathbf{R}^{-1} \circ \mathbf{R} = \text{id}$. In particular, we have that $\mathbf{R}(\varphi) = \mathbf{R} \circ \mathbf{R}^{-1}\big(h(\alpha)\big) = h(\alpha)$, i.e. , for all U, we obtain the following commutative diagram:

$$\begin{tikzcd}
\pi_0(f^{-1}(U)) \arrow[r, "\pi_0(\varphi)"] \arrow[d, "="] & \pi_0(g_\epsilon^{-1}(U)) \arrow[d, "="] \\
\pi_0(f^{-1}(U)) \arrow[r, "h(\alpha)"] \arrow[d, "="] & \pi_0(g_\epsilon^{-1}(U)) \arrow[d, "h(U)^{-1}"] \\
\pi_0(f^{-1}(U)) \arrow[r, "\alpha"] & \pi_0(g^{-1}(U^\epsilon))
\end{tikzcd} \tag{10}$$

using the inverse of the isomorphism $h(U)$ in Equation (7). By the proof of Lemma 15 in (Bauer et al., 2015), there exist $\Phi\colon R \to S$ and $\Psi\colon S \to R$ sucht that

$$\begin{aligned}
\sup_{\substack{(r,r'),(s,s') \\ \in C(\Phi,\Psi)}} \frac{1}{2}|d_f(r,r') - d_g(s,s')| \leq 3(\epsilon+\delta) \\
||f - g \circ \Phi||_\infty \leq \epsilon + \delta \\
||g - f \circ \Psi||_\infty \leq \epsilon + \delta
\end{aligned} \tag{11}$$

for all sufficiently small $\delta > 0$. For $r \in R$, we now show that $BF(r)$ is close to $BG \circ \Phi(r)$ in the interleaving distance.

Let $\kappa > 0$, $t \in \mathbb{R}_{\geq 0}$ and $B(f(r),t) \subseteq \mathbb{R}$ be an open ball of radius t around $f(r)$. Since $|f(r) - g \circ \Phi(r)| \leq \epsilon + \delta$, if $\kappa > \epsilon + \delta$, we get:

$$\begin{aligned}
B\big(f(r),t\big) &\subseteq B\big(g \circ \Phi(r), t+\kappa\big) \\
&\subseteq B\big(g \circ \Phi(r), t+\kappa+2\epsilon\big) \\
&\subseteq B\big(f(r), t+2(\kappa+\epsilon)\big).
\end{aligned} \tag{12}$$

Therefore, by functoriality of $\mathcal{F}$ and the ϵ-interleaving between $\mathcal{F}$ and $\mathcal{G}$ in Equation (4) we obtain:

$$\begin{tikzcd}
\mathcal{F}\Big(B\big(f(r),t\big)\Big) \arrow[rr, "\mathcal{F}(\iota)"] \arrow[d, "\mathcal{F}(\iota)"'] & & \mathcal{F}\Big(B\big(f(r),t+2(\kappa+\epsilon)\big)\Big) \\
\mathcal{F}\Big(B\big(g \circ \Phi(r),t+\kappa\big)\Big) \arrow[rr, "\mathcal{F}(\iota)"] \arrow[dr, "(\alpha_{t+\kappa},\eta_{t+\kappa})"'] & & \Big(B\big(g \circ \Phi(r),t+\kappa+2\epsilon\big)\Big) \arrow[u, "\mathcal{F}(\iota)"'] \\
& \mathcal{G}\Big(B\big(g \circ \Phi(r),t+\kappa+\epsilon\big)\Big) \arrow[ur, "(\beta_{t+\kappa+\epsilon},\rho_{t+\kappa+\epsilon})"'] &
\end{tikzcd} \tag{13}$$

If we apply $\mathbf{dom}$ to Equation (13) we obtain:

$$\begin{tikzcd}
\pi_0\Big(f^{-1}\big(B(f(r),t)\big)\Big) \arrow[dr, "\pi_0(\iota)"] \arrow[dd, "\pi_0(\iota)"'] & \\
& \pi_0\Big(f^{-1}\big(B(f(r),t+2(\kappa+\epsilon))\big)\Big) \\
\pi_0\Big(f^{-1}\big(B(g \circ \Phi(r),t+\kappa)\big)\Big) \arrow[dr, "\pi_0(\iota)"] \arrow[dd, "\alpha_{t+\kappa}"'] & \\
& \pi_0\Big(f^{-1}\big(B(g \circ \Phi(r),t+\kappa+2\epsilon)\big)\Big) \arrow[uu, "\pi_0(\iota)"'] \\
\pi_0\Big(g^{-1}\big(B(g \circ \Phi(r),t+\kappa+\epsilon)\big)\Big) \arrow[ur, "\beta_{t+\kappa+\epsilon}"'] &
\end{tikzcd} \tag{14}$$

Let $B_{d_f}(r,t)$ be the open ball of radius t around r in R. Since, $B_{d_f}(r,t) \subseteq f^{-1}(B(f(r),t))$ is by definition path-connected, $B_{d_f}(r,t) \in \pi_0\Big(f^{-1}\big(B(f(r),t)\big)\Big)$ and, since $r \in B_{d_f}(r,t)$, we have $B_{d_f}(r,t) = [r]$ the path-component of r in $f^{-1}(B(f(r),t))$. By the same argument, $B_{d_g}(\Phi(r),t+\kappa+\epsilon) = [\Phi(r)] \in \pi_0\Big(g^{-1}\big(B(g \circ \Phi(r),t+\kappa+\epsilon)\big)\Big)$. Moreover, $\pi_0(\iota)([r]) = [\iota(r)] = [r] \in \pi_0\Big(f^{-1}\big(B(g \circ \Phi(r),t+\kappa)\big)\Big)$. By using Equation (10) for $U = B(g \circ \Phi(r),t+\kappa)$ we obtain:

$$\begin{tikzcd}
\pi_0\big(f^{-1}(B(g \circ \Phi(r),t+\kappa))\big) \arrow[dr, "\pi_0(\varphi)"] \arrow[dd, "\alpha"'] & \\
& \pi_0\big(g_\epsilon^{-1}(B(g \circ \Phi(r),t+\kappa))\big) \arrow[dl, "h^{-1}"] \\
\pi_0\big(g^{-1}(B(g \circ \Phi(r),t+\kappa+\epsilon))\big) &
\end{tikzcd} \tag{15}$$

By Equation (7), $h^{-1} := \pi_0(p_1) \circ \pi_0(q)^{-1}$ and, by (Bauer et al., 2015) Section 3.2, $\Phi := p_1 \circ \tilde{\varphi}_\delta$. Since $\tilde{\varphi}_\delta(r) \in \overline{\varphi}(B_{d_f}(r,\delta)) = q^{-1}\big(\varphi(B_{d_f}(r,\delta))\big)$, $\varphi(B_{d_f}(r,\delta))$ is path-connected and $\varphi(r) \in \varphi(B_{d_f}(r,\delta))$, we have that $[\varphi(r)] = [q(\tilde{\varphi}_\delta(r))] = \pi_0(q)([\tilde{\varphi}_\delta(r)])$. Hence, $\pi_0(q)^{-1}([\varphi(r)]) = [\tilde{\varphi}_\delta(r)]$. By definition of Φ, we have $[\Phi(r)] = [p_1 \circ \tilde{\varphi}_\delta(r)] = \pi_0(p_1)([\tilde{\varphi}_\delta(r)])$. Therefore, $h^{-1} \circ \pi_0(\varphi)([r]) = h^{-1}([\varphi(r)]) = \pi_0(p_1) \circ \pi_0(q)^{-1}([\varphi(r)]) = \pi_0(p_1)([\tilde{\varphi}_\delta(r)]) = [\Phi(r)] = \alpha_{t+\kappa}([r])$. By commutativity of Equation (14), $\beta_{t+\kappa+\epsilon} \circ \alpha_{t+\kappa}([r]) = \beta_{t+\kappa+\epsilon}([\Phi(r)]) = \pi_0(\iota)([r]) = [r]$. As a consequence, $\alpha_{t+\kappa} \circ \pi_0(\iota)(B_{d_f}(r,t)) = B_{d_g}(\Phi(r),t+\kappa+\epsilon)$ and $\pi_0(\iota) \circ \beta_{t+\kappa+\epsilon}(B_{d_g}(\Phi(r),t+\kappa+\epsilon)) = B_{d_f}(r,t+2(\kappa+\epsilon))$. Thus, the interleaving in Equation (13) yields the following commutative diagram in $\mathbf{Vec}$:

$$\begin{tikzcd}
F\big(B_{d_f}(r,t)\big) \arrow[rr, "F(\iota)"] \arrow[d, "F(\iota)"'] & & F\big(B_{d_f}(r,t+2(\kappa+\epsilon))\big) \\
F([r]) \arrow[rr, "F(\iota)"] \arrow[dr, "\eta_{t+\kappa}([r])"'] & & F([r]) \arrow[u, "F(\iota)"'] \\
& G\big(B_{d_g}(\Phi(r),t+\kappa+\epsilon)\big) \arrow[ur, "\rho_{t+\kappa+\epsilon}([\Phi(r)])"'] &
\end{tikzcd} \tag{16}$$

where $[r]$ and $[\Phi(r)]$ denote the topological path-components in the respective preimages.

We now start with $\Phi(r)$. Similar to Equation (12) we obtain the following inclusions of open intervals in $\mathbb{R}$:

$$\begin{aligned}
B\big(g\circ\Phi(r),t\big) &\subseteq B\big(f(r),t+\kappa\big) \\
&\subseteq B\big(f(r),t+\kappa+2\epsilon\big) \\
&\subseteq B\big(g\circ\Phi(r),t+2(\kappa+\epsilon)\big)
\end{aligned} \tag{17}$$

for every $\kappa > \epsilon+\delta$. Therefore, by functoriality of $\mathcal{G}$ and the ϵ-interleaving between $\mathcal{F}$ and $\mathcal{G}$ in Equation (4) we obtain:

$$\begin{tikzcd}
\mathcal{G}\Big(B\big(g\circ\Phi(r),t\big)\Big) \arrow[rr, "\mathcal{G}(\iota)"] \arrow[d, "\mathcal{G}(\iota)"'] & & \mathcal{G}\Big(B\big(g\circ\Phi(r),t+2(\kappa+\epsilon)\big)\Big) \\
\mathcal{G}\Big(B\big(f(r),t+\kappa\big)\Big) \arrow[rr, "\mathcal{G}(\iota)"] \arrow[dr, "(\beta_{t+\kappa},\rho_{t+\kappa})"'] & & \mathcal{G}\Big(B\big(f(r),t+\kappa+2\epsilon\big)\Big) \arrow[u, "\mathcal{G}(\iota)"'] \\
& \mathcal{F}\Big(B\big(f(r),t+\kappa+\epsilon\big)\Big) \arrow[ur, "(\alpha_{t+\kappa+\epsilon},\eta_{t+\kappa+\epsilon})"'] &
\end{tikzcd} \tag{18}$$

and, by applying $\mathbf{dom}$, we get:

$$\begin{tikzcd}
\pi_0\Big(g^{-1}\big(B(g\circ\Phi(r),t)\big)\Big) \arrow[dr, "\pi_0(\iota)"] \arrow[dd, "\pi_0(\iota)"'] & \\
& \pi_0\Big(g^{-1}\big(B(g\circ\Phi(r),t+2(\kappa+\epsilon))\big)\Big) \\
\pi_0\Big(g^{-1}\big(B(f(r),t+\kappa)\big)\Big) \arrow[dr, "\pi_0(\iota)"] \arrow[dd, "\beta_{t+\kappa}"'] & \\
& \pi_0\Big(g^{-1}\big(B(f(r),t+\kappa+2\epsilon)\big)\Big) \arrow[uu, "\pi_0(\iota)"'] \\
\pi_0\Big(f^{-1}\big(B(f(r),t+\kappa+\epsilon)\big)\Big) \arrow[ur, "\alpha_{t+\kappa+\epsilon}"'] &
\end{tikzcd} \tag{19}$$

As in the previous case, we have that $B_{d_g}(\Phi(r),t) \subseteq g^{-1}\big(B(g\circ\Phi(r),t)\big)$ is the path-component of $\Phi(r)$, i.e. $[\Phi(r)] = B_{d_g}(\Phi(r),t) \in \pi_0\Big(g^{-1}\big(B(g\circ\Phi(r),t)\big)\Big)$ and, analogously, $[r] = B_{d_f}(r,t+\kappa+\epsilon) \in \pi_0\Big(f^{-1}\big(B(f(r),t+\kappa+\epsilon)\big)\Big)$. We now use the analog of Equation (10) for $U = B(f(r),t+\kappa)$, ψ from the interleaving in Equation (9) and β to obtain:

$$\begin{tikzcd}
\pi_0\big(g^{-1}(B(f(r),t+\kappa))\big) \arrow[dr, "\pi_0(\psi)"] \arrow[dd, "\beta"'] & \\
& \pi_0\big(f_\epsilon^{-1}(B(f(r),t+\kappa))\big) \arrow[dl, "h^{-1}"] \\
\pi_0\big(f^{-1}(B(f(r),t+\kappa+\epsilon))\big) &
\end{tikzcd} \tag{20}$$

By Equation (7), $h^{-1} := \pi_0(p_1)\circ\pi_0(q)^{-1}$ and, by (Bauer et al., 2015) Section 3.2, $\Psi := p_1\circ\tilde{\psi}_\delta$. Since $\tilde{\psi}_\delta(\Phi(r)) \in \overline{\psi}(B_{d_g}(\Phi(r),\delta)) = q^{-1}\big(\psi(B_{d_g}(\Phi(r),\delta))\big)$, $\psi(B_{d_g}(\Phi(r),\delta))$ is path-connected and $\psi(\Phi(r)) \in \psi(B_{d_g}(\Phi(r),\delta))$, we have that $[\psi(\Phi(r))] = [q(\tilde{\psi}_\delta(\Phi(r)))] = \pi_0(q)([\tilde{\psi}_\delta(\Phi(r))])$. Hence, $\pi_0(q)^{-1}([\psi(\Phi(r))]) = [\tilde{\psi}_\delta(\Phi(r))]$. By definition of Ψ, we have $[\Psi(\Phi(r))] = [p_1\circ\tilde{\psi}_\delta(\Phi(r))] = \pi_0(p_1)([\tilde{\psi}_\delta(\Phi(r))])$. Therefore, $h^{-1}\circ\pi_0(\psi)([\Phi(r)]) = h^{-1}([\psi(\Phi(r))]) = \pi_0(p_1)\circ\pi_0(q)^{-1}([\psi(\Phi(r))]) = \pi_0(p_1)([\tilde{\psi}_\delta(\Phi(r))]) = [\Psi(\Phi(r))] = \beta_{t+\kappa}([\Phi(r)])$.

From Equation (11) we get $\frac{1}{2}|d_f(r,\Psi\circ\Phi(r))| \leq 3(\epsilon+\delta)$. If $\kappa+\epsilon > 6(\epsilon+\delta)$, then $B_{d_f}(r,6(\epsilon+\delta)) \subseteq B_{d_f}(r,t+\kappa+\epsilon) \subseteq f^{-1}\big(B(f(r),t+\kappa+\epsilon)\big)$. Hence, since r and $\Psi\circ\Phi(r) \in B_{d_f}(r,t+\kappa+\epsilon)$ and $B_{d_f}(r,t+\kappa+\epsilon)$ is path-connected, $[r] = [\Psi\circ\Phi(r)] \in \pi_0\big(f^{-1}(B(f(r),t+\kappa+\epsilon))\big)$.

Therefore, starting with $B_{d_g}(\Phi(r),t) = [r]$, we obtain $\beta_{t+\kappa}\circ\pi_0(\iota)([\Phi(r)]) = \beta_{t+\kappa}([\Phi(r)]) = [\Psi\circ\Phi(r)] = [r] = B_{d_f}(r,t+\kappa+\epsilon)$. This implies that we can extract the following commutative diagram from Equation (18):

$$\begin{tikzcd}
G\big(B_{d_g}(\Phi(r),t)\big) \arrow[rr, "G(\iota)"] \arrow[d, "G(\iota)"'] & & G\big(B_{d_g}(\Phi(r),t+2(\kappa+\epsilon))\big) \\
G([\Phi(r)]) \arrow[rr, "G(\iota)"] \arrow[dr, "\rho_{t+\kappa}([\Phi(r)])"'] & & G([\Phi(r)]) \arrow[u, "G(\iota)"'] \\
& F\big(B_{d_f}(r,t+\kappa+\epsilon)\big) \arrow[ur, "\eta_{t+\kappa+\epsilon}([r])"'] &
\end{tikzcd} \tag{21}$$

Now we define

$$\begin{aligned}
&\mu_t\colon F\big(B_{d_f}(r,t)\big) \to G\big(B_{d_g}(\Phi(r),t+\kappa+\epsilon)\big) \\
&\mu_t := \eta_{t+\kappa}\circ F(\iota) \\
&\nu_t\colon G\big(B_{d_g}(\Phi(r),t)\big) \to F\big(B_{d_f}(r,t+\kappa+\epsilon)\big) \\
&\nu_t := \rho_{t+\kappa}\circ G(\iota)
\end{aligned} \tag{22}$$

Since $\mathcal{F}$ and $\mathcal{G}$ are ϵ-interleaved we have the following commutative diagram

$$\begin{tikzcd}
\mathcal{F}\big(B(f(r),t)\big) \arrow[r, "\mathcal{F}(\iota)"] \arrow[d, "(\alpha_t,\eta_t)"'] & \mathcal{F}\big(B(g\circ\Phi(r),t+\kappa)\big) \arrow[d, "(\alpha_{t+\kappa},\eta_{t+\kappa})"] \\
\mathcal{G}\big(B(f(r),t+\epsilon)\big) \arrow[r, "\mathcal{G}(\iota)"] & \mathcal{G}\big(B(g\circ\Phi(r),t+\kappa+\epsilon)\big)
\end{tikzcd} \tag{23}$$

Following the component $B_{d_f}(r,t)$ we get

$$\begin{tikzcd}
F\big(B_{d_f}(r,t)\big) \arrow[r, "F(\iota)"] \arrow[d, "\eta_t"'] & F\big([r]\big) \arrow[d, "\eta_{t+\kappa}"] \\
G\big(\dots\big) \arrow[r, "\mathcal{G}(\iota)"] & G\big(B_{d_g}(\Phi(r),t+\kappa+\epsilon)\big)
\end{tikzcd} \tag{24}$$

This implies that the map $\mu_t = \eta_{t+\kappa}\circ F(\iota) = G(\iota)\circ\eta_t$. Analogously we obtain that $\nu_t = \rho_{t+\kappa}\circ G(\iota) = F(\iota)\circ\rho_t$. Moreover, for $s < t \in \mathbb{R}_{\geq 0}$, the following diagram and its analog for ν obviously commute:

$$\begin{tikzcd}
F\big(B_{d_f}(r,s)\big) \arrow[r, "F(\iota)"] \arrow[d, "\mu_s"'] & F\big(B_{d_f}(r,t)\big) \arrow[d, "\mu_t"] \\
G\big(B_{d_g}(\Phi(r),s+\kappa+\epsilon)\big) \arrow[r, "G(\iota)"] & G\big(B_{d_g}(\Phi(r),t+\kappa+\epsilon)\big)
\end{tikzcd} \tag{25}$$

Combining these results with Equation (16) and Equation (21), we obtain the following $(\kappa+\epsilon)$-interleaving:

$$\begin{tikzcd}
F\big(B_{d_f}(r,t)\big) \arrow[d, "F(\iota)"'] \arrow[dr, "\mu_t"] & G\big(B_{d_g}(\Phi(r),t)\big) \arrow[d, "G(\iota)"] \arrow[dl, "\nu_t"'] \\
F\big(B_{d_f}(r,t+\kappa+\epsilon)\big) \arrow[d, "F(\iota)"'] \arrow[dr, "\mu_{t+\kappa+\epsilon}"] & G\big(B_{d_g}(\Phi(r),t+\kappa+\epsilon)\big) \arrow[d, "G(\iota)"] \arrow[dl, "\nu_{t+\kappa+\epsilon}"'] \\
F\big(B_{d_f}(r,t+2(\kappa+\epsilon))\big) & G\big(B_{d_g}(\Phi(r),t+2(\kappa+\epsilon))\big)
\end{tikzcd} \tag{26}$$

Hence, $BF(r)$ and $BG(\Phi(r))$ are $(\kappa+\epsilon)$-interleaved for every $\kappa > 5\epsilon+6\delta$. Since $\inf\{\kappa+\epsilon \mid \kappa > 5\epsilon+6\delta \text{ and } \delta > 0\} = 6\epsilon$, we finally obtain $d_{\mathrm{I}}\big(BF(r), BG(\Phi(r))\big) \leq 6\epsilon$. By symmetry, we analogously obtain $d_{\mathrm{I}}\big(BF(\Psi(s)), BG(s)\big) \leq 6\epsilon$. Together with Equation (11), these bounds imply the theorem.