

Stability and Approximations for Decorated Reeb Spaces

Justin Curry ✉ 

University at Albany, State University of New York, NY, USA

Washington Mio ✉

Florida State University, Tallahassee, FL, USA

Tom Needham ✉ 

Florida State University, Tallahassee, FL, USA

Osman Berat Okutan ✉

Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

Florian Russold ✉

Graz University of Technology, Austria

Abstract

Given a map $f : X \rightarrow M$ from a topological space X to a metric space M , a *decorated Reeb space* consists of the Reeb space, together with an attribution function whose values recover geometric information lost during the construction of the Reeb space. For example, when $M = \mathbb{R}$ is the real line, the Reeb space is the well-known Reeb graph, and the attributions may consist of persistence diagrams summarizing the level set topology of f . In this paper, we introduce decorated Reeb spaces in various flavors and prove that our constructions are Gromov-Hausdorff stable. We also provide results on approximating decorated Reeb spaces from finite samples and leverage these to develop a computational framework for applying these constructions to point cloud data.

2012 ACM Subject Classification Mathematics of computing $\rightarrow$ Algebraic topology; Theory of computation $\rightarrow$ Computational geometry

Keywords and phrases Reeb spaces, Gromov-Hausdorff distance, Persistent homology

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.44

Related Version *Full Version*: <https://arxiv.org/abs/2312.01982> [12]

Funding *Justin Curry*: Supported by NSF CCF-1850052 and NASA 80GRC020C0016.

Washington Mio: Supported by NSF DMS-1722995.

Tom Needham: Supported by NSF DMS-2107808 and DMS-2324962.

Florian Russold: Supported by the Austrian Science Fund (FWF): W1230.

1 Introduction

Graphical summaries of topological spaces have a long history in pure mathematics [20, 29] and have more recently served as popular tools for shape and data analysis; prominent examples include Reeb graphs [19], Mapper graphs [33] and merge trees [26]. The related concept of persistent homology is another ubiquitous tool for topological summarization in modern data science [6]. Graphical and persistence-based summaries frequently capture complementary information about a dataset, and there is a recently developing body of work which studies methods for enriching graphical descriptors with additional geometric or topological data [10, 11, 13, 23, 24].

In this paper we prove a Gromov-Hausdorff stability result for an enriched topological summary called the **Decorated Reeb Space**. The Reeb graph summarizes connectivity of level sets of a function $f : X \rightarrow \mathbb{R}$ defined on some topological space X ; the Reeb space



© Justin Curry, Washington Mio, Tom Needham, Osman Berat Okutan, and Florian Russold; licensed under Creative Commons License CC-BY 4.0


40th International Symposium on Computational Geometry (SoCG 2024).


Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 44; pp. 44:1–44:17



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** The Decorated Reeb Graph pipeline. (a) A **point cloud**, noisily sampled from a space homotopy equivalent to $(S^1 \times S^1) \vee S^1$. (b) The (estimated) **Reeb graph** associated to the point cloud, with respect to the height function. (c) To each node of the Reeb graph, we associate a **filtered space**, which is a sequence of subsets of the original point cloud. (d) Taking persistent homology of the filtered space associated to each node yields a **decorated Reeb graph**; each node is attributed with a persistence diagram.

generalizes the Reeb graph to summarize maps $f : X \rightarrow M$ valued in some metric space M – such considerations are natural when dealing with multi-variate data [15]. The enrichment presented in this paper synthesizes persistent homology with the Reeb space to provide a single topological descriptor more powerful than persistence or Reeb spaces alone. Before summarizing the content of the paper, we first provide an illustration of the decorated Reeb space pipeline in order to shed light on the nature and importance of our results.

Illustration of the Pipeline. Consider Figure 1. Part (a) shows a point cloud in $\mathbb{R}^3$, sampled (with additive noise) from a space that is homotopy equivalent to $(S^1 \times S^1) \vee S^1$. Using height along the z -axis as a function on the point cloud, we obtain an estimation of the Reeb graph of the space, as shown in (b). To each node of the Reeb graph, we assign a filtration of the original point cloud, as illustrated in (c). For a fixed anchor node in the Reeb graph (highlighted in red) and for each scale r we include those points in the original point cloud that correspond to nodes in the Reeb graph within Reeb radius r (see Definition 1) of the anchor node. The top row of (c) shows subsets of the Reeb graph within two different Reeb radii and the bottom row shows the original point clouds within those radii of the anchor node. To this filtration we obtain a bi-filtered simplicial complex by considering the Vietoris-Rips complex (with proximity parameter s) at each Reeb radius r . Taking a one-dimensional slice of the (r, s) -parameter space, we compute the (one-dimensional) persistent homology and add this as an attribution to the anchor node in the Reeb graph. By letting the anchor node vary across the Reeb graph, and repeating this process, we obtain a novel structure in TDA called the **Decorated Reeb Graph**. Part (d) of Figure 1 shows the resulting persistence diagram for three different anchor nodes. We note that the persistent homology decoration at each node captures the local topology of the original point cloud, and this construction is more sensitive to topological features that exist near each anchor node. In our example, the attributions on the right side of the Reeb graph capture the topology of the torus (there are two off-diagonal points in the H_1 persistence diagram), whereas the attribution on the left hand side captures only the cycle coming from the larger circle.

Summary of Main Results. The output signature of the Decorated Reeb Graph pipeline described above – that is, a graph whose nodes are attributed with persistence diagrams – is the same as that of the construction considered in [13]. However, as noted in [13, Remark 5.1], the computational pipeline in that paper is largely divorced from the theory, and the abstract stability results there have little relevance to the algorithmic implementation. This paper gives

a novel theoretical construction of decorated Reeb spaces arising from functions $f : X \rightarrow M$ defined on connected spaces, borrowing tools from metric geometry. Our main contributions are:

- We introduce the concept of the **Reeb radius** associated to a map $f : X \rightarrow M$ and show that this is closely tied to Reeb graph smoothings considered in [14]; see Theorem 2.
- We show that, under (quantifiable) tameness assumptions on the spaces and functions, if $f : X \rightarrow M$ and $g : Y \rightarrow M$ are close in a certain Gromov-Hausdorff sense, then so are their decorated Reeb spaces; see Theorem 22.
- We show that a tame function $f : X \rightarrow M$ can be approximated by a finite graph so that the (continuous) decorated Reeb space of X is well approximated by the (discrete) decorated Reeb space of the graph; see Theorem 25.
- Finally, we illustrate our constructions via computational examples in Section 4.4.

2 Decorated Reeb Spaces

This section introduces the main constructions that we study throughout the paper. We begin with some preliminary definitions, which allows us to set conventions and notation.

2.1 Reeb Spaces and Smoothings

Let X be a topological space, (M, d_M) be a metric space, and $f : X \rightarrow M$ be a map. In this paper, every topological space is assumed to be compact, connected, and locally path connected, and every map is assumed to be continuous. We refer to the data (X, f) as an **M -field**, and denote the class of all M -fields as $\mathcal{F}_M$.

The **Reeb space** associated to $(X, f) \in \mathcal{F}_M$ is defined to be the quotient space of X under the equivalence relation $\sim$ on X defined by $x \sim x'$ if and only if x and x' lie in the same connected component of a level set of f . This construction was independently conceived by Georges Reeb [29] and Aleksandr Kronrod [20], but the attribution to Kronrod is frequently dropped. Keeping with this convention, we denote the Reeb space by $R_f = X / \sim$ and the quotient map by $\pi_f : X \rightarrow R_f$. For $x \in X$, we denote the equivalence class $\pi_f(x)$ by $[x]$. By definition, the map $f : X \rightarrow M$ factors through R_f to define an induced map $\bar{f} : R_f \rightarrow M$, i.e., $\bar{f}([x]) = f(x)$. When $M = \mathbb{R}$, the resulting quotient R_f is typically called the **Reeb graph**, although additional hypotheses are required to ensure a true graph structure [32]. These assumptions are automatically satisfied in most computational settings and one should consult [19, 36] for a sample of the applications to data analysis.


Both the Reeb graph and Reeb space are sensitive to perturbations of the map $f : X \rightarrow M$. One way of remedying this is to consider analogs of the smoothing operation first introduced in [14]. Given $\epsilon > 0$, the **ϵ -smoothing** of the Reeb space R_f , denoted R_f^ϵ , is defined in a two-step process. First one constructs the space T_f^ϵ and map $f^\epsilon : T_f^\epsilon \rightarrow M$ defined by

$$T_f^\epsilon := \{(x, v) \in X \times M : d_M(f(x), v) \leq \epsilon\}, \quad f^\epsilon(x, v) := v. \quad (1)$$

If R_{f^ϵ} is the Reeb space of $(T_f^\epsilon, f^\epsilon)$, then R_f^ϵ is the image of X under the composition,

$$\pi_f^\epsilon : X \hookrightarrow T_f^\epsilon \rightarrow R_{f^\epsilon} \quad \text{where} \quad x \mapsto (x, f(x)) \mapsto [(x, f(x))], \quad (2)$$

of the inclusion and quotient map. We note that R_f^0 can be identified with R_f , so this smoothing operation produces a 1-parameter family of Reeb spaces starting with R_f .



■ **Figure 2** The graph of a generic function $f : [0, 1] \rightarrow \mathbb{R}$ is its own Reeb graph. Notice that the Reeb radius is not symmetric in x and y .

2.2 Reeb Radius

A key concept introduced in this paper is the *Reeb radius*, which contains all information about which points in X are identified in the Reeb space and its associated smoothings. Later, we will use this to metrize our Reeb space and to induce decorations.

► **Definition 1** (Reeb radius). For $(X, f) \in \mathcal{F}_M$, define the **Reeb radius** $\rho_f : X \times X \rightarrow \mathbb{R}$ by

$$\rho_f(x, y) := \inf_{\gamma} \sup_t d_M(f(x), f(\gamma(t))), \quad (3)$$

where the infimum is taken over all paths $\gamma : [0, 1] \rightarrow X$ with $\gamma(0) = x$, $\gamma(1) = y$.

The Reeb radius is very similar to the **Reeb distance** of [3]. This is the map

$$\partial_f : X \times X \rightarrow \mathbb{R}, \quad \partial_f(x, y) := \inf_{\gamma} \text{diam}(f(\gamma([0, 1]))), \quad (4)$$

where the infimum is also taken over continuous paths from x to y . It is straightforward to show that $\rho_f \leq \partial_f \leq 2\rho_f$, but one clear difference is that the Reeb radius is not necessarily symmetric; see Figure 2 for an example.

We now show how the Reeb radius characterizes Reeb smoothings.

► **Theorem 2.** For $(X, f) \in \mathcal{F}_M$ and $x, y \in X$, $\pi_f^\epsilon(x) = \pi_f^\epsilon(y)$ if and only if $f(x) = f(y)$ and $\rho_f(x, y) \leq \epsilon$.

► **Lemma 3.** For $x, y \in X$, y is contained in the connected component of x in $f^{-1}(\mathbb{D}^r(f(x)))$ if and only if $\rho_f(x, y) \leq r$, where $\mathbb{D}^r(v)$ denotes the closed ball with radius r and center v .

Proof. Assume $\rho_f(x, y) \leq r$. For each integer $n > 0$, let C_n denote the connected component of x in $f^{-1}(\mathbb{D}^{r+1/n}(f(x)))$. Note that $y \in C_n$ for all n . The collection $\{C_n\}$ forms a decreasing family of compact connected sets. By [16, Corollary 6.1.19], $C := \cap C_n$ is connected and $x, y \in C \subseteq f^{-1}(\mathbb{D}^r(f(x)))$.

Now, assume that y is contained in the connected component of x in $f^{-1}(\mathbb{D}^r(f(x)))$. For an integer $n > 0$, let U_n denote the connected component of x in $f^{-1}(\mathbb{B}^{r+1/n}(f(x)))$, where $\mathbb{B}^{r+1/n}$ denotes the open ball with radius $r + 1/n$. We have $y \in U_n$, which is path connected, since it is open, connected, and X is locally path connected. Thus, there exists a path in U_n from x to y , so that $\rho_f(x, y) \leq r + 1/n$. Since $n > 0$ was arbitrary, we have $\rho_f(x, y) \leq r$. ◀

Proof of Theorem 2. Let $x \in X$ and $f(x) = v$. By Equation (1) we have that

$$(f^\epsilon)^{-1}(v) = \{(x', v) : d_M(f(x), v) \leq \epsilon\} = f^{-1}(\mathbb{D}^\epsilon(v)) \times \{v\},$$

where we continue to use $\mathbb{D}^\epsilon$ for a closed ϵ -ball in M . Hence, by Equation (2), $\pi_f^\epsilon(x) = \pi_f^\epsilon(y)$ if and only if $f(y) = f(x)$ and y is in the connected component of x in $f^{-1}(\mathbb{D}^\epsilon(f(x)))$. This happens if and only if $\rho_f(x, y) \leq \epsilon$ by Lemma 3. $\blacktriangleleft$

2.3 Metrizing the Reeb Space

One of our main goals is to establish Gromov-Hausdorff stability of our constructions. In order to formulate these results, we first need to metrize the Reeb space R_f . We do this using the Reeb radius, but first we need to establish some of its properties.

► **Proposition 4.** Let $(X, f) \in \mathcal{F}_M$. For all $x, y, z \in X$, we have:

- i) $\rho_f(x, z) \leq \rho_f(x, y) + \rho_f(y, z)$.
- ii) $\rho_f(x, y) = 0$ if and only if $[x] = [y]$.
- iii) If $[x] = [x']$ and $[y] = [y']$, then $\rho_f(x, y) = \rho_f(x', y')$.

► **Remark 5.** With these properties, ρ_f induces a well defined *quasimetric* on R_f , meaning that it only lacks the symmetry axiom of a metric; recall Figure 2.

Proof.

- i) Given paths α from x to y and β from y to z , let γ denote the path from x to z obtained by concatenating α and β . We have

$$\begin{aligned} \rho_f(x, z) &\leq \sup_t d_M(f(x), f(\gamma(t))) = \max \left(\sup_t d_M(f(x), f(\alpha(t))), \sup_t d_M(f(x), f(\beta(t))) \right) \\ &\leq \max \left(\sup_t d_M(f(x), f(\alpha(t))), d_M(f(x) + f(y)) + \sup_t d_M(f(y), f(\beta(t))) \right) \\ &\leq \sup_t d_M(f(x), f(\alpha(t))) + \sup_t d_M(f(y), f(\beta(t))). \end{aligned}$$

Since this holds for arbitrary paths α and β , the triangle inequality follows.

- ii) This is the $\epsilon = 0$ case of Theorem 2.
- iii) Using part i), $\rho_f(x, y) \leq \rho_f(x, x') + \rho_f(x', y') + \rho_f(y', y) = \rho_f(x', y')$. Repeating this argument with the roles of x, y and x', y' reversed proves $\rho_f(x', y') \leq \rho_f(x, y)$. $\blacktriangleleft$

We now obtain a metric on R_f by symmetrizing the Reeb radius.

► **Definition 6.** For $(X, f) \in \mathcal{F}_M$, we define $d_f : R_f \times R_f \rightarrow [0, \infty)$ by

$$d_f([x], [y]) := 2 \max(\rho_f(x, y), \rho_f(y, x)).$$

As already remarked, it is not hard to show that d_f is bi-Lipschitz equivalent to the metric induced by the Reeb distance (Equation (4)), but they are not equal in general.

2.4 Decorations

We wish to endow the Reeb space R_f of a map $f : X \rightarrow M$ with an additional function $R_f \rightarrow Y$, for some (perhaps extended- and/or pseudo-) metric space Y ; this function will be our attribution or **decoration** function. The metric space Y will typically be the space of all (multi-parameter) persistence modules or barcodes; the decoration will capture topological information about X that is lost in the quotient R_f .

In more detail, recall that Topological Data Analysis (TDA) typically views a function $f : X \rightarrow \mathbb{R}$ as a stepping stone to defining a **filtration** – an increasing chain of closed subsets that exhausts X . As such, an $\mathbb{R}$ -field (X, f) can be equivalently regarded as a **filtered space** – a space along with its exhaustion $\{f^{-1}(-\infty, t]\}_{t \in \mathbb{R}}$. In this sense, our first decoration function assigns a filtration to each point of the Reeb space.

► **Definition 7** (Filtration decoration for Reeb Spaces). *Let $(X, f) \in \mathcal{F}_M$. The **filtration decoration** for R_f is the map $D_f : R_f \rightarrow \mathcal{F}_{\mathbb{R}}$ defined by*

$$D_f([x]) := (X, \rho_f(x, \cdot)),$$

where ρ_f is the Reeb radius; see Figure 1(c) for an example.

Filtration or filtered space decorations are somewhat unwieldy to work with computationally, but these structures can be summarized using persistent homology. To this end, we assume the reader is familiar with the basic notions of persistent homology: the **Vietoris-Rips complex**, the **persistent homology barcode**, and the **bottleneck distance** between barcodes; see [6] for a reminder. In the following, we consider a **simplicial filtration**; in general, this is a family $\{K^r\}_{r \in P}$ of simplicial complexes indexed by a poset $(P, \leq)$ such that K^r is a subcomplex of K^s whenever $r \leq s$. In this paper, we consider **1-parameter** and **2-parameter simplicial filtrations**, where $P = \mathbb{R}_{\geq 0}$ and $P = \mathbb{R}_{\geq 0}^2$ (with its product poset structure), respectively. In particular, we use $\{\text{VR}^r(Y)\}_{r \geq 0}$ for the Vietoris-Rips complex of a metric space Y . We now introduce our main computational object of study, the barcode decoration, which is an attribution valued in $\mathcal{BC}$ – the space of barcodes with the bottleneck distance.

► **Definition 8** (Barcode decoration for Reeb Spaces). *Let $(X, f) \in \mathcal{F}_M$ and further assume that X is a metric space. Let $\lambda, c \geq 0$ and let k be a non-negative integer. The **barcode decoration** for R_f, λ, c, k is the map $B_{f,k}^{\lambda,c} : R_f \rightarrow \mathcal{BC}$, where $B_{f,k}^{\lambda,c}([x])$ is the k -dimensional persistent homology barcode of the Vietoris-Rips 1-parameter simplicial filtration*

$$\{\text{VR}^r(\{y \in X : \rho_f(x, y) \leq \lambda r + c\})\}_{r \geq 0}.$$

See Figures 1 and 7 for examples.

► **Remark 9.** In fact, we should only consider $B_{f,k}^{\lambda,c}([x])$ to be a *persistence module*, as the tameness conditions for the guaranteed existence of a barcode decomposition [9] may not hold. With a view toward computation, we will abuse terminology and still refer to the decorations as barcodes, but our results hold when decorations are considered instead as persistence modules and the interleaving distance is used in place of bottleneck distance.

3 Gromov-Hausdorff Stability

In this section, we establish stability of the decorated Reeb space constructions defined above.

3.1 Metric Fields and Multiscale Comparisons

As Definition 8 intimates, filtration-attributed Reeb spaces can be compared using the Vietoris-Rips filtration if and only if X has a metric. This will be essential as we move from a connected topological space X to a point-sampling of X along with its map $f : X \rightarrow M$. To emphasize this extra structure, we call the pair (X, f) an **M -metric field** when X has a metric (otherwise, $f : X \rightarrow M$ is just an M -field). Recall that the collection of all M -fields

is denoted $\mathcal{F}_M$. By contrast, we will denote the collection of all M -metric fields by $\mathcal{MF}_M$. Methods for metrizing $\mathcal{MF}_M$ can be found in [2], but here we will focus on developing variations on the Gromov-Hausdorff distance to compare metric fields.

Recall that a **correspondence** between sets X and Y is a subset $\mathcal{R} \subset X \times Y$ such that the coordinate projections are surjective when restricted to $\mathcal{R}$.

► **Definition 10** ((r,s)-Correspondences between metric fields). *Let $(X, f), (Y, g) \in \mathcal{MF}_M$, let $\mathcal{R}$ be a correspondence between X, Y , and let $r, s \geq 0$. We call $\mathcal{R}$ an (r, s) -**correspondence** between (X, f) and (Y, g) if for all $(x, y), (x', y') \in \mathcal{R}$ we have*

$$|d_X(x, x') - d_Y(y, y')| \leq 2r \quad \text{and} \quad d_M(f(x), g(y)) \leq s.$$

We use the following metrization of $\mathcal{MF}_M$ in our stability results below.

► **Definition 11.** *The **Gromov-Hausdorff distance** between M -metric fields (X, f) and (Y, g) , denoted $d_{\text{GH}}((X, f), (Y, g))$, is defined to be the infimum over r such that there is an (r, r) -correspondence between (X, f) and (Y, g) .*

When M is a single point, the above definition reduces to the regular Gromov-Hausdorff distance between metric spaces. The case $M = \mathbb{R}$ reduces to the definition given in [7, Definition 2.4]. This definition gives a pseudometric on metric fields.

The filtration decoration of Definition 7 produces a metric field valued in metric fields. That is, for $(X, f) \in \mathcal{MF}_M$ and $[x] \in R_f$, the filtration decoration $D_f([x]) = (X, \rho_f(x, \cdot))$ is naturally an $\mathbb{R}$ -metric field, i.e., D_f is a map $D_f : R_f \rightarrow \mathcal{MF}_{\mathbb{R}}$. The class $\mathcal{MF}_{\mathbb{R}}$ can be endowed with the GH pseudometric above, so that $(X, D_f) \in \mathcal{MF}_{\mathcal{MF}_{\mathbb{R}}}$ (where we consider D_f as a function on X by composing with the Reeb quotient map). The following definition shows how we can compare these decorations.

► **Definition 12** ((r,s,t)-Correspondences between metric field-valued functions). *Let (M, d_M) be a metric space. Let (X, d_X) and (Y, d_Y) be metric spaces endowed with functions $D_X : X \rightarrow \mathcal{MF}_M$ and $D_Y : Y \rightarrow \mathcal{MF}_M$ – that is, $(X, D_X), (Y, D_Y) \in \mathcal{MF}_{\mathcal{MF}_M}$. Let $\mathcal{R}$ be a correspondence between X and Y . We call $\mathcal{R}$ an (r, s, t) -**correspondence** between (X, D_X) and (Y, D_Y) if, for all $(x, y), (x', y') \in \mathcal{R}$, $|d_X(x, x') - d_Y(y, y')| \leq 2r$ and $\mathcal{R}$ is an (s, t) -correspondence between $D_X(x)$ and $D_Y(y)$.*

We define a variant of the **Gromov-Hausdorff distance** between (X, D_X) and (Y, D_Y) , denoted $\hat{d}_{\text{GH}}((X, D_X), (Y, D_Y))$, to be the infimum over r such that there is an (r, r, r) -correspondence between them.

3.2 Stability of Filtration Decorated Reeb Spaces

To study the stability properties of the decorated Reeb space construction, we need to restrict ourselves to a certain subclass of M -metric fields.

► **Definition 13** ((L, ϵ)-Connectivity). *We say $(X, f) \in \mathcal{MF}_M$ has (L, ϵ) -**connectivity** if for all $x, y \in X$, $\rho_f(x, y) \leq Ld_X(x, y) + 2\epsilon$. We denote the subspace of $\mathcal{MF}_M$ consisting of M -metric fields with (L, ϵ) -connectivity by $\mathcal{MF}_M^{L, \epsilon}$.*

► **Example 14.** Suppose $(X, f) \in \mathcal{MF}_M$ satisfies $\sup_{x \in X} d_M(f(x), g(x)) \leq \epsilon$ for some L -Lipschitz $g : X \rightarrow M$. If X is a geodesic space, then f has (L, ϵ) -connectivity. Indeed, let γ be a geodesic in X from x to y . Then $d_M(f(x), f(\gamma(t))) \leq d_M(g(x), g(\gamma(t))) + 2\epsilon \leq Ld_X(x, y) + 2\epsilon$. This implies that $\rho_f(x, y) \leq Ld_X(x, y) + 2\epsilon$.

► **Theorem 15** (Stability of Decorated Reeb Spaces). *Let $(X, f), (Y, g) \in \mathcal{MF}_M^{L, \epsilon}$. Let $\mathcal{R}$ be an (r, s) -correspondence between (X, f) and (Y, g) and let $\mathcal{S}$ be the induced correspondence between the Reeb spaces R_f and R_g given by $([x], [y]) \in \mathcal{S}$ whenever $(x, y) \in \mathcal{R}$. Then, $\mathcal{S}$ is a $(2Lr + 2s + 2\epsilon, r, 2Lr + 2s + 2\epsilon)$ -correspondence between (R_f, D_f) and (R_g, D_g) .*

► **Corollary 16.** *Let $(X, f), (Y, g) \in \mathcal{MF}_M^{L, \epsilon}$. Then,*

$$\hat{d}_{\text{GH}}((R_f, D_f), (R_g, D_g)) \leq 2(L + 1)d_{\text{GH}}((X, f), (Y, g)) + 2\epsilon.$$

Proof of Theorem 15. Let $(x, y), (x', y') \in \mathcal{R}$. Let $\gamma : [0, 1] \rightarrow Y$ be a path from y to y' . For every $\delta > 0$ there is a partition $0 = t_0 < t_1 < \dots < t_n = 1$ such that $\text{diam}(g(\gamma([t_i, t_{i+1}])) \leq \delta$ for all i . Let $y_i = \gamma(t_i)$ and $(x_i, y_i) \in \mathcal{R}$ such that $x_0 = x$ and $x_n = x'$. Note that

$$\rho_f(x_i, x_{i+1}) \leq Ld_f(x_i, x_{i+1}) + \epsilon \leq L(2r + \delta) + 2\epsilon, \quad \forall i.$$

By definition of ρ_f , there exist paths α_i from x_i to x_{i+1} such that $\rho_f(x_i, \alpha_i(t)) \leq L(2r + \delta) + 2\epsilon$ for all t . Hence, we have

$$\begin{aligned} d_M(f(x), f(\alpha_i(t))) &\leq d_M(f(x), f(x_i)) + \rho_f(x_i, \alpha_i(t)) \\ &\leq \sup_t d_M(g(x), g(\gamma(t))) + L(2r + \delta) + 2s + 2\epsilon. \end{aligned}$$

By concatenating α_i 's, we get a path from x to x' . Hence, the inequality above implies

$$\rho_f(x, x') \leq \sup_t d_M(g(x), g(\gamma(t))) + L(2r + \delta) + 2s + 2\epsilon.$$

Since $\delta > 0$ and γ were arbitrary, we get $\rho_f(x, x') \leq \rho_g(y, y') + 2Lr + 2s + 2\epsilon$. One can similarly show that $\rho_g(y, y') \leq \rho_f(x, x') + 2Lr + 2s + 2\epsilon$. Therefore $|\rho_f(x, x') - \rho_g(y, y')| \leq 2Lr + 2s + 2\epsilon$.

This implies that the metric distortion of $\mathcal{S}$ between R_f and R_g is less than or equal to $(4Lr + 4s + 4\epsilon)$, and $\mathcal{R}$ is an $(r, 2Lr + 2s + 2\epsilon)$ -correspondence between $(X, D_f([x]))$ and $(Y, D_g([y]))$ for all $(x, y) \in \mathcal{R}$. Combining these two, we see that $\mathcal{S}$ is an $(2Lr + 2s + 2\epsilon, r, 2Lr + 2s + 2\epsilon)$ -correspondence between (R_f, D_f) and (R_g, D_g) . ◀

3.3 Stability of Barcode Decorated Reeb Spaces

To establish the stability of barcode decorations (Definition 8), let us analyze the stability of the intermediate steps going from a metric filtration to a barcode.

► **Definition 17.** *For any $\mathbb{R}$ -metric field $f : X \rightarrow \mathbb{R}$, we have a two-parameter simplicial filtration $\{S_f^{t,u}\}_{t,u \geq 0}$ defined by $S_f^{t,u} := \text{VR}^t(\{x \in X \mid f(x) \leq u\})$.*

The one parameter version of this, where $t = u$, was shown to be stable in [7]. To establish a multi-parameter stability result, we utilize a multiparameter version of the homotopy interleaving distance (in the sense of [18] and *not* in the sense of [4]) below.

► **Definition 18** (cf. [18]). *Given two 2-parameter simplicial filtrations $S^{*,*}$ and $T^{*,*}$, an (r, s) -homotopy interleaving between them is a pair of 2-parameter families of simplicial maps $\phi^{t,u} : S^{t,u} \rightarrow T^{t+r, u+s}$ and $\psi^{t,u} : T^{t,u} \rightarrow S^{t+r, s+u}$ that commute with the structure maps, and whose composition is homotopy equivalent to the inclusion maps $S^{*,*} \hookrightarrow S^{*+2r, *+2s}$ and $T^{*,*} \hookrightarrow T^{*+2r, *+2s}$.*

We adapt the standard contiguity proof of Vietoris-Rips stability [8] to this multi-parameter setting.

► **Proposition 19.** *Let $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ be $\mathbb{R}$ -metric fields and let $\mathcal{R}$ be an (r, s) -correspondence between them. Then $S_f^{*,*}$ and $S_g^{*,*}$ are $(2r, s)$ -homotopy interleaved.*

Proof. Let $F : X \rightarrow Y$ and $G : Y \rightarrow X$ be functions such that $(x, F(x)), (G(y), y) \in \mathcal{R}$. Note that F induces a simplicial map from $S_f^{t,u}$ to $S_g^{t+2r,u+s}$, and similarly G induces a simplicial map. As in [8], composition of these maps are contiguous to the inclusion, hence $S_f^{*,*}$ and $S_g^{*,*}$ are $(2r, s)$ -homotopy interleaved. ◀

As we already defined (r, s) -homotopy interleavings of simplicial filtrations, we can, by analogy, also define (r, s) -interleavings of persistence modules [22] by simply replacing each simplicial complex with its homology. Homotopy invariance of homology [27] implies that if two simplicial filtrations are (r, s) -homotopy interleaved, then their persistent homology modules are (r, s) -interleaved. Assuming our modules are point-wise finite dimensional [9], we can induce barcodes by taking 1-parameter slices of a 2-parameter module, as in [21].

► **Definition 20.** *Given a two parameter persistence module $P^{*,*}$ and $\lambda, c \geq 0$, we denote by $B_P^{\lambda,c}$ the persistence barcode of the 1-parameter filtration $\{P^{t,\lambda t+c}\}_{t \geq 0}$.*

The following result is an adaptation of [21, Lemma 1] to our multiparameter setting.

► **Proposition 21.** *Let $P^{*,*}$ and $Q^{*,*}$ be (r, s) -interleaved persistence modules. Given $\lambda, c \geq 0$,*

$$d_b(B_P^{\lambda,c}, B_Q^{\lambda,c}) \leq \max(r, s/\lambda),$$

where d_b denotes the bottleneck distance.

Proof. Let $\epsilon = \max(r, s/\lambda)$. It is enough to show that the 1-parameter filtrations $V^* := \{P^{t,\lambda t+c}\}_{t \geq 0}$ and $W^* := \{Q^{t,\lambda t+c}\}_{t \geq 0}$ are ϵ -interleaved. Let (ϕ, ψ) be an (r, s) -interleaving between P and Q . Note that $t + r \leq t + \epsilon$ and $\lambda t + c + s \leq \lambda(t + \epsilon) + c$. Hence, composing ϕ with the structure maps of Q , we get a map from $V^t \rightarrow W^{t+\epsilon}$ for all $t \geq 0$:

$$V^t = P^{t,\lambda t+c} \xrightarrow{\phi} Q^{t+r,\lambda t+c+s} \longrightarrow Q^{t+\epsilon,\lambda(t+\epsilon)+c} = W^{t+\epsilon}$$

Similarly we get maps from W^* to $V^{*+\epsilon}$. Since (ϕ, ψ) is an interleaving, V^*, W^* are ϵ -interleaved. ◀

With the notation established above, for $[x] \in R_f$, we have $B_{f,k}^{\lambda,c}([x]) := B_{PH_k(S_{D_f}([x]))}^{\lambda,c}$.

► **Theorem 22.** *For $(X, f), (Y, g) \in \mathcal{MF}_M^{L,\epsilon}$, we have*

$$d_{\text{GH}}((R_f, B_{f,k}^{\lambda,c}), (R_g, B_{g,k}^{\lambda,c})) \leq 2 \max(1, 1/\lambda)((L+1)d_{\text{GH}}((X, f), (Y, g)) + \epsilon).$$

Proof. Let $\mathcal{R}$ be an (r, r) -correspondence between (X, f) and (Y, g) . By Theorem 15, Proposition 19 and Proposition 21, $\mathcal{R}$ is an $((2L+2)r + 2\epsilon, \max(2r, \frac{(2L+2)r+2\epsilon}{\lambda}))$ -correspondence between $(R_f, B_{f,k}^{\lambda,c}), (R_g, B_{g,k}^{\lambda,c})$. The maximum of $(2L+2)r + 2\epsilon$ and $\frac{(2L+2)r+2\epsilon}{\lambda}$ is equal to $2 \max(1, 1/\lambda)((L+1)r + \epsilon)$. ◀

► **Remark 23.** In the case that M is a point, R_f, R_g become singletons, L, ϵ can be taken to be zero, and λ, c do not effect the barcode assigned to the single point, which becomes the barcode of the persistent homology of the Vietoris-Rips filtration of the corresponding metric space. So, taking $\lambda = 1$, Theorem 22 reduces to the classical result on the stability of Vietoris-Rips barcodes [7].

4 Decorated Reeb Spaces for Combinatorial Graphs

Having established the stability of the decorated Reeb space in the continuous setting, we now show how we can reliably approximate it in the finite setting. First, notice that the concept of Reeb radius makes sense for finite combinatorial graphs, where paths in a topological space are replaced by edge paths in a graph. More precisely, given a simple finite graph $G = (V, E)$ and a function $g : V \rightarrow M$ to a metric space M , we define $\rho_g : V \times V \rightarrow \mathbb{R}$ by

$$\rho_g(v, w) := \min \left\{ \max_{i=1, \dots, n} d_M(g(v), g(v_i)) \mid v = v_0, \dots, v_n = w, [v_i, v_{i+1}] \in E \forall i \right\}.$$

The Reeb space R_g , metric d_g , and decorations D_g and $B_{g,k}^{\lambda,c}$ are defined similarly.

4.1 An Algorithm for the Reeb Radius Function


Here we present an efficient algorithm for computing the Reeb radius function of a discrete graph by modifying Dijkstra's shortest path algorithm. See Figure 3 for an example.

■ **Algorithm 1** Input: Connected graph $G = (V, E)$, $g : V \rightarrow M$, $x \in V$. Output: $\rho_g(x, \cdot) : V \rightarrow \mathbb{R}$.

```

for  $v$  in  $V$  do:
     $\rho_g(x, v) \leftarrow \infty$ 
end for
 $\rho_g(x, x) \leftarrow 0$ 
 $Q \leftarrow [x]$ 
while  $Q$  is non-empty do:
     $v \leftarrow \operatorname{argmin}_{w \in Q} \rho_g(x, w)$ 
     $Q \leftarrow Q \setminus \{v\}$ 
    for  $w$  in the set of neighbours of  $v$  do:
        if  $\rho_g(x, w) = \infty$  then:
             $\rho_g(x, w) \leftarrow \max(\rho_g(x, v), d_M(g(x), g(w)))$ 
             $Q \leftarrow Q \cup \{w\}$ 
        end if
    end for
end while
return  $\rho_g(x, \cdot)$ 

```



■ **Figure 3** On the left is a graph endowed with an $\mathbb{R}$ -valued function g given by node height. The other two graphs are the spanning trees generated by Algorithm 1 with respect to the root node x , with a directed edge $[w, v]$ for each v, w that is processed inside the conditional statement. At each node, the value of $\rho_g(x, v)$ is written. The directed path from v to x realizes $\rho_g(x, v)$.

4.2 Finite Approximations of Decorated Reeb Spaces

The following result shows that, with a carefully constructed graph, we can approximate the Reeb radius function of a metric field.

► **Proposition 24.** *Let $(X, f) \in \mathcal{MF}_M^{L, \epsilon}$, let $\delta > 0$, and let $V \subset X$ be a finite subspace such that for all $x \in X$, there exists $v \in V$ such that $d_X(x, v) < \delta$ and $d_M(f(x), f(v)) \leq \delta$. Let G be the simple metric graph with node set V and the edges given by $[v, w] \in E$ if $d_X(v, w) \leq 3\delta$. Let $g : V \rightarrow M$ be the restriction of f to V . Then, for all $v, w \in V$,*

$$|\rho_f(v, w) - \rho_g(v, w)| \leq (3L + 1)\delta + \epsilon.$$

Proof. Let $\gamma : [0, 1] \rightarrow X$ be a path from v to w . Take $0 = t_0 < \dots < t_n = 1$ such that $d_X(\gamma(t_i), \gamma(t_{i+1})) \leq \delta$ for all i . Let $x_i = \gamma(t_i)$, and $v_i \in V$ such that $v_0 = v$, $v_n = w$, $d_X(x_i, v_i) \leq \delta$, and $d_X(f(x_i), f(v_i)) \leq \delta$. Then $d_X(v_i, v_{i+1}) \leq 3\delta$, which means that $(v_0, \dots, v_n)$ is an edge path in G . We have

$$\max_i d_M(g(v), g(v_i)) \leq \sup_t d_M(f(v), f(\gamma(t))) + \delta.$$

Infimizing over γ , we get $\rho_g(v, w) \leq \rho_f(v, w) + \delta$.

Let $(v = v_0, \dots, v_n = w)$ be an edge path in G . We have $\rho_f(v_i, v_{i+1}) \leq 3L\delta + \epsilon$. Hence, there is a path α_i from x_i to x_{i+1} in X such that for all x over the image of α_i , we have

$$d_M(f(v), f(x)) \leq d_M(f(v), f(x_i)) + \rho_f(x_i, x) \leq d_M(f(v), f(x_i)) + 3L\delta + \epsilon.$$

By concatenating α_i 's, we get a path from v to w in X , hence

$$\rho_f(v, w) \leq \max_i d_M(f(v), f(x_i)) + 3L\delta + \epsilon.$$

Infimizing over edge paths, we get $\rho_f(v, w) \leq \rho_g(v, w) + 3L\delta + \epsilon$. ◀

► **Theorem 25.** *Let (X, f) and $G = (V, E)$ with $g : V \rightarrow M$ be as in Proposition 24. Then there is a $((5L + 1)\delta + 3\epsilon, \delta, (5L + 1)\delta + 3\epsilon)$ -correspondence between (R_f, D_f) and (R_g, D_g) .*

Proof. Let $\mathcal{R}$ be the correspondence between X, V given by $(x, v) \in \mathcal{R}$ if $d_X(x, v) \leq \delta$. Let $(x, v), (y, w) \in \mathcal{R}$.

$$\rho_f(x, y) - \rho_f(v, w) \leq \rho_f(x, v) + \rho_f(w, y) \leq Ld_X(x, v) + \epsilon + Ld_X(w, y) + \epsilon \leq 2(L\delta + \epsilon).$$

By using the triangle inequality for ρ_f , one can see that $|\rho_f(v, w) - \rho_f(x, y)| \leq 2(L\delta + \epsilon)$. Hence, by Proposition 24,


$$|\rho_f(x, y) - \rho_g(v, w)| \leq (5L + 1)\delta + 3\epsilon.$$

This shows that $\mathcal{R}$ is a $(\delta, (5L + 1)\delta + 3\epsilon)$ -correspondence between $(X, \rho_f(x, \cdot))$ and $(V, \rho_g(v, \cdot))$. Then, if $\mathcal{S}$ denotes the correspondence between R_f and R_g given by $([x], [y]) \in \mathcal{S}$ if $(x, y) \in \mathcal{R}$, then $\mathcal{S}$ is a $((5L + 1)\delta + 3\epsilon, \delta, (5L + 1)\delta + 3\epsilon)$ -correspondence between (R_f, D_f) and (R_g, D_g) . ◀

► **Corollary 26.** *Let $f : X \rightarrow M$ and $G = (V, E)$ with $g : V \rightarrow M$ be as in Proposition 24. Then,*

$$\hat{d}_{\text{GH}}((R_f, B_{f,k}^{\lambda, c}), (R_g, B_{g,k}^{\lambda, c})) \leq \max(1, 1/\lambda)((5L + 2)\delta + 3\epsilon)$$

Proof. Follows from Theorem 25 and Proposition 19 and Proposition 21. ◀



■ **Figure 4** A graph with function g given by node height, its Reeb graph and two smoothings.

4.3 Simplifying Graphs by Smoothing

As we have seen above, the Reeb space construction can be adapted to the setting of (combinatorial) graphs by changing continuous paths to edge paths. As Reeb smoothing requires considering a certain subspace of $X \times M$, this approach does not directly generalize to graphs – as M does not carry a graph structure, there is no natural way of seeing $G \times M$ as a finite graph. However, Theorem 2 gives a characterization of Reeb smoothings in terms of the Reeb radius, which does make sense in the graph setting.


► **Definition 27** (Reeb smoothing of graphs). Let $G = (V, E)$ be a finite simple graph and $g : V \rightarrow M$, where M is a metric space. We define R_g^ϵ as the finite graph whose set of nodes V_g^ϵ is the quotient set $V / \sim$, where $v \sim w$ if $f(v) = f(w)$ and $\rho_g(v, w) \leq \epsilon$. Denote the quotient map by $\pi_g^\epsilon : V \rightarrow V_g^\epsilon$. The set of edges E_g^ϵ are induced by E in the sense that if $[v, w] \in E$, then $[\pi_g^\epsilon(v), \pi_g^\epsilon(w)] \in E_g^\epsilon$. We metrize V_g^ϵ by choosing a representative from each class and using d_g distance between them. Let us denote such a metric by d_g^ϵ . If V has a metric structure, we also use the barcode decorations of the selected representatives to get a barcode decoration V_g^ϵ , which we denote by $D_{g^\epsilon, k}^{\lambda, c}$.

► **Remark 28.** Note that the method of metrization and decoration described in Definition 27 depends on the representatives we choose, but it saves us from doing new computations, and choosing different representatives changes the metric at most by ϵ .

► **Example 29.** Figure 4 shows a simple graph endowed with a function (height along the y -axis), its associated Reeb graph and two Reeb smoothings of the Reeb graph.

The main advantage of using ϵ -smoothings is that while the metric structure compared to the Reeb graph does not change much, the graph structure can be much simpler even for small ϵ . For example, if the original function $g : V \rightarrow M$ on the node set takes distinct values for all vertices, then the Reeb graph will have the same graph structure as the original graph. However, it is possible that by changing values of g by at most a small δ , there could be many vertices taking the same value under g . The ϵ -smoothing will then produce a much simpler graph. This is an alternative to the Mapper approach [33], where a similar effect can be achieved by adjusting the bin size and connectivity parameters.

► **Example 30.** Figure 5 shows a Reeb smoothing of the USAir97 graph [31]. The nodes represent 332 US airports, with edges representing direct flights between them. We smooth the graph using PageRank [5] as the function g via the operation described in Definition 27, and allowing small adjustments in g , as described above. The resulting smoothed graph clarifies the hub structure of the airport system. These results are similar to those of [30], where the graph was simplified using the Mapper algorithm.



■ **Figure 5** The USAir97 graph of airports and connections, endowed with the PageRank function, together with a smoothed Reeb graph illustrating hub structure.

► **Proposition 31.** *Let $G = (V, E)$ be a finite simple graph and $g : V \rightarrow M$, where V and M are metric spaces. Let $h : V \rightarrow M$ such that $d_M(g(v), h(v)) \leq \delta$ for all $v \in V$. Then,*

$$d_{\text{GH}}((R_g, B_{g,k}^{\lambda,c}), (R_h^\epsilon, D_{h^\epsilon,k}^{\lambda,c})) \leq 2 \max(1, 1/\lambda)(\delta + \epsilon).$$

Proof. Let $\mathcal{R}$ be the correspondence between R_g, R_h^ϵ consisting of elements of the form $(\pi_g(v), \pi_h^\epsilon(v))$ for $v \in V$. Let us show that $\mathcal{R}$ is an $(\delta + 2\epsilon, 0, (2\delta + \epsilon)/\lambda)$ correspondence between the corresponding decorated Reeb graphs.

Let $v, w \in \mathcal{R}$. Let $v', w' \in V$ such that $\pi_h^\epsilon(v') = \pi_h^\epsilon(v)$ and $\pi_h^\epsilon(w') = \pi_h^\epsilon(w)$. Note that since $h(v) = h(v')$, we have $d_h([v], [v']) = 2\rho_h(v, v') \leq 2\epsilon$. Similarly $d_h([w], [w']) \leq 2\epsilon$. Since $d_M(g(\cdot), h(\cdot)) \leq \delta$, we have $|d_g(v, w) - d_h(v, w)| \leq 2\delta$. Hence $|d_g([v], [w]) - d_h([v'], [w'])| \leq 2\delta + 4\epsilon$. For $x \in V$, we have


$$|\rho_g(v, x) - \rho_h(v', x)| \leq |\rho_g(v, x) - \rho_h(v, x)| + |\rho_h(v, x) - \rho_h(v', x)| \leq 2\delta + \rho_h(v, v') \leq 2\delta + \epsilon.$$

Hence, the identity correspondence is a $(0, 2\delta + \epsilon)$ correspondence between $\rho_g(v, \cdot) : V \rightarrow \mathbb{R}$ and $\rho_h(v', \cdot) : V \rightarrow \mathbb{R}$. By Proposition 21, corresponding decorations have bottleneck distance at most $(2\delta + \epsilon)/\lambda$. ◀


4.4 Computational Examples

In this subsection, we present some simple examples which illustrate our computational pipeline. For the sake of simplicity, we focus on fields with $M = \mathbb{R}$, so that our examples are all Decorated Reeb Graphs.

Synthetic data and the effect of the λ -parameter. Figure 6 shows the DRG pipeline for a synthetic planar point cloud; the function f is height along the y -axis. The Reeb graph in (b) is obtained by fitting a k -nearest neighbors graph to the point cloud ($k = 10$, with respect to Euclidean distance), then smoothing the resulting (true) Reeb graph by rounding function values and applying the smoothing procedure described in Section 4.3. The decorations are degree-1 persistent homology diagrams $B_{f,1}^{\lambda,0}$ (Definition 8), for various values of λ . We see from this simple example that the decorations pick up the homology of the underlying point cloud from a localized perspective, where homological features that are closer to a node are accentuated in the diagram. Part (c) shows the effect of the λ parameter on the decoration for a particular node; as λ is increased, homological features appearing at larger Reeb radii from the node become more apparent.



■ **Figure 6** DRGs for synthetic data, illustrating dependence on the λ parameter. (a) Synthetic 2D point cloud endowed with function given by height along the y -axis. (b) Smoothed Reeb graph, decorated with persistence diagrams $B_{f,1}^{1,0}$, using the notation of Definition 8. Barcodes are displayed for several nodes. (c) Barcode decorations of the form $B_{f,1}^{\lambda,0}$ for the indicated node, with varying λ .



■ **Figure 7** Illustration of DRGs and the effect of the function. **Top Row:** (a) Point cloud sampled from a surface mesh of a camel shape from the Sumner database, endowed with function f given by height along the z -axis. (b) The Reeb graph, decorated with persistence diagrams $B_{f,1}^{1,0}$, using the notation of Definition 8. Barcodes are displayed for the indicated nodes of the Reeb graph. (c) Shows the first steps in the filtered space associated to the indicated node. The barcode is computed by taking the Vietoris-Rips persistent homology with radius equal to the Reeb radius. **Bottom Row:** Similar, but with function g given by p -eccentricity with $p = 2$. For the DRG in (b), all nodes except the indicated one have empty diagrams.

Shape data and the effect of the function. Figure 7 applies the DRG to a “camel” shape; this is a point cloud sampled from a shape in the classic mesh database from [34]. The figure illustrates DRGs for two different functions on the shape: the first is height along the z -axis and the second is the p -eccentricity function, which for $p = 2$ and a finite metric space (X, d) is the function $x \mapsto (\sum_{y \in X} d(x, y)^2)^{1/2}$. This function gives a higher value to those points in the metric space that are far “on average” from other points in the space. In each case, the Reeb graph is approximated using a Mapper-like construction from [13], and decorations are of the form $B_{f,1}^{1,0}$, with respect to the appropriate function f . As expected, the resulting DRG depends heavily on the function, both in the shape of the Reeb graph and in the structure of the decorations.




Figure 8 Point cloud matching results. (a) Example point clouds from each shape class. (b) The distance matrix and MDS embedding for decorated Reeb graph representations of the point clouds. (c) The distance matrix and MDS embedding using only Reeb graphs.

Distances between point clouds. To illustrate the usage of the DRG pipeline in a metric-based analysis task, we perform a simple point cloud comparison experiment. The data consists of 10 samples each from 4 classes of meshes from the database of [34]. Each mesh is sampled to form a 3D point cloud, and each point cloud is endowed with a function f given by height along the z -axis. A decorated Reeb graph is computed for each shape, first using the barcode decorations $B_{f,1}^{2,0}$. Each barcode decoration is then vectorized as a persistence image [1], to improve computational efficiency, resulting in a Reeb graph with each node decorated by a vector in $\mathbb{R}^{25 \times 25}$. Actual computation of the field Gromov-Hausdorff distance (Definition 11) is intractable, so we use Fused Gromov-Wasserstein (FGW) distance [35] as a more computationally convenient proxy. This distance is a variant of the Gromov-Wasserstein (GW) distance [25], which can be viewed, roughly, as an L^p -relaxation of Gromov-Hausdorff distance. The GW and FGW distances can be approximated efficiently via gradient descent and entropic regularization [17, 28], and have been similarly used to compare topological signatures of datasets in several recent papers [10, 13, 23, 24].

Our results are shown in Figure 8. We show the pairwise distance matrix between DRGs with respect to FGW distance, as well as a multidimensional scaling plot. Observe that the classes are well-separated by their topological signatures. For comparison, we also compute GW distances between the underlying Reeb graphs; here, the classes are much less distinct, indicating that the decorations recover valuable distinguishing information about the shapes which is lost in the construction of their Reeb graphs.

We note that this is only an illustrative example. To fully utilize the power of this framework, it is likely that feature optimization is necessary; for example, via graph neural networks, as in [13]. Further exploration of this technique will be the subject of future work.

References

- 1 Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18, 2017.
- 2 Soheil Anbouhi, Washington Mio, and Osman Berat Okutan. On metrics for analysis of functional data on geometric domains, 2024. [arXiv:2309.10907](https://arxiv.org/abs/2309.10907).
- 3 Ulrich Bauer, Xiaoyin Ge, and Yusu Wang. Measuring distance between Reeb graphs. In *Proceedings of the thirtieth annual symposium on Computational geometry*, pages 464–473, 2014.
- 4 Andrew Blumberg and Michael Lesnick. Universality of the homotopy interleaving distance. *Transactions of the American Mathematical Society*, 376(12):8269–8307, 2023.
- 5 Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1-7):107–117, 1998.

- 6 Gunnar Carlsson. Topological pattern recognition for point cloud data. *Acta Numerica*, 23:289–368, 2014.
- 7 Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Mémoli, and Steve Y Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum*, 28(5):1393–1403, 2009.
- 8 Frédéric Chazal, Vin De Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014.
- 9 William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.
- 10 Justin Curry, Haibin Hang, Washington Mio, Tom Needham, and Osman Berat Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, 6(3):371–428, 2022.
- 11 Justin Curry, Washington Mio, Tom Needham, Osman Berat Okutan, and Florian Russold. Convergence of leray cosheaves for decorated mapper graphs, 2023. [arXiv:2303.00130](#).
- 12 Justin Curry, Washington Mio, Tom Needham, Osman Berat Okutan, and Florian Russold. Stability and approximations for decorated reeb spaces, 2023. [arXiv:2312.01982](#).
- 13 Justin Curry, Washington Mio, Tom Needham, Osman Berat Okutan, and Florian Russold. Topologically attributed graphs for shape discrimination. In *Topological, Algebraic and Geometric Learning Workshops 2023*, pages 87–101. PMLR, 2023.
- 14 Vin De Silva, Elizabeth Munch, and Amit Patel. Categorified Reeb graphs. *Discrete & Computational Geometry*, 55(4):854–906, 2016.
- 15 Herbert Edelsbrunner, John Harer, and Amit K Patel. Reeb spaces of piecewise linear mappings. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 242–250, 2008.
- 16 Ryszard Engelking. *General Topology*. Heldermann Verlag Berlin, 1989.
- 17 Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *The Journal of Machine Learning Research*, 22(1):3571–3578, 2021.
- 18 Patrizio Frosini, Claudia Landi, and Facundo Mémoli. The persistent homotopy type distance. *Homology, Homotopy and Applications*, 21(2):231–259, 2019.
- 19 Xiaoyin Ge, Issam Safa, Mikhail Belkin, and Yusu Wang. Data skeletonization via Reeb graphs. *Advances in Neural Information Processing Systems*, 24, 2011.
- 20 Aleksandr Semenovich Kronrod. On functions of two variables. *Uspekhi matematicheskikh nauk*, 5(1):24–134, 1950.
- 21 Claudia Landi. The rank invariant stability via interleavings. *Research in computational topology*, pages 1–10, 2018.
- 22 Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015.
- 23 Mingzhe Li, Carson Storm, Austin Yang Li, Tom Needham, and Bei Wang. Comparing morse complexes using optimal transport: An experimental study, 2023. [arXiv:2309.04681](#).
- 24 Mingzhe Li, Xinyuan Yan, Lin Yan, Tom Needham, and Bei Wang. Flexible and probabilistic topology tracking with partial optimal transport, 2023. [arXiv:2302.02895](#).
- 25 Facundo Mémoli. Gromov-Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.
- 26 Dmitriy Morozov, Kenes Beketayev, and Gunther Weber. Interleaving distance between merge trees. *Discrete and Computational Geometry*, 49(22-45):52, 2013.
- 27 James R Munkres. *Elements of algebraic topology*. CRC press, 2018.
- 28 Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In *International conference on machine learning*, pages 2664–2672. PMLR, 2016.

- 29 Georges Reeb. Sur les points singuliers d'une forme de pfaff complètement 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.
- 30 Paul Rosen, Mustafa Hajij, and Bei Wang. Homology-preserving multi-scale graph skeletonization using mapper on graphs, 2023. [arXiv:1804.11242](https://arxiv.org/abs/1804.11242).
- 31 Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. URL: <https://networkrepository.com>.
- 32 Vladimir Vasilievich Sharko. About Kronrod-Reeb graph of a function on a manifold. *Methods of Functional Analysis and Topology*, 12(04):389–396, 2006.
- 33 Gurjeet Singh, Facundo Mémoli, Gunnar E Carlsson, et al. Topological methods for the analysis of high dimensional data sets and 3d object recognition. *PBG@ Eurographics*, 2:091–100, 2007.
- 34 Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Transactions on graphics (TOG)*, 23(3):399–405, 2004.
- 35 Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused Gromov-Wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020.
- 36 Weiming Wang, Yang You, Wenhai Liu, and Cewu Lu. Point cloud classification with deep normalized Reeb graph convolution. *Image and Vision Computing*, 106:104092, 2021.