

Labeled Interleaving Distance for Reeb Graphs

Fangfei Lan^{1*}, Salman Parsa² and Bei Wang¹

^{1*}University of Utah, Salt Lake City, Utah, USA.

²DePaul University, Chicago, Illinois, USA.

*Corresponding author(s). E-mail(s): fangfei.lan@sci.utah.edu;
Contributing authors: s.parsa@depaul.edu; beiwang@sci.utah.edu;

Abstract

Merge trees, contour trees, and Reeb graphs are graph-based topological descriptors that capture topological changes of (sub)level sets of scalar fields. Comparing scalar fields using their topological descriptors has many applications in topological data analysis and visualization of scientific data. Recently, Munch and Stefanou introduced a labeled interleaving distance for comparing two labeled merge trees, which enjoys a number of theoretical and algorithmic properties. In particular, the labeled interleaving distance between merge trees can be computed in polynomial time. In this work, we define the labeled interleaving distance for labeled Reeb graphs. We then prove that the (ordinary) interleaving distance between Reeb graphs equals the minimum of the labeled interleaving distance over all labelings. We also provide an efficient algorithm for computing the labeled interleaving distance between two labeled contour trees (which are special types of Reeb graphs that arise from simply-connected domains). In the case of merge trees, the notion of the labeled interleaving distance was used by Gasparovic et al. to prove that the (ordinary) interleaving distance on the set of (unlabeled) merge trees is intrinsic. As our final contribution, we present counterexamples showing that, on the contrary, the (ordinary) interleaving distance on (unlabeled) Reeb graphs (and contour trees) is not intrinsic. It turns out that, under mild conditions on the labelings, the labeled interleaving distance is a metric on isomorphism classes of Reeb graphs, analogous to the ordinary interleaving distance. This provides new metrics on large classes of Reeb graphs.

Keywords: Reeb graphs, merge trees, interleaving distance, topological data analysis

1 Introduction

Topological descriptors such as merge trees, contour trees, and Reeb graphs capture topological changes of (sub)level sets of scalar fields. Comparing scalar fields using their topological descriptors has a number of applications in topological data analysis (TDA) and visualization of scientific datasets such as combustion and molecular dynamics simulations, including symmetry detection [1–4], shape matching and retrieval [3, 5–8]; feature tracking [9–11] and event detection [1, 3, 12]; clustering and classification [7, 8], summarization [13, 14], uncertainty visualization [14–16], and interactive exploration [14, 17]; see [18] for a survey.




Fig. 1 An example of a Reeb graph.

Given a continuous function $\tilde{f} : \mathbb{M} \rightarrow \mathbb{R}$ defined on a connected domain $\mathbb{M}$, a Reeb graph records the connectivity of its *level sets*. Two points $x, y \in \mathbb{M}$ are considered *equivalent* (w.r.t. $\tilde{f}$), denoted as $x \sim y$, if $\tilde{f}(x) = \tilde{f}(y) = t$ and x and y belong to the same connected component of the level set $\tilde{f}^{-1}(t)$, for some $t \in \mathbb{R}$. The *Reeb graph* $\mathbb{M}/\sim$ is the quotient space obtained by identifying equivalent points. For well-behaved data, e.g., a Morse function on a manifold, $\mathbb{M}/\sim$ is a graph G that inherits a function f from the original input data $(\mathbb{M}, \tilde{f})$. See Figure 1 for an example. A *contour tree* is a special type of Reeb graph when the domain $\mathbb{M}$ is simply connected. A merge tree, on the other hand, relies on equivalence relations among points in the *sublevel sets* of $\tilde{f}$. That is, $x \sim y$, if $\tilde{f}(x) = \tilde{f}(y) = t$, and they belong to the same connected component of the sublevel set $\tilde{f}^{-1}(\infty, t]$, for some $t \in \mathbb{R}$.

Since Reeb graphs generalize contour trees, and to some extent, merge trees, we mainly focus on comparative measures for Reeb graphs. A number of distances have been proposed for Reeb graphs and their variants, such as interleaving distance [19–24], functional distortion distance [25, 26], functional contortion distance [27], edit distance [3, 28–30], Gromov-Hausdorff distance [31, 32], distances based on branch decompositions and matching [1, 33], and metrics for phylogenetic trees [34]. The bottleneck distance [35] is also defined for Reeb graphs [31] by computing the distance between their

extended persistence diagrams [36]. The stability of distances between Reeb graphs, their equivalence and the inequalities between them has been the topic of much research [21, 25, 26, 30, 37, 38]. Recently, it has been shown that the edit distance is universal, that is, it is stable and upper bounds all other stable distances [30], and hence is more discriminating. There is now a vast literature on the subject of distances for Reeb graphs, see [18, 39] for surveys.

From the above distances, the bottleneck distance is polynomial-time computable. The interleaving distance, the functional distortion distance and the edit distance are actually metrics on the isomorphism classes of Reeb graphs, hence most discriminatory. However, a major drawback of these metrics is that they are hard to compute. Reeb graph 0-interleaving is shown to be graph isomorphism complete [21, 40]. And interleaving distances between multiparameter persistence modules are also NP-hard [40, 41]. As noted in [32], the reduction of [42] shows that it is NP-hard to approximate the interleaving distance for merge trees [22] within a factor of 3. Therefore, it is also hard for Reeb graphs, since merge trees are special cases of Reeb graphs.

Nevertheless, not all hope is lost. Recently, Munch and Stefanou [23] showed that the l^∞ -cophenetic metric originally defined on phylogenetic trees is an example of an interleaving distance on labeled merge trees, where nodes of a merge tree are given labels from a fixed set. Gasparovic et al. [43] extended this work on the labeled interleaving distance of merge trees and proved that the (ordinary) interleaving distance is obtained as the minimum of the labeled distance over all possible labelings. The advantage of the labeled interleaving distance is that it can be computed efficiently in $O(n^2)$ (n being the number of critical points of f). Such a notion makes the distance computation feasible in real-world applications [14, 44], and it appears to be a reasonable replacement for the (ordinary) interleaving distance, which is quite desirable.

A labeled interleaving distance is applicable when there is a natural labeling for the nodes, or when a labeling may be inferred from the data. For instance, a climate simulation may give rise to an ensemble of scalar fields (e.g., temperature and pressure) simulated with different parameter settings. Each ensemble member is a scalar field defined on the same underlying mesh and gives rise to a slightly different merge tree. We may use the indices of mesh nodes as the labeling or infer correspondences between nodes of the merge tree based on similarities among features of their underlying scalar fields.

Apart from computational efficiency, another desirable property for a distance between Reeb graphs is for the distance to be *intrinsic*, i.e., the distance between any two Reeb graphs can be realized or approximated arbitrarily closely by a geodesic (a shortest continuous path) in the space of Reeb graphs [31, 43]. An intrinsic distance is desirable for not only discrimination but also interpolation between a pair of Reeb graphs. Instead of studying the intrinsic-ness of a distance directly, Carriere and Oudot [31] studied the *intrinsic metrics* induced by a number of distances between Reeb graphs, and showed that the intrinsic versions of Gromov-Hausdorff distance, interleaving distance, functional distortion distance and the bottleneck distance are all equivalent.

To the best of our knowledge, it remains unknown whether most of the above distances proposed for Reeb graphs are intrinsic or not [31]. Gasparovic et al. [43] gave the first positive answer in the setting of merge trees. They used the labeled interleaving distance to prove that the ordinary interleaving distance on merge trees is strictly intrinsic. This means that one can always find geodesics in the space of merge trees connecting two given merge trees, and such that their length approximates the distance arbitrarily closely. Therefore, there exists an average, albeit not necessarily unique, of two merge trees, i.e., a merge tree that lies halfway between them in the sense of the interleaving distance. Thus, for instance, in the climate simulation application, we may compute an average merge tree from an ensemble of merge trees. It appears hopeful that we could perform statistics on the space of contour trees and Reeb graphs, similar to merge trees, using the interleaving distance. However, we show that the (ordinary, unlabeled) interleaving distance is not intrinsic. We arrived at this result by trying to prove the opposite, that is, to use the labeled interleaving distance to deduce, as in [43], that there are intermediary Reeb graphs. Hence the connection to the first part of our results.

Very recently, Bauer et al. [45] proved that the interleaving distance of sheaves on the real line is not intrinsic, which implies that the interleaving distance of Reeb graph is not intrinsic either. We were not aware of this result. In contrast, our counterexamples are simple and we have a direct argument that proves that the interleaving distance of Reeb graphs is not intrinsic.

We aim to generalize the results of [43] to the setting of Reeb graphs, succeeding partly:

- Our main contribution is the definition of a labeled interleaving distance for Reeb graphs (Section 3) and Theorem 1, which states that the ordinary interleaving distance between a pair of Reeb graphs can be obtained by taking the minimum of the labeled interleaving distance over all suitable labelings. Moreover, we show that under reasonable assumptions on the labelings, the resulting distance enjoys the same metric properties as the (unlabeled) interleaving distance (Theorem 2 and Appendix B).
- We provide a simple $\tilde{O}(n^2)$ algorithm to compute the labeled interleaving distance between contour trees, where n is the number of nodes in the contour tree (Section 4).
- In the negative direction, we provide simple but not immediate counterexamples and prove that the interleaving distance on Reeb graphs (and contour trees) is not intrinsic (Section 5). It follows that computing averages of Reeb graphs is challenging at least with respect to the interleaving distance.

2 Background

2.1 Reeb Graphs

Assume we are given a well-behaved function $f : \mathbb{M} \rightarrow \mathbb{R}$, i.e., a Morse function on a manifold, a constructible $\mathbb{R}$ -space [21], or a levelset-tame function [46] on a

topological space. Following Figure 1, a *Reeb graph* $\mathbb{M}/\sim$ arises from a quotient map $\pi : (\mathbb{M}, \tilde{f}) \rightarrow (\mathbb{M}/\sim, f)$. In our context, a Reeb graph is denoted as a pair $R = (G, f)$ where G is a combinatorial finite multi-graph and $f : |G| \rightarrow \mathbb{R}$ a function inherited from the input data defined as $f(\pi(x)) = \tilde{f}(x)$. Here, G is a combinatorial object and $|G|$ is its underlying topological space; we sometimes use G in place of $|G|$ for simplicity. In fact, our results apply to a more general setting, where a Reeb graph is defined to be a graph G with function values associated with each node such that the function value on the edges is strictly monotone (with the maximum and minimum determined by the function values at the nodes) [21]. For the rest of the paper, we work with this definition. Therefore, f is strictly monotone on any edge of G . We also assume that f is injective on the nodes.

The *up-degree* d^+ of a node v of G is the number of edges uv such that $f(u) > f(v)$. The *down-degree* d^- is defined analogously. A node is called a *split node* if its $d^+ > 1$ and $d^- = 1$. A node is called a *join node* if $d^+ = 1$ and $d^- > 1$. In Figure 1, u is a split node and v is a join node. A node with $d^+ = d^- = 1$ is a *regular node*. A node with $d^+ = 0$ and $d^- = 1$ is a *(local) maximum*. A *(local) minimum* is defined analogously. A non-regular node is *critical* and corresponds to a critical point of $\tilde{f}$ (i.e., for a smooth $\tilde{f}$).

In degenerate scenarios, a Reeb graph might contain a degenerate node, for instance, with $d^+ = 0$ and $d^- = 2$. For technical reasons and for simplifying our proofs, we consider such a node a superposition of two nodes, one with $d^+ = 1$ and $d^- = 2$, and another with $d^+ = 0$ and $d^- = 1$. We apply similar operations for nodes with $d^+ \geq 2$ and $d^- \geq 2$. This way, a single degenerate node of the Reeb graph consists of several superposed simpler nodes. We consider the nodes in superposition to be connected by edges of length 0.

When it is clear from the context, we abuse the notations and denote by R some geometric realization of the Reeb graph $R = (G, f)$ (e.g., embedded in $\mathbb{R}^3$ for visualization) such that f is shown as a height function in the vertical direction; see Figure 1. In this way, we could talk about an *arc* being a subset of R that does not necessarily start or end at a node. In addition, we denote a geometric realization of the graph G also by G .

2.2 Interleaving Distance Between Reeb Graphs

We recall the (ordinary) interleaving distance in a geometric way following notations in [21, 24]. We refrain from using categorical language as long as it is not needed in our arguments; see [21] for a categorical interpretation of Reeb graphs.

Let $R_1 = (G_1, f_1)$ and $R_2 = (G_2, f_2)$ be two Reeb graphs. A *morphism* of Reeb graphs, $h : R_1 \rightarrow R_2$, is a continuous function $h : G_1 \rightarrow G_2$ that is *function-preserving*, i.e., $f_2(h(x)) = f_1(x)$.

For a Reeb graph $R = (G, f)$ and $\varepsilon \geq 0$, we define the ε -*thickening* of R , denoted $\mathcal{T}^\varepsilon(R)$, to be $G \times [-\varepsilon, \varepsilon]$ with the product topology, together with the function $F^\varepsilon : G \times [-\varepsilon, \varepsilon] \rightarrow \mathbb{R}$, $F^\varepsilon(x, t) \mapsto f(x) + t$. We then define the ε -*smoothing* $\mathcal{S}^\varepsilon(R)$ to be the Reeb graph of $\mathcal{T}^\varepsilon(R)$ with respect to F^ε . $\mathcal{S}^\varepsilon(R)$ is




Fig. 2 Examples for a contour tree (top) and a Reeb graph (bottom). From left to right: a Reeb graph $R = (G, f)$, its ε -thickening $G \times [-\varepsilon, \varepsilon]$ with a function F^ε , and the Reeb graph $\mathcal{S}^\varepsilon(R)$ of the ε -thickening.

equipped with a function f^ε that is inherited from F^ε . See Figure 2 for examples involving a contour tree (top) and a Reeb graph (bottom) respectively. The ε -thickening of the Reeb graph in Figure 2 (bottom) is tilted slightly to reveal its structure. For simplicity, we abbreviate $R^\varepsilon = \mathcal{S}^\varepsilon(R)$.

By construction, there is an inclusion $i : G \rightarrow G \times [-\varepsilon, \varepsilon]$, $x \mapsto (x, 0)$, as well as a quotient map $\pi : G \times [-\varepsilon, \varepsilon] \rightarrow R^\varepsilon$; their composition gives rise to an ε -shift morphism, $\eta := \pi \circ i : R \rightarrow R^\varepsilon$, which assigns to the point $x \in G$ the connected component of $(x, 0) \in G \times [-\varepsilon, \varepsilon]$ in the pre-image of F^ε , see [21] for details.

If $\phi : R_1 \rightarrow R_2^\varepsilon$ is a morphism, then, for every $\delta \geq 0$ including $\delta = \varepsilon$, there exists a morphism $\phi^\delta : R_1^\delta \rightarrow R_2^{\delta+\varepsilon}$.¹ Analogously, if $\psi : R_2 \rightarrow R_1^\varepsilon$ is a morphism, then there exists a morphism $\psi^\varepsilon : R_2^\varepsilon \rightarrow R_1^{2\varepsilon}$. Similarly, given a morphism $\eta_1 : R_1 \rightarrow R_1^\varepsilon$, there exists a morphism $\eta_1^\varepsilon : R_1^\varepsilon \rightarrow R_1^{2\varepsilon}$; η_2^ε is defined analogously.

Definition 1 Let R_1 and R_2 be a pair of Reeb graphs. An ε -interleaving between R_1 and R_2 is given by two morphisms, $\phi : R_1 \rightarrow R_2^\varepsilon$, $\psi : R_2 \rightarrow R_1^\varepsilon$, such that the

¹This map is defined first as a function-preserving map from the thickening $T^\delta(R_1)$ into the thickening $T^\delta(R_2^\varepsilon)$ as $\Phi^\delta(x, t) = (\phi, t)$. This map will induce the map ϕ^δ on the Reeb graphs.

following diagram commutes.

$$\begin{array}{ccccc}
 R_1 & \xrightarrow{\eta_1} & R_1^\varepsilon & \xrightarrow{\eta_1^\varepsilon} & R_1^{2\varepsilon} \\
 & \searrow \phi & & \searrow \phi^\varepsilon & \\
 & & & & \\
 R_2 & \xrightarrow{\eta_2} & R_2^\varepsilon & \xrightarrow{\eta_2^\varepsilon} & R_2^{2\varepsilon} \\
 & \swarrow \psi & & \swarrow \psi^\varepsilon & \\
 & & & &
 \end{array} \tag{1}$$

Equivalently, see [21, Definition 4.38], we require that the following two relations hold,

$$\begin{aligned}
 \phi^\varepsilon \circ \psi &= \eta_2^{2\varepsilon} \\
 \psi^\varepsilon \circ \phi &= \eta_1^{2\varepsilon},
 \end{aligned} \tag{2}$$

where $\eta_1^{2\varepsilon} = \eta_1^\varepsilon \circ \eta_1 : R_1 \rightarrow R_1^{2\varepsilon}$ and $\eta_2^{2\varepsilon} = \eta_2^\varepsilon \circ \eta_2 : R_2 \rightarrow R_2^{2\varepsilon}$ are the 2ε -shift morphisms.

The *interleaving distance* is defined to be

$$d_I(R_1, R_2) := \inf\{\varepsilon \geq 0 \mid \text{there exists an } \varepsilon\text{-interleaving between } R_1 \text{ and } R_2\}. \tag{3}$$

Smoothing by 2ε is isomorphic to smoothing twice by ε . It is known that if we consider ε -smoothings of a Reeb graph R for increasing values of $\varepsilon \geq 0$, the minima and join nodes move downwards (i.e., they move to points with lower function values in the smoothed Reeb graph), and maxima and split nodes move upwards, see [26, 47] for a proof. Recall that we consider a degenerate node (e.g. with $d^+ > 1$ and $d^- > 1$) as a superposition of a join node and a split node. When smoothed by ε these nodes give rise to two nodes, one split and the other join, that are separated by 2ε .

If $T := R$ is a contour tree, every node v of T corresponds to a node in T^ε which lies ε above or below in function value. We denote this correspondence by $s : V(T) \rightarrow V(T^\varepsilon)$. For nodes in a superposition, each superposed node has its image under s . This is the reason behind the idea of superposition, since each superposed node moves differently when smoothed.

If R is a Reeb graph, a split node v and a join node w might cancel each other in the smoothed graph, and a loop might disappear. In this case, there are no node correspondences in the smoothed graph. For a split vertex v , we define $s(v) = \pi((v, \varepsilon)) \in R^\varepsilon$, and, for a merge vertex w , we define $s(w) = \pi((w, -\varepsilon)) \in R^\varepsilon$. In this way, $s : V(R) \rightarrow R^\varepsilon$ is defined on all nodes.

2.3 Intrinsic Metrics and Geodesic Spaces

Two Reeb graphs R_1 and R_2 are *isomorphic* if there are function-preserving continuous maps $h : R_1 \rightarrow R_2$ and $g : R_2 \rightarrow R_1$ such that $f \circ g$ and $g \circ f$ are identity maps. This is equivalent to $d_I(R_1, R_2) = 0$ since by definition, h and g define a 0-interleaving.

The interleaving distance is a metric if we identify isomorphic Reeb graphs, see [21, Proposition 4.8]. Given any metric d over the set of (isomorphism classes) of Reeb graphs, we obtain the *space of Reeb graphs*. Given a continuous

path γ (i.e., continuous w.r.t. d) into the space of Reeb graphs, where $\gamma(0) = R_1$ and $\gamma(1) = R_2$, the length $\ell_d(\gamma)$ induced by d is defined as

$$\ell_d(\gamma) = \sup_{n, \Gamma} \sum_{i=1}^n d(\gamma(\Gamma(i)), \gamma(\Gamma(i+1))),$$

where n ranges over all $\mathbb{N}$ (natural numbers) and Γ ranges over all partitions of $[0, 1]$, i.e. all ordered sets of n points $\Gamma = (\Gamma(1), \dots, \Gamma(n))$ such that $\Gamma(i) \in [0, 1]$; see [43, 48].

Given a metric d , the *intrinsic metric* $\hat{d}$ induced by d is defined as $\hat{d}(R_1, R_2) = \inf_{\gamma} \ell_d(\gamma)$; it is the infimum of the lengths of all paths from one point to another. It is known that $d \leq \hat{d}$. A metric d is *intrinsic* if it is always equal to its intrinsic version $\hat{d}$; in this case, the metric space is said to be a *length space*. A metric space is a *geodesic space* if any two points in the space can be connected by a curve of length equal to the distance between the two points; the metric is then called *strictly intrinsic* [31, 43, 48]. A strictly intrinsic metric is of course intrinsic, thus a geodesic space is necessarily a length space.

3 Labeled Interleaving Distance Between Reeb Graphs

In this section, we first define the labeled interleaving distance between a pair of labeled Reeb graphs and prove its properties. Then we show that such a distance generalizes the labeled interleaving distance between merge trees.

3.1 Labeled Reeb Graphs and Their Distance

Definition 2 Let L be a finite set of *labels*. For simplicity, $L = [N] := \{1, \dots, N\}$. Let $V = V(R)$ denote the set of nodes of a Reeb graph $R = (G, f)$. A *labeling* of R is a function $\lambda : L \rightarrow V$. We call the triple $R_\lambda = (G, f, \lambda)$ an *L -labeled Reeb graph*, or simply a *labeled Reeb graph* when the set of labels L is clear from context. R_λ is *fully-labeled* if λ is surjective. Note that λ is not necessarily injective. A *morphism* between labeled Reeb graphs is defined to be the morphism of the underlying unlabeled Reeb graphs.

Definition 3 Let $v \in V(R)$ be a node in the Reeb graph $R = (G, f)$. The ε -*path-neighborhood* of v , denoted $P^\varepsilon(v)$, is $\pi(\{v\} \times [-\varepsilon, \varepsilon]) \subset R^\varepsilon$ in the ε -smoothed Reeb graph R^ε . Here $\pi : G \times [-\varepsilon, \varepsilon] \rightarrow R^\varepsilon$ is the quotient map. For any $x \in R$, $P^\varepsilon(x)$ is defined similarly.

As illustrated in Figure 3, we observe that for any point $x \in R$ and any $\varepsilon \geq 0$, $P^\varepsilon(x)$ is a monotonic path in R^ε such that $f^\varepsilon(P^\varepsilon(x)) = [f(x) - \varepsilon, f(x) + \varepsilon]$.




Fig. 3 The ε -path-neighborhoods of node a and b are highlighted in green and yellow in the right-most figure, respectively.

In other words, $P^\varepsilon(x)$ is a monotonic path centered on the image of x in R^ε , i.e., centered on $\eta(x) := \pi \circ i(x)$.

Definition 4 Let $R_{1,\lambda_1} = (G_1, f_1, \lambda_1)$ and $R_{2,\lambda_2} = (G_2, f_2, \lambda_2)$ be two L -labeled Reeb graphs, and let $\varepsilon \geq 0$. We say a pair of morphisms $\phi : R_1 \rightarrow R_2^\varepsilon$ and $\psi : R_2 \rightarrow R_1^\varepsilon$ define a *labeled ε -interleaving* between R_{1,λ_1} and R_{2,λ_2} if the following hold:

1. ϕ and ψ define an ε -interleaving between R_1 and R_2 .
2. For each $\ell \in L$, we have the following label-preserving properties,

$$\begin{aligned} \phi^\varepsilon(s(\lambda_1(\ell))) &\in P^\varepsilon(s(\lambda_2(\ell))), \\ \psi^\varepsilon(s(\lambda_2(\ell))) &\in P^\varepsilon(s(\lambda_1(\ell))). \end{aligned} \tag{4}$$

In the the above formulae, $P^\varepsilon(s(\lambda_1(\ell))) \subset R_1^{2\varepsilon}$ and $P^\varepsilon(s(\lambda_2(\ell))) \subset R_2^{2\varepsilon}$. The *labeled interleaving distance* between R_{1,λ_1} and R_{2,λ_2} is defined as

$$\begin{aligned} d_I^L(R_{1,\lambda_1}, R_{2,\lambda_2}) \\ = \inf\{\varepsilon \geq 0 \mid \text{there exists a labeled } \varepsilon\text{-interleaving between } R_{1,\lambda_1} \text{ and } R_{2,\lambda_2} \}. \end{aligned} \tag{5}$$

Recall that the function s maps the nodes of a Reeb graph to the points of the ε -smoothed Reeb graph (see the end of [Section 3](#)). Intuitively, it outputs the point where the original node would arrive at after moving up or down by ε . For simplified notation, we refer to all such functions by s , instead of naming them differently based on their domain or co-domain.

From [Definition 4](#), it follows easily that for all R_1, R_2, λ_1 , and λ_2 , we have,

$$d_I^L(R_{1,\lambda_1}, R_{2,\lambda_2}) \geq d_I(R_1, R_2).$$

If the label set L is empty, then the second condition in [Definition 4](#) is vacuous and the labeled interleaving distance equals the unlabeled one.

Let R_1 and R_2 be two Reeb graphs. It turns out that the existence of maps defining an ε -interleaving between R_1 and R_2 does not depend on all of the nodes and edges of R_1 and R_2 . Intuitively, it depends only on the topological




Fig. 4 ε -essential (left) and ε -inessential nodes (right).

features that are significant enough w.r.t. to ε . We make this precise by defining the notion of ε -essential and ε -inessential nodes.

Definition 5 Let $R = (G, f)$ be a Reeb graph and let $\varepsilon > 0$ be given. A split node v is called ε -inessential if one of the following holds:

1. There exists a join node w and two paths Q_1 and Q_2 such that each Q_i ($i \in \{1, 2\}$) joins v to w with $f(v) \leq f(x) \leq f(v) + 4\varepsilon$, for all $x \in Q_i$.
2. Let $\mathcal{Q}$ be the set of all paths $Q = vu$ such that $f(u) \geq f(v) + 2\varepsilon$ and $\forall x \in Q, f(x) \geq f(v)$. Then all paths in $\mathcal{Q}$ use the same outgoing edge from v .

An ε -inessential join node is defined analogously. Then ε -essential nodes are all maxima, minima, join and split nodes that are not ε -inessential.

The first condition of the ε -inessential definition implies that there is a loop of height at most 4ε such that v is the lowest point of the loop. The second condition indicates that the parts of the graph reachable from v via paths that do not go below $f(v)$ either have heights at most 2ε , or they are reached using the same edge incident on v ; see Figure 4.

We call a labeling an ε -essential labeling if every ε -essential node is labeled. The following main theorem shows that we obtain the (ordinary) interleaving distance between two Reeb graphs if we consider all possible ε -essential labelings.

Theorem 1 Let R_1 and R_2 be two Reeb graphs and set $\varepsilon = d_I(R_1, R_2)$. There exist a set L of labels and ε -essential labelings λ_1 and λ_2 such that for L -labeled Reeb graphs R_{1,λ_1} and R_{2,λ_2} , we have

$$d_I(R_1, R_2) = d_I^L(R_{1,\lambda_1}, R_{2,\lambda_2}).$$

The reason for considering ε -essential labelings in Theorem 1 is to present a label set L and labelings that cover a good portion of the nodes; the theorem would be trivial for $L = \emptyset$.

We first sketch the proof and then present it in full detail. Let $\varepsilon = d_I(R_1, R_2)$. We consider the case of a maximum or an essential split node v in




Fig. 5 An illustration for the proof of [Theorem 1](#).

R_1 , as illustrated in [Figure 5](#). Then $v'' = s(s(v))$ is also a maximum or split node in $R_1^{2\varepsilon}$. Let $P = P^\varepsilon(s(v)) \subset R_1^{2\varepsilon}$ be the ε -path neighborhood of $s(v)$. We then consider the space $W = (\psi^\varepsilon)^{-1}(P) \subset R_2^\varepsilon$. From the properties of ε -interleaving, we deduce that $w = \phi(v) \in W$. Next, we walk upward in W : starting from w until reaching a split node or a maximum. We argue that one of these nodes (i.e., a split node or a maximum) is reachable inside W . We denote such a maximum or split node by $s(y)$ for some $y \in R_2$, and pair v with y . These pairs provide the desired labeling.

Proof Let $\varepsilon = d_I(R_1, R_2)$ be realized by morphisms $\phi : R_1 \rightarrow R_2^\varepsilon$, $\psi : R_2 \rightarrow R_1^\varepsilon$. Let $V_1 = V(R_1)$ and $V_2 = V(R_2)$ be the node sets. We will construct a subset $B \subset V_1 \times V_2$ and use B to define a labeling on R_1 and R_2 .

As illustrated in [Figure 5](#), let $v \in V_1$ be either a maximum or an essential split node. Then there is also a maximum or a split node $v'' = s(s(v))$ in $R_1^{2\varepsilon}$ such that $f_1^{2\varepsilon}(v'') = f_1(v) + 2\varepsilon$. Let $w = \phi(v)$ in R_2^ε . Also, let $v' = s(v)$ and let $P = P^\varepsilon(v')$ be the ε -path neighborhood of v' in $R_1^{2\varepsilon}$. P is a path with endpoints v'' from above and $v^{2\varepsilon}$ from below. For simplicity, we denote the image $\eta_1(v_i)$ for any node $v_i \in R_1$ by v_i^ε . Then, $\eta_1^\varepsilon(\eta_1(v)) = v^{2\varepsilon}$. Let $W \subset R_2^\varepsilon$ be the pre-image $W = (\psi^\varepsilon)^{-1}(P)$. We have $w \in W$, since $\psi^\varepsilon(w) = \psi^\varepsilon(\phi(v)) = \eta_1^\varepsilon(\eta_1(v))$.

We walk up (i.e. in the direction of increasing function values) from $w \in W$ until we reach a split node or a maximum. We claim that one of these two possibilities must occur while the function value is at most $f_1(v) + 2\varepsilon = f_1^{2\varepsilon}(v'')$. If v is a maximum then the point with the maximum value in W must be a maximum node with function value at most $f_1^{2\varepsilon}(v'')$. It follows that in this case we either reach a split node or a maximum. If v is an essential split node, we claim that for a sufficiently small $\delta > 0$, there are two points a and b in R_1 with $f_1(a) = f_1(b) = f_1(v) + 2\varepsilon + \delta$ such that $a^{2\varepsilon} \neq b^{2\varepsilon}$. To see this, we argue as follows. We move up along the two different branches of the split node v . If the branches join before advancing by value

$2\varepsilon + \delta$, then this satisfies the first condition of inessentiality (Definition 5), violating our assumption that v is an essential split node. Thus, $a \neq b$ at function value $f(v) + 2\varepsilon + \delta$. Since the split node moves up at most 2ε in $R_1^{2\varepsilon}$, $a^{2\varepsilon} = b^{2\varepsilon}$ is only possible if there exists a join node $j \in R_1$ that moves past the value of a and b from a higher function value. Since j moves down by 2ε , j'' will merge with the split node v'' if it moves at most δ more. Thus, there must be a loop with height at most $4\varepsilon + \delta$. Since δ is small, the loop has height at most 4ε , which again satisfies the first condition of inessentiality. We reach a contradiction. Therefore, we have $a^{2\varepsilon} \neq b^{2\varepsilon}$.

Since $\psi^\varepsilon \circ \phi = \eta_1^\varepsilon \circ \eta_1$ and $a^{2\varepsilon} \neq b^{2\varepsilon}$, then $\phi(a) \neq \phi(b)$. This implies that we will see a split node when moving up from w before passing the function value $f_1(v) + 2\varepsilon$ (in fact, before passing $f_1(v) + 2\varepsilon + \delta$ by the above argument for all small $\delta > 0$, which implies before passing $f_1(v) + 2\varepsilon$). It might be that the split node has value $f_1(v) + 2\varepsilon$. This proves our claim for an essential split node.

Regardless of the type of the node v , let y' be the split node or the maximum we reach when moving up from w . Note that it could be that $y' = w$. Then $f_2^\varepsilon(y') \leq f_2^\varepsilon(v'') = f_1(v) + 2\varepsilon$, meaning that we travel by at most 2ε in function values. y' corresponds to a split or a maximum y in R_2 such that $y' = s(y)$ (recall our convention on superposition of nodes). We add (v, y) to B , indicating that $v \in R_1$ and $y \in R_2$ are assigned the same label. We have $\psi^\varepsilon(y') \in P^\varepsilon(v')$ by construction. If y is an inessential split mode, we have completed the proof. Otherwise, we need to show that $\phi^\varepsilon(v') \in P^\varepsilon(y')$ to stop considering the node y .

Let Q be the path that we have traversed connecting w to y' . Since Q is a monotone path with height at most 2ε , applying Lemma 1 below with v and $z = y'$ implies that $\phi^\varepsilon(v') \in P^\varepsilon(y')$.

We analogously pair all minima and essential join nodes of R_1 with minima and join nodes of R_2 . After pairing all the essential nodes of R_1 with nodes of R_2 , we repeat the above process for the essential nodes of R_2 that are not paired yet. In the end, we obtain a set B containing all paired nodes. We choose B to be our label set L . By construction the morphisms ϕ and ψ define a labeled ε -interleaving w.r.t. this choice of labels. $\square$

Lemma 1 *Let ϕ and ψ define an ε -interleaving between R_1 and R_2 . Let v be a split node or a maximum in the Reeb graph R_1 and let z be a split node or a maximum in R_2^ε . Let $M_\varepsilon^-(z) \subset R_2^\varepsilon$ be the union of all points x that are connected to z by a monotone path Q_x of height at most 2ε and such that $f_2(q) \leq f_2(z)$, for all $q \in Q_x$. If $\phi(v) \in M_\varepsilon^-(z)$ then $\phi^\varepsilon(v') \in P^\varepsilon(z)$. The analogous statement holds for a join node or a minimum.*

Proof Set $x = \phi(v)$, then there is a monotone path Q connecting x to z that lies below z . Let F_2^ε be the function on $\mathcal{T}^\varepsilon(R_2^\varepsilon)$. Since the height of Q is at most 2ε , the point $(x, \varepsilon) \in \mathcal{T}^\varepsilon(R_2^\varepsilon)$ and the point $(z, \varepsilon - Q)$ (where $\varepsilon - Q \geq -\varepsilon$) are connected in the pre-image $(F_2^\varepsilon)^{-1}(F_2^\varepsilon(x, \varepsilon)) \subset \mathcal{T}^\varepsilon(R_2^\varepsilon)$, see Figure 6. By the definition of ϕ^ε , we have $\phi^\varepsilon(v') = \pi_2^{2\varepsilon}(\phi(v), \varepsilon) = \pi_2^{2\varepsilon}(x, \varepsilon) = \pi_2^{2\varepsilon}(z, \varepsilon - \|Q\|)$. It follows that $\phi^\varepsilon(v') \in P^\varepsilon(z) \subset R_2^{2\varepsilon}$. $\square$

Remark 1 Based on Theorem 1, observe that the nodes with the same label move in the same direction when smoothed. We call this *consistency* between the two




Fig. 6 An illustration for the proof of [Lemma 1](#).

labelings. This condition is necessary for the distance to be a metric; see [Theorem 2](#) and its proof in [Appendix B](#).

Theorem 2 Let $\mathcal{R}_c^L$ be a maximal set of isomorphism classes of consistently L -labeled Reeb graphs. The labeled interleaving distance is an extended metric on $\mathcal{R}_c^L$.

We have made additional observations regarding the local neighborhood of a node u under the mapping of η w.r.t. its ε -path-neighborhood, see [Appendix A](#).

3.2 Generalizing Labeled Interleaving Between Merge Trees

In this section, we show that our definition of the labeled interleaving distance between Reeb graphs generalizes the labeled interleaving distance defined in [\[43\]](#) for merge trees. Viewed as graphs with functions defined on them, merge trees may be considered as special cases of Reeb graphs. In [\[43\]](#), the maximum node of a merge tree T is defined to have value ∞ and a labeling λ is assumed to be surjective on the set of non-root leaves. A matrix $U(T_\lambda)$ called the *induced matrix* is assigned to a labeled merge tree T_λ . It is defined as

$$U(T_\lambda)[i, j] = f(\text{LCA}(\lambda(i), \lambda(j))),$$

where $L = [n]$ and LCA denotes the lowest common ancestor of nodes labeled i and j in T_λ . Given two labeled merge trees T_{1, λ_1} and T_{2, λ_2} with the same label set L , the labeled interleaving distance between the two merge trees (in

the sense of [43]), denoted $\bar{d}_I^L$, is defined as

$$\bar{d}_I^L(T_{1,\lambda_1}, T_{2,\lambda_2}) = \|U(T_1, \lambda_1) - U(T_2, \lambda_2)\|_{\max}.$$

where $\|\cdot\|_{\max}$ is the element-wise maximum.

A labeled merge tree as defined in [43] is also a labeled merge tree based on our definition (as a special case of Reeb graphs), if we replace ∞ by a large number depending on the function values of the tree. We assume this modification for the following proposition.

Proposition 1 *For any two labeled merge trees (in the sense of [43]), we have*

$$d_I^L(T_{1,\lambda_1}, T_{2,\lambda_2}) = \bar{d}_I^L(T_{1,\lambda_1}, T_{2,\lambda_2}).$$

The main observation for the proof of this proposition is that for a merge tree with root at ∞ , the smoothing operation does not change the structure of the tree, but it pushes the entire tree down uniformly.

Proof Since the ∞ nodes in the two trees have large and equal function values, the maps ϕ, ψ realizing $d_I^L(T_{1,\lambda_1}, T_{2,\lambda_2})$ map these nodes to the path-neighborhoods of each other. Since they do not cause any restriction, we ignore these nodes.

We first show that $d_I^L(T_{1,\lambda_1}, T_{2,\lambda_2}) \leq \bar{d}_I^L(T_{1,\lambda_1}, T_{2,\lambda_2})$. Let $\bar{\varepsilon} = \bar{d}_I^L(T_{1,\lambda_1}, T_{2,\lambda_2})$. Since every leaf is labeled, every node is the LCA of two leaves and must appear in the matrix. Therefore, the right hand side is the maximum of the differences between nodes with the same label, and nodes that are LCA for the same pair of labeled leaves. We claim that there exist ϕ, ψ defining a $\bar{\varepsilon}$ -labeled interleaving between the two trees. We can define ϕ by mapping a leaf v_1 of T_1 to the $\bar{\varepsilon}$ -path-neighborhood of the node with the same label, say v_2 in T_2 . This is possible because v_2 has moved down by at least the height-difference between these nodes. With similar reasoning, after mapping two leaves v_1 and w_1 , we can also map the join node which is their LCA to the LCA of nodes with the same label in T_2 . The paths connecting v_1 and w_1 to their individual LCA are therefore mapped to the smoothed tree $T_2^{\bar{\varepsilon}}$. The map ϕ defined here satisfies a condition stronger than the label-preserving property (4), see Remark 2 in Section 3.2. These paths cover the whole tree, hence we have defined the morphism on all of T_1 . Since we can consider T_1 as a subset of $\mathcal{T}^{\bar{\varepsilon}}(T_2)$, and analogously T_2 as a subset of $\mathcal{T}^{\bar{\varepsilon}}(T_1)$, it is easily checked that the commutativity relations (2) hold.

Next, we show that $\bar{d}_I^L(T_{1,\lambda_1}, T_{2,\lambda_2}) \leq d_I^L(T_{1,\lambda_1}, T_{2,\lambda_2})$. Let $\varepsilon = d_I^L(T_{1,\lambda_1}, T_{2,\lambda_2})$. Let ϕ and ψ realize d_I^L . Let $v_1 \in T_1$ and $v_2 \in T_2$ be leaves with the same label. Since the second condition of the labeled interleaving distance is satisfied by ϕ , we see that the height-difference between $v'_1 = s(v_1)$ and $v'_2 = s(v_2)$ is at most ε . Since all nodes move in the same direction, the height difference between v_1 and v_2 is also at most ε .

We also need to show that the height difference between unlabeled join nodes is at most ε . Take two leaves $v_1, w_1 \in T_1$ and set $u_1 = LCA(v_1, w_1)$. Let $A_1 \subset T$ be the union of two paths that connect v_1 and w_1 to u_1 . Let v_2 and w_2 be the nodes in T_2 with the same labels as v_1 and w_1 , respectively. Set $u_2 = LCA(v_2, w_2)$ and let $A_2 \subset T_2$ be the union of the two paths connecting u_2 to v_2 and w_2 . Now observe that since we are working with merge trees, the part of $T_1^{\bar{\varepsilon}}$ other than the edge connecting

to the ∞ node is exactly the same as T_1 but shifted down by ε . Therefore, there exist paths $A_1^\varepsilon \subset T_1^\varepsilon$ and $A_2^\varepsilon \subset T_2^\varepsilon$ corresponding to A_1 and A_2 , respectively.

Now assume the height difference between u_1 and u_2 is larger than ε . Without loss of generality, assume u_1 is higher than u_2 as in Figure 7. Observe that now the subtree rooted at $u'_2 = s(u_2)$ must be mapped by ψ^ε to a single branch going down from $u''_1 = s(s(u_1))$, and assume it is the branch containing $w''_1 = s(s(w_1))$. However, then $\psi^\varepsilon(v'_2)$ is also on the same branch and cannot lie on $P^\varepsilon(v'_1)$, as shown in Figure 7. This is a contradiction. $\square$




Fig. 7 An illustration for the proof of Proposition 1.

Remark 2 Observe that in the case of merge trees, we could replace the second condition of the labeled interleaving distance for Reeb graphs (Definition 4) with the following stronger condition that for all labels ℓ in L ,

$$\begin{aligned} \phi(\lambda_1(\ell)) &\in P^\varepsilon(\lambda_2(\ell)), \\ \psi(\lambda_2(\ell)) &\in P^\varepsilon(\lambda_1(\ell)); \end{aligned}$$

where ϕ and ψ are maps that realize an ε -interleaving. However, we believe that with this definition, the analogous Theorem 1 will not hold for Reeb graphs. Nevertheless, we do not make a claim at this moment and leave this question open. Lemma 5 implies that the above conditions are stronger than our label-preserving conditions (4).

Remark 3 The algorithm presented in this paper can be applied to the case of merge trees. However, the algorithm of [43] is linear and therefore more efficient for merge trees.

4 Algorithms

We first present a general framework (Section 4.1) for computing the labeled interleaving distance between Reeb graphs. The framework consists of two subproblems: the *existence subproblem* (Section 4.2) and the *commutativity*

subproblem (Section 4.3). We solve the *existence subproblem* for the general case of Reeb graphs. For the *commutativity subproblem*, we consider the simpler case of contour trees and give a solution that depends on the existence of unique paths between pairs of nodes. We prove that our specialized efficient algorithm for contour trees has a run-time in $\tilde{O}(n^2)$.




Fig. 8 A spanning (left) and a non-spanning labeling (right). Labeled nodes are in blue.

When the label set L is empty, the labeled interleaving distance equals the unlabeled one. Therefore, we impose a restriction on L to obtain a polynomial algorithm in the case of contour trees. We say that a labeling λ is *spanning*, if:

1. For any pair of labeled nodes u, v , there is at most a single path connecting them without labeled nodes in the interior of the path;
2. All such paths cover the entire graph.

See Figure 8 for examples. By a *path* we mean a simple, possibly non-monotone path in the graph. For instance, the labeling of a fully-labeled graph is spanning. Furthermore, if the labeling is surjective on all leaves of a contour tree, the resulting labeling is spanning.

If the labelings λ_1 and λ_2 are arbitrary, it is often true that the labeled interleaving distance is ∞ . For simplicity, we assume that the labelings are *consistent*, meaning that two nodes with the same label move in the same direction when smoothed. It is therefore not allowed that a minimum and a maximum have the same label. See Appendix B for details.

4.1 Algorithmic Framework

In order to compute the labeled interleaving distance we need to define two morphisms ϕ and ψ satisfying the definition of labeled ε -interleaving (Definition 4), such that ε is smallest with this property. If for some ε there is a labeled ε -interleaving, then for any $\varepsilon' > \varepsilon$, there is also a labeled ε' -interleaving. Our general approach is to start with a large enough ε for which a labeled interleaving exists, then to compute the infimum by a binary search over the possible values. We explain in the following how we determine whether a labeled ε -interleaving exists for a fixed ε . We break this into two subproblems. Let R_{1,λ_1} and R_{2,λ_2} be the two input labeled Reeb graphs; let R_1 and R_2 be their underlying unlabeled Reeb graphs. The first subproblem is to compute for a given

$\varepsilon \geq 0$ all morphisms $\phi : R_1 \rightarrow R_2^{\varepsilon_1}$ and $\psi : R_2 \rightarrow R_1^{\varepsilon_1}$ satisfying (4). We call this subproblem the *existence subproblem*. The second subproblem is to decide if the relations in (2) are satisfied for maps computed in the existence subproblem. We call this the *commutativity subproblem*.

The distance we are looking for depends on the function values. We follow the convention that the complexity required to store the function values is bounded by a fixed polynomial in the complexity of the graph. We also assume that the nodes of our Reeb graphs have distinct function values, which can be achieved by perturbation if necessary. In this section, n represents the complexity of the input, which is the total number of nodes and edges in R_1 and R_2 .

4.2 Existence Subproblem

We focus on computing $\phi : R_1 \rightarrow R_2^\varepsilon$, and we compute ψ analogously. Let $V_L(R_1) \subset V(R_1)$ denote the set of nodes in R_1 with labels, i.e., the image of λ_1 . Consider $v \in V_L(R_1)$ and a label $l \in L$ such that $\lambda_1(l) = v$. Let $v' = \lambda_2(l)$. Then for all labels $l \in L$, if $v = \lambda_1(l)$ and $v' = \lambda_2(l)$, ϕ^ε must satisfy,

$$f_2^\varepsilon(s(v')) - \varepsilon \leq f_2^{2\varepsilon}(\phi^\varepsilon(s(v))) \leq f_2^\varepsilon(s(v')) + \varepsilon. \quad (6)$$

Since the path-neighborhoods are strictly monotone sub-paths of the corresponding Reeb graph, if (6) is satisfied for l , v , and v' , then the point $\phi^\varepsilon(s(v))$ is the point $x \in P_2^\varepsilon(s(v'))$ such that $f_2^{2\varepsilon}(x) = f_1^\varepsilon(s(v))$. It follows that for each labeled node v , the image of $s(v)$ under ϕ^ε is uniquely determined. Next, we determine all the possible points for the image of v under ϕ ; there might be more than one possibility. Recall that ϕ^ε is the map induced on the Reeb graphs by a map $\Phi^\varepsilon : \mathcal{T}^\varepsilon(R_1) \rightarrow \mathcal{T}^\varepsilon(R_2^\varepsilon)$ defined as $\Phi^\varepsilon(x, t) = (\phi(x), t)$. Also recall that $s(v)$ is the projection of the point $(v, \pm\varepsilon)$ to the Reeb graph (sign of ε depends on the type of v). Let $\pi_2^{2\varepsilon} : \mathcal{T}^\varepsilon(R_2^\varepsilon) \rightarrow R_2^{2\varepsilon}$ and $y = \phi^\varepsilon(s(v)) \in R_2^{2\varepsilon}$. The pre-image of y in $\mathcal{T}^\varepsilon(R_2^\varepsilon)$, denoted $(\pi_2^{2\varepsilon})^{-1}(y)$, is a collection of horizontal line segments that intersect with vertical line segments of the form $z \times [-\varepsilon, \varepsilon]$, where $z \in R_2^\varepsilon$, see Figure 9. W.l.o.g., we assume v is a split node or a maximum. From $\Phi^\varepsilon(v, \varepsilon) = (\phi(v), \varepsilon)$, we know that the value of $\phi(v)$ is the midpoint of *some* vertical segment that intersects the pre-image of y at its highest point. In brief, the set of possible values of $\phi(v)$, denoted $C(v)$, are determined by the condition that $\phi(v) \times [-\varepsilon, \varepsilon] \in \mathcal{T}^\varepsilon(R_2^\varepsilon)$ intersects the pre-image of y at $(\phi(v), \varepsilon)$. We say that a choice of value in $C(v)$ for each v result in a *choice* of $\phi|_{V_L}$. Given any ε , these collections can be obtained easily in the same run time as computing a Reeb graph R , because we can compute a smoothing by computing the Reeb graph of $\mathcal{T}^\varepsilon(R)$.

In the next step, we need to decide if there is a choice of images in $C(v)$ such that the resulting map can be extended to edges of R_1 . Since we assume that the labeling λ_1 is spanning, it follows that the graph $R_1 = Q_1 \cup Q_2 \dots \cup Q_k$, where Q_i is a simple path that starts and ends with labeled nodes u_i and w_i .




Fig. 9 An illustration of the choices of $\phi(v)$.

Moreover, for $i \neq j$, Q_i and Q_j intersect at the endpoints, if at all. Our goal now is to extend the morphism ϕ to these paths.

Lemma 2 *For a given $\varepsilon > 0$, a choice of $\phi|_{V_L}$ can be extended to a morphism $\phi : R_1 \rightarrow R_2^\varepsilon$ if and only if for all i , Q_i can be mapped into a path connecting $\phi^\varepsilon(u_i)$ to $\phi^\varepsilon(w_i)$ by a continuous, function-preserving map.*

We call the problem defined in the above lemma the *path extension* problem. Apart from choices of $\phi|_{V_L}$, this path extension subproblem also becomes harder or easier depending on the restrictions we impose on the structure of the Reeb graphs and on the labelings. Let us denote the computational complexity of the path extension problem by $c(n)$.

Proposition 2 *Let R_{1,λ_1} and R_{2,λ_2} be labeled Reeb graphs such that the labelings are spanning. Let $p(n)$ be an upper bound on the number of choices for $\phi|_{V_L}$ for all ε . There is an algorithm that computes the smallest ε for which the existence problem has a solution in time $O(n c(n) p(n) \log n)$, assuming $c(n) = \Omega(\log n)$; and in time $O(np(n) \log^2 n)$, otherwise.*

Proof We perform a binary search on the values of ε . First we choose ε to be large enough such that the existence of the morphism ϕ is guaranteed. Using (6), we can compute a value for ε such that the morphism ϕ is defined on nodes. It is not difficult to determine here whether the distance is ∞ . Next we increase this ε to a value ε' such that every split node in R_2^ε is above the maximum value of f_1 , denoted $\max(f_1)$, and every join node is below the minimum value of f_1 , denoted $\min(f_1)$. For such an ε , $(f_2^\varepsilon)^{-1}([\min(f_1), \max(f_1)])$ is a monotonic arc. It is then possible to extend ϕ to all paths connecting labeled nodes.

Starting from ε' we execute a binary search starting from the interval $[0, \varepsilon']$. During each iteration, let $[\alpha, \beta]$ be the interval at hand. We compute the midpoint $\varepsilon = (\alpha + \beta)/2$ of the interval. Theoretically, we can compute the Reeb graph R_2^ε using the algorithm of [49] on the space $\mathcal{T}^\varepsilon(R_2) = G_2 \times [-\varepsilon, \varepsilon]$ in $O(n \log n)$ time, since this latter space can be turned into a simplicial complex with $O(n)$ simplices. After computing the smoothed Reeb graph R_2^ε we use the path extension subroutine and

spend $O(np(n))$ time to check if each choice of the morphism ϕ defined on labeled nodes can be extended to all paths connecting labeled nodes. If the extension is possible for all paths for some choice, we set $\beta = \varepsilon$; otherwise, we set $\alpha = \varepsilon$. This step takes $O(nc(n)p(n) + n \log n)$ time to execute.

To finish the proof, we need to argue that the binary search stops after $O(\log n)$ iterations. Observe that as we smooth the Reeb graphs, the merge and split nodes of both input graphs move down and up, respectively. Therefore, we need to take careful consideration when the function values of the two sets of nodes cross each other. Then the value of ε we are looking for has the same asymptotic complexity as the input function values, which is at most $O(n)$ by our assumptions. Hence, we need at most $O(\log n)$ steps in the binary search. $\square$

4.3 Commutativity Subproblem

In [Section 4.2](#), we have computed ε such that certain morphisms ϕ and ψ exist. However, there might not be a unique way of mapping a path and the choice might be important for satisfying the commutativity requirement. If the first Betti numbers $\beta_1(G_i)$, for $i = 1, 2$, is $O(\log n)$, then there are polynomially many paths to choose from. We consider a simpler setting when the Reeb graphs are trees, i.e., contour trees. In this case, there is a unique path connecting any two vertices. Using this property we can prove that for a given ε there is at most a single label-preserving morphism. We then check this morphism if it satisfies the commutativity relations. We thus have the following theorem.

Theorem 3 *Let R_{1,λ_1} and R_{2,λ_2} be contour trees. Assume λ_1, λ_2 are consistent labelings that cover the leaves of the contour trees. Then there is an algorithm for computing the labeled interleaving distance in time $O(n^2 \log^2 n)$.*

Proof First, we claim that if the Reeb graph is a tree then $p(n) = 1$, and thus there is at most one choice for $\phi|_{V_L}$. Let w be the global maximum of R_2 . Since w is a labeled node, some node $v \in R_1$ has the same label as w . We know the point $\phi^\varepsilon(s(v))$ and the point $\phi(v)$ is one of the points in $C(v)$. These choices are caused by the existence of a join node below $s(w)$ as shown in [Figure 9](#), where $s(w)$ could be the maximum of R_2^ε . Since a join node moves down during smoothing, it can only have moved down from above w , which is not possible. Hence, there must be a single choice in $C(v)$. We thus defined ϕ on a single node uniquely. We show now that this morphism extends uniquely to all of R_1 , if at all.

We define ϕ on the sub-tree T rooted at v that consists only of join nodes. We move down from v inductively, and at each step, we add a monotone path P that has either only split nodes in the interior, or no node at all. Assume that we have already extended our map to some sub-tree of T . Let P connect nodes v_1 and v_2 such that the morphism is uniquely defined on v_1 . We want to extend the map uniquely to P or deduce that such an extension is not possible. There could be multiple choices for v_2 in $C(v_2)$. Since v_2 moves down when smoothed, these choices must be caused by some split node in R_2^ε . For the sake of contradiction, suppose there are two points $q_1, q_2 \in C(v_2)$ such that there is a monotone path from $\phi(v_1)$ to q_i , $i = 1, 2$. These paths have higher function values than $C(v_1)$, i.e., $\phi(v_1)$. There is also a split node

that connects q_1 and q_2 and has function value less than $C(v_2)$. This is however impossible, since R_2^ε is a tree. Therefore, there is at most one point in $C(v_2)$ that we can connect to. It follows that we can extend to T uniquely, if at all.

We now have nodes in T with ϕ defined on them. We can now start moving up from these nodes and define the function using a similar argument as above. These are split nodes in the interior of the path P and can possibly include v itself. We continue extending the function ϕ as long as there are split and join sub-trees remaining whose roots have been already uniquely mapped. This completes the proof of our claim that $p(n) = 1$.

Since the labelings are spanning, we can solve the existence subproblem as in [Proposition 2](#). To solve the path extension subproblem, we only need to check the possibility of mapping a path to another unique path, which takes $O(n)$ time per edge. As there is a unique way of mapping any edge, we need to perform a binary search as in [Proposition 2](#) to find the smallest ε such that the commutativity relations hold for the same maps found from existence subproblem. We can do this simultaneously as we solve the existence problem. Checking the commutativity can be done in linear time per edge, and therefore in time $O(n^2)$ for every step of the binary search. The entire algorithm takes $O(n^2 \log n) = \tilde{O}(n^2)$ time. $\square$

5 Nonexistence of Geodesics

As explained in [Section 1](#), it is important to know whether a given metric on Reeb graphs (or contour trees) is intrinsic, i.e., whether the metric space has geodesics. In [\[21\]](#), the labeled interleaving distance between merge trees is used to argue that the (ordinary) interleaving distance between merge trees is strictly intrinsic. Therefore, the space of merge trees under this metric is intrinsic. Here, we continue our discussion in [Section 3.2](#) on generalizing labeled interleaving between merge trees to Reeb graphs. Specifically, we consider this intrinsic property for the interleaving distance for Reeb graphs and contour trees. We show that the (ordinary) interleaving distance is not strictly intrinsic in these more general settings. Therefore, the labeled interleaving distance must not be strictly intrinsic as well. Our counterexamples are simple, however, the fact that they are counterexamples is not immediate.

If a space is intrinsic, there is an “average” for any two objects in the space. That is, for all pairs of Reeb graphs R_1 and R_2 , there exists an average Reeb graph $M = M(R_1, R_2)$ such that $d_I(M, R_1) = d_I(M, R_2) = d_I(R_1, R_2)/2$. We show below that the statement is not true for Reeb graphs or contour trees.

5.1 Counterexample for Reeb Graphs

Let R_1 be a line segment of height α and let R_2 be a simple loop of height α . The morphisms drawn in [Figure 10](#) show that $d_I(R_1, R_2) = \alpha/4$, because of the loop of height α in R_2 .

We assume for the sake of contradiction that there is a Reeb graph M such that $d_I(M, R_1) = d_I(M, R_2) = \alpha/8 = \varepsilon$. We have the following commutative diagram,




Fig. 10 A counterexample for Reeb graphs.

$$\begin{array}{ccccc}
 M & \longrightarrow & M^{\alpha/8} & \longrightarrow & M^{\alpha/4} \\
 & \searrow & \nearrow & & \nearrow \\
 R_2 & \longrightarrow & R_2^{\alpha/8} & \longrightarrow & R_2^{\alpha/4}
 \end{array}$$

Since R_2 and $R_2^{\alpha/4}$ each contains a loop of height α and $\alpha/2$, respectively, and the above diagram is commutative, $M^{\alpha/8}$ must contain a loop of height at least $\alpha/2$. Then by Lemma 3, M must have a loop of height at least $\alpha/2 + \alpha/4 = 3\alpha/4$. Now consider the relationship between R_1 and M . Since R_1 is a line segment and M has a loop of height at least $3\alpha/4$, the interleaving distance $d_I(M, R_1)$ is at least $(3\alpha/4)/4 = 3\alpha/16$. This is because that the loop in M must disappear after smoothing by 2ε , similar to the example in Figure 10. However, $3\alpha/16 > \alpha/8 = \varepsilon$, contradicting our assumption that $d_I(M, R_1) = d_I(M, R_2) = \alpha/8$.

Lemma 3 Let R be a Reeb graph. If R^ε has a loop of height a , then R must contain a loop of height $a + 2\varepsilon$.

Proof Let r be an injective, function-preserving map $r : R^\varepsilon \rightarrow \mathcal{T}^\varepsilon(R)$. r exists by our assumption that nodes have distinct function values. Let O be the loop of R^ε with height a , and let $x \in R^\varepsilon$ be a point of O . Let $y = r(x)$ and $J = z \times [-\varepsilon, \varepsilon] \subset \mathcal{T}^\varepsilon(R)$ be the vertical interval containing y . We move y up or down until it coincides with $z \times \{0\} \in R \times \{0\}$. We perform this move for all $x \in O$ such that x is a vertex of R^ε or z is a vertex of R .

The highest point of O , $\max(O)$, is a join node and its image must move ε up to lie in $R \times \{0\}$. Similarly, the lowest point of O , $\min(O)$, is a split node and its image

must move down ε to lie in $R \times \{0\}$. Other nodes move at most ε . Now, $\max(O)$ and $\min(O)$ are the highest and lowest nodes of a new loop with height $a + 2\varepsilon$. $\square$

5.2 Counterexample for Contour Trees

Since the loop in [Section 5.1](#) is essential, one might conjecture that when there is no loop in a Reeb graph (i.e., for a contour tree), the geodesic property holds. In this section, we present a counterexample to this conjecture.




Fig. 11 A counterexample for contour trees.

Let R_1 be an X-shaped contour tree in [Figure 11](#) (top), where we consider the crossing point to be the superposition of a join node and a split node that are arbitrarily close. Let R_2 be a line segment in [Figure 11](#) (bottom). The total height of the Reeb graphs are equal and we denote it by α . Let the lowest point of these graphs be the origin of the height function and let the line segment be centered at $\alpha/2$. Observe that $d_I(R_1, R_2) = \alpha/4$. The maps $\phi : R_1 \rightarrow R_2^\varepsilon$ and $\psi : R_2 \rightarrow R_1^\varepsilon$ defining an $\alpha/4$ -interleaving are drawn in [Figure 11](#). We again consider the following commutative diagram:

$$\begin{array}{ccccc}
 M & \longrightarrow & M^{\alpha/8} & \longrightarrow & M^{\alpha/4} \\
 & \searrow & \nearrow & & \nearrow \\
 R_1 & \longrightarrow & R_1^{\alpha/8} & \longrightarrow & R_1^{\alpha/4}
 \end{array}$$

For the sake of contradiction, we assume there is a contour tree M such that $d_I(M, R_1) = d_I(M, R_2) = \alpha/8$. We call a (part of a) contour tree a *join-split structure*, or JS-structure², if it consists of a split and a join node connected by a *connecting arc*, see [Figure 12](#) (left). We call the function (height) value difference between the split and the join node the *spread* of the JS-structure.

²This term is not intended to be related to the one in [\[50\]](#).

Note that the spread can be negative. For instance, R_1 in Figure 11 is a JS-structure where the connecting arc is trivial and the spread is 0.




Fig. 12 Left: A joint-split structure. Right: additional illustration for the counterexample involving contour trees.

As illustrated in Figure 12 (right), the tree $R_1^{\alpha/4}$ is a JS-structure with spread $\alpha/2$. Let j and s denote the join and split nodes in $R_1^{\alpha/4}$. Since the above diagram commutes, there must be a split node s' in $M^{\alpha/8}$ lying at or below s and a join node j' lying at or above j . Then, there is a JS-structure in $M^{\alpha/8}$ of spread at most $\alpha/2$. Since there are corresponding join and split nodes $j_0, s_0 \in M$ with function values $\alpha/8$ above j' and below s' , respectively, there must be a JS-structure in M of spread at most $\alpha/4$. Since R_2 is a line segment, if $d_I(M, R_2) = \alpha/8 = \varepsilon$, then after moving 2ε down, the join node j_0 must have a function value of at most 0; and after moving 2ε up, the split node s_0 must have a function value of at least α . Then s_0 must lie at most $\alpha/4$ below α and j_0 must lie at most $\alpha/4$ above 0. It follows that the spread is at least $\alpha/2$, which contradicts our earlier assumption.

Remark 4 In the above example, the join and split nodes in R_1 coincide. However, this is not essential and they can be separated by δ where δ is sufficiently small but positive. For example, it may be a Reeb graph of a Morse function defined on a 2-manifold.

5.3 More on Metric Properties of the Space of Reeb Graphs

The counterexamples above show that the spaces of Reeb graphs and contour trees are not geodesic w.r.t. d_I , implying that d_I is not strictly intrinsic. We can also deduce the following,

Proposition 3 *The interleaving distance d_I is not intrinsic.*

Proof If the interleaving distance was intrinsic, then there would exist paths connecting the two Reeb graphs R_1 and R_2 whose lengths, computed using the metric d_I ,

approach the interleaving distance arbitrarily closely. Then there would exist Reeb graphs which are arbitrarily close to a midpoint of R_1 and R_2 w.r.t. d_I . However, this cannot be true by the above example. $\square$

In the examples presented above, the parameter α is arbitrary. It follows that there are counterexamples that are arbitrarily close to each other w.r.t. d_I . In [21], the authors asked if $\hat{d}_I$ (and other intrinsic versions of distances on Reeb graphs) is strictly intrinsic. We conjecture that $\hat{d}_I$ is not strictly intrinsic, and more generally we conjecture the following,

Conjecture 1 *The space of isomorphism classes of Reeb graphs with any metric d is not a geodesic space.*

6 Conclusion and Discussion

In this paper, we define a labeled interleaving distance for Reeb graphs and prove its properties. The labeled interleaving distance does not approximate the ordinary interleaving distance, nor does it help with its computation. Rather, the labeled interleaving distance may be a substitute for the ordinary interleaving distance in some applications. We also show that the ordinary interleaving distance between Reeb graphs is not intrinsic. Therefore, there is no interpolation or average Reeb graph w.r.t. this distance between two close enough Reeb graphs. There are at least two ways to tackle this issue. First, we could try to find a Reeb graph which approximates an average. This would require a careful determination of the criteria for this average, and whether we should give more importance to one Reeb graph compared to the other one. Second, we could try to come up with distances with intrinsic properties. Even if [Conjecture 1](#) is true, there might be useful pseudo-metrics.

We present a first algorithm for computing the labeled interleaving distance between contour trees in polynomial time. We do not know whether a more efficient algorithm exists. It would also be useful to know if a polynomial algorithm exists for labeled Reeb graphs. Chambers et al. [24] recently defined a family of metrics related to the interleaving distance on Reeb graphs, using the concept of smoothing and truncation. Our current labeled interleaving distance only uses the concept of smoothing. Understanding the intrinsic properties of these truncated interleaving distances could also be an interesting continuation of the present work.

Acknowledgments. This work was partially supported by a grant from the Department of Energy (DOE) DE-SC0021015, and grants from the National Science Foundation (NSF) IIS-1910733 and IIS-2145499.

Appendix A An Observation on Path Neighborhoods

The following is an observation which we do not use but it might find application elsewhere.

Lemma 4 *Let u be a node in a Reeb graph R and let $N_\varepsilon(u)$ be the union of all points that are connected to u by a path Q of height at most ε and such that $f(q) \geq f(u)$ for all $q \in Q$, or $f(q) \leq f(u)$ for all $q \in Q$. Then $\eta(N_\varepsilon(u)) \subset P^\varepsilon(u)$.*




Fig. A1 An illustration for the proof of Lemma 4.

Proof Take $x \in N_\varepsilon(u)$. By assumption, the point x is connected by some path $Q_x \subset Q$ to u such that Q_x lies completely above or below u in the function values. This is because for a path to attain the value $f(u)$, it must be crossing $f(u)$ or be incident to u . Consider the thickening $\mathcal{T}^\varepsilon(R)$ and the function $f^\varepsilon : \mathcal{T}^\varepsilon(R) \rightarrow \mathbb{R}$. Since $Q_x \leq \varepsilon$, there is a point $u_x \in u \times [-\varepsilon, \varepsilon]$ such that $f^\varepsilon(u_x) = f^\varepsilon((x, 0)) = f(x)$, see Figure A1. We claim that x and u_x are connected in the pre-image $(f^\varepsilon)^{-1}(f(x))$. Consider the strip $Q_x \times [-\varepsilon, \varepsilon]$. The range $[f(u), f(u_x)]$ is always covered by lines $q \times [-\varepsilon, \varepsilon]$, for $q \in Q_x$. $\square$

Appendix B Metric Properties of the Labeled Interleaving Distance

It is easy to see that the labeled interleaving distance is not necessarily finite. This happens when the nodes that are labeled by the same label are originally apart and also move in opposite directions when smoothed, see Figure B2 for an example. Consequently, it could potentially only be an extended metric on the isomorphism class of labeled Reeb graphs, with the same label set L .




Fig. B2 A pair of inconsistently labeled Reeb graphs

In this section, we first put a condition on possible labelings and show that under this condition the labeled interleaving distance becomes a metric. Since the label set L is arbitrary, this provides a large number of metrics on large subsets of Reeb graphs that could be of interest in many applications, both in theory and in practice.

For a label set L , we denote the set of all L -labeled Reeb graphs by $\mathcal{R}^L$. We say that two labelled Reeb graphs $R_{1,\lambda_1}, R_{2,\lambda_2}$ are *consistently labeled* if for each label $\ell \in L$, the nodes $v_1 = \lambda_1(\ell)$ and $v_2 = \lambda_2(\ell)$ move in the same direction when smoothed. In other words, if v_1 is a maximum or split, then v_2 is also a maximum or a split, and if v_1 is a minimum or a join, then v_2 is a minimum or a join. A set of labeled Reeb graph $\mathcal{R} \subset \mathcal{R}^L$ is called *consistently labeled* if each pair of labeled Reeb graphs in it are consistently labeled.

The following lemma will be used in the proof of [Theorem 2](#).

Lemma 5 *Let R_1 and R_2 be two Reeb graphs, and $v_1 \in R_1, v_2 \in R_2$ be two nodes such that they move in the same direction when smoothed. Let $\phi : R_1 \rightarrow R_2^\varepsilon$ be a morphism satisfying $\phi(v_1) \in P^\varepsilon(v_2)$. Then, for any $\delta \geq 0$,*

$$\phi^\delta(s_\delta(v_1)) \in P^\varepsilon(s_\delta(v_2)),$$

where s_δ denotes the correspondence $s_\delta : V(R) \rightarrow V(R^\delta)$ for any pair of Reeb graphs.

Proof Without loss of generality, assume that both v_1 and v_2 are maxima or split nodes. The condition $\phi(v_1) \in P^\varepsilon(v_2)$ implies that $|f_1(v_1) - f_2(v_2)| \leq \varepsilon$. Since the nodes move in the same direction we have $|f_1^\delta(s_\delta(v_1)) - f_2^\delta(s_\delta(v_2))| \leq \varepsilon$.

On the other hand, we know that $y = \phi^\delta(s(v_1)) \in P^{\varepsilon+\delta}(v_2)$, and y is at the function distance at most ε from $s_\delta(v_2)$. Hence, y has to be in $P^\varepsilon(s_\delta(v_2))$, since this path neighborhood is the upper part of $P^{\varepsilon+\delta}(v_2)$. See [Figure B3](#) for an example. $\square$

Proof of Theorem 2 All conditions of R_c^L being a metric are trivial to check except for the triangle inequality, for which we provide a proof. Let $R_1, R_2, R_3 \in \mathcal{R}_c^L$ be Reeb graphs such that $d_I^L(R_1, R_2) = \varepsilon_1$ and $d_I^L(R_2, R_3) = \varepsilon_2$. We need to prove that $d_I^L(R_1, R_3) \leq \varepsilon_1 + \varepsilon_2$. For this purpose, it is enough to show that there is a labeled $\varepsilon_1 + \varepsilon_2$ -interleaving between R_1 and R_3 . Let ϕ_1, ψ_1 define a labeled ε -interleaving between R_1 and R_2 , and ϕ_2, ψ_2 be analogous for R_2 and R_3 . We need to define $\phi : R_1 \rightarrow R_3^{\varepsilon_1+\varepsilon_2}$, and $\psi : R_3 \rightarrow R_1^{\varepsilon_1+\varepsilon_2}$. We take $\phi = \phi_2^{\varepsilon_1} \circ \phi_1$ and $\psi = \psi_1^{\varepsilon_2} \circ \psi_2$ to be the compositions. As in [\[21\]](#), Proposition 4.5, this pair of morphisms define an $\varepsilon_1 + \varepsilon_2$ -interleaving. Here we show that these morphisms satisfy the additional labeling requirement, namely,

$$\phi^{\varepsilon_1+\varepsilon_2}(s(v_1)) \in P^{\varepsilon_1+\varepsilon_2}(s(v_3))$$

when v_1 and v_3 have the same label, and analogously for ψ . We focus on ϕ .

Let $v_1 \in R_1, v_2 \in R_2$ and $v_3 \in R_3$ be nodes with the same label. First, note that

$$\phi^{\varepsilon_1+\varepsilon_2} = (\phi_2^{\varepsilon_1} \phi_1)^{\varepsilon_1+\varepsilon_2} = \phi_2^{2\varepsilon_1+\varepsilon_2} \circ \phi_1^{\varepsilon_1+\varepsilon_2}.$$

Second, using [Lemma 5](#) with $\delta = \varepsilon_2$,

$$\phi_1^{\varepsilon_1}(s_{\varepsilon_1}(v_1)) \in P^{\varepsilon_1}(s_{\varepsilon_1}(v_2))$$




Fig. B3 An illustration for the proof of Lemma 5.

implies

$$y = \phi_1^{\varepsilon_1 + \varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_1)) \in P^{\varepsilon_1}(s_{\varepsilon_1 + \varepsilon_2}(v_2)).$$

Similarly, $\phi_2^{\varepsilon_2}(s_{\varepsilon_2}(v_2)) \in P^{\varepsilon_2}(s_{\varepsilon_2}(v_3))$ implies

$$\phi_2^{\varepsilon_1 + \varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_2)) \in P^{\varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_3)).$$

By "thickening" the two sides of the above we obtain,

$$\phi_2^{2\varepsilon_1 + \varepsilon_2}(P^{\varepsilon_1}(s_{\varepsilon_1 + \varepsilon_2}(v_2))) \subset P^{\varepsilon_1 + \varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_3)).$$

Putting these together,

$$\phi^{\varepsilon_1 + \varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_1)) = \phi_2^{2\varepsilon_1 + \varepsilon_2}(y) \in \phi_2^{2\varepsilon_1 + \varepsilon_2}(P_1^{\varepsilon}(s_{\varepsilon_1 + \varepsilon_2}(v_2))) \subset P^{\varepsilon_1 + \varepsilon_2}(s_{\varepsilon_1 + \varepsilon_2}(v_3)),$$

which is what we wanted to prove. $\square$

References

- [1] Saikia, H., Seidel, H.P., Weinkauff, T.: Extended branch decomposition graphs: Structural comparison of scalar data. *Computer Graphics Forum (CGF)* **33**(3), 41–50 (2014). <https://doi.org/10.1111/cgf.12360>
- [2] Saikia, H., Seidel, H.-P., Weinkauff, T.: Fast similarity search in scalar fields using merging histograms. In: Carr, H., Garth, C., Weinkauff, T. (eds.) *Topological Methods in Data Analysis and Visualization IV*, pp. 121–134. Springer, Cham (2017)
- [3] Sridharamurthy, R., Masood, T.B., Kamakshidasan, A., Natarajan, V.: Edit distance between merge trees. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* **26**(3), 1518–1531 (2020). <https://doi.org/10.1109/TVCG.2018.2873612>

- [4] Thomas, D.M., Natarajan, V.: Symmetry in scalar field topology. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* **17**(12), 2035–2044 (2011). <https://doi.org/10.1109/TVCG.2011.236>
- [5] Barra, V., Biasotti, S.: 3D shape retrieval using kernels on extended Reeb graphs. *Pattern Recognition* **46**(11), 2985–2999 (2013). <https://doi.org/10.1016/j.patcog.2013.03.019>
- [6] Biasotti, S., Marini, S., Mortara, M., Patane, G., Spagnuolo, M., Falcidieno, B.: 3D shape matching through topological structures. *International Conference on Discrete Geometry for Computer Imagery*, 194–203 (2003). https://doi.org/10.1007/978-3-540-39966-7_18
- [7] Hilaga, M., Shinagawa, Y., Kohmura, T., Kunii, T.L.: Topology matching for fully automatic similarity estimation of 3D shapes. In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 2001)*, pp. 203–212 (2001). <https://doi.org/10.1145/383259.383282>
- [8] Zhang, X., Bajaj, C.L., Baker, N.: Fast matching of volumetric functions using multi-resolution dual contour trees. Technical report, Texas Institute for Computational and Applied Mathematics, Austin, Texas (2004)
- [9] Edelsbrunner, H., Harer, J., Mascarenhas, A., Pascucci, V.: Time-varying Reeb graphs for continuous space-time data. *Computational Geometry*, 149–166 (2008). <https://doi.org/10.1016/j.comgeo.2007.11.001>
- [10] Saikia, H., Weinkauff, T.: Global feature tracking and similarity estimation in time-dependent scalar fields. *Computer Graphics Forum (CGF)* **36**(3), 1–11 (2017). <https://doi.org/10.1111/cgf.13163>
- [11] Sohn, B.-S., Bajaj, C.: Time-varying contour topology. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* **12**(1), 14–25 (2006). <https://doi.org/10.1109/TVCG.2006.16>
- [12] Li, M., Palande, S., Yan, L., Wang, B.: Sketching merge trees for scientific data visualization (2021) [arXiv:2101.03196](https://arxiv.org/abs/2101.03196) [cs.CG]
- [13] Lohfink, A.-P., Wetzels, F., Lukasczyk, J., Weber, G.H., Garth, C.: Fuzzy contour trees: alignment and joint layout of multiple contour trees. *Computer Graphics Forum (CGF)* **39**(3), 343–355 (2020). <https://doi.org/10.1111/cgf.13985>
- [14] Yan, L., Wang, Y., Munch, E., Gasparovic, E., Wang, B.: A structural average of labeled merge trees for uncertainty visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* **26**(1), 832–842

- (2020). <https://doi.org/10.1109/TVCG.2019.2934242>
- [15] Günther, D., Salmon, J., Tierny, J.: Mandatory critical points of 2D uncertain scalar fields. *Computer Graphics Forum (CGF)* **33**(3), 31–40 (2014). <https://doi.org/10.1111/cgf.12359>
 - [16] Wu, K., Zhang, S.: A contour tree based visualization for exploring data with uncertainty. *International Journal for Uncertainty Quantification* **3**(3), 203–223 (2013). <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2012003956>
 - [17] Poco, J., Doraiswamy, H., Talbert, M., Morisette, J., Silva, C.T.: Using maximum topology matching to explore differences in species distribution models. In: 2015 IEEE Scientific Visualization Conference (SciVis), pp. 9–16 (2015). <https://doi.org/10.1109/SciVis.2015.7429486>
 - [18] Yan, L., Masood, T.B., Sridharamurthy, R., Rasheed, F., Natarajan, V., Hotz, I., Wang, B.: Scalar field comparison with topological descriptors: properties and applications for scientific visualization. *Computer Graphics Forum (CGF)* **40**(3), 599–633 (2021). <https://doi.org/10.1111/cgf.14331>
 - [19] Bubenik, P., de Silva, V., Scott, J.: Interleaving and gromov-hausdorff distance (2018) [arXiv:1707.06288](https://arxiv.org/abs/1707.06288) [math.CT]
 - [20] Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J., Oudot, S.Y.: Proximity of persistence modules and their diagrams. In: 25th Annual Symposium on Computational Geometry (SoCG 2009), pp. 237–246. Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1542362.1542407>
 - [21] De Silva, V., Munch, E., Patel, A.: Categorified Reeb graphs. *Discrete & Computational Geometry* **55**(4), 854–906 (2016). <https://doi.org/10.1007/s00454-016-9763-9>
 - [22] Morozov, D., Beketayev, K., Weber, G.: Interleaving distance between merge trees. *Proceedings of Topology-Based Methods in Visualization (TopoInVis)* (2013)
 - [23] Munch, E., Stefanou, A.: The ℓ^∞ -cophenetic metric for phylogenetic trees as an interleaving distance. *Association for Women in Mathematics Series*, vol. 17, pp. 109–127. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11566-1_5
 - [24] Chambers, E.W., Munch, E., Ophelders, T.: A family of metrics from the truncated smoothing of Reeb graphs. 37th International Symposium on Computational Geometry (SoCG 2021) **189**, 22–12217 (2021). <https://doi.org/10.4230/LIPIcs.SoCG.2021.22>

- [25] Bauer, U., Ge, X., Wang, Y.: Measuring distance between Reeb graphs. In: 30th International Symposium on Computational Geometry (SoCG 2014), pp. 464–474 (2014). <https://doi.org/10.1145/2582112.2582169>
- [26] Bauer, U., Munch, E., Wang, Y.: Strong equivalence of the interleaving and functional distortion metrics for Reeb graphs. In: 31st International Symposium on Computational Geometry (SoCG 2015), vol. 34, pp. 461–475 (2015). <https://doi.org/10.4230/LIPIcs.SOCG.2015.461>
- [27] Bauer, U., Bjerkevik, H.B., Fluhr, B.: Quasi-universality of Reeb graph distances. In: 38th International Symposium on Computational Geometry (SoCG 2022). Leibniz International Proceedings in Informatics (LIPIcs), vol. 224, pp. 14–11418. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2022). <https://doi.org/10.4230/LIPIcs.SoCG.2022.14>
- [28] Di Fabio, B., Landi, C.: The edit distance for Reeb graphs of surfaces. *Discrete & Computational Geometry* **55**(2), 423–461 (2016). <https://doi.org/10.1007/s00454-016-9758-6>
- [29] Bauer, U., Di Fabio, B., Landi, C.: An edit distance for Reeb graphs. In: Ferreira, A., Giachetti, A., Giorgi, D. (eds.) *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, Eindhoven, The Netherlands (2016). <https://doi.org/10.2312/3dor.20161084>
- [30] Bauer, U., Landi, C., Memoli, F.: The Reeb graph edit distance is universal. *Foundations of Computational Mathematics* **21**(5), 1441–1464 (2020). <https://doi.org/10.1007/s10208-020-09488-3>
- [31] Carrière, M., Oudot, S.: Local equivalence and intrinsic metrics between Reeb graphs. In: 33rd International Symposium on Computational Geometry (SoCG 2017), vol. 77, pp. 25–12515. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017). <https://doi.org/10.4230/LIPIcs.SoCG.2017.25>
- [32] Touli, E.F., Wang, Y.: FPT-algorithms for computing Gromov-Hausdorff and interleaving distances between trees. *Proceedings of the 27th Annual European Symposium on Algorithms*, 83–18314 (2019). <https://doi.org/10.4230/LIPIcs.ESA.2019.83>
- [33] Beketayev, K., Yeliussizov, D., Morozov, D., Weber, G., Hamann, B.: Measuring the distance between merge trees. *Topological Methods in Data Analysis and Visualization III*, 151–165 (2014). https://doi.org/10.1007/978-3-319-04099-8_10
- [34] Cardona, G., Mir, A., Rosselló, F., Rotger, L., Sánchez, D.: Cophenetic metrics for phylogenetic trees, after Sokal and Rohlf. *BMC Bioinformatics*

- 14**(1) (2013). <https://doi.org/10.1186/1471-2105-14-3>
- [35] Edelsbrunner, H., Harer, J.: Computational Topology: an Introduction. American Mathematical Society, Providence, Rhode Island (2010)
 - [36] Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics* **9**(1), 79–103 (2009). <https://doi.org/10.1007/s10208-008-9027-z>
 - [37] Bakke Bjerkevik, H.: On the stability of interval decomposable persistence modules. *Discrete & Computational Geometry* **66**(1), 92–121 (2021). <https://doi.org/10.1007/s00454-021-00298-0>
 - [38] Botnan, M., Lesnick, M.: Algebraic stability of zigzag persistence modules. *Algebraic & Geometric Topology* **18**(6), 3133–3204 (2018). <https://doi.org/10.2140/agt.2018.18.313>
 - [39] Bollen, B., Chambers, E., Levine, J.A., Munch, E.: Reeb graph metrics from the ground up (2022) [arXiv:2110.05631](https://arxiv.org/abs/2110.05631) [cs.CG]
 - [40] Bjerkevik, H., Botnan, M.: Computational complexity of the interleaving distance. In: Speckmann, B., Tóth, C.D. (eds.) 34th International Symposium on Computational Geometry (SoCG 2018), pp. 13–11315. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2018). <https://doi.org/10.4230/LIPIcs.SocG.2018.13>
 - [41] Bjerkevik, H., Botnan, M., Kerber, M.: Computing the interleaving distance is NP-hard. *Foundations of Computational Mathematics* **20**(5), 1237–1271 (2020). <https://doi.org/10.1007/s10208-019-09442-y>
 - [42] Agarwal, P.K., Fox, K., Nath, A., Sidiropoulos, A., Wang, Y.: Computing the Gromov-Hausdorff distance for metric trees. *ACM Transactions on Algorithms* **14**(2) (2018). <https://doi.org/10.1145/3185466>
 - [43] Gasparovic, E., Munch, E., Oudot, S., Turner, K., Wang, B., Wang, Y.: Intrinsic interleaving distance for merge trees (2022) [arXiv:1908.00063](https://arxiv.org/abs/1908.00063) [cs.CG]
 - [44] Yan, L., Masood, T.B., Rasheed, F., Hotz, I., Wang, B.: Geometry-aware merge tree comparisons for time-varying data with interleaving distances. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2022). <https://doi.org/10.1109/TVCG.2022.3163349>
 - [45] Bauer, U., Botnan, M.B., Fluhr, B.: Universal distances for extended persistence (2022) [arXiv:2007.01834](https://arxiv.org/abs/2007.01834) [math.AT]

- [46] Dey, T.K., Wang, Y.: Reeb graphs: approximation and persistence. *Discrete & Computational Geometry* **49**(1), 46–73 (2013). <https://doi.org/10.1007/s00454-012-9463-z>
- [47] Alharbi, R., Chambers, E.W., Munch, E.: Realizable piecewise linear paths of persistence diagrams with reeb graphs (2021) [arXiv:2107.04654](https://arxiv.org/abs/2107.04654) [cs.CG]
- [48] Burago, D., Burago, Y., Ivanov, S.: *A Course in Metric Geometry* vol. 33. American Mathematical Society, Providence, Rhode Island (2001)
- [49] Parsa, S.: A deterministic $O(m \log m)$ time algorithm for the Reeb graph. *Discrete & Computational Geometry* **49**(4), 864–878 (2013). <https://doi.org/10.1007/s00454-013-9511-3>
- [50] Wang, S., Wang, Y., Wenger, R.: The JS-graphs of join and split trees. *Proceedings of the 30th Annual Symposium on Computational Geometry*, 539–548 (2014). <https://doi.org/10.1145/2582112.2582162>