Measuring Distance between Reeb Graphs

Ulrich Bauer*

Xiaoyin Ge[†]

Yusu Wang[†]

Abstract

One of the prevailing ideas in geometric and topological data analysis is to provide descriptors that encode useful information about hidden objects from observed data. The Reeb graph is one such descriptor for a given scalar function. The Reeb graph provides a simple yet meaningful abstraction of the input domain, and can also be computed efficiently. Given the popularity of the Reeb graph in applications, it is important to understand its stability and robustness with respect to changes in the input function, as well as to be able to compare the Reeb graphs resulting from different functions.

In this paper, we propose a metric for Reeb graphs, called the functional distortion distance. Under this distance measure, the Reeb graph is stable against small changes of input functions. At the same time, it remains discriminative at differentiating input functions. In particular, the main result is that the functional distortion distance between two Reeb graphs is bounded from below by (and thus more discriminative than) the bottleneck distance between both the ordinary and extended persistence diagrams for appropriate dimensions.

As an application of our results, we analyze a natural simplification scheme for Reeb graphs, and show that persistent features in Reeb graph remains persistent under simplification. Understanding the stability of important features of the Reeb graph under simplification is an interesting problem on its own right, and critical to the practical usage of Reeb graphs.

^{*}IST Austria (Institute of Science and Technology Austria), A-3400 Klosterneuburg. Email: ulrich.bauer@ist.ac.at.

[†]Computer Science and Engineering Department, The Ohio State University, Columbus, OH 43221. Emails: gex, yusu@cse.ohio-state.edu.

1 Introduction

One of the prevailing ideas in geometric and topological data analysis is to provide descriptors that encode useful information about hidden objects from observed data. The Reeb graph is one such descriptor for a given function. Specifically, given a continuous scalar function $f : X \to \mathbb{R}$ defined on a domain X, the levelset of f at value a is the set $f^{-1}(a) = \{x \in X \mid f(x) = a\}$. As the scalar value a increases, connected components appear, disappear, split and merge in the level set, and the Reeb graph of f tracks such changes. It provides a simple yet meaningful abstraction of the input domain. The Reeb graph was first termed and used for shape understanding by Shinagawa et al. [30]. Since then, it has been used in a variety of applications in fields such as graphics and scientific visualization [2, 4, 19, 20, 22, 28, 29, 30, 32, 33, 35, 36]; also see [5] for a survey.

The Reeb graph can be computed efficiently in $O(m \log m)$ time for a piecewise-linear function defined on an arbitrary simplicial complex domain with *m* vertices, edges and triangles [26] (a randomized algorithm was given in [18]). This is in contrast to, for example, the $O(m^3)$ time (or matrix multiplication time) needed to compute even just the first-dimensional homology information for the same simplicial complex. Being a graph structure, it is also simple to represent and manipulate. These properties make the Reeb graph also appealing for analyzing high-dimensional points data. For example, a variant of the Reeb graph is proposed in [31] for analyzing high dimensional data, and in [17], the Reeb graph is used to recover a hidden geometric graph from its point samples. Very recently in [11], it is shown that a certain Reeb graph can reconstruct a metric graph with respect to the Gromov-Hausdorff distance.

Given the popularity of the Reeb graph in both low and high dimensional data analysis, it is important to understand its stability and robustness with respect to changes in the input function (both in function values and in the domain). To measure the stability, we need to first define distance between two Reeb graphs. Furthermore, an important usage of the Reeb graph is as a summary for the hidden function that is easier to process than the original domain or the function. For example, one can compare shapes by comparing Reeb graphs of appropriate functions defined on them, or compute the average of a collection of input function in the space of Reeb graphs. Again, a central problem involved is to have a meaningful distance measure between Reeb graphs.

In this paper, we propose a metric for Reeb graphs, called the *functional distortion distance*, drawing intuition from the Gromov-Hausdorff distance for measuring metric distortion. Under this distance, the Reeb graph is stable against small changes of input functions; at the same time it remains discriminative at differentiating functions (these statements will be made precise in Section 4). In particular, the main result is that the functional distortion distance between two Reeb graphs is bounded from below by (and thus more discriminative than) the *bottleneck distance* between the (ordinary and extended) persistence diagrams for appropriate dimensions. On the other hand, the functional distortion distance yields the same type of L_{∞} stability that persistence diagrams enjoy [12, 9, 10]. The persistence diagram has been a popular topological summary of shapes and functions, and the bottleneck distance is introduced in [12] as a natural distance measure for persistence diagrams. However, as the simple example in Figure 1 (a) shows, the Reeb graph takes the structure of the domain (and the function) into account, and can be strictly more discriminative than the persistence diagrams in differentiating functions.

As an application of our results, we show in Section 5 that "important" features in the Reeb graph remains important under the simplification of the Reeb graph. Understanding the stability of Reeb graph features under simplification is an interesting problem on its own right: In practice, one often collapses small branches and loops in the Reeb graph to remove noise; e.g [14, 17, 27]. It is crucial that by collapsing a collection of small features, there is no cascading effect that causes larger features to be destroyed, and our results in Section 5 confirm that.

Very recently, there has been some interesting work on developing a distance measure for the *merge trees*, which are closely related to loop-free Reeb graphs (i.e, the contour trees). In particular, the ε -interleaving distance based on the idea of interleaving between 0-dimensional persistence modules is proposed in [24]. Similar stability and discriminative results relating this measure to the *ordinary* persistent homology are obtained. However, it is not clear how to extend these results to loops, an important family of features encoded in the Reeb graph. The proposed functional distortion distance is similar to this ε -interleaving distance when applied to merge trees.

However, the main technical challenge involved in our work is to relate the proposed distance to the 1st *extended* persistent homology, which describes loops in the Reeb graph and whose handling relies on a very different approach. Finally, another distance measure based on the *branch decomposition* of contour trees is proposed in [3], and a polynomial time algorithm to compute it for two merge trees is given. This measure, however, is not stable with respect to changes in the function. It is also not clear how to extend this measure to Reeb graphs.

2 Preliminaries and Problem Definition

Reeb graphs. Given a continuous function $f : X \to \mathbb{R}$ defined on a triangulable topological space X, for each scalar value $\alpha \in \mathbb{R}$, a *level set* of f w.r.t α is a set $f^{-1}(\alpha) = \{x \in X : f(x) = \alpha\}$. A level set may contain several connected components. Two points $x, y \in X$ are in the same equivalent class *iff* they are in the same connected component, denoted by $x \sim y$. The *Reeb graph* of the function $f : X \to \mathbb{R}$, denoted by \mathbb{R}_f , is the quotient space X/ \sim , which is the set of equivalent classes equipped with the quotient topology induced by this equivalent relation.

Intuitively, \mathbb{R}_f is obtained by continuously collapsing every component of each level set into a single point, and there is a natural surjection $\mu : X \to \mathbb{R}_f$ (the *quotient map*) associated with this process, satisfying $\mu(x) = \mu(y)$ if and only if $x \sim y$. The input function f on X also induces a continuous function $\tilde{f} : \mathbb{R}_f \to \mathbb{R}$ defined as $\tilde{f}(z) = f(x)$ for any preimage $x \in \mu^{-1}(z)$ of z. To simplify notation, we often write $f : \mathbb{R}_f \to \mathbb{R}$ instead of \tilde{f} when there is no ambiguity. In all illustrations of this paper, we plot the Reeb graph with the height of a point z corresponding to f(z).



Given a point $x \in R_f$, we use the term *up-degree* (resp. *down-degree*) of x to denote the number of branches incident to x that are higher (resp. lower) than x w.r.t. values of f. A point is *regular* if both of its up-degree and down-degree equal to 1, and *critical* otherwise. A critical point is a minimum (maximum) if it has down-degree 0 (up-degree 0), and a down-fork saddle (up-fork saddle) if it has down-degree (up-degree) larger than 1. A critical point can be degenerate, having more than one types of criticality (such as both a minimum and an up-fork saddle if it has down-degree 0 and up-degree 2). From now on, we use the term *node* to refer to a critical point in the Reeb graph. For simplicity of exposition, we assume that all nodes of the Reeb graph have distinct f function values. Note that because of the monotonicity of f at regular points, the Reeb graph together with its associated function is completely described by the function values on the nodes.

Persistent homology and persistence diagrams. The notion of persistence was originally introduced by Edelsbrunner et al. in [16]. There has since been a great amount of development both in theory and in applications; see e.g [37, 7, 10]. This paper does not concern the theory of persistence, hence below we only provide a simple description so as to introduce the notion of *persistence diagrams*, which will be used later. We refer the readers to [25] for a more detailed treatment of homology groups in general and [15] for persistent homology.

Given a function $f : X \to \mathbb{R}$ defined on a topological space, we call $X_{\leq a} = \{x \in X \mid f(x) \leq a\}$ the *sublevel* set of X w.r.t. f. Let $H_p(Y)$ denote the p-th homology group of a topological space Y. In this paper, we always consider homology with \mathbb{Z}_2 coefficients, so $H_p(Y)$ is a vector space. As we sweep through X in increasing values of a we inspect the changes in $H_p(X_{\leq a})$. Sometimes, new homology classes (such as a family of loops, which gives 1-dimensional homology classes) will be *created*. Sometimes, existing homology classes will be *destroyed* (such as a loop now bounds a surface patch and thus becomes null-homologous). More specifically, consider the following sequence of vector spaces,

$$0 = \mathsf{H}_p(\mathsf{X}_{\leq a_0}) \to \mathsf{H}_p(\mathsf{X}_{\leq a_1}) \to \dots \to \mathsf{H}_p(\mathsf{X}_{\leq a_N}) = \mathsf{H}_p(\mathsf{X}), \tag{1}$$

where each homomorphism μ_i^j : $H_p(X_{\leq a_i}) \to H_p(X_{\leq a_j})$ is induced by the canonical inclusion $X_{\leq a_i} \hookrightarrow X_{\leq a_j}$. Throughout this paper, we assume that *f* is *tame*, i.e., $N < \infty$ and $H_p(X_{\leq a_i}) < \infty$ for all *i*. A homology class



Figure 1: (a) The two trees have the same persistence diagrams (thus the bottleneck distance between their persistence diagrams is 0), but their tree structures are different. The functional distortion distance measure will differentiate these two cases. In (b), solid dots are minimum and maximum, empty dots are essential saddles, and crossed-dots are ordinary saddles. The ordinary saddle x_6 merges components C_1 and C_2 in the sublevel set below it, represented by minima x_1 and x_2 respectively. The resulting critical pair (x_2, x_6) gives rise to the point (a_2, x_6) in Dg₀(R_f) in (c), where $a_i = f(x_i)$ for $i \in [1, 12]$. The essential saddle x_9 is paired with the up-fork saddle x_4 , corresponding to the thin loop $x_4x_8x_6x_5x_9x_4$ created at x_9 . This gives rise to the point (a_4, a_9) in the extended persistence diagram ExDg₁(R_f) in (d).

is created at time a_i if it exists in $H_p(X_{\leq a_i})$ but does not have preimage in $H_p(X_{\leq a_{i-1}})$ under the map μ_{i-1}^i . A homology class is destroyed at time a_j if it exists in $H_p(X_{\leq a_j})$ but its image is zero in $H_p(X_{\leq a_{j+1}})$. Persistent homology records such birth and death events, and one can encode the information in the *ordinary persistence diagram*. In particular, the *p*-th ordinary persistence diagram of *f*, denoted by $Dg_p(f)$, consists of a multi-set of points in the 2-dimensional plane, where each point (b, d) corresponds uniquely to the birth time *b* and death time *d* of some *p*-dimensional homology class. See Figure 1 (c) for an example of the 0-th dimensional persistence diagram. In general, since $H_p(X)$ may not be trivial, any nontrivial homology class of $H_p(X)$, referred to as an *essential homology class*, will never die during the sequence (1). For example, there is a point (a_1, ∞) in Figure 1 (b) indicating a 0-dimensional homology class that was created at a_1 but never dies.

One can use the following sequence, where an extended sequence of relative homology groups is appended at the end of Sequence (1), to obtain a pairing of the essential homology classes (i.e, homology classes of $H_p(X)$):

$$0 = \mathsf{H}_p(\mathsf{X}_{\leq a_0}) \to \dots \to \mathsf{H}_p(\mathsf{X}) \to \mathsf{H}_p(\mathsf{X}, \mathsf{X}_{\geq a_N}) \to \mathsf{H}_p(\mathsf{X}, \mathsf{X}_{\geq a_{N-1}}) \to \dots \to \mathsf{H}_p(\mathsf{X}, \mathsf{X}_{\geq a_0}) = 0.$$
(2)

Here $X_{\geq a}$ denotes the *super-level set* $X_{\geq a} = \{x \in X \mid f(x) \geq a\}$. Since the last vector space $H_p(X, X_{\geq a_0}) = 0$, each essential homology class will necessarily die in the *extended part* of the above sequence at some relative homology group $H_p(X, X_{\geq a_j})$. We refer to the multi-set of points encoding the birth and death time of *p*-dimensional homology classes born in the original part and die in the extended part of Sequence (2), as the *p*-th extended persistence diagram of *f*, denoted by $ExDg_p(f)$. In particular, each point (b, d) in $ExDg_p(f)$ uniquely refers to a *p*-dimensional (essential) homology class that is born in $H_p(X_{\leq b})$ and died after $H_p(X, X_{\geq d})$. See Figure 1 (d) for an example.

Relation of Reeb graph and persistent homology. There is a natural way to define "features" of the Reeb graph and their "importance", which turns out to be consistent with information encoded in the diagrams $Dg_0(R_f)$ and $ExDg_1(R_f)$ of the function $f : R_f \to \mathbb{R}$. Since R_f is a graph, we only need to consider the 0- and 1-dimensional (persistent) homology. We provide an intuitive treatment below.

Imagine that we sweep through R_f in increasing values of *a* and inspect the changes in $H_0((R_f)_{\leq a})$ (i.e, the connected components of the sublevel set $(R_f)_{\leq a}$). New components in the sublevel sets are created at local minima of R_f . At any moment, associate each connected component *C* in the sublevel set with the lowest local minimum *m* contained in *C*: intuitively, *C* is created by *m*.

Consider a down-fork saddle node s with a = f(s); for now, assume for the sake of simplicity that its down degree is 2. By the construction of Reeb graph, s merges two connected components from the level set just below

it into a single one in the level set $(R_f)_a$. If these two connected components are also disjoint within the open sublevel set $(R_f)_{<a}$, for reasons that will become obvious soon we call *s* an *ordinary saddle*; otherwise, it is an *essential saddle*. In other words, an ordinary saddle not only merges two components in the *level set*, it also merges two components, say C_1 and C_2 , in the *sublevel set*. Let x_1 and x_2 be the global minimum of C_1 and C_2 , respectively. After the merging in the sublevel set $(R_f)_{\le a}$, the new component is represented by the lower of x_1 and x_2 . We say that the component represented (created) by the higher minimum, say x_2 , is killed by the ordinary saddle *s*, and we obtain a critical pair (x_2, s) . The critical pair (x_2, s) gives rise to a unique point $(f(x_2), f(s))$ in the 0-th ordinary persistence diagram $Dg_0(R_f)$, representing a connected component created at time $f(x_2)$ and killed at f(s). Indeed, there is a one-to-one correspondence between the set of such pairs of minima and ordinary down-fork saddles and points in the 0th persistence diagram $Dg_0(R_f)$ with finite coordinates; see Figure 1 (b) and (c). A symmetric procedure with -f will produce critical pairs between maxima and ordinary up-fork saddles, corresponding to points in the 0th persistence diagram $Dg_0(R_{-f})$. These critical pairs capture the *branching features* of a Reeb graph.

For an essential saddle *s* with down-degree 2, its two lower branches are connected in the *sublevel set*; see Figure 1 (b) and (d). Hence, the introduction of s will create a family of cycles Γ_s in the sublevel set $(R_f)_{\leq f(s)}$, all with *s* being the highest point. Since R_f is a graph, these cycles are non-trivial in R_f , and their corresponding homology classes will not be destroyed in ordinary persistent homology. Consider the cycle γ with largest minimum value of *f* among all cycles in Γ_s . We use γ to represent the *cycle feature* (1-cycle) created by *s* and pair *s* with the point *s'* achieving the minimum on γ . The persistence of this feature is f(s) - f(s'), which measures the size (height) of this cycle w.r.t. *f*. It turns out that *s'* is necessarily an essential up-fork saddle [1], and we call such a pair (s', s) an *essential pair*. The essential pair (s', s) gives rise to a unique point (f(s'), f(s)) in the 1st extended persistence diagram of *f*. Indeed, the collection of essential pairs has a one-to-one correspondence to points in $\text{ExDg}_1(R_f)$. (The extended persistence diagram $\text{ExDg}_1(R_{-f})$ is the reflection of $\text{ExDg}_1(R_f)$ and thus encodes the same information as $\text{ExDg}_1(R_f)$.) These essential pairs capture the *cycle features* of a Reeb graph.

In short, the branching features and cycle features of a Reeb graph give rise to points in the 0th ordinary and 1st extended persistence diagrams, respectively. However, the persistence diagram captures only the lifetime of features, but not how these features are connected; see Figure 1 (a). In this paper we aim to develop a way of measuring distance between Reeb graphs which also takes into account the graph structure.

3 A Distance Measure for Reeb Graphs

From now on, consider two Reeb graphs R_f and R_g , with functions $f : R_f \to \mathbb{R}$ and $g : R_g \to \mathbb{R}$ defined on them. While topologically each Reeb graph is simply a graph, it is important to note that it also has a function define on it (induced from the input scalar field). Hence the distance measure should depend on both the graph structures and the functions f and g. Approaching the problem through graph isomorphisms does not seem viable, as small perturbation of the function f can create arbitrary number of new branches and loops in the graph. To this end, we first put the following metric structure on a Reeb graph R_f to capture the function f information.

Specifically, for any two points (not necessarily nodes) $u, v \in R_f$, let π be a path between u and v. The *range* of this path is the interval range(π) := [min_{$x \in \pi$} f(x), max_{$x \in \pi$} f(x)], and its *height* is simply the length of this interval, denoted by height(π) = max_{$x \in \pi$} $f(x) - min_{x \in \pi} f(x)$. We define the distance

$$d_f(u, v) = \min_{\pi: u \rightsquigarrow v} \operatorname{height}(\pi), \tag{3}$$

where π ranges over all paths from *u* to *v*. Note that this is in fact a metric, since on Reeb graphs there is no path of constant function value between two points $u \neq v$. We put *f* in the subscript to emphasize the dependency on the input function. Intuitively, $d_f(u, v)$ is the minimal function difference one has to overcome to move from *u* to *v*.

To define a distance between R_f and R_g , we need to connect the spaces R_f and R_g , which is achieved by

continuous maps ϕ_{\rightarrow} : $R_f \rightarrow R_g$ and ϕ_{\leftarrow} : $R_g \rightarrow R_f$. The functional distortion distance is defined as:

$$d_{FD}(\mathbf{R}_f, \mathbf{R}_g) := \inf_{\phi_{\to}: \mathbf{R}_f \to \mathbf{R}_g, \ \phi_{\leftarrow}: \mathbf{R}_g \to \mathbf{R}_f} \max\{\max_{x \in \mathbf{R}_f} |f(x) - g(\phi_{\to}(x))|, \ \max_{y \in \mathbf{R}_g} |g(y) - f(\phi_{\leftarrow}(y))|, \\ \max_{x \in \mathbf{R}_f} d_f(x, \phi_{\leftarrow} \circ \phi_{\to}(x)), \ \max_{y \in \mathbf{R}_g} d_g(y, \phi_{\to} \circ \phi_{\leftarrow}(y))\},$$
(4)

where ϕ_{\rightarrow} and ϕ_{\leftarrow} range over all continuous maps between R_f and R_g . The first two terms measure the distortion between the function values of f and g under these two maps, respectively. The last two terms measure intuitively how far these two maps are to the "inverse" of each other. In particular, the point $x' = \phi_{\leftarrow} \circ \phi_{\rightarrow}(x)$ and x are connected within some interval level set { $x \in R_f \mid f(x) \in I$ } for an interval I with height $d_f(x, \phi_{\leftarrow} \circ \phi_{\rightarrow}(x)) \le d_{FD}(R_f, R_g)$.

Remarks. Interestingly, it turns out that by treating R_f and R_g as metric graphs (equipped with the metrics d_f and d_g , respectively), $\delta = d_{FD}(R_f, R_g)$ approximates some variant of the *Gromov-Hausdorff distance* between these metric graphs within a constant factor; see Appendix A for details. Roughly speaking, this means that there exists correspondences of points in R_f and in R_g such that for any points $x, y \in R_f$, the distance $d_g(\tilde{x}, \tilde{y})$ between (any of) their correspondences $\tilde{x}, \tilde{y} \in R_g$ is close to the distance $d_f(x, y)$ between x and y with $|d_f(x, y) - d_g(\tilde{x}, \tilde{y})| = O(\delta)$. Furthermore, for any possible set of correspondences between R_f and R_g , this metric distortion $|d_f(x, y) - d_g(\tilde{x}, \tilde{y})|$ can be $\Theta(\delta)$ for some corresponding pairs (x, \tilde{x}) and (y, \tilde{y}) . That is, δ corresponds to the minimal metric distortion possible. Intuitively, the distance distortion between the two trees in Figure 1 (a) is large (no matter how we identify correspondences between points from them). Thus the functional distortion distance between them is also large, making it more discriminative than the bottleneck distance between persistence diagrams.

4 Properties of the functional distortion distance

In this section, we show that the functional distortion distance is both stable and discriminative. Note that it is somewhat meaningless to discuss the stability of a distance measure alone without understanding its discriminative power – It is easy to have a distance measure that is stable, such as simply always returning the value 0. The goal is to have a distance measure that can reflect the information encoded in the Reeb graph while being stable.

4.1 Stability and Relation to Ordinary Persistence Diagram

Suppose that the domains X, Y where f and g are defined are homeomorphic. (Note that this homeomorphism condition is only required for the stability result so that we can define distance between input scalar fields f and g. It is not necessary for Theorem 4.2 and 4.3.) Define the L_{∞} -distance between the functions f and g as $D_{\infty}(f,g) = \inf_{h:X \to Y} ||f - g \circ h||_{\infty}$, where h ranges over all homeomorphisms from X to Y. Then we have the following stability result for the metric d_{FD} for Reeb graphs:

Theorem 4.1 $d_{FD}(\mathbf{R}_f, \mathbf{R}_g) \leq D_{\infty}(f, g)$.

Proof: Simply consider the homeomorphism h that gives rise to $D_{\infty}(f,g)$. Composed with the quotient map between a domain and its Reeb graph, h and h^{-1} induce two maps ϕ^h_{\rightarrow} and ϕ^h_{\leftarrow} between R_f and R_g . Under ϕ^h_{\rightarrow} and ϕ^h_{\leftarrow} , the maximum of the four terms in the right-hand-side of Eqn (4) is $D_{\infty}(f,g)$.

The above result is similar to the stability result obtained for the bottleneck distance between persistence diagrams [13], as well as for the ε -interleaving distance between merge trees[24].

The main part of this section is devoted to discussing the discriminative power of the functional distortion distance for Reeb graphs. In particular, we relate this distance with the bottleneck distance between persistence diagrams. We have already seen in Figure 1 (a) that there are cases where the functional distortion distance

is strictly larger (more discriminative) than the bottleneck distance between persistence diagrams of according dimensions (0th ordinary and 1st extended persistence diagrams). Below we show that the functional distortion distance is always at least as large as (as discriminative as) the bottleneck distance. We will discuss separately the branching features (ordinary persistence diagram) and the cycle features (extended persistence diagram). For the former, we have the following main result. The proof is rather standard, and similar to the result on ε -interleaving distance between merge trees in [24]. It can be found in Appendix B.

Theorem 4.2 $d_B(Dg_0(R_f), Dg_0(R_g)) \le d_{FD}(R_f, R_g)$. Similarly, we have $d_B(Dg_0(R_{-f}), Dg_0(R_{-g})) \le d_{FD}(R_f, R_g)$.

4.2 Relation to Extended Persistence Diagram

Recall that the range of cycle features in the Reeb graph correspond to points in the 1st extended persistence diagram. In what follows we will show the follow main theorem, which states that $d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$ is at least as discriminative as the bottleneck distance between the 1st extended persistence diagrams for \mathbf{R}_f and \mathbf{R}_g , denoted by $\text{ExDg}_1(\mathbf{R}_f)$ and $\text{ExDg}_1(\mathbf{R}_g)$, respectively.¹

Theorem 4.3 $d_B(\operatorname{ExDg}_1(\mathbf{R}_f), \operatorname{ExDg}_1(\mathbf{R}_g)) \leq 3d_{FD}(\mathbf{R}_f, \mathbf{R}_g).$

For simplicity of exposition, we assume that $d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$ can be achieved by optimal continuous maps $\phi_{\rightarrow} : \mathbf{R}_f \rightarrow \mathbf{R}_g$ and $\phi_{\leftarrow} : \mathbf{R}_g \rightarrow \mathbf{R}_f$. The case where $d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$ is achieved in the limit can be handled by taking the limit of a sequence of continuous maps. Let $\delta = d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$.

Thin bases. Let $Z_1(R_f)$ be the 1-dimensional cycle group of R_f with \mathbb{Z}_2 coefficients. Since the Reeb graph has a graph structure, every 1-cycle of R_f represents a unique homology class in $H_1(R_f)$; that is, $H_1(R_f) \cong Z_1(R_f)$. Let range $(\gamma) = [\min_{x \in \gamma} f(x), \max_{x \in \gamma} f(x)]$ denote the *range* of a cycle γ , and let height (γ) be the length of this interval. A cycle is *thinner* than another one if its height is strictly smaller. A cycle γ is *thin* if there is no other cycle with smaller height than γ but having the same minimum or maximum value in f as γ . See Figure 1 (b), where the cycle $x_4x_8x_6x_5x_9x_4$ is thin, while the cycle $x_3x_4x_9x_5x_6x_3$ is not. Given a basis of $Z_1(R_f)$, consider the sequence of the heights of the cycles in it ordered in non-decreasing order. A basis for $Z_1(R_f)$ is a *thin* basis if its height sequence is smaller than or equal in lexicographic order to that of any other basis of $Z_1(R_f)$.

From now on, we fix an arbitrary thin basis $\mathcal{G}_f = \{\gamma_1, \dots, \gamma_n\}$ of $Z_1(R_f)$ and $\mathcal{G}_g = \{\zeta_1, \dots, \zeta_m\}$ of $Z_1(R_g)$, with n and m being the rank of $Z_1(R_f)$ and $Z_1(R_g)$, respectively. It is known [13] that every cycle in a thin basis of R_f is necessarily a thin cycle, and the ranges [b, d] of cycles in \mathcal{G}_f (resp. in \mathcal{G}_g) correspond one-to-one to the points (b, d) in the 1st extended persistence diagram $\text{ExDg}_1(R_f)$ (resp. in $\text{ExDg}_1(R_g)$). For example, in Figure 1 (b), the two cycles $x_3x_4x_8x_6x_3$ and $x_4x_8x_6x_5x_9x_4$ form a thin basis, corresponding to points $(f(x_3), f(x_8))$, and $(f(x_4), f(x_9))$ in $\text{ExDg}_1(R_f)$ in (d).

Given any cycle ℓ of R_f (resp. of R_g), we can represent ℓ uniquely as a linear combination of cycles in \mathcal{G}_f (resp. \mathcal{G}_g), which we call the *thin basis decomposition of* ℓ *with respect to* \mathcal{G}_f (*resp.* \mathcal{G}_g); we omit the reference to \mathcal{G}_f and \mathcal{G}_g since they will be fixed from now on. The thin cycle with the largest height from the thin basis decomposition of ℓ is called the *dominating cycle of* ℓ , denoted by dom(ℓ). If there are multiple cycles with the same maximal height, then by convention we choose the one with smallest index in \mathcal{G}_f (resp. in \mathcal{G}_g) as the dominating cycle. A cycle ℓ is α -stable if its dominating cycle has height strictly larger than 2α ; that is, height(dom(ℓ)) > 2α . The following property of the dominating cycle is easy to see (a proof is given in Appendix C.1).

Lemma 4.4 A set of cycles ℓ_1, \ldots, ℓ_k in R_f with distinct dominating cycles is linearly independent.

¹Some readers may wonder why we have not employed an approach similar to the one above handling branching features. A discussion can be found in Appendix C.6.

 α -matching. The main use of thin cycles is that we will prove Theorem 4.3 by showing the existence of an α -matching between \mathcal{G}_f and \mathcal{G}_g . Specifically, two thin cycles ℓ_1 and ℓ_2 are α -close if their ranges [a, b] and [c, d] are within Hausdorff distance α , i.e., $|c - a| \leq \alpha$ and $|d - b| \leq \alpha$. (Note that two α -close cycles can differ in height by at most 2α .) An α -matching for \mathcal{G}_f and \mathcal{G}_g is a set of pairs $\mathcal{M} = \{\langle \gamma_i, \zeta_j \rangle \mid \gamma_i \in \mathcal{G}_f, \zeta_j \in \mathcal{G}_g\}$ such that:

(I) For each pair $\langle \gamma_i, \zeta_j \rangle \in \mathcal{M}$, γ_i and ζ_j are α -close; and

(II) Every α -stable cycle in \mathcal{G}_f and \mathcal{G}_g (i.e., with height larger than 2α) appears in *exactly one* pair in \mathcal{M} .

Since each point (b, d) in the extended persistence diagram corresponds to the range [b, d] of a unique cycle in a given thin basis, $d_B(\text{ExDg}_1(\mathbb{R}_f), \text{ExDg}_1(\mathbb{R}_g)) \leq \alpha$ if and only if there is an α -matching for \mathcal{G}_f and \mathcal{G}_g . Hence our goal now is to prove that there exists a 3δ -matching for \mathcal{G}_f and \mathcal{G}_g , which will then imply Theorem 4.3.

Properties of ϕ_{\rightarrow} **and** ϕ_{\leftarrow} . Recall that $\phi_{\rightarrow} : \mathbb{R}_f \to \mathbb{R}_g$ and $\phi_{\leftarrow} : \mathbb{R}_g \to \mathbb{R}_f$ are the optimal continuous maps that achieve $\delta = d_{FD}(\mathbb{R}_f, \mathbb{R}_g)$. Below we first provide several results on the effect of these maps on thin cycles. Let $Z_1^{\delta}(\mathbb{R}_f)$ denote the subgroup of $Z_1(\mathbb{R}_f)$ generated by cycles with height at most 2δ . Note that a cycle is in $Z_1^{\delta}(\mathbb{R}_f)$ if and only if this cycle is *not* δ -stable. Lemma 4.5 below states that ϕ_{\leftarrow} is "close" to the inverse of ϕ_{\rightarrow} , and Lemma 4.6 relates the range of $\phi_{\rightarrow}(\gamma)$ with the range of γ . Their proofs are in Appendix C.2 and C.3, respectively.

Lemma 4.5 Given any cycle ℓ of \mathbb{R}_f , we have $\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell) \in \ell + \mathbb{Z}_1^{\delta}(\mathbb{R}_f)$. That is, $\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell) = \ell + \ell'$, where ℓ' is not δ -stable. A symmetric statement holds for any cycle in \mathbb{R}_g .

Lemma 4.6 Given any thin cycle γ of \mathbb{R}_f with range [b, d], we have that the range of any cycle in the thin basis decomposition of $\phi_{\rightarrow}(\gamma)$ must be contained in the interval $[b - \delta, d + \delta]$.

Lemma 4.7 For any cycle ℓ of R_f , we have height $(dom(\phi_{\rightarrow}(\ell))) \ge height(dom(\ell)) - 2\delta$. A symmetric statement holds for any cycle of R_g .

Proof: Let $\gamma_s = \operatorname{dom}(\ell)$. If height $(\gamma_s) \le 2\delta$, the claim follows directly from non negativity of the height. Now assume that height $(\gamma_s) > 2\delta$. First, we claim that $\operatorname{dom}(\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell)) = \gamma_s$. This is because by Lemma 4.5, $\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell) \in \ell + Z_1^{\delta}(\mathbb{R}_f)$. Since height $(\gamma_s) > 2\delta$, γ_s still belongs to the thin basis decomposition of $\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell)$ and still has the largest height.

Now set $\tilde{\zeta} = \phi_{\rightarrow}(\ell)$ with $\zeta_{i_1} + \cdots + \zeta_{i_a}$ being its thin basis decomposition. Observe that for any loop ℓ' in \mathbb{R}_f , we have that height $(\phi_{\rightarrow}(\ell')) \leq \text{height}(\ell') + 2\delta$, which follows from the fact that for any $x \in \mathbb{R}_f$, $|f(x) - g(\phi_{\rightarrow}(x))| \leq \delta$; recall Eqn (4). A symmetric statement holds for a loop from \mathbb{R}_g . We thus have:

$$\operatorname{height}(\operatorname{dom}(\phi_{\leftarrow}(\tilde{\zeta}))) = \operatorname{height}(\operatorname{dom}(\sum_{j=1}^{a} \phi_{\leftarrow}(\zeta_{i_{j}}))) \le \max_{j=1}^{a} \operatorname{height}(\phi_{\leftarrow}(\zeta_{i_{j}})) \le \max_{j=1}^{a} \operatorname{height}(\zeta_{i_{j}}) + 2\delta] = \max_{j=1}^{a} \operatorname{height}(\zeta_{i_{j}}) + 2\delta = \operatorname{height}(\operatorname{dom}(\tilde{\zeta})) + 2\delta.$$
(5)

Since we have shown earlier that $dom(\phi_{\leftarrow} \circ \phi_{\rightarrow}(\gamma)) = \gamma_s$, it follows that $dom(\phi_{\leftarrow}(\tilde{\zeta})) = dom(\phi_{\leftarrow} \circ \phi_{\rightarrow}(\gamma)) = \gamma_s = dom(\ell)$. Combining this with Eqn (5), we have that

$$\operatorname{height}(\operatorname{dom}(\phi_{\rightarrow}(\ell)) = \operatorname{height}(\operatorname{dom}(\tilde{\zeta})) \geq \operatorname{height}(\operatorname{dom}(\phi_{\leftarrow}(\tilde{\zeta}))) - 2\delta = \operatorname{height}(\operatorname{dom}(\ell)) - 2\delta.$$

In fact, while the above result is sufficient for our later argument, a more careful argument shows that $dom(\phi_{\rightarrow}(\gamma))$ is δ -close to γ if γ is a thin cycle. Intuitively, this already provides some mapping of base cycles from \mathcal{G}_f to cycles from \mathcal{G}_g such that each pair of corresponding cycles are δ -close. However, the main challenge is to show that there exists a *one-to-one* correspondence for all 3δ -stable cycles (recall the definition of a 3δ -matching of \mathcal{G}_f and \mathcal{G}_g). For this, we need to take a slight detour to relate cycles in \mathcal{G}_f with those in \mathcal{G}_g in a stronger sense:

Proposition 4.8 For any thin cycle γ_k from \mathcal{G}_f , one can compute a (not necessarily thin) cycle $\widehat{\gamma}_k$ such that $\gamma_k = \operatorname{dom}(\widehat{\gamma}_k)$ and:

$$\phi_{\rightarrow}(\widehat{\gamma}_k) \in \sum_{j=1}^r \zeta_{k_j} + \mathsf{Z}_1^{\delta}(\mathsf{R}_g),$$

where each ζ_{k_i} is 3 δ -close to γ_k , for any $j \in [1, r]$.

Proof: Assume the thin basis decomposition of $\phi_{\rightarrow}(\gamma_k)$ has the form:

$$\phi_{\rightarrow}(\gamma_k) \in \sum_{j=1}^r \zeta_{k_j} + \sum_{j=1}^s \tilde{\zeta}_j + \mathsf{Z}_1^{\delta}(\mathsf{R}_g),$$

where the first term contains the "good" thin cycles whose range is 3δ -Hausdorff close to range(γ_k), the last term contains the "small" thin cycles that are δ -stable, and the middle term contains those "bad" thin cycles $\tilde{\zeta}_j$ that are neither 3δ -close to γ_k nor small. (If a "good" thin cycle 3δ -close to γ_k is also δ -stable, i.e, "small", we put it in either category.) We wish to remove the middle term $\tilde{\zeta} = \sum_{j=1}^{s} \tilde{\zeta}_j$. Set $\gamma' = \phi_{\leftarrow}(\tilde{\zeta})$. It then follows from Lemma 4.5 that $\phi_{\rightarrow}(\gamma') = \phi_{\rightarrow} \circ \phi_{\leftarrow}(\tilde{\zeta}) \in \tilde{\zeta} + Z_1^{\delta}(R_g)$. Set $\widehat{\gamma}_k := \gamma_k + \gamma'$. It is then easy to verify that $\phi_{\rightarrow}(\widehat{\gamma}_k) = \phi_{\rightarrow}(\gamma_k) + \phi_{\rightarrow}(\gamma') \in \sum_{j=1}^{r} \zeta_{k_j} + Z_1^{\delta}(R_g)$.

What remains is to show that dom($\widehat{\gamma}_k$) = γ_k . Set [b, d] := range(γ_k). By Lemma 4.6, range($\widetilde{\zeta}_j$) $\subseteq [b - \delta, d + \delta]$, for any $j \in [1, s]$. Since the each $\widetilde{\zeta}_j$ is not 3 δ -close to γ_k , it is then necessary that either range($\widetilde{\zeta}_j$) $\subseteq [b - \delta, d - 3\delta)$ or range($\widetilde{\zeta}_j$) $\subseteq (b + 3\delta, d + \delta]$, for any $j \in [1, s]$. This means that height($\widetilde{\zeta}_j$) $< d - b - 2\delta$. Apply Lemma 4.6 to the loop $\widetilde{\zeta}_j$. We have that the height of any loop in the thin basis decomposition of $\phi_{\leftarrow}(\widetilde{\zeta}_j)$ is *strictly smaller* than $d - b - 2\delta + 2\delta = d - b$ for any $j \in [1, s]$. Hence all loops in the thin basis decomposition of $\gamma' = \phi_{\leftarrow}(\widetilde{\zeta})$ have a height strictly smaller than d - b (= height(γ_k)). This means γ_k has the largest height among the loops of the thin basis decomposition of $\widehat{\gamma}_k = \gamma_k + \gamma'$, and it is the only loop with this largest height. It then follows that $\gamma_k = \text{dom}(\widehat{\gamma}_k)$.

Corollary 4.9 Let $\widehat{\mathcal{G}}_f$ denote the set of cycles $\{\widehat{\gamma}_k\}_{k=1}^n$, where each $\widehat{\gamma}_k$ is as specified in Proposition 4.8. $\widehat{\mathcal{G}}_f$ forms a basis for $Z_1(R_f)$.

Proof: Since the dominating cycles for cycles in $\widehat{\mathcal{G}}_f$ are all distinct, it follows from Lemma 4.4 that all cycles in $\widehat{\mathcal{G}}_f$ are linearly independent. Hence $\widehat{\mathcal{G}}_f$ also forms a (not necessarily thin) basis for $Z_1(R_f)$.

Let Φ denote the matrix of the mapping from the base cycles in $\widehat{\mathcal{G}}_f$ (columns, domain) to those in \mathcal{G}_g (rows, range) as induced by ϕ_{\rightarrow} , i.e., the *i*th column of Φ specifies the representation of $\phi_{\rightarrow}(\widehat{\gamma}_i)$ using basis elements from \mathcal{G}_g , with $\Phi[i][j] = 1$ if ζ_j is in the thin basis decomposition of $\phi_{\rightarrow}(\widehat{\gamma}_i)$. Let $\overline{\Phi}$ be the submatrix of Φ with columns corresponding to basis elements $\widehat{\gamma}_i$ that are 3δ -stable, and rows corresponding to basis elements ζ_j that are δ -stable. See Figure 2 (a). By Proposition 4.8, $\overline{\Phi}[i][j] = 1$ implies that the basis element $\zeta_j \in \mathcal{G}_g$ is 3δ -close to the basis element $\gamma_i \in \mathcal{G}_f$. Recall that our goal is to show that there is a 3δ -matching for \mathcal{G}_f and \mathcal{G}_g . Intuitively, non-zero entries in $\overline{\Phi}$ will provide potential matchings for basis elements in \mathcal{G}_f to establish a 3δ -matching between \mathcal{G}_f and \mathcal{G}_g that we need. The proofs of the following two results can be found in Appendix C.4 and C.5, respectively.

Lemma 4.10 The columns of $\widetilde{\Phi}$ are linearly independent.

Corollary 4.11 We can identify a unique row index *j* for each column index *i* in $\tilde{\Phi}$ such that $\tilde{\Phi}[i][j] = 1$.

Proof of Theorem 4.3. Recall that by Proposition 4.8, $\overline{\Phi}[i][j] = 1$ implies that the cycle γ_i and ζ_j are 3δ -close. It follows from Corollary 4.11 that there is an injective map F from the set of 3δ -stable cycles in \mathcal{G}_f to cycles in \mathcal{G}_g such that each pair of corresponding cycles are 3δ -close. By a symmetric argument (switching the role of R_f

and R_g), there is also an injective map G from the 3δ -stable cycles in \mathcal{G}_g to cycles in \mathcal{G}_f where each corresponding pair of cycles are 3δ -close. However, F and G may not be consistent and do not directly give rise to a 3δ -matching of \mathcal{G}_f and \mathcal{G}_g yet. In what follows, we will modify F to obtain another injective map \widehat{F} such that (i) for any cycle of \mathcal{G}_f , its image under \widehat{F} is 3δ -close to it; (ii) any 3δ -stable cycle in \mathcal{G}_f is mapped by \widehat{F} to a cycle in \mathcal{G}_g that is 3δ -close, and (iii) all 3δ -stable cycles in \mathcal{G}_g are contained in im \widehat{F} , the image of \widehat{F} . Note that \widehat{F} provides exactly the set of correspondences necessary in a 3δ -matching between \mathcal{G}_f and \mathcal{G}_g . As discussed at the beginning of this section, this then means that $d_B(\text{ExDg}_1(R_f), \text{ExDg}_1(R_g)) \leq 3\delta = 3d_{FD}(R_f, R_g)$, proving Theorem 4.3.



Figure 2: (a). The *i*th column in the matrix Φ specifies the representation of $\phi_{\rightarrow}(\widehat{\gamma}_i)$ using the basis elements in \mathcal{G}_g . The shaded submatrix represents $\overline{\Phi}$. (b). A bipartite graph view of the augmenting process. Left: Thick path alternates between an *F*-induced (solid) and an *G*-induced (dash-dotted) edge. Right: Thick solid edges are induced by the modified injective map \widetilde{F} (which are used to be *G*-induced edges in the left figure).

It remains to show how to construct $\widehat{\mathsf{F}}$ satisfying conditions (i), (ii) and (iii) above. Conditions (i) and (ii) already hold for F , so the main task is to establish condition (iii) while maintaining (i) and (ii). Start with $\widehat{\mathsf{F}} = \mathsf{F}$. Let $y_0 \notin \inf \widehat{\mathsf{F}}$ be any 3δ -stable cycle in \mathcal{G}_g that is not yet in $\inf \widehat{\mathsf{F}}$. Let $x_1 = \mathsf{G}(y_0) \in \mathcal{G}_f$. We continue with $y_i = \widehat{\mathsf{F}}(x_i)$ if x_i is 3δ -stable. Next, if y_i is 3δ -stable, then set $x_{i+1} = \mathsf{G}(y_i)$. We repeat this process, until we reach x_k or y_k which is not 3δ -stable any more. At this time, we modify $\widehat{\mathsf{F}}(x_i)$ to be y_{i-1} for each $j \in [1, k]$ (originally, $\widehat{\mathsf{F}}(x_i) = y_i$). Throughout this process, all y_i other than y_0 are already in $\inf \widehat{\mathsf{F}}$. After the modification of $\widehat{\mathsf{F}}$, we have $y_0 \in \inf \widehat{\mathsf{F}}$, while all other y_i remain in $\inf \widehat{\mathsf{F}}$. The only exception is when the above process terminates by reaching some y_k which is not 3δ -stable (the termination condition), in which case y_k will not be in $\inf \widehat{\mathsf{F}}$ after the modification of $\widehat{\mathsf{F}}$. Nevertheless, the number of 3δ -stable cycles contained in $\inf \widehat{\mathsf{F}}$ increases by one (i.e., y_0) by the above process. An alternative way to view this is to consider the specific bipartite graph $G' = (P \cup Q, E')$, where nodes in P and Q correspond to basis cycles in \mathcal{G}_f and \mathcal{G}_g , respectively, and edges E' are those corresponding to a cycle and its image under either the map $\widehat{\mathsf{F}}$ or G . The sequence y_0, x_1, y_1, \ldots specifies a path with edges alternating between the $\widehat{\mathsf{F}$ -induced matching along this path, much similar to the use of augmenting paths to obtain maximum bipartite matching. See Figure 2 (b) for an illustration.

We repeat the above path augmentation process for any remaining 3δ -stable cycle in $\mathcal{G}_g \setminus \operatorname{im} \widehat{\mathsf{F}}$. This process will terminate because after each augmentation process, the number of 3δ -stable cycles contained in $\operatorname{im} \widehat{\mathsf{F}}$ strictly increases. In the end, we obtain an injective map $\widehat{\mathsf{F}}$ from the set of 3δ -stable cycles of \mathcal{G}_f to cycles in \mathcal{G}_g such that $\operatorname{im} \widehat{\mathsf{F}}$ contains all 3δ -stable cycles of \mathcal{G}_g . Putting everything together proves Theorem 4.3.

5 Simplification of Reeb Graphs

Simplifying a Reeb graph can help to remove noise and to create multi-resolution representation of input scalar fields; see e.g, [14, 17, 27]. As described in Section 2, there is a natural way to define "importance" of branching



Figure 3: (a) Removing a branching feature spanned by (m, s) by merging paths π_1 and π_2 . This removes the point (f(m), f(s)) from the 0-th ordinary persistence diagram. (b) Removing a cycle feature spanned by (s_1, s_2) ; this removes the point $(f(s_1), f(s_2))$ from the 1st extended persistence diagram. There are other branches potentially connected to nodes such as a, b, c, d, represented as dashed un-finished curves, which will not be affected by the merging process.

and cycle features, which is equivalent to ordinary and extended persistence in according dimensions. It is common in practice to simplify the Reeb graph by removing all features with persistence smaller than a certain threshold δ . In Figure 3, we illustrate a natural merging strategy (see e.g, [17, 27]) to remove branching and cycle features; a detailed description of this simplification scheme is in Appendix D.1. The main result is the following.

Theorem 5.1 Suppose we simplify a Reeb graph R by removing features of persistence $\leq \delta$ using the strategy shown in Figure 3 and detailed in Appendix D.1. The bottleneck distance between the (ordinary and extended) persistence diagram for R and for its simplification \widetilde{R} is at most $\delta\delta$.

This result is obtained by showing that the functional distortion distance $d_{FD}(\mathbf{R}, \mathbf{R})$ between a Reeb graph and its simplification is bounded by 2δ , where δ is the largest persistence of features removed. See Appendix D for details. Interestingly, we remark that a similar approach can also be used to show that the distance between the standard Reeb graph for a function f and the α -Reeb graph as introduced in [11] is bounded by 3α . We omit the rather straightforward details here.

6 Concluding Remarks

In this paper, we propose a distance measure for Reeb graphs, under which the Reeb graph is stable with respect to changes in the input function under the L_{∞} norm. More importantly, we show that this distance measure is bounded from below by and is strictly more discriminative at differentiating scalar fields than the bottleneck distance between both 0-th ordinary and 1st extended persistence diagrams. Similar to the use of the Gromov-Hausdorff distance for metric spaces, having the Reeb graph distance metric provides a formal language for describing and studying various properties of the Reeb graphs. Indeed, we have already shown that, by bounding the functional distortion distance between a Reeb graph and its simplification, we can prove that important (persistent) features are preserved under simplification, which addresses a key practical issue.

A natural question is how to compute the functional distortion distance for Reeb graphs. We believe that there is an exponential time algorithm to approximate $d_{FD}(R_f, R_g)$, similar to what is known for the ε -interleaving distance for merge trees [24]. However, it remains an open prolbem to develop more efficient algorithms. We remark that comparing unlabeled trees is a computationally hard problem in general: The commonly used tree edit distance and tree alignment distance are NP-hard to compute (and sometimes even to approximate) [6, 21]. Similarly, no polynomial time algorithms yet exist for the Gromov-Hausdorff distance in general. It will be interesting to see whether by using the extra scalar field information associated with merge trees and Reeb graphs, more efficient exact or approximation algorithms for computing functional distortion distance can be developed.

References

- [1] P. K. Agarwal, H. Edelsbrunner, J. Harer, and Y. Wang. Extreme elevation on a 2-manifold. *Discrete and Computational Geometry (DCG)*, 36(4):553–572, 2006.
- [2] M. Attene, S. Biasotti, and M. Spagnuolo. Shape understanding by contour driven retiling. *The Visual Computer*, 19(2-3):127–138, 2003.
- [3] K. Beketayev, D. Yeliussizov, D. Morozov, G. H. Weber, and B. Hamann. Measuring the distance between merge trees. In Workshop on Topological Methods in Data Analysis and Visualization: Theory, Algorithms and Applications (TopoInVis'13), 2013.
- [4] S. Biasotti, B. Falcidieno, and M. Spagnuolo. Extended Reeb graphs for surface understanding and description. In Proc. 9th Internat. Conf. Discrete Geom. for Computer Imagery, pages 185–197, 2000.
- [5] S. Biasotti, D. Giorgi, M. Spagnuolo, and B. Falcidieno. Reeb graphs for shape analysis and applications. *Theor. Comput. Sci.*, 392(1-3):5–22, 2008.
- [6] P. Bille. A survey on tree edit distance and related problems. Theor. Comput. Sci., 337(1-3):217-239, 2005.
- [7] G. Carlsson and V. de Silva. Zigzag persistence. Foundations of Computational Mathematics, 10(4):367–405, 2010.
- [8] G. Carlsson, V. de Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. In Proc. 25th Annu. ACM Sympos. Comput. Geom., pages 247–256, 2009.
- [9] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Oudot. Proximity of persistence modules and their diagrams. In Proc. 25th ACM Sympos. on Comput. Geom., pages 237–246, 2009.
- [10] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. The structure and stability of persistence modules, 2013. arXiv:1207.3674.
- [11] F. Chazal and J. Sun. Gromov-Hausdorff approximation of metric spaces with linear structure, 2013. arXiv:1305.1172.
- [12] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. Discrete & Computational Geometry, 37(1):103–120, 2007.
- [13] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.
- [14] H. Doraiswamy and V. Natarajan. Output-sensitive construction of Reeb graphs. *IEEE Trans. Vis. Comput. Graph.*, 18(1):146–159, 2012.
- [15] H. Edelsbrunner and J. Harer. Computational Topology: An Introduction.
- [16] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
- [17] X. Ge, I. Safa, M. Belkin, and Y. Wang. Data skeletonization via Reeb graphs. In Proc. 25th Annu. Conf. Neural Information Processing Systems (NIPS), pages 837–845, 2011.
- [18] W. Harvey, R. Wenger, and Y. Wang. A randomized O(m log m) time algorithm for computing Reeb graph of arbitrary simplicial complexes. In Proc. 25th Annu. ACM Sympos. Compu. Geom., pages 267–276, 2010.

- [19] F. Hétroy and D. Attali. Topological quadrangulations of closed triangulated surfaces using the Reeb graph. *Graph. Models*, 65(1-3):131–148, 2003.
- [20] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proc. SIGGRAPH '01*, pages 203–212, 2001.
- [21] T. Jiang, L. Wang, and K. Zhang. Alignment of trees an alternative to tree edit. *Theor. Comput. Sci.*, 143(1):137–148, 1995.
- [22] P. Kanongchaiyos and Y. Shinagawa. Articulated Reeb graphs for interactive skeleton animation. In S. Hashimoto, editor, *Multimedia Modeling: Modeling Multimedia Information and System*, pages 451–467. World Scientific, 2000.
- [23] F. Mémoli. On the use of Gromov-Hausdorff distances for shape comparison. In *Sympos. Point Based Graphics*, pages 81–90, 2007.
- [24] D. Morozov, K. Beketayev, and G. Weber. Interleaving distance between merge trees. In Workshop on Topological Methods in Data Analysis and Visualization: Theory, Algorithms and Applications (TopoInVis'13), 2013.
- [25] J. R. Munkres. Elements of Algebraic Topology. Westview Press, 1996.
- [26] S. Parsa. A deterministic O(m log m) time algorithm for the Reeb graph. In ACM Sympos. Comput. Geom. (SoCG), pages 269–276, 2012.
- [27] V. Pascucci, G. Scorzelli, P.-T. Bremer, and A. Mascarenhas. Robust on-line computation of Reeb graphs: simplicity and speed. ACM Trans. Graph. (SIGGRAPH 2007), 26(3), 2007.
- [28] G. Patanè, M. Spagnuolo, and B. Falcidieno. Para-graph: Graph-based parameterization of triangle meshes with arbitrary genus. *Comput. Graph. Forum*, 23(4):783–797, 2004.
- [29] Y. Shi, R. Lai, S. Krishna, N. Sicotte, I. Dinov, and A. W. Toga. Anisotropic Laplace-Beltrami eigenmaps: Bridging Reeb graphs and skeletons. *Computer Vision and Pattern Recognition Workshop*, 0:1–7, 2008.
- [30] Y. Shinagawa, T. Kunii, and Y. L. Kergosien. Surface coding based on Morse theory. *IEEE Comput. Graph. Appl.*, 11(5):66–78, 1991.
- [31] G. Singh, F. Mémoli, and G. Carlsson. Topological methods for the analysis of high dimensional data sets and 3D object recognition. In *Eurograph. Sympos. Point-Based Graphics*, pages 91–100, 2007.
- [32] J. Tierny. Reeb graph based 3D shape modeling and applications. PhD thesis, Université des Sciences et Technologies de Lille, 2008.
- [33] T. Tung and F. Schmitt. The augmented multiresolution Reeb graph approach for content-based retrieval of 3D shapes. *Internat. J. Shape Modeling*, 11(1):91–120, 2005.
- [34] J. H. van Lint and R. M. Wilson. A course in Combinatorics. Cambridge University Press, 1992.
- [35] Z. Wood, H. Hoppe, M. Desbrun, and P. Schröder. Removing excess topology from isosurfaces. ACM Trans. Graph., 23(2):190–208, 2004.
- [36] E. Zhang, K. Mischaikow, and G. Turk. Feature-based surface parameterization and texture mapping. *ACM Trans. Graph.*, 24(1):1–27, 2005.
- [37] A. Zomorodian and G. E. Carlsson. Computing persistent homology. Discrete Comput. Geom., 33(2):249– 274, 2005.

A Relation to Gromov-Hausdorff Distance

We can view the Reeb graphs R_f and R_g as metric graphs $\mathcal{R}_f = (R_f, d_f)$ and $\mathcal{R}_g = (R_f, d_g)$, equipped with metrics d_f and d_g , respectively. A natural distance for metric spaces is the *Gromov-Hausdorff* distance, which we define below, following the definitions and results from [23].

Definition A.1 (Correspondence [23]) For sets A and B, a subset $C \subset A \times B$ is a correspondence (between A and B) if and only if:

(1) $\forall a \in A$, there exists $b \in B$ s.t. $(a, b) \in R$;

(2) $\forall b \in B$, there exists $a \in A$ s.t. $(a, b) \in R$.

Let $\Pi(A, B)$ denote the set of all possible correspondences between sets A and B.

Definition A.2 (GH-distance, Proposition 5 of [23]) *The Gromov-Hausdorff distance between two metric spaces* $X = (X, d_X)$ and $\mathcal{Y} = (Y, d_Y)$ is

$$d_{GH}(\mathcal{X}, \mathcal{Y}) = \frac{1}{2} \inf_{C \in \Pi(A,B)} \max_{(x,y), (x',y') \in C} |d_X(x,x') - d_Y(y,y')|.$$

We can now talk about the Gromov-Hausdorff distance between R_f and R_g . However, note that translation f + c and negation -f do not change the metric structures of the Reeb graph R_f . However, such changes in the functions do affect our functional distortion distance. To remove the effect of the difference in the function values, we now define the following *functional GH distance* between R_f and R_g , which measures not only the distance distortion, but also the changes in the function value between corresponding points.

Definition A.3 (functional GH distance) The functional GH distance between two metric Reeb graphs $\mathcal{R}_f = (\mathcal{R}_f, d_f)$ and $\mathcal{R}_g = (\mathcal{R}_f, d_g)$ is defined as:

$$d_{fGH}(\mathcal{R}_f, \mathcal{R}_g) = \inf_{C \in \Pi(\mathcal{R}_f, \mathcal{R}_g)} \max\{\max_{(x, y) \in C} |f(x) - g(y)|, \max_{(x, y), (x', y') \in C} |d_f(x, x') - d_g(y, y')\}.$$

It turns out that we have the following relations, which imply that our functional distortion distance roughly measures the minimum distortion in both function values (between f and g) and in their induced metrics (between d_f to d_g). We note that this relation does not generalize to spaces with dimension higher than 1.

Theorem A.4 $\frac{1}{3}d_{FD}(\mathbf{R}_f,\mathbf{R}_g) \leq d_{fGH}(\mathcal{R}_f,\mathcal{R}_g) \leq 4d_{FD}(\mathbf{R}_f,\mathbf{R}_g).$

Proof: **Part I:** First, we prove the left inequality $\frac{1}{3}d_{FD}(R_f, R_g) \leq d_{fGH}(\mathcal{R}_f, \mathcal{R}_g)$. Fix an arbitrary positive real value ε . Let C denote an ε -optimal correspondence, i.e., the maximum of the two terms in Def. A.3 is less than or equal to $d_{fGH}(\mathcal{R}_f, \mathcal{R}_g) + \varepsilon$. Set $\beta = d_{fGH}(\mathcal{R}_f, \mathcal{R}_g)$. Our final goal is to show that $d_{FD}(R_f, R_g) \leq 3\beta$. We do this by constructing *continuous* maps $\phi_{\rightarrow}^{\varepsilon} : R_f \rightarrow R_g$ and $\phi_{\leftarrow}^{\varepsilon} : R_g \rightarrow R_f$ based on the ε -optimal correspondence C such that each of the four terms in Eqn (4) can be bounded by $3\beta + \Theta(\varepsilon)$.

We now show how to construct a certain continuous map $\phi_{\rightarrow}^{\varepsilon}$: $\mathbf{R}_f \rightarrow \mathbf{R}_g$. First, construct an ε -subdivision of \mathbf{R}_f as described in Section D.2 (see also Figure 4): the graph \mathbf{R}_f is augmented with a set of nodes $V_{\varepsilon} = \{v_1, \ldots, v_N\}$ such that f is monotonic on each resulting arc, and the height of an arc $v_i v_j$ (which is $|f(v_i) - f(v_j)|$ since f is monotonic on $v_i v_j$) is at most ε . We set $\phi_{\rightarrow}^{\varepsilon}(v_i)$ to be an arbitrary but fixed \tilde{v}_i such that $(v_i, \tilde{v}_i) \in \mathbf{C}$ (this choice can be made canonical by, saying, taking the \tilde{v}_i that has the smallest g-function value among all points in \mathbf{R}_g that form a corresponding pair with v_i).

Next, we extend such mappings over nodes in V_{ε} to a continuous map over the entire R_g . In particular, consider an arc $v_i v_j$ and assume w.l.o.g. that $f(v_i) \le f(v_j)$. Consider $\tilde{v}_i = \phi_{\rightarrow}^{\varepsilon}(v_i)$ and $\tilde{v}_j = \phi_{\rightarrow}^{\varepsilon}(v_j)$. Since both $(v_i, \tilde{v}_i), (v_j, \tilde{v}_j) \in C$ and C is ε -optimal, we know that $|d_f(v_i, v_j) - d_g(\tilde{v}_i, \tilde{v}_j)| \le \beta + \varepsilon$, thus $d_g(\tilde{v}_i, \tilde{v}_j) \le d_f(v_i, v_j) + \beta \le \varepsilon$.

 $\beta + 2\varepsilon$. This means that there is a path $\pi(\tilde{v}_i, \tilde{v}_j)$ in R_g connecting \tilde{v}_i to \tilde{v}_j whose height is at most $\beta + 2\varepsilon$. We now extend $\phi_{\rightarrow}^{\varepsilon}$ to any homeomorphism from the arc $v_i v_j$ of R_f to this path $\pi(\tilde{v}_i, \tilde{v}_j)$ with $\phi_{\rightarrow}^{\varepsilon}(v_i) = \tilde{v}_i$ and $\phi_{\rightarrow}^{\varepsilon}(v_j) = \tilde{v}_j$. For example, one can take the arc-length parametrizations of the arc $v_i v_j$ and of the path $\pi(\tilde{v}_i, \tilde{v}_j)$, and then map up points proportionally. The concatenation of all these pieces of $\phi_{\rightarrow}^{\varepsilon}$ on each arc of R_f gives rise to the continuous map $\phi_{\rightarrow}^{\varepsilon} : R_f \to R_g$.

Given any point $x \in \mathbb{R}_f$, assume that x is from arc $v_i v_j$. Then $\tilde{x} := \phi_{\rightarrow}^{\varepsilon}(x)$ is mapped to some point in $\pi(\tilde{v}_i, \tilde{v}_j)$. Since C is an ε -optimal correspondence, by definition in Def. A.3, $g(\tilde{v}_i) \in [f(v_i) - \beta - \varepsilon, f(v_j) + \beta + \varepsilon]$ and $g(\tilde{v}_j) \in [f(v_i) - \beta - \varepsilon, f(v_j) + \beta + \varepsilon]$. Since the path $\pi(\tilde{v}_i, \tilde{v}_j)$ has height at most $\beta + 2\varepsilon$, we then have range($\pi(\tilde{v}_i, \tilde{v}_j)$) $\in [f(v_i) - 2\beta - 3\varepsilon, f(v_j) + 2\beta + 3\varepsilon]$. Since $x \in v_i v_j$ and $\tilde{x} \in \pi(\tilde{v}_i, \tilde{v}_j)$, it then follows that $g(\tilde{x}) \in [f(v_i) - 2\beta - 3\varepsilon, f(v_j) + 2\beta + 3\varepsilon]$ and thus $|g(\tilde{x}) - f(x)| \le 2\beta + 4\varepsilon$ for any $x \in \mathbb{R}_f$.

Symmetrically, we can take an ε -subdivision of R_g with nodes $U_{\varepsilon} = {\tilde{u}_1, \ldots, \tilde{u}_M}$, and construct a continuous map $\phi_{\leftarrow}^{\varepsilon} : R_g \to R_f$. Using the same argument as above, we have that $\max_{\tilde{y} \in R_g} |g(\tilde{y}) - f(\phi_{\leftarrow}^{\varepsilon}(\tilde{y}))| \le 2\beta + 4\varepsilon$.

We now bound $d_f(x, \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}^{\varepsilon}(x))$ for any $x \in \mathbb{R}_f$. Set $\tilde{x} = \phi_{\rightarrow}^{\varepsilon}(x)$ and $x' = \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}^{\varepsilon}(x)$. Assume that *x* is on some arc $v_i v_j$ and \tilde{x} is on some arc $\tilde{u}_a \tilde{u}_b$; see the right figure for an illustration. Since *x* is from arc $v_i v_j$, we have that $d_f(x, v_i) \le \varepsilon$. By construction of $\phi_{\rightarrow}^{\varepsilon}$, there is a path $\pi(\tilde{v}_i, \tilde{v}_j)$ with height at most $\beta + 2\varepsilon$ between $\tilde{v}_i = \phi_{\rightarrow}^{\varepsilon}(v_i)$ and $\tilde{v}_j = \phi_{\rightarrow}^{\varepsilon}(v_j)$, as both $(v_i, \tilde{v}_i), (v_i, \tilde{v}_i) \in \mathbb{C}$. Since \tilde{x} is in $\pi(\tilde{v}_i, \tilde{v}_j)$ and \tilde{x} is also in arc \tilde{u}_a, \tilde{u}_b , we



thus have that $d_g(\tilde{u}_a, \tilde{v}_i) \leq \beta + 3\varepsilon$. On the other hand, by construction of $\phi_{\leftarrow}^{\varepsilon}$, there is a path $\pi(u_a, u_b)$ with height at most $\beta + 2\varepsilon$ between $u_a = \phi_{\leftarrow}^{\varepsilon}(\tilde{u}_a)$ and $u_b = \phi_{\leftarrow}^{\varepsilon}(\tilde{u}_b)$. Since $x' \in \pi(u_a, u_b)$, $d_f(x', u_a) \leq \beta + 2\varepsilon$. Finally, since $(u_a, \tilde{u}_a), (v_i, \tilde{v}_i)$ are both from the ε -optimal correspondence C, we have that $|d_f(u_a, v_i) - d_g(\tilde{u}_a, \tilde{v}_i)| \leq \beta + \varepsilon$; in particular, $d_f(u_a, v_i) \leq d_g(\tilde{u}_a, \tilde{v}_i) + \beta \leq 2\beta + 4\varepsilon$. Hence $d_f(x, x') \leq d_f(x, v_i) + d_f(v_i, u_a) + d_f(u_a, x') \leq \varepsilon + (2\beta + 4\varepsilon) + (\beta + 2\varepsilon) = 3\beta + 7\varepsilon$.

In other words, $\max_{x \in \mathbb{R}_f} d_f(x, \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}^{\varepsilon}(x)) \le 3\beta + 7\varepsilon$. A symmetric argument shows that $\max_{\tilde{y} \in \mathbb{R}_g} d_g(\tilde{y}, \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}^{\varepsilon}(\tilde{y})) \le 3\beta + 7\varepsilon$. Hence the following distance is bounded from above by $3\beta + 7\varepsilon$:

$$D_{\phi_{\rightarrow}^{\varepsilon},\phi_{\leftarrow}^{\varepsilon}} := \max\{\max_{x \in \mathbb{R}_{f}} \{|f(x) - g(\phi_{\rightarrow}^{\varepsilon}(x))|, d_{f}(x,\phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}^{\varepsilon}(x))\}, \max_{\tilde{y} \in \mathbb{R}_{g}} \{|g(\tilde{y}) - f(\phi_{\leftarrow}^{\varepsilon}(\tilde{y}))|, d_{g}(\tilde{y},\phi_{\rightarrow}^{\varepsilon} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{y}))\}\}.$$

Putting everything together, we have that $d_{FD}(\mathbf{R}_f, \mathbf{R}_g) \leq \lim_{\varepsilon \to 0} D_{\phi_{\to}^{\varepsilon}, \phi_{\leftarrow}^{\varepsilon}} = 3\beta = 3d_{fGH}(\mathcal{R}_f, \mathcal{R}_g).$

Part II: We now prove the right inequality. Set $\delta = d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$, and let $\phi_{\rightarrow} : \mathbf{R}_f \rightarrow \mathbf{R}_g$ and $\phi_{\leftarrow} : \mathbf{R}_g \rightarrow \mathbf{R}_f$ be the optimal continuous maps that achieve δ^2 . Let C denote the correspondence generated by ϕ_{\rightarrow} and ϕ_{\leftarrow} , i.e., C consists of corresponding pairs of the form $(x, \phi_{\rightarrow}(x))$ or $(\phi_{\leftarrow}(\tilde{y}), \tilde{y})$. It is easy to see that the first term in Def. A.3 satisfies $\max_{(x,y)\in C} |f(x) - g(y)| \le \delta$. We now bound the second term from above. Specifically, consider two arbitrary pairs (x, \tilde{x}) and (y, \tilde{y}) from C. There are four cases: (i) $\tilde{x} = \phi_{\rightarrow}(x)$ and $\tilde{y} = \phi_{\rightarrow}(y)$; (ii) $x = \phi_{\leftarrow}(\tilde{x})$, $y = \phi_{\leftarrow}(\tilde{y})$; (iii) $x = \phi_{\leftarrow}(\tilde{x})$ but $\tilde{y} = \phi_{\rightarrow}(y)$; or (iv) $\tilde{x} = \phi_{\rightarrow}(x)$ but $y = \phi_{\leftarrow}(\tilde{y})$. Below we only show the proof for case (i). The handling of other cases is similar.

Consider the path $\pi(x, y)$ from x to y in \mathbb{R}_f whose height equals to $d_f(x, y)$. Let the path $\pi' = \phi_{\rightarrow}(\pi(x, y))$. Since $\max_{x' \in \mathbb{R}_f} |f(x') - g(\phi_{\rightarrow}(x'))| \le \delta$, we have

$$d_g(\tilde{x}, \tilde{y}) \le \text{height}(\pi') \le \text{height}(\pi(x, y)) + 2\delta = d_f(x, y) + 2\delta.$$
(6)

On the other hand, consider the optimal path $\tilde{\pi}(\tilde{x}, \tilde{y})$ between \tilde{x} and \tilde{y} that realizes $d_g(\tilde{x}, \tilde{y})$. Set $x' = \phi_{\leftarrow}(\tilde{x})$ and $y' = \phi_{\leftarrow}(\tilde{y})$. Using a similar argument as above, the path $\tilde{\pi}' = \phi_{\leftarrow}(\tilde{\pi}(\tilde{x}, \tilde{y}))$ connecting x' and y' satisfies $d_f(x', y') \leq \text{height}(\tilde{\pi}') \leq d_g(\tilde{x}, \tilde{y}) + 2\delta$. Furthermore, note that $x' = \phi_{\leftarrow}(\tilde{x}) = \phi_{\leftarrow} \circ \phi_{\rightarrow}(x)$. Hence by definition of $d_{FD}(\mathbb{R}_f, \mathbb{R}_g)$, we have that $d_f(x, x') \leq \delta$. Similarly, $d_f(y, y') \leq \delta$. Hence

$$d_f(x, y) \le d_f(x, x') + d_f(x', y') + d_f(y', y) \le \delta + (d_g(\tilde{x}, \tilde{y}) + 2\delta) + \delta = d_g(\tilde{x}, \tilde{y}) + 4\delta.$$
(7)

²The case where $d_{FD}(\mathbf{R}_f, \mathbf{R}_g)$ is achieved in the limit can be handled by taking a sequence of continuous maps.

Putting Eqns (6) and (7) together, we have that $\max_{(x,\tilde{x}),(y,\tilde{y})\in\mathbb{C}} |d_f(x,y) - d_g(\tilde{x},\tilde{y})| \le 4\delta$, which in turn proves the right inequality.

B Relation to Ordinary Persistence Diagram

B.1 Proof of Theorem 4.2

Recall that ϕ_{\rightarrow} : $R_f \rightarrow R_g$ and ϕ_{\leftarrow} : $R_g \rightarrow R_f$ are the optimal continuous maps that achieve $\delta = d_{FD}(R_f, R_g)$. We now show that maps in following diagrams are well-defined for an arbitrary real values $\alpha \le \alpha'$, and the induced homomorphisms commute at the 0th homology level, where *i*, *j* stand for canonical inclusion maps.

$$(\mathbf{R}_{f})_{\leq \alpha} \xrightarrow{i} (\mathbf{R}_{f})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{f})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{f})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{f})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{g})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{g})_{\leq \alpha'} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{g})_{\leq \alpha'} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{g})_{\leq \alpha'+\delta} \xrightarrow{\phi_{\leftarrow}} (\mathbf{R}_{g})_{< \alpha'+\delta} \xrightarrow{\phi_{$$

By Eqn (4), $\max_{x \in \mathbb{R}_f} |f(x) - g(\phi_{\rightarrow}(x))| \le \delta$. Hence $\phi_{\rightarrow} : (\mathbb{R}_f)_{\le \alpha} \longrightarrow (\mathbb{R}_g)_{\le \alpha + \delta}$ is well defined for any $\alpha \in \mathbb{R}$. Similarly, $\phi_{\leftarrow} : (\mathbb{R}_g)_{\le \beta} \longrightarrow (\mathbb{R}_f)_{\le \beta + \delta}$ is well defined for any $\beta \in \mathbb{R}$. It is easy to see that the second diagram above commutes. We now show that the first diagram commutes at the homology level for the 0th dimension. We need to show that for any 0-cycle c in $(\mathbb{R}_f)_{\le \alpha}$, $[i(c)] = [\phi_{\leftarrow} \circ j \circ \phi_{\rightarrow}(c)]$, where [c'] is the homology class represented by a cycle c'. Assume w.l.o.g that the 0-cycle $c = x_1 + x_2$ contains only two points x_1, x_2 from $(\mathbb{R}_f)_{\le \alpha}$; the argument easily extends to the case where c contains an arbitrary even number of points. Let $x'_1 = \phi_{\leftarrow} \circ \phi_{\rightarrow}(x_1)$ and $x'_2 = \phi_{\leftarrow} \circ \phi_{\rightarrow}(x_2)$. Since $d_f(x_1, x'_1) \le \delta$, we know that there is a path (1-chain) $\pi(x_1, x'_1)$ with height at most δ connecting x_1 and x'_1 . In other words, x_1 and x'_1 are connected in $(\mathbb{R}_f)_{\le \alpha+\delta} \subseteq (\mathbb{R}_f)_{\le \alpha'+2\delta}$. Similarly, x_2 and x'_2 are connected in $(\mathbb{R}_f)_{\le \alpha'+2\delta}$. Hence the new 0-cycle $c' = x'_1 + x'_2 = \phi_{\leftarrow} \circ j \circ \phi_{\rightarrow}(c)$ is homologous to c in $(\mathbb{R}_f)_{\le \alpha+2\delta}$. Thus, $[i(c)] = [c'] = [\phi_{\leftarrow} \circ j \circ \phi_{\rightarrow}(c)]$.

A similar argument also shows that the symmetric versions of the diagrams in Eqn (8) (by switching the roles of R_f and R_g) also commute at the 0th homology level. This means that the two persistence modules $\{H_0((R_f)_{\leq \alpha})\}_{\alpha}$ and $\{H_0((R_g)_{\leq \beta})\}_{\beta}$ are strongly δ -interleaved. The first half of Theorem 4.2 then follows from Theorem 4.8 of [9].

The same argument works for the scalar fields $-f : \mathbb{R}_f \to \mathbb{R}$ and $-g : \mathbb{R}_g \to \mathbb{R}$, which proves the second half of Theorem 4.2. Recall that while $Dg_0(\mathbb{R}_f)$ captures (minimum, down-saddle) persistence pairs, $Dg_0(\mathbb{R}_{-f})$ captures (up-saddle,maximum) persistence pairs.

We remark that in the proof of Theorem 4.2, we considered the ordinary persistence diagrams for the functions $\tilde{f} : \mathbb{R}_f \to \mathbb{R}$ and $\tilde{g} : \mathbb{R}_g \to \mathbb{R}$ induced on the Reeb graphs. However, the same argument holds for the ordinary persistence diagrams for the functions $f : X \to \mathbb{R}$ and $g : Y \to \mathbb{R}$ as induced by the persistence modules $\{H_0(X_{\leq \alpha})\}_{\alpha}$ and $\{H_0(Y_{\leq \beta})\}_{\beta}$. Indeed, $Dg_0(f)$ is identical to $Dg_0(\tilde{f})$, and $Dg_0(-f)$ is identical to $Dg_0(-\tilde{f})$. Similar statements hold for g and \tilde{g} . Furthermore, it is easy to check that Eqns. (3) and (4) can be extended for the original functions $f : X \to \mathbb{R}$ and $g : Y \to \mathbb{R}$ that produce the Reeb graphs. The distance $d_{FD}(X, Y)$ defined this way is fully captured by the distance between the corresponding Reeb graphs: $d_{FD}(\mathbb{R}_f, \mathbb{R}_g) = d_{FD}(X, Y)$. Putting these together, we have that Theorem 4.2 still holds if we replace \tilde{f} by f and \tilde{g} by g.

C Relation to Extended Persistence Diagram

C.1 Proof of Lemma 4.4

We show that $\sum_{j=1}^{a} \ell_{i_j} \neq 0$ for any subset $\{i_1, \ldots, i_a\} \subseteq \{1, 2, \ldots, k\}$. Specifically, consider the maximum height of dominating cycles of any cycle in $\{\ell_{i_j}\}_{i=1}^{a}$. First, assume that there is only a unique cycle, say ℓ_{i_a} , whose

dominating cycle dom(ℓ_{i_a}) has this maximum height. It then follows that this thin cycle dom(ℓ_{i_a}) is not in the thin basis decomposition of any other loop ℓ_{i_j} , $j \neq a$. Since the thin basis decomposition of the cycle $\sum_{j=1}^{a} \ell_{i_j}$ is simply the sum (modulo 2) of thin basis decomposition of each ℓ_{i_j} , dom(ℓ_{i_a}) must exist in the thin basis decomposition of the cycle $\sum_{j=1}^{a} \ell_{i_j}$; in fact, dom($\sum_{j=1}^{a} \ell_{i_j}$) = dom(ℓ_{i_a}). Therefore $\sum_{j=1}^{a} \ell_{i_j} \neq 0$.

If there are multiple cycles whose dominating cycle has the maximal height, then we consider the one whose dominating cycle has the smallest index among all of them. The same argument as above shows that this cycle will present in the thin basis decomposition of the loop $\sum_{i=1}^{a} \ell_{i_i}$, implying that $\sum_{i=1}^{a} \ell_{i_i} \neq 0$.

C.2 Proof of Lemma 4.5

The input Reeb graph R_f is a finite graph, and there are only a finite number of cycles. Hence there are only a finite number of height values (reals) that these cycles can have. Let α denote the lowest height of any cycle whose height is strictly larger than 2δ — if the height of a cycle is strictly smaller than this value α , then it must be smaller than or equal to 2δ . We set ε to be any positive constant between 0 and $\alpha/2 - \delta$; that is, $2\delta + 2\varepsilon < \alpha$. Thus if a cycle γ satisfies height(γ) $\leq 2\delta + 2\varepsilon < \alpha$, then necessarily height(γ) $\leq 2\delta$.

Now for the given cycle ℓ , let $\tilde{\ell}$ denote $\phi_{\leftarrow} \circ \phi_{\rightarrow}(\ell)$, and for any point $x \in \ell$, let $\tilde{x} \in \tilde{\ell}$ denote $\phi_{\leftarrow} \circ \phi_{\rightarrow}(x)$. By Eqn (4), there is a path $\pi(x, \tilde{x})$ connecting x to \tilde{x} with height($\pi(x, \tilde{x})$) $\leq \delta$. Assume that the cycle ℓ has only a single connected component – the case with multiple components can be handled in a component-wise manner. This means that we can consider ℓ as a closed curve on \mathbb{R}_f . Fix an arbitrary orientation on ℓ and consider the induced orientation on $\tilde{\ell}$ by the map $\phi_{\leftarrow} \circ \phi_{\rightarrow}$. Let $\ell[x, x'] \subseteq \ell$ denote the subcurve from x to x' under this orientation; and similarly for $\ell'[x, x']$. Start with an arbitrary point $x = x_0 \in \ell$, and consider $\tilde{x}_0 \in \tilde{\ell}$. Since $\phi_{\leftarrow} \circ \phi_{\rightarrow}$ is a continuous map, as we move x along ℓ continuously, \tilde{x} moves continuously. In step i, as we move along ℓ (starting from x_i), we set x_{i+1} to be the first point such that the loop $c_i = \ell[x_i, x_{i+1}] \circ \pi(x_{i+1}, \tilde{x}_{i+1}) \circ (-\tilde{\ell}[\tilde{x}_i, \tilde{x}_{i+1}]) \circ \pi(\tilde{x}_i, x_i)$ satisfies height(c_i) = $2\delta + 2\varepsilon$. If no such c_i exists, then we set x_{i+1} to be x_0 , and the process terminates. It is easy to see that by construction, we have that $\ell = \tilde{\ell} + \sum_i c_i$. Since each c_i satisfies height(c_i) $\leq 2\delta + 2\varepsilon < \alpha$, as discussed earlier it then follows that height(c_i) $\leq 2\delta$. Hence $c_i \in Z_1^{\delta}(\mathbb{R}_f)$ for each i and $\ell' = \sum_i c_i \in Z_1^{\delta}(\mathbb{R}_f)$, implying that $\tilde{\ell} = \ell + \ell' \in \ell + Z_1^{\delta}(\mathbb{R}_f)$.

C.3 Proof of Lemma 4.6

First, by Eqn (4), we have $\max_{x \in \mathbb{R}_f} |f(x) - g(\phi_{\rightarrow}(x))| \le \delta$. Hence $\operatorname{range}(\phi_{\rightarrow}(\gamma)) \subseteq [b - \delta, d + \delta]$. Now let b' be the smallest left endpoint of the range of any loop in the thin basis decomposition of $\phi_{\rightarrow}(\gamma)$. We will prove that $b' \ge b - \delta$.

Suppose this is not case and $b' < b - \delta$. Then let $\zeta_{i_1}, \ldots, \zeta_{i_a} \in \mathcal{G}_g$, $a \ge 1$, be the set of cycles from the thin basis decomposition of $\phi_{\rightarrow}(\gamma)$ that have b' as the left endpoint of their range. Set $\rho = \zeta_{i_1} + \cdots + \zeta_{i_a}$. Assume ζ_{i_a} has the largest height among these thin cycles. Note that range(ζ_{i_j}) \subseteq range(ζ_{i_a}) for any j < a, as all these ranges share the same left endpoint b'. Hence range(ρ) \subseteq range(ζ_{i_a}).

On the other hand, recall that the left endpoint of range($\phi_{\rightarrow}(\gamma)$) is at least $b - \delta$, which is strictly larger than b'. This means that range(ρ) is a proper subset of range(ζ_{i_a}); i.e., range(ρ) \subset range(ζ_{i_a}). Hence ρ has strictly smaller height than ζ_{i_a} . This however contradicts that \mathcal{G}_g is a thin basis, because we can replace ζ_{i_a} in \mathcal{G}_g with ρ and obtain a basis element with smaller height (the resulting set of cycles remain independent). Therefore it is not possible that $b' < b - \delta$.

A symmetric argument shows that the largest right endpoint of the range of any loop in the thin basis decomposition of $\phi_{\rightarrow}(\gamma)$ is at most $d + \delta$. Hence the range of any loop in the thin basis decomposition of $\phi_{\rightarrow}(\gamma)$ is a subset of $[b - \delta, d + \delta]$.

C.4 Proof of Lemma 4.10

Consider an arbitrary subset of indices i_1, \ldots, i_s whose corresponding columns are in $\widetilde{\Phi}$ (i.e, each $\widehat{\gamma}_{i_a}$ is 3δ -stable), and let $\widehat{\gamma} = \widehat{\gamma}_{i_1} + \cdots + \widehat{\gamma}_{i_s}$. We will show that $\phi_{\rightarrow}(\widehat{\gamma})$ is δ -stable; that is, dom $(\phi_{\rightarrow}(\widehat{\gamma}))$ has height at least 2δ . This means that the linear combination of the corresponding columns in $\widetilde{\Phi}$ contains a non-zero element. Since this holds for any subset of columns from $\widetilde{\Phi}$, we have that the columns in $\widetilde{\Phi}$ are linearly independent.

It remains to show that $\widehat{\gamma}$ is δ -stable. Recall that for any index a, γ_a is the dominating cycle of $\widehat{\gamma}_a$. Assume w.l.o.g. that γ_{i_1} has the largest height among all γ_{i_j} , for $j \in [1, s]$. (If there are multiple cycles from $\{\gamma_{i_j}\}_{j=1}^s$ having this largest height, let γ_{i_s} be the one with smallest index.) From the end of the proof of Proposition 4.8, we note that for any index a, γ_a is the *only* cycle with maximum height among all cycles in the thin basis decomposition of $\widehat{\gamma}_a$. In other words, all other cycles in the thin basis decomposition of $\widehat{\gamma}_a$ have strictly smaller height than γ_a . Putting these together, it follows that γ_{i_1} must exist in the thin basis decomposition of $\widehat{\gamma}$ w.r.t. the original thinnest basis \mathcal{G}_f and in fact, $\gamma_{i_1} = \operatorname{dom}(\widehat{\gamma})$. Since γ_{i_1} is 3δ -stable (thus its height at least 6δ), it then follows from Lemma 4.7 that height($\operatorname{dom}(\phi_{\rightarrow}(\widehat{\gamma}))$) \geq height($\operatorname{dom}(\widehat{\gamma})$) $- 2\delta \geq 4\delta > 2\delta$. So $\phi_{\rightarrow}(\widehat{\gamma})$ is δ -stable.

C.5 Proof of Corollory 4.11

View the matrix $\overline{\Phi}$ as the adjacency graph for the following bipartite graph $G = (P \cup Q, E)$, where P are the columns of $\overline{\Phi}$, Q are the rows, and there is an edge between $p \in P$ and $q \in Q$ iff the corresponding entry in the matrix $\overline{\Phi}$ is 1. We now claim that there is a *P*-saturated matching for G: that is, there is a matching of G such that every node in P is matched exactly once, and each node in Q is matched at most once. Note that this immediately implies the lemma. Specifically, for any subset of nodes $P' \subseteq P$, let $Q' \subseteq Q$ be the union of neighbors of nodes from P'. In other words, Q' is the set of rows with at least one non-zero entry in the columns P'. If |Q'| < |P'|, then these columns of $\overline{\Phi}$ will be linearly dependent, which violates Lemma 4.10. Hence we have $|Q'| \ge |P'|$. Now by Hall's Theorem (e.g, Page 35 of [34]), a P-satuated matching exist for G.

C.6 Further Discussion

The proof of Theorem 4.2 leverages a beautiful stability result for a pair of interleaving persistence modules, originally introduced in [9] and recently generalized in [10]. This approach however, is not immediately applicable to bound the relation to the extended persistence diagrams in our case (i.e., to prove Theorem 4.3), since the extended persistence module is not parametrized over the reals, but over two copies of the real line. Therefore, the notion of interleaving distance does not apply directly. A better choice to capture the extended persistence diagrams as we define in this paper seems to be the *level-set zigzag* persistence module [7, 8] (specifically see the discussion on relation of the extended persistence diagrams and level-set zigzag persistence modules in [8]). Unfortunately, it is not yet clear how to extend the interleaving concept to zigzag persistence modules and obtain similar type of stability results as given in [9].

Our approach in Section 4.2 instead aims to provide explicit correspondences between certain basis cycles of $H_1(R_f)$ and of $H_1(R_g)$, by intuitively treating the correspondences as edges in a specific bipartite graph, and establishing an 3δ -matching in this graph. Interestingly, it appears that this approach can be potentially generalized to establish matching between certain basis cycles generating two input *p*-th levelset zigzag persistence modules for *arbitrary dimension p* (instead of p = 1 in our case). This will then lead to stability results for two general levelset zigzag persistence modules, and it will be very interesting to investigate this direction. We note, however, that the stability result obtained in Theorem 4.3 does not seem to be tight. We conjecture that the tight bound is $d_B(\text{ExDg}_1(R_f), \text{ExDg}_1(R_g)) \le \delta = d_{FD}(R_f, R_g)$ instead of $d_B(\text{ExDg}_1(R_f), \text{ExDg}_1(R_g)) \le 3\delta$. It will be interesting to see whether the upper bound can be improved.

D Simplification of Reeb graphs

Reeb graphs have been used as a meaningful summary of the input functions. Simplifying a Reeb graph can help to remove noise or single out major features, and to create multi-resolution representation of input domains; see e.g. [14, 17, 27]. As we described in Section 2, there is a natural way to define "importance" of both branching and loop features, which turns out to be equivalent to ordinary and extended persistence in the according dimensions. Indeed, it is common practice to simplify the Reeb graph by removing all features with persistence smaller than a threshold, say δ . In this section, we prove that by removing small features using a natural merging strategy, (branching and loop) features with large persistence will not be killed, and will roughly maintain their persistence ("importance").

D.1 A Natural Simplification Scheme for Reeb graphs

We now introduce a natural simplification scheme for Reeb graphs (see, e.g., [17, 27]). See Figure 3 for an illustration.

Given an ordinary persistence pair (m, s), assume that m is a minimum and s a down-saddle point. Recall that the down-fork saddle s merges two disjoint lower branches C_1 and C_2 , and m is the higher global minimum of the two. To remove the feature (m, s), we wish to merge the branch containing m, say C_1 , into the other branch C_2 , so that afterwards, m and s become regular points (i.e., with up-degree and down-degree both being 1). In particular, we perform the following operation (see Figure 3 (a)). Let m_2 denote the global minimum of C_2 . We choose an arbitrary path $\pi_1 \subseteq C_1$ from s to m, and an arbitrary $\pi' \subseteq C_2$ from s to m_2 . Now imagine we traverse π' starting from s. We stop when we encounter the first point $x \in \pi'$ such that f(x) = f(m), and set π_2 be the subcurve of π' from s to x. We now merge π_1 and π_2 into a new monotonic arc π_3 between s and x such that any point $p \in \pi_1 \cup \pi_2$ is mapped to $q \in \pi_3$ such that f(q) = f(p). Pairs of type (up-saddle, maximum) are treated in a symmetric way.

Given an extended persistence pair (s_1, s_2) between an up-saddle s_1 and a down-saddle s_2 , let γ be a thin cycle that spans it. W.l.o.g. assume that γ consists only of a single connected component: if γ has multiple connected components, then there must exist one that contains both s_1 and s_2 . That loop is necessarily thin and thus we can simply set γ to be that loop. Let π_1 and π_2 denote the two disjoint sub-curves of γ connecting s_1 and s_2 . To cancel the this feature, intuitively, we wish to merge π_1 and π_2 to kill the cycle γ . Note that π_1 and π_2 may not be monotonic (w.r.t. the input function f); however, all points in π_1 and π_2 have function values within the range $[f(s_1), f(s_2)]$. The merging of π_1 and π_2 results in a new monotonic arc π_3 from s_1 and s_2 , such that for every point $p \in \gamma$ is mapped to $q \in \pi_3$ such that f(q) = f(p). See Figure 3 (b) for an illustration.

Note that since a critical pair (m, s) (resp. an essential pair (s_1, s_2)) corresponds uniquely to a persistence pair (f(m), f(s)) (resp. $(f(s_1), f(s_2))$) in the ordinary (resp. extended) persistence diagram, the above process also removes a point from the respective persistence diagram.

Let R (induced by function f) and R' denote the Reeb graph before and after the simplification of a persistence pair $\tau = (\mathbf{b}, \mathbf{d})$ by collapsing its corresponding branching or loop feature. Let π_1^{τ} and π_2^{τ} be as introduced above³. Call $\gamma^{\tau} = \pi_1^{\tau} \cup \pi_2^{\tau}$ the *merging path w.r.t.* τ : γ^{τ} is a closed curve corresponding to a thin cycle spanning (\mathbf{b}, \mathbf{d}) when it is an extended persistence pair, and a connected path with \mathbf{b} and \mathbf{d} being the respective minimum and maximum function values on it otherwise. The merging path γ^{τ} will be collapsed into a single monotonic arc to kill the persistence pair τ . We can view the removal of τ in a more formal way as follows: We say that two points $x, y \in \mathbf{R}$ are τ -equivalent, denoted by $x \sim_{\tau} y$, if f(x) = f(y) and $x, y \in \gamma^{\tau}$. The simplified Reeb graph R' is the quotient space \mathbf{R}/\sim_{τ} ; the corresponding quotient map $\mu_{\tau} : \mathbf{R} \to \mathbf{R}'$ satisfies $\mu_{\tau}(x) = \mu_{\tau}(y)$ if and only if $x \sim_{\tau} y$. The function $f : \mathbf{R} \to \mathbf{R}$ induces a function $f' : \mathbf{R}' \to \mathbf{R}$ such that for any $x' \in \mathbf{R}'$, f'(x') = f(x) for any $x \in \mu_{\tau}^{-1}(x')$.

³Note that the choices of π_1^{τ} and π_2^{τ} are not canonical. However, these choices do not affect the final results. Alternatively, one can make a canonical choice by say choosing π_1^{τ} and π_2^{τ} (or the thin cycle γ for a loop feature) as the shortest possible among all such choices.

Now given an input Reeb graph R, suppose we wish to eliminate a set of persistence pairs { $\tau_1 = (\mathbf{b}_1, \mathbf{d}_1), \tau_2 = (\mathbf{b}_2, \mathbf{d}_2), \dots, \tau_k = (\mathbf{b}_k, \mathbf{d}_k)$ }. Compute the merging path γ^{τ_i} for each persistence pair τ_i in R. Similar to above, two points x, y are τ_i -equivalent, denoted by \sim_{τ_i} if $x, y \in \gamma^{\tau_i}$ and f(x) = f(x). We now define an equivalence relation \sim as the transitive closure of all \sim_{τ_i} s for $i \in [1, k]$. This is equivalent to collapsing γ^{τ_i} s for all $i \in [1, k]$ in an arbitrary order to kill the persistence pairs τ_1, \dots, τ_k ⁴. The final simplified Reeb graph \widetilde{R} is obtained as the quotient space R/\sim . The quotient map $\mu : \mathbb{R} \to \widetilde{\mathbb{R}}$ satisfies $\mu(x) = \mu(y)$ if and only if $x \sim y$, and we have a well-defined function $g : \widetilde{\mathbb{R}} \to \mathbb{R}$ induced by function $f : \mathbb{R} \to \mathbb{R}$ such that $g(\mu(x)) = f(x)$ for any $x \in \mathbb{R}$. Let δ denote the largest persistence of τ_1, \dots, τ_k . We have the following properties of $\widetilde{\mathbb{R}}$.

Observation D.1 (*i*) Given any two points $x, y \in \mathbb{R}$, we have that $d_g(\mu(x), \mu(y)) \leq d_f(x, y)$. (*ii*) Given a point $\tilde{x} \in \widetilde{\mathbb{R}}$, for any two points $x_0, x_1 \in \mu^{-1}(\tilde{x})$, we have $d_f(x_0, x_1) \leq 2\delta$.

Proof: Claim (i) follows easily since the continuous map μ is surjective. We now prove (ii). Since $\mu(x_0) = \mu(x_1) = \tilde{x}$, by the definition of μ , there exists a set of equivalent relations $\sim_{\tau_{j_1}}, \ldots, \sim_{\tau_{j_a}}$ with the index set $\{j_1, \ldots, j_a\} \subseteq \{1, \ldots, k\}$ such that $y_0 := x_0 \sim_{\tau_{j_1}} y_1 \sim_{\tau_{j_2}} y_2 \cdots \sim_{\tau_{j_a}} y_a := x_1$. Set $\alpha = f(x_0) = f(x_1)$. All y_i have the same function value α . For each $i \in [1, a]$, we have that $y_{i-1} \sim_{\tau_{j_i}} y_i$, which is induced by the merging path $\gamma^{\tau_{j_i}}$ with height $(\gamma^{\tau_{j_i}}) \leq \delta$. In other words, there is a path $\pi_i \subseteq \gamma^{\tau_{j_i}}$ connecting y_{i-1} to y_i and range $(\pi_i) \subseteq [\alpha - \delta, \alpha + \delta]$. The concatenation of these π_i s gives rise to a path π connecting $y_0 = x_0$ and $y_a = x_1$, and range $(\pi) \subseteq [\alpha - \delta, \alpha + \delta]$. Thus proves Claim (ii).



Figure 4: Left: An ε -subdivision of the simplified Reeb graph \widetilde{R} such that each arc is monotonic and of height at most ε . Right: For an arc $\tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$, we map it under $\phi_{\leftarrow}^{\varepsilon}$ to a certain path $\pi(v_i, v_j)$ such that $\mu(v_i) = \tilde{v}_i, \mu(v_j) = \tilde{v}_j$, and the range of the path $\pi(v_i, v_j) \subseteq \mathbb{R}$ is contained within the interval $[f(v_i) - \delta, f(v_j) + \delta]$. Note that the paths $\pi(v_i, v_j)$ and $\pi(v_i, v_j)$ for different arcs $\tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$ and $\tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$ from \widetilde{R} are not necessarily disjoint.

D.2 Distance between R and R

Note that while the simplification scheme removes persistence pairs τ_1, \ldots, τ_k , it is not clear how other points in the persistence diagram for the original Reeb graph R are affected. In this section, we bound the bottleneck distance between the persistence diagrams for R and for \widetilde{R} . Specifically, we bound the functional distortion distance $d_{FD}(R, \widetilde{R})$, where we have $f : R \to \mathbb{R}$ and $g : \widetilde{R} :\to \mathbb{R}$ (defined at the end of Section D.1). We do so by constructing continuous maps $\phi_{\to} : R \to \widetilde{R}$ and $\phi_{\leftarrow} : \widetilde{R} \to R$, and bounding the four terms in Eqn (4), which in turn provides an upper bound for $d_{FD}(R, \widetilde{R})$.

The continuous map $\phi_{\rightarrow} : \mathbb{R} \to \widetilde{\mathbb{R}}$ can simply be taken as the surjective map $\mu : \mathbb{R} \to \widetilde{\mathbb{R}}$. For the opposite continuous map $\phi_{\leftarrow} : \widetilde{\mathbb{R}} \to \mathbb{R}$, we will in fact construct a sequence of them. First, we need the following result, which is a slight generalization of Observation D.1, and whose proof is similar but more tedious.

⁴For example, assume w.l.o.g. we merge in increasing order of *i*. Let R_{i-1} denote the simplified Reeb graph after merging the first *i* – 1 persistence pairs. In the *i*th stage, we will collapse the *image of* γ^{τ_i} in R_{i-1} .

Lemma D.2 Let $\tilde{x}, \tilde{y} \in \widetilde{R}$ be two points in \widetilde{R} with a monotonic path $\tilde{\pi}$ between \tilde{x} and \tilde{y} such that $d_g(\tilde{x}, \tilde{y}) = g(\tilde{y}) - g(\tilde{x}) = \varepsilon$. Let x and y be arbitrary preimages for \tilde{x} and \tilde{y} , respectively. Then $d_f(x, y) \le 2\delta + \varepsilon$.

In fact, there is a path π from x to y such that the highest point t in π satisfies $f(t) \le f(y) + \delta$, and the lowest point b in π satisfies $f(b) \ge f(x) - \delta$.

Now for a fixed positive real ε , we use the following procedure to construct a continuous map $\phi_{\leftarrow}^{\varepsilon} : \widetilde{R} \to R$ (see Figure 4 for an illustration). First, we subdivide the simplified Reeb graph \widetilde{R} by adding a set of nodes, so that for the resulting graph (still denoted by \widetilde{R}) with nodes $V_{\varepsilon} = \{\widetilde{v}_1, \ldots, \widetilde{v}_m\}$, every arc is monotonic and has height at most ε : note the height of an monotonic path from x to y is simply the difference in the function values of x and y. We refer to the resulting augmented graph \widetilde{R} as an ε -subdivision of \widetilde{R} . Now for each $\widetilde{v}_i \in V_{\varepsilon}$, we set $\phi_{\leftarrow}^{\varepsilon}(\widetilde{v}_i)$ to be an arbitrary but fixed pre-image $v_i \in \mu^{-1}(\widetilde{v}_i)$. For each arc $\widetilde{\pi}(\widetilde{v}_i, \widetilde{v}_j)$ from \widetilde{R} , consider the path $\pi(v_i, v_j)$ connecting the two pre-image points v_i and v_j as stated in Lemma D.2. We set the restriction of $\phi_{\leftarrow}^{\varepsilon}$ over the arc $\widetilde{\pi}(\widetilde{v}_i, \widetilde{v}_j)$ to be any homeomorphism from $\widetilde{\pi}(\widetilde{v}_i, \widetilde{v}_j)$ to $\pi(v_i, v_j)$ with $\phi_{\leftarrow}^{\varepsilon}(\widetilde{v}_i) = v_i$ and $\phi_{\leftarrow}^{\varepsilon}(\widetilde{v}_j) = v_j$. The union of $\phi_{\leftarrow}^{\varepsilon}(\widetilde{\pi}(\widetilde{v}_i, \widetilde{v}_j))$ s for all arcs $\widetilde{\pi}(\widetilde{v}_i, \widetilde{v}_j)$ from \widetilde{R} gives rise to the continuous map $\phi_{\leftarrow}^{\varepsilon} : \widetilde{R} \to R$.

By definition, for any $x \in \mathbb{R}$, $|f(x) - g(\phi_{\rightarrow}(x))| = |f(x) - g(\mu(x))| = 0$; hence $\max_{x \in \mathbb{R}} |f(x) - g \circ \phi_{\rightarrow}(x)| = 0$. Furthermore, let $\tilde{x} = \mu(x) = \phi_{\rightarrow}(x)$ and $x' = \phi_{\leftarrow}^{\varepsilon}(\tilde{x}) = \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}(x)$. Assume w.l.o.g. that \tilde{x} is contained in arc $\tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$. Obviously, we have $g(\tilde{x}) = f(x) \in [f(v_i), f(v_j)] = [g(\tilde{v}_i), g(\tilde{v}_j)]$. Since the sub-arc $\tilde{\pi}(\tilde{v}_i, \tilde{x}) \subseteq \tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$ is also monotonic with height at most ε , by Lemma D.2, we have that x and v_i is connected by a path π_1 whose range is contained within $[f(v_i) - \delta, f(x) + \delta] \subseteq [f(v_i) - \delta, f(v_j) + \delta]$. At the same time, since $\tilde{x} \in \tilde{\pi}(\tilde{v}_i, \tilde{v}_j), x'$ is contained in the path $\pi(v_i, v_j)$ by the construction of $\phi_{\leftarrow}^{\varepsilon}$ above. Hence x' and v_i is connected by a path π_2 whose range is contained within $[f(v_i) - \delta, f(v_j) + \delta]$. By concatenating π_1 and π_2 , we have that $d_f(x, x') = d_f(x, \phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}(x)) \le 2\delta + \varepsilon$.

Next, consider any point $\tilde{x} \in \widetilde{R}$. Assume that \tilde{x} is from the arc $\tilde{\pi}(\tilde{v}_i, \tilde{v}_j)$ in \widetilde{R} , and assume w.l.o.g. that $g(\tilde{v}_i) \leq g(\tilde{v}_j)$. By construction, $g(\tilde{x}) \in [g(\tilde{v}_i), g(\tilde{v}_j)] = [f(v_i), f(v_j)]$. Set $x = \phi_{\leftarrow}^{\varepsilon}(\tilde{x})$. By the construction of $\phi_{\leftarrow}^{\varepsilon}$, $f(x) \in [f(v_i) - \delta, f(v_j) + \delta]$. Since $|f(v_j) - f(v_i)| \leq \varepsilon$, we thus have that $|f(x) - g(\tilde{x})| \leq \delta + \varepsilon$. Thus $\max_{\tilde{x}\in\widetilde{R}} |f(\phi_{\leftarrow}^{\varepsilon}(\tilde{x})) - g(\tilde{x})| \leq \delta + \varepsilon$. Furthermore, by Lemma D.2, $d_f(x, v_i) \leq 2\delta + \varepsilon$. By Observation D.1, this means that $d_g(\mu(x), \mu(v_i)) = d_g(\phi_{\rightarrow} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x}), \tilde{v}_i) \leq 2\delta + \varepsilon$. Since $d_g(\tilde{x}, \tilde{v}_i) \leq d_g(\tilde{v}_i, \tilde{v}_j) \leq \varepsilon$, it then follows that $d_g(\tilde{x}, \phi_{\rightarrow} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x})) \leq d_g(\tilde{x}, \phi_{\rightarrow} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x})) \leq 2\delta + 2\varepsilon$. Hence $\max_{\tilde{x}\in\widetilde{R}} d_g(\tilde{x}, \phi_{\rightarrow} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x})) \leq 2\delta + 2\varepsilon$. Putting everything together, we have that under the maps ϕ_{\rightarrow} and $\phi_{\leftarrow}^{\varepsilon}$, we have that

$$\begin{split} d_{\phi_{\rightarrow},\phi_{\leftarrow}^{\varepsilon}} &:= \max\{\max_{x\in\mathbb{R}} |f(x) - g \circ \phi_{\rightarrow}(x)|, \quad \max_{x\in\mathbb{R}} d_f(x,\phi_{\leftarrow}^{\varepsilon} \circ \phi_{\rightarrow}(x)) \\ \max_{\tilde{x}\in\widetilde{\mathbb{R}}} |g(\tilde{x}) - f \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x})|, \quad \max_{\tilde{x}\in\widetilde{\mathbb{R}}} d_g(\tilde{x},\phi_{\rightarrow} \circ \phi_{\leftarrow}^{\varepsilon}(\tilde{x}))\} \le 2\delta + 2\varepsilon. \end{split}$$

By Eqn (4), we have that $d_{FD}(\mathbb{R}, \widetilde{\mathbb{R}}) \leq \lim_{\varepsilon \to 0} d_{\phi_{\rightarrow}, \phi_{\leftarrow}^{\varepsilon}} = 2\delta$. Combining this with Theorems 4.2 and 4.3, we thus have that $d_B(Dg_0(\mathbb{R}_f), Dg_0(\mathbb{R}_g)) \leq d_{FD}(\mathbb{R}, \widetilde{\mathbb{R}}) = 2\delta$, $d_B(Dg_0(\mathbb{R}_{-f}), Dg_0(\mathbb{R}_{-g})) \leq 2\delta$, and $d_B(ExDg_1(\mathbb{R}_f), ExDg_1(\mathbb{R}_g)) \leq 3d_{FD}(\mathbb{R}, \widetilde{\mathbb{R}}) = 6\delta$.