Intrinsic Interleaving Distance for Merge Trees

Ellen Gasparovic[*], Elizabeth Munch[†], Steve Oudot[‡],
Katharine Turner[§], Bei Wang[¶], Yusu Wang[||]

Abstract

Merge trees are a type of graph-based topological summary that tracks the evolution of connected components in the sublevel sets of scalar functions. They enjoy widespread applications in data analysis and scientific visualization. In this paper, we consider the problem of comparing two merge trees via the notion of interleaving distance in the metric space setting. We investigate various theoretical properties of such a metric. In particular, we show that the interleaving distance is intrinsic on the space of labeled merge trees and provide an algorithm to construct metric 1-centers for collections of labeled merge trees. We further prove that the intrinsic property of the interleaving distance also holds for the space of unlabeled merge trees. Our results are a first step toward performing statistics on graph-based topological summaries.

1 Introduction

Many applications in science and engineering use scalar functions to describe and model their data. For example, atmospheric scientists compare the Weather Research and Forecasting model with daily surface observations in weather forecasts, where both simulated and observed parameters (such as surface temperature, pressure, precipitation, and wind speed) can be modeled as scalar functions. We are interested in comparing scalar functions by comparing their topological summaries. There are several types of summaries constructed from topological methods, including vector-based summaries such as persistence diagrams [29] and barcodes [33], as well as graph-based summaries such as merge trees, contour trees [13], and Reeb graphs [44].

The merge tree (sometimes referred to as a barrier tree [31] or a join tree [13]) for a given space $\mathbb{X}$ equipped with a scalar function is a construction that tracks the evolution of sublevel sets. For a given function $f : \mathbb{X} \to \mathbb{R}$, the merge tree encodes the connected components of the sublevel sets $f^{-1}(-\infty, a]$ for $a \in \mathbb{R}$. This construction is closely related to that of the Reeb graph [44], which analogously encodes connected components of the level sets $f^{-1}(a)$. The contour tree [13] is a special type of Reeb graph for a simply connected domain. Both merge trees and Reeb graphs are related to the level-set topology through critical points of the scalar functions [39]. Furthermore, the mapper graph [48], which has found considerable success in applications, can be viewed as an approximation of a Reeb graph [16, 42, 14]. These constructions are referred to as graph-based summaries as the output object of study is always a graph G equipped with an induced real-valued

[*]Union College, gasparoe@union.edu
[†]Michigan State University, muncheli@msu.edu
[‡]Inria Saclay, steve.oudot@inria.fr
[§]Australian National University, katharine.turner@anu.edu.au
[¶]University of Utah, beiwang@sci.utah.edu
[||]Ohio State University, yusu@cse.ohio-state.edu

function $f : G \to \mathbb{R}$. They have appeared in many contexts and applications over the last few decades [53, 43, 54].

Considerable recent effort has gone into understanding how to perform proper statistics on graph-based summaries. For instance, how does one define the mean of a collection of these objects? The first step toward answering this question is to determine a metric for the comparison of two summaries. This has been extensively studied recently with the creation of a veritable zoo of metric options for Reeb graphs and merge trees [40, 23, 2, 4, 1, 25, 3, 50, 15, 5]; see Section 2.3 for a discussion of some of these metrics. In particular, Carriére and Oudot [15] have investigated whether some of these metrics are *intrinsic* in the more general case of Reeb graphs; i.e., that the distance between two (close enough) graphs can be realized by a geodesic.

In this paper, we continue the investigation into the intrinsic-ness of these metrics with the more narrow view of merge trees. One of the main distances we study is the interleaving distance. This metric was originally given in the context of persistence modules [17, 18] as a generalization of the bottleneck distance, and has been ported to merge trees [40, 51] and Reeb graphs [23, 20] via a category-theoretic viewpoint [9, 24]. When restricting ourselves from Reeb graphs to merge trees, we can actually work in an even more restrictive setting that has desirable theoretical properties, namely, labeled merge trees. In this case, we study a data triple: a merge tree T with its function $f : T \to \mathbb{R}$, and a labeling $\pi : \{1, \cdots, n\} \to V(T)$ of its vertices, which at a minimum encompasses the leaves of T. The interleaving distance for labeled merge trees has been investigated in [41], where it is shown that the metric can be naturally realized as the L_∞-distance for a particular matrix construction. This construction has already been discovered in the context of dendrograms [49] and phylogenetic trees [12], where the objects of interest are closely related to merge trees. The phylogenetic tree literature, in particular, provides a wealth of other options for metrics [45, 46, 22, 21, 26, 6, 8, 7, 30, 11, 19]. There has also been interest in that community for creating summaries of collections of phylogenetic trees [38, 37, 32].

These ideas are also closely related to those of ultrametrics, a strengthening of the triangle inequality for a metric into a requirement that $d(x, y) \leq \max\{d(x, z), d(z, y)\}$. Independent of the phylogenetic tree work, there has been extensive interest in what is known as Gelfand's Problem from the ultrametric literature, that is, to describe all finite ultrametric spaces up to isometry using graph theory. The answer to this question is exactly a restriction of the labeled merge tree, although their literature never calls it such [34, 35, 36, 28, 27].

Our contributions. In Section 2, we provide the necessary background on labeled merge trees and establish a correspondence between labeled merge trees and a particular class of matrices known as ultra matrices (Lemma 2.9). We also discuss several metrics on labeled and unlabeled merge trees, how to obtain their intrinsic counterparts, and the known relationships between these metrics. Then, in Section 3, we prove a stability result for the labeled interleaving distance d_I^L (Lemma 3.2), which we use both to show that d_I^L is strictly intrinsic on the space of labeled merge trees (Corollary 3.3) as well as to construct 1-centers for collections of labeled merge trees in Section 3.3.

Section 4 focuses on unlabeled merge trees and the interleaving distance. In particular, given two unlabeled merge trees, we show that the unlabeled interleaving distance between them is equal to the infimum over all finite labelings for the two trees of the labeled interleaving distance between them (Proposition 4.1). In fact, we show in Corollary 4.4 that the infimum is always achieved. Section 4 concludes with the result that the interleaving distance is intrinsic on the space of unlabeled merge trees (Theorem 4.5). We conclude with a discussion of open problems and future work in Section 5.

2 Background

In this section, we give the basic definitions for our constructions of interest. We refer to Section 1 for an overview of notation. For the entirety of the section, we fix n, and denote $\{1, \cdots, n\}$ by $[n]$.

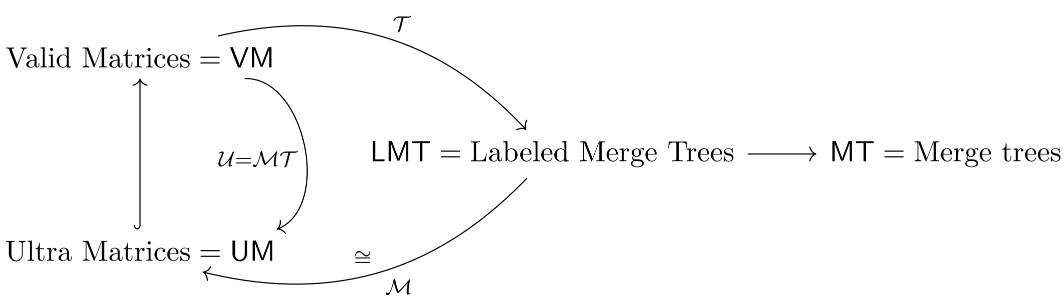

Figure 1: A roadmap of key notation.

2.1 Labeled Merge Trees

First, we give the definition of a *merge tree* and related notions arising from the phylogenetic tree literature that we will make use of shortly.

Definition 2.1. *A* ***merge tree*** *is a pair (T, f) of a finite rooted tree T with vertex set $V(T)$ and a function $f : V(T) \to \mathbb{R} \cup \infty$ such that adjacent vertices do not have equal function value, every non-root vertex has exactly one neighbor with higher function value, and the root (a degree one node) is the only point with the value ∞. The space of merge trees is denoted* MT*.*

We commonly call the function f a *height function*, the non-root vertices with degree 1 are called *leaves*, and we let *depth*(u) denote the largest height difference between the vertex u in T and any node in the subtree rooted at u. All merge trees under consideration in this paper are assumed to be *finite*. MT also denote the space of unlabeled merge trees.

In some sense, replacing a merge tree edge $e = (u, v)$ with $f(u) < f(v)$ by a subdivision of that edge where the interior vertex w satisfies $f(u) < f(w) < f(v)$ does not change the inherent structure of the tree (sometimes such a tree is referred to as an *argumented merge tree*). We consider two merge trees to be the same if one can be obtained from the other by a sequence of such subdivisions or the inverse operation.

Furthermore, the merge tree structure induces a poset relation on the vertices of T. We say v is an *ancestor* of w and write $v \succ w$ if the unique path from v to w strictly decreases in f. This occurs if and only if w is in the subtree of v. We use $\mathrm{LCA}(v, w) \in T$ to mean the lowest common ancestor of v and w (or $\mathrm{LCA}_f(v, w)$ if the function needs to be emphasized), and $f(\mathrm{LCA}(v, w))$ for its function value. We abuse notation and write $\mathrm{LCA}(S)$ for the lowest common ancestor of any finite set $S \subset V(T)$.

The merge tree structure provides a method for inducing a metric on the underlying tree vertices via the metric given by the length of the unique path between two points. Note that there is a canonical weighting associated to any merge tree (T, f), namely, $\omega(u, v) = |f(u) - f(v)|$ for any two adjacent vertices u and v in the tree. Further, as paths are unique in a tree, we can define a metric for any pair of vertices by $\delta_T(u, v) = \sum \omega(e)$ for the edges in the path from u to v.

In Section 3, we will be focusing on *labeled merge trees*, defined as follows.

Definition 2.2. *A* ***labeled merge tree*** *is a triple* (T, f, π) *consisting of a merge tree* (T, f) *along with a map* $\pi : [n] \to V(T)$ *that is surjective on the set of leaves. When additional data are unnecessary or clear from context, we sometimes write* T *for* (T, f, π)*. The space of labeled merge trees is denoted* LMT*.*

Analogously to the unlabeled case, we consider two labeled merge trees to be the same if one can be obtained via edge contractions or insertions that respect the function values and existing labels.

This definition is closely related to that of a *weighted, rooted* X*-tree* from the phylogenetic literature [47]. Specifically, given a set X, an X-tree is a pair (T, ϕ) where T is a tree and $\phi : X \to V(T)$ is a map so that every vertex of degree at most 2 is in the image. The difference is that such weighted graphs do not keep track of function values, so that two different labeled merge trees that induce the same weighting might be considered to be the same X-tree. Thus, a labeled merge tree can be thought of as a weighted, labeled X-tree (where $X = [n]$) with $f(u)$ specified for a subset of vertices u that includes all leaves, as function values for the remaining vertices can be deduced from the weights.

As with X-trees, the labels for our merge tree are allowed to go to vertices that are not leaves; we essentially think of these as degenerate labeled leaves. Further, we do allow for non-injectivity of π, so a vertex can have multiple labels. See Section 2 for an example with labels on degenerate leaves and vertices with more than one label.

2.2 Relating Merge Trees and Matrices

In this section, we give the relationship between labeled merge trees and a particular class of matrices. Again, see Section 1 for an overview of notation.

We begin with the traditional notion of an *ultrametric* and our variant of it that relaxes one of the conditions, which will be closely related to our labeled merge trees.

Definition 2.3. *An* ***ultrametric*** *is a function* $d : X \times X \to \mathbb{R}$ *such that for any* $x, y, z \in X$*,*

- $d(x, y) \geq 0$ *and is equal to 0 if and only if* $x = y$*,*
- $d(x, y) = d(y, x)$*, and*
- $d(x, y) \leq \max\{d(x, z), d(z, y)\}$*.*

Definition 2.4. *A* ***relaxed ultrametric*** *is a function* $d : X \times X \to \mathbb{R}$ *such that for any* $x, y, z \in X$*,*

- $d(x, y) = d(y, x)$*, and*
- $d(x, y) \leq \max\{d(x, z), d(z, y)\}$*.*

It is well known that ultrametrics satisfy the isosceles triangle property. That is, for any triple x, y, z, at least two of $d(x, y)$, $d(y, z)$, and $d(x, z)$ must be equal. Otherwise, assume without loss of generality that $d(x, y) < d(y, z) < d(x, z)$, and then $d(x, z) \not\leq \max\{d(x, y), d(y, z)\}$. Note that this further implies that the pair that are equal must be at least as big as the third value since $d(x, y) = d(y, z) < d(x, z)$ still violates the ultrametric property. Note that relaxed ultrametrics still satisfy the isosceles triangle property.

When we have a set $X \cong [n]$, the information in a relaxed ultrametric can be stored as follows.

Definition 2.5. *A symmetric matrix* $M \in \mathbb{R}^{n \times n}$ *is called* ***valid*** *if* $M_{ii} \leq M_{ij}$ *for all* $1 \leq i, j \leq n$*. A valid matrix* M *is called* ***ultra*** *if* $M_{ij} \leq \max\{M_{ik}, M_{kj}\}$*. The spaces of valid and ultra matrices are denoted* VM *and* UM*, respectively.*

In particular, a relaxed ultrametric on $[n]$ is represented by an ultra matrix. Inspired by the cophenetic matrix construction of [12] that is studied in relation to merge trees in [41], there is a natural way to associate a matrix to a labeled merge tree as follows.

Definition 2.6. *The* ***induced matrix of the labeled merge tree*** (T, f, π)*, denoted* $\mathcal{M}(T, f, \pi) \in \mathbb{R}^{n\times n}$*, is the matrix*

$$\mathcal{M}(T, f, \pi)_{ij} = f(\mathrm{LCA}(\pi(i), \pi(j))).$$

See Section 2 for an example.

Lemma 2.7. *The induced matrix of a labeled merge tree is an ultra matrix. That is,* $\mathcal{M}(T, f, \pi) \in \mathsf{UM}$ *for* $(T, f, \pi) \in \mathsf{LMT}$*.*

Proof. Let $M = \mathcal{M}(T, f, \pi)$ for $(T, f, \pi) \in \mathsf{LMT}$. First, to check that it is a valid matrix, we see that M_{ii} is simply the function value $f(\pi(i))$. So, as $f(u) \leq f(LCA(u, v))$ by definition, we have

$$M_{ii} = f(\pi(i)) \leq f(\mathrm{LCA}(\pi(i), \pi(j))) = M_{ij}.$$

To check that it is an ultra matrix, let $u = \mathrm{LCA}(\pi(i), \pi(k))$, $v = \mathrm{LCA}(\pi(j), \pi(k))$, and $w = \mathrm{LCA}(\pi(i), \pi(j), \pi(k))$. This means that $u \preceq w$ and $v \preceq w$. If u and v are not comparable, then there are two distinct paths from $\pi(j)$ to each of them, and thus we have a loop $\pi(j) \to u \to w \to v \to \pi(j)$, contradicting the tree property of T. If u and v are comparable, assume without loss of generality that $u \preceq v$; then v is a common ancestor for $\pi(i)$, $\pi(j)$, and $\pi(k)$, and thus $w \preceq v$. This implies $f(w) \leq f(v)$, and so

$$M_{ij} \leq f(w) \leq f(v) = \max\{f(u), f(v)\} = \max\{M_{ik}, M_{jk}\}.$$

□

A valid matrix may be viewed as representing a function on the complete graph of n vertices, with function value M_{ii} on vertex i and function value M_{ij} on edge ij. Note that because M is a valid matrix, any sublevel set of the resulting function is a simplicial complex as every edge has equal or higher function value than either of its vertices. Given a valid matrix, one may obtain a labeled merge tree and subsequently an ultra matrix in the following way.

Definition 2.8. *The* ***labeled merge tree of a valid matrix*** $M \in \mathbb{R}^{n\times n}$*, denoted* $\mathcal{T}(M)$*, is the labeled merge tree (in the sense of* $a \mapsto \pi_0(f^{-1}(-\infty, a])$*) of the complete graph with the induced function.*

Note that the labeling is inherited by including internally labeled vertices if there is any pair $i \neq j$ for which $M_{ii} = M_{ij}$. See Section 2 for a labeled tree containing an example of when $M_{ii} = M_{ij} = M_{jj}$ creating a leaf with two labels, as well as an example where $M_{ii} = M_{ij} > M_{jj}$ creating an internal labeled vertex.

Lemma 2.9. $\mathcal{M}$ *induces a 1-1 correspondence between labeled merge trees and ultra matrices.*

Proof. We start with injectivity of $\mathcal{M}$. From [47], a metric δ is called a tree metric if there exists a weighted $[n]$-tree (i.e., a weighted X-tree with $X = [n]$) (T, f, π, ω) for which $\delta(i, j) = \sum_{e\in\gamma} \omega(e)$ for γ the unique path from $\pi(i)$ to $\pi(j)$ if $\pi(i) \neq \pi(j)$, and is 0 otherwise. By [47, Thm. 7.1.8], such a weighted $[n]$-tree representation is unique. For any $(T, f, \pi) \in \mathsf{LMT}$, we can construct a tree metric $\delta_T : [n] \times [n] \to \mathbb{R}_{\geq 0}$ uniquely from $\mathcal{M}(T, f, \pi)$ by setting

$$\delta_T(i, j) = 2\mathcal{M}(T)_{ij} - \mathcal{M}(T)_{ii} - \mathcal{M}(T)_{jj};$$

$$\begin{pmatrix} a_1 & a_1 & a_4 & a_4 \\ \cdot & a_1 & a_4 & a_4 \\ \cdot & \cdot & a_3 & a_3 \\ \cdot & \cdot & \cdot & a_2 \end{pmatrix}$$

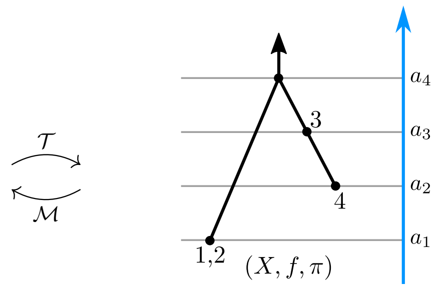

Figure 2: An example of a labeled merge tree with two types of degenerate labels. As all matrices used are symmetric, we only show the upper triangular portion.

this is the length of the path from $\pi(i)$ to $\pi(j)$. So, given any $(T, f, \pi), (T', f', \pi') \in \mathsf{LMT}$, we construct the two tree metrics δ_T and δ'_T. However, these two tree metrics are equivalent as $\delta'_T(x, y) \leq \delta_T(x, y) \leq 2\delta'_T(x, y)$. This implies that any continuity condition is the same under either choice of metric.

For ease of notation, denote $\mathcal{M}(T, f, \pi)$ by $\mathcal{M}(T)$; similarly for $\mathcal{M}(T')$. If $\mathcal{M}(T) = \mathcal{M}(T')$, then $(T, f, \pi, \omega) = (T', f', \pi', \omega')$ as weighted $[n]$-trees by keeping the weighting but forgetting the function values and which vertex is the root. Since the function value of any labeled vertex can be determined by $\mathcal{M}(T)_{ii}$, this implies that $T = T'$ as labeled merge trees.

Next, we tackle surjectivity of $\mathcal{M}$. Given any ultra matrix M, we want a labeled merge tree T for which $\mathcal{M}(T) = M$. In particular, we will show that $T = \mathcal{T}(M)$ satisfies this requirement, which further gives that $\mathcal{T}$ is the inverse of $\mathcal{M}$. To construct $\mathcal{T}(M)$, let K be the complete graph on n vertices with vertices labeled $v_1, \cdots, v_n$. Define the map $s : K \to \mathbb{R}$ on the simplicial complex K by $s(v_i) = M_{ii}$ (vertex map) and $s(v_i, v_j) = M_{ij}$ (edge map). Because M is a valid matrix, this gives a well-defined map; in particular, $s(v_i) \leq s(v_i, v_j)$ for any $i \neq j$.

First, we check the diagonal entries of the matrix. By definition of the construction of $\mathcal{T}(M)$, there is a vertex $\pi(i)$ in the resulting tree with function value $f(\pi(i)) = s(v_i) = M_{ii}$, so clearly $\mathcal{M}(\mathcal{T}(M))_{ii} = f(\mathrm{LCA}(\pi(i), \pi(i))) = f(\pi(i)) = M_{ii}$.

Finally, we check the non-diagonal entries, so assume $i \neq j$ and consider M_{ij}. Note that $\mathcal{M}(\mathcal{T}(M))_{i,j} = f(\mathrm{LCA}(\pi(i), \pi(j)))$ is exactly the function value for which the components containing v_i and v_j merge in the sublevel set persistence of $s : K \to \mathbb{R}$. (See the end of Section 2.3 for a very brief discussion of sublevel set persistence.) Because $s(v_i, v_j) = M_{ij}$, this means that $\mathcal{M}(\mathcal{T}(M))_{i,j} \leq M_{ij}$. Seeking a contradiction, assume that $\mathcal{M}(\mathcal{T}(M))_{i,j} < M_{ij}$. In order for the components with v_i and v_j to have merged before M_{ij}, there must be a path $\gamma = v_i u_1 u_2 \cdots u_k v_j$ for which every internal edge e has $s(e) < M_{ij}$. By the isosceles property using the triangle $v_i v_j u_1$, we know that $s(v_i, v_j) = M_{ij}$ and $s(v_i, u_1) < M_{ij}$, so $s(v_j, u_1) = M_{ij}$. The same logic for triangle $u_1 u_2 v_j$ implies that $s(v_j, u_2) = M_{ij}$. Repeating this process for the entire path, we conclude finally that $s(v_j, u_{k-1}) = M_{ij}$. However, then the triangle $v_j u_{k-1} u_k$ has both $s(u_{k-1}, u_k)$ and $s(u_k, v_j)$ strictly less than $s(u_{k-1}, v_j)$, contradicting the isosceles triangle property. Thus, we conclude that no such path exists, and therefore $\mathcal{M}(\mathcal{T}(M))_{ij} = M_{ij}$. □

In the course of the above proof, we have showed that $\mathcal{M}\mathcal{T}$ is the identity when restricted to ultra matrices, but this is not the case when extending to only valid matrices. However, this construction does offer a method for turning a valid matrix into an ultra matrix.

Definition 2.10. *The* ***ultra matrix of a valid matrix*** $M \in \mathsf{VM}$*, denoted* $\mathcal{U}(M)$*, is defined to be the induced matrix of* $\mathcal{T}(M)$*. That is,* $\mathcal{U} = \mathcal{M}\mathcal{T}$*.*

Basically, given a valid matrix M, we can consider M to induce weights of a complete graph on n vertices. We then compute a minimal spanning tree $\mathcal{T}(M)$ of this complete graph based on the weights. The resulting tree $\mathcal{T}(M)$ gives rise to an induced relaxed ultra matrix $\mathcal{M}\mathcal{T}(M)$.

2.3 Available Metrics

There are a number of metrics that may be defined on the space of (labeled) merge trees. Note that any metric defined on unlabeled merge trees can be extended to labeled merge trees by simply forgetting the labeling information (while likely turning the metric into a pseudometric). We now introduce several of these metrics. In this paper, we focus on interleaving distance d_I and labeled interleaving distance d_I^L. Other popular distances include functional distortion distance d_{FD} and bottleneck distance d_B. We note that $d_I = d_F$ in the case of merge trees (Theorem 6.2 of the arXiv version of [2]); while they are strongly equivalent for general Reeb graphs [4].

Interleaving distance. The interleaving distance is an idea arising from the generalization of the bottleneck distance for persistence diagrams to arbitrary persistence modules [17]. Generalizations abound [9, 42, 24], but the analog for merge trees was first given in [40]. We give a modified definition here, which was shown to be equivalent to the original in [51].

Definition 2.11. *Given two merge trees $(T, f), (T', f')$, a δ-**good** map $\alpha : (T, f) \to (T', f')$ is a continuous map on the metric trees such that the following properties hold:*

(i) For any $x \in |T|$, $f'(\alpha(x)) - f(x) = \delta$;

(ii) For any $w \in \mathrm{Im}(\alpha)$ with $x' := LCA(\alpha^{-1}(w))$, $f(x') - f(u) \leq 2\delta$ for all $u \in \alpha^{-1}(w)$; and

(iii) For any $w \notin \mathrm{Im}(\alpha)$, $depth(w) \leq 2\delta$.

*The **interleaving distance** is then defined to be*

$$d_I((T, f), (T', f')) = \inf\{\delta \mid \exists\, \delta\text{-good } \alpha : (T, f) \to (T', f')\}.$$

One particularly useful property that we will use later is the following.

Lemma 2.12. *Let $\alpha : (T, f) \to (T', f')$ be a continuous map such that $f'(\alpha(x)) = f(x) + \delta$ for any $x \in |T|$. Assume $u \preceq v$. Then*

- *$\alpha(u) \preceq \alpha(v)$, and*
- *if w is the unique ancestor of $\alpha(u)$ with $f'(w) = f(v) + \delta$, then $w = \alpha(v)$.*

Proof. Note that $u \preceq v$ implies that $f(u) \leq f(v)$ and further that the unique path γ from u to v in T is monotone increasing in f. Then the image of γ in T', $\alpha(\gamma)$, satisfies $f'(\alpha(\gamma(t))) = f(\gamma(t)) + \delta$ and thus is monotone increasing in f'. Thus, by definition, we have that $\alpha(u) \preceq \alpha(v)$. Further, the uniqueness of paths implies that if w is the unique ancestor with $f'(w) = f(v) + \delta$, then it must be the endpoint of γ, and so $w = \alpha(v)$. □

Labeled interleaving distance. The following related metric is closely related to one originally defined in [12] for comparing phylogenetic trees.

Definition 2.13. *Given two labeled merge trees, the **labeled interleaving distance** is*

$$d_I^L((T, f, \pi), (T', f', \pi')) = \|\mathcal{M}(T, f, \pi) - \mathcal{M}(T', f', \pi')\|_\infty.$$

The reason for calling such a distance an interleaving distance comes from [41] where it is shown that this metric arises as an interleaving distance on a particular category with a flow [24]. Note that because we need the labels in order to be able to have a well-defined matrix, this metric only works on labeled merge trees.

Functional distortion distance. The next distance, introduced in [2], takes inspiration from the Gromov–Hausdorff distance. Anytime we have a correspondence between points in two different metric spaces X and Y, we can ask what is the maximum difference between the distance between a pair of points in X and the distance between a corresponding pair in Y. This is called the distortion of the correspondence, and the Gromov-Hausdorff distance between X and Y is the infimum of all correspondence distortions. A natural generalization is to consider the distortion of functions that may not be metrics, and a natural class of correspondences are those that arise from a pair of continuous maps.

Definition 2.14. *Given two merge trees (T, f) and (T', f'), let $\phi : T \to T'$ and $\psi : T' \to T$ be continuous maps. If*

$$G(\phi, \psi) = \{(x, \phi(x)) : x \in T\} \cup \{(\psi(y), y)) : y \in T'\}$$

and

$$D(\phi, \psi) = \sup_{(x,y),(\tilde{x},\tilde{y}) \in G(\phi,\psi)} \frac{1}{2}|d_f(x, \tilde{x}) - d_{f'}(y, \tilde{y})|,$$

then the ***functional distortion distance*** *d_{FD} is defined to be*

$$d_{FD}((T, f), (T', f')) = \inf_{\phi,\psi} \max\{D(\phi, \psi), ||f - f' \circ \phi||_\infty, ||f \circ \psi - f'||_\infty\}.$$

Bottleneck distance. In order to define the next distance, we require a bit of background. In *persistent homology*, colloquially known as *persistence*, one is interested in the changing homology of an increasing sequence of subspaces associated to a topological space X. Given a smooth function $f : X \to \mathbb{R}$, one may define a *sublevel set filtration* $f^{-1}(-\infty, a_1) \subseteq f^{-1}(-\infty, a_2) \subseteq \ldots \subseteq f^{-1}(-\infty, \infty)$, and these inclusions induce a *persistence module* in any homological dimension k by applying the kth homology functor with coefficients in a field. A homology class is said to be *born* in $H_k(f^{-1}(-\infty, a_i))$ if its class does not exist in any previous $H_k(f^{-1}(-\infty, a_i - \epsilon))$, and it is said to *die* in $H_k(f^{-1}(-\infty, a_j))$ if its homology class merges with another class in some $H_k(f^{-1}(-\infty, a_k))$ with $a_k < a_i$. In this way, elements of the kth homology group may be tracked through the filtration and recorded in the kth *persistence diagram*, which is a multiset of points in the extended plane lying above the diagonal $\Delta = \{(x, x) : x \in \mathbb{R}\}$ (since a class dies after it is born). Given persistence diagrams $D = \mathrm{Dg}(f)$ and $D' = \mathrm{Dg}(f')$, a *partial matching* (or *partial correspondence*) between the two diagrams is a subset $\Gamma \subset D \times D'$, where to every $p \in D$ there is associated at most one $p' \in D'$ so that $(p, p') \in \Gamma$, and the analog holds for every $p' \in D'$. The *bottleneck distance* between D and D' is given by

$$d_B(D, D') = \inf_\Gamma \max\left\{\max_{p \in D} \delta_D(p), \max_{p' \in D'} \delta_{D'}(p')\right\},$$

where Γ ranges over all partial matchings between D and D', and where $\delta_D(p) = ||p - p'||_\infty$ if p is matched to some $p' \in D'$ and $\delta_D(p) = d_\infty(p, \Delta)$ if p is unmatched (similarly for $\delta_{D'}(p')$).

Definition 2.15. *The* ***bottleneck distance*** *between two merge trees (T, f) and (T', f') is*

$$d_B((T, f), (T', f')) = d_B(\mathrm{Dg}(f), \mathrm{Dg}(f')).$$

2.4 Intrinsic Metrics

The *induced intrinsic metric* on a metric space is the infimum of the lengths of all paths from one point to another. A metric space is said to be a *length space* if the original metric d coincides with the intrinsic metric $\hat{d}$. It is known that d is always less than or equal to $\hat{d}$. Given a metric d on merge trees, we may define its intrinsic version as follows.

Definition 2.16. *Given two merge trees, let $\gamma : [0,1] \to \mathsf{MT}$ be a continuous path in d such that $\gamma(0) = T$ and $\gamma(1) = T'$. The* ***length of γ induced by the distance*** *d is defined as*

$$L_d(\gamma) = \sup_{n,\sum} \sum_{i=0}^{n-1} d(\gamma(t_i), \gamma(t_{i+1})),$$

where n ranges over $\mathbb{N}$ and $\sum$ ranges over all partitions $0 = t_0 \leq t_1 \leq \ldots \leq t_n = 1$ of $[0,1]$. The ***intrinsic metric $\hat{d}$ induced by the distance*** *d is*

$$\hat{d}(T, T') = \inf_{\gamma} L_d(\gamma).$$

Recall that a metric space is said to be 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. In this case, the metric is said to be *strictly intrinsic*. Note that a geodesic space is necessarily a length space. The following theorem provides a sufficient characterization for a complete metric space to be a geodesic space.

Theorem 2.17 ([10, Thm. 2.1.16])**.** *Let X be a complete metric space. If for any two points a and b from X there exists a midpoint between a and b, then X is a geodesic metric space.*

3 Geodesics and 1-Centers for Labeled Merge Trees

In this section, we prove two results involving the labeled interleaving distance and provide methods for constructing geodesics and 1-centers for collections of labeled merge trees.

3.1 More on the Labeled Interleaving Distance

We begin by pointing out a completeness result for LMT (Section 3.1) and proving a stability result for d_I^L (Section 3.2). As a consequence, we show in the next subsection in Corollary 3.3 that the metric d_I^L is strictly intrinsic on LMT.

Lemma 3.1. *The space of n-labeled merge trees, LMT, equipped with the labeled interleaving distance, d_I^L, is complete.*

Proof. By Section 2.9, $\mathsf{UM} \cong \mathsf{LMT}$, and the space of ultra matrices, UM, is a closed subspace of $(\mathbb{R}^{n\times n}, d_\infty)$, which is complete. □

Lemma 3.2. *For any pair of valid matrices $M, M' \in \mathsf{VM}$,*

$$d_I^L(\mathcal{T}(M), \mathcal{T}(M')) \leq \|M - M'\|_\infty.$$

Proof. Since, by definition, $d_I^L(\mathcal{T}(M), \mathcal{T}(M')) = \|\mathcal{U}(M) - \mathcal{U}(M')\|_\infty$, we will actually establish the inequality $\|\mathcal{U}(M) - \mathcal{U}(M')\|_\infty \leq \|M - M'\|_\infty$.

Let $\delta = \|M - M'\|_\infty$. Set $T = \mathcal{M}(M)$ and $T' = \mathcal{M}(M')$ to be the associated merge trees, and $\widetilde{M} = \mathcal{U}(M)$ and $\widetilde{M}' = \mathcal{U}(M')$ the induced ultra matrices. Consider any pair of (possibly equal) labels i and j with $1 \leq i \leq j \leq n$. We consider the vertices v_i and v_j in the complete graph K with $s, s' : K \to \mathbb{R}$ the maps on K induced by M and M', respectively. As v_i and v_j are in the same component of the M_{ij}-sublevel set of s, there is a path γ in K with $s(e) \leq \widetilde{M_{ij}}$ for all edges e in the path. Because $\|M - M'\|_\infty \leq \delta$, we have that

$$s'(e) \leq s(e) + \delta \leq \widetilde{M_{ij}} + \delta$$

for every $e \in \gamma$. So, v_i and v_j are in the same component of the $(\widetilde{M_{ij}} + \delta)$-sublevel set of s' and thus $\widetilde{M'}_{ij} \leq \widetilde{M_{ij}} + \delta$.

Symmetrically, for any $t < \widetilde{M_{ij}} - \delta$, v_i and v_j do not lie in the same connected component of the t-sublevel set of s'. Otherwise, by the same argument as above, v_i and v_j would belong to the same connected component of the $(t + \delta)$-sublevel set of T with $t + \delta < \widetilde{M_{ij}}$, a contradiction. Hence, $\widetilde{M'_{ij}} \geq \widetilde{M_{ij}} - \delta$. It follows that $|\widetilde{M'_{ij}} - \widetilde{M_{ij}}| \leq \delta$, and since this is true for all labels $1 \leq i \leq j \leq n$, the symmetric matrices $\widetilde{M}, \widetilde{M'}$ satisfy $\|\widetilde{M} - \widetilde{M'}\|_\infty \leq \delta$. Hence, $d_I^L(T, T') = \|\widetilde{M} - \widetilde{M'}\|_\infty \leq \delta$. □

3.2 Geodesics in LMT

The next corollary looks at the straight line between the matrices associated to two labeled merge trees. Specifically, given any two labeled merge trees $T, T' \in \mathsf{LMT}$, we know that their associated matrices $M = \mathcal{M}(T), M' = \mathcal{M}(T')$ are ultra matrices. We can define the line between them by setting $M^\lambda := (1 - \lambda)M + \lambda M'$ for $\lambda \in [0, 1]$. While not necessarily ultra matrices, it is easy to check that $M^\lambda \in \mathsf{VM}$ for all $\lambda \in [0, 1]$. We can then pull this back to a path of labeled merge trees by setting $T^\lambda = \mathcal{T}(M^\lambda)$.

Corollary 3.3 (LMT Geodesics)**.** *Given any two labeled merge trees $T, T' \in \mathsf{LMT}$, and their corresponding ultra matrices $M = \mathcal{M}(T), M' = \mathcal{M}(T')$, the family of merge trees $\left\{T^\lambda := \mathcal{T}\left(M^\lambda\right)\right\}_{\lambda \in [0,1]}$ defines a geodesic between T and T' in the metric d_I^L. As a consequence, on the space of labeled merge trees, the metric d_I^L is strictly intrinsic.*

Proof. Let δ denote the distance $d_I^L(T, T') = \|M - M'\|_\infty$. For any $0 \leq \lambda \leq \lambda' \leq 1$, the linearly interpolating matrices $M^\lambda, M^{\lambda'}$ satisfy $\|M^\lambda - M^{\lambda'}\|_\infty \leq (\lambda' - \lambda)\,\delta$. Hence, by Section 3.2, we have $d_I^L(T^\lambda, T^{\lambda'}) \leq (\lambda' - \lambda)\,\delta$. Since this is true for all $0 \leq \lambda \leq \lambda' \leq 1$, the triangle inequality implies that the family $\{T^\lambda\}_{\lambda \in [0,1]}$ defines a geodesic between T and T'. □

See the example of Section 3. Setting $\lambda = 1/2$, M^λ is the matrix (labeled M) shown in the middle green circle, and T^λ (labeled $\mathcal{T}(M)$) is the tree shown at the far right.

3.3 1-centers in LMT

A metric 1-center of a finite set of labeled merge trees is one that minimizes the maximum distance to any other tree in the set.

Definition 3.4. *Given a metric space (X, d), the 1-**center** $c \in X$ of a finite point set $P = \{p_1, \cdots, p_m\} \subset X$ is*

$$c = \arg\min_{x \in X} \max_{p \in P} d(x, p).$$

That is, c is the center of the minimum enclosing ball of P.

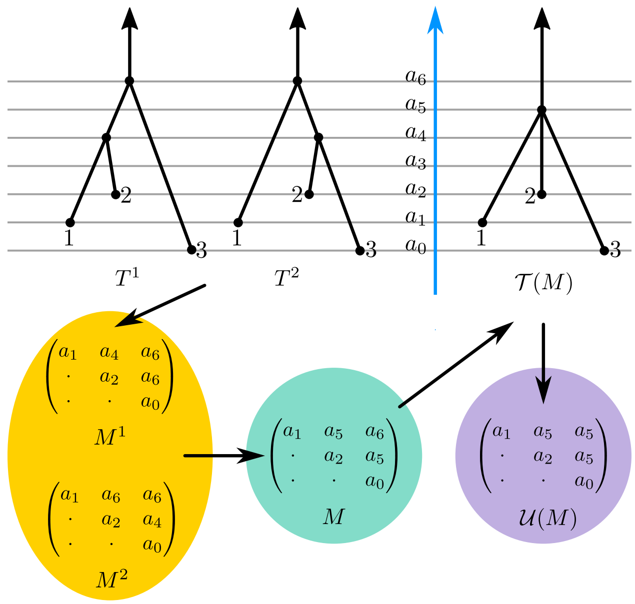

Figure 3: An example of the averaging process for labeled merge trees. T^1 and T^2 are labeled merge trees with induced matrices M^1 and M^2. M is the pointwise average of M^1 and M^2, but is not an ultra matrix. The labeled merge tree $\mathcal{T}(M)$ is shown, whose induced matrix is the ultra matrix $\mathcal{U}(M)$.

In the case of a finite collection of numbers χ in $\mathbb{R}$, the 1-center is simply the midpoint of the enclosing interval, $(\max(\chi) + \min(\chi))/2$. For a collection of matrices $M^1, \cdots, M^N$ with L_∞-distance, an option for a 1-center in the space of all matrices is the entry-wise 1-center of the matrices.

We study the following procedure in order to obtain a 1-center in the space of valid matrices for a collection of labeled merge trees under the labeled interleaving distance d_I^L. First, map the collection of labeled merge trees to their corresponding induced ultra matrices, and then take a 1-center in the space of valid matrices. This coincides with the entrywise defined 1-center in the space of all matrices mentioned previously. This 1-center is a valid matrix since $M_{ii}^k \le M_{ij}^k$ for all k implies that

$$(\max_k(M_{ii}^k) + \min_k(M_{ii}^k))/2 \le (\max_k(M_{ij}^k) + \min_k(M_{ij}^k))/2.$$

This 1-center (while being valid) may not be an ultra matrix, so we can replace it by its labeled merge tree and take its corresponding ultra matrix. See Section 3 for an example of the process. The following proposition establishes the fact that this construction yields a 1-center.

Proposition 3.5 (LMT 1-Center)**.** *Let $M^1, \cdots, M^N$ be a set of ultra matrices, and M their 1-center in the space* VM *of valid matrices. Let $\mathcal{T}(M)$ denote the labeled merge tree of M, and $U = \mathcal{M}(\mathcal{T}(M))$ its corresponding ultra matrix. Then $\mathcal{T}(M)$ is a 1-center for the labeled merge trees $\{\mathcal{T}(M^i)\}_{i=1}^N$; while U is a 1-center for the set of ultra matrices $\{M^1, \cdots, M^N\}$.*

Proof. For ease of notation, set $T = \mathcal{T}(M)$ and $T^i = \mathcal{T}(M^i)$. Let $\delta = \max_i \|M - M^i\|_\infty$. Then $d_I^L(T, T^i) \le \|M - M^i\|_\infty$ by Section 3.2. It then follows that

$$\max_i d_I^L(T, T^i) \le \max_i \|M - M^i\|_\infty \le \delta.$$

Thus $\{T_i\}_{i=1}^N$ is contained in a ball of radius δ centered at T.

To show that this is in fact a *minimum enclosing ball*, assume there exists a $\widetilde{T}$ such that $\max_i d_I^L(\widetilde{T}, T^i) < \delta$. Set $\widetilde{M} = \mathcal{M}(\widetilde{T})$. Then for any i,

$$\|\widetilde{M} - M^i\|_\infty = d_I^L(\mathcal{T}(\widetilde{M}), \mathcal{T}(M^i)) = d_I^L(\widetilde{T}, T^k) < \delta.$$

This contradicts the assumption that M is a 1-center, and thus T is a 1-center for $\{T^1, \cdots, T^N\}$. By the relation between distance for ultra matrices and for their corresponding labeled merge trees, U is a 1-center for $\{M^1, \cdots, M^N\}$, as well. □

4 Interleaving Distances for Unlabeled Merge Trees

Moving to the unlabeled setting, we establish the existence of a certain labeling for a pair of merge trees that allows us to show that the interleaving distance for unlabeled merge trees is intrinsic.

Proposition 4.1. *Given two merge trees (T, f) and (T', f'), let L and L' be the respective leaf sets. Then*

$$d_I((T, f), (T', f')) = \inf_{\pi, \pi'} d_I^L((T, f, \pi), (T', f', \pi')) \tag{1}$$

where the infimum is taken over all finite labelings of the two given merge trees, π and π', using at most $|L| + |L'|$ labels.

Prior to proving the proposition, we will investigate the following construction of a labeling when given a δ-good map. First, note that given two labeled merge trees (T, f, π) and (T', f', π'), where $\pi : [n] \to V(T)$ and $\pi' : [n] \to V(T')$, the labeling information can be equivalently stored as an ordered collection of pairs $\Pi = \{(\pi(i), \pi'(i)) \mid i \in [n]\} \subseteq V(T) \times V(T')$. Since the order of the labels does not matter for this particular application, we will build Π iteratively and assign the integers at the end.

Let L and L' denote the leaf sets for T and T', respectively. Assume we are given a δ-good map α as described in Section 2.11. While this map is defined on the underlying metric trees, note that we can subdivide the trees so that $\alpha(v)$ is a vertex in T' for any vertex in T, and further that every point in the set $\alpha^{-1}(w)$ is a vertex in T if w is a vertex in T'.

Then, we construct the labeling Π as follows.

(S-1) Fix some $v \in L$, and let $w = \alpha(v)$. Then for every $u \in \alpha^{-1}(w)$, add (u, w) to Π. Repeat this for every vertex in L.

(S-2) For any leaf node $w \in L' \setminus \mathrm{Im}(\alpha)$, let x be its lowest ancestor contained in $\mathrm{Im}(\alpha)$. Let $u \in \alpha^{-1}(x)$ be an arbitrary preimage of x from $|T|$. Add (u, w) to Π. Repeat for all leaves in L'.

(S-3) Fix an ordering on the pairs in $\Pi = \{(u_i, w_i) \mid i \in [n]\}$ and define $\pi(i) = u_i \in T$ and $\pi'(i) = w_i \in T'$.

Observe that since the preimage of any leaf node $w \in L' \cap \mathrm{Im}(\alpha)$ must be some vertex (or vertices) in L, any $w \in L' \cap \mathrm{Im}(\alpha)$ will be paired with some $u \in L$ by the process in (S-1), so this procedure does not miss any leaves in T'. See Section 4 for an example.

To use this construction to prove Section 4.1, we will use the following two lemmas.

Lemma 4.2. *For any $(u, w) \in \Pi$, $|f(u) - f'(w)| \leq \delta$.*

$$\mathcal{M}(T,f,\pi) = \begin{pmatrix} a_1 & a_4 & a_7 & a_7 & a_1 & a_4 & a_7 \\ \cdot & a_2 & a_7 & a_7 & a_4 & a_2 & a_7 \\ \cdot & \cdot & a_4 & a_5 & a_7 & a_7 & a_5 \\ \cdot & \cdot & \cdot & a_0 & a_7 & a_7 & a_4 \\ \cdot & \cdot & \cdot & \cdot & a_1 & a_4 & a_7 \\ \cdot & \cdot & \cdot & \cdot & \cdot & a_2 & a_7 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & a_4 \end{pmatrix}$$

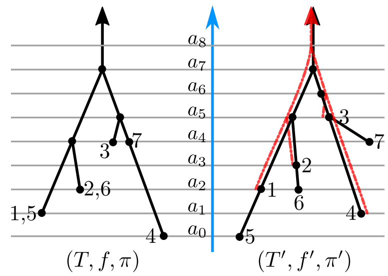

$$\mathcal{M}(T',f',\pi') = \begin{pmatrix} a_2 & a_5 & a_7 & a_7 & a_2 & a_5 & a_7 \\ \cdot & a_3 & a_7 & a_7 & a_5 & a_3 & a_7 \\ \cdot & \cdot & a_5 & a_5 & a_7 & a_7 & a_5 \\ \cdot & \cdot & \cdot & a_1 & a_7 & a_7 & a_5 \\ \cdot & \cdot & \cdot & \cdot & a_0 & a_5 & a_7 \\ \cdot & \cdot & \cdot & \cdot & \cdot & a_2 & a_7 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & a_4 \end{pmatrix}$$

Figure 4: Given $\alpha : (T, f) \to (T', f')$, this is an example of the labeling induced by the procedure discussed after Section 4.1. The image of the map α is given by the red dashed lines, and α is δ-good for $\delta = a_{i+1} - a_i$. Labels 1-4 were generated in (S-1), the rest in (S-2). Note that there were two options for the location of label 7 in T. The other choice would be the same as the vertex labeled 3, and would only change the red entries in $\mathcal{M}(T, f, \pi)$.

Proof. If (u, w) is generated from (S-1) above, then the lemma holds by property (i) in the definition of the δ-good map α (see Definition 2.11). If (u, w) is generated from (S-2), then the lemma follows from property (iii) of the δ-good map α. Indeed, let x be the lowest ancestor of w contained in $\mathrm{Im}(\alpha)$, so that $\alpha(u) = x$. Then $0 \le f'(x) - f'(w) \le 2\delta$ and $f'(x) - f(u) = \delta$, implying that $|f'(w) - f(u)| \le \delta$. ☐

Lemma 4.3. *For any* $(u_1, w_1), (u_2, w_2) \in \Pi$, $|f(LCA(u_1, u_2)) - f'(LCA(w_1, w_2))| \le \delta$.

Proof. Assume we are given α, a δ-good map. If (u_i, w_i) is generated from (S-1), set $w_i' = w_i$. If (u_i, w_i) is generated via (S-2), then let w_i' be the lowest ancestor of w_i in $\mathrm{Im}(\alpha)$. In both cases, we have that $\alpha(u_i) = w_i'$ and $w_i \preceq w_i'$.

Set $u_0 = LCA(u_1, u_2)$, $w_0 = LCA(w_1, w_2)$ and $w_0' = LCA(w_1', w_2')$. We will first show that $w_0 = w_0'$. If both pairs come from (S-1), then $w_i = w_i'$ and the claim is obvious. So, assume that at least one, say (u_1, w_1), comes from (S-2) and thus $w_1 \ne w_1'$. As $w_i \preceq w_i' \preceq w_0'$ for each i, the least common ancestor property implies $w_0 \preceq w_0'$. Seeking a contradiction, assume that w_0 is not a common ancestor of both w_i'; without loss of generality, say w_0 is not an ancestor of w_1'. Let $z = \mathrm{LCA}(w_0, w_1', w_2')$. Then there are two paths in T' from w_1 to z: one through w_0 and one through w_1'. This contradicts the tree assumption of T'. Therefore, w_0 is a common ancestor of w_i', implying $w_0' \preceq w_0$, and so $w_0 = w_0'$.

We will now prove the main claim, namely, that $|f(u_0) - f'(w_0)| \le \delta$. To see that this is the case, assume that the claim does not hold; that is, either $f(u_0) - f'(w_0) > \delta$ or $f'(w_0) - f(u_0) > \delta$. Suppose first that $f'(w_0) - f(u_0) > \delta$, and consider $\alpha(u_0)$. Because $u_i \preceq u_0$ for $i = 1, 2$, by Section 2.12 we must have that $w_i' = \alpha(u_i) \preceq \alpha(u_0)$ for $i = 1, 2$. However, then $\alpha(u_0)$ is an ancestor of both w_1' and w_2' with

$$f'(\alpha(u_0)) = f(u_0) + \delta < f'(w_0),$$

contradicting the least common ancestor assumption of w_0.

Next, suppose $f(u_0) - f'(w_0) > \delta$ and consider $\alpha^{-1}(w_0)$. We claim that any point in $\alpha^{-1}(w_0)$ is a descendant of u_0; i.e., $v \preceq u_0$ for all $v \in \alpha^{-1}(w_0)$. Otherwise, we have that

$$f(LCA(\alpha^{-1}(w_0))) > f(u_0) > f'(w_0) + \delta = f(v) + 2\delta$$

for any $v \in \alpha^{-1}(w_0)$, contradicting property (ii) of Section 2.11. For $i = 1, 2$, let v_i be the unique ancestor of u_i with $f(v_i) = f'(w_0) - \delta$. By Section 2.12, since $\alpha(u_i) = w_i'$ and w_0 is the unique ancestor of w_i' with $f'(w_0) = f(v_i) + \delta$, this implies that $\alpha(v_i) = w_0$. That is, $v_i \in \alpha^{-1}(w_0)$. Further, $v_1 \neq v_2$. Otherwise if $v := v_1 = v_2$, then

$$f(v) = f'(w_0) - \delta < f(u_0) - 2\delta < f(u_0)$$

and thus v is a lower common ancestor of u_1 and u_2 than u_0, a contradiction. Hence, $\mathrm{LCA}(v_1, v_2) = u_0$. However,

$$f(u_0) - f(v_i) = f(u_0) - f'(w_0) + \delta > 2\delta.$$

This also contradicts property (ii) of Section 2.11, finishing the proof of Section 4.3. □

Proof of Section 4.1. Say we have a δ-good map α for some $\delta \geq d_I((T,f),(T',f'))$. We construct the labelings π, π' as described above. Then Lemmas 4.2 and 4.3 imply that

$$d_I^L((T,f,\pi),(T',f',\pi')) \leq \delta.$$

As this is true for any δ, $\inf_{\Pi} d_I^L((T,f,\pi),(T',f',\pi')) \leq d_I((T,f),(T',f'))$.

To show the other inequality, assume we are given any pair of labelings π, π' and assume

$$d_I^L((T,f,\pi),(T',f',\pi')) = \delta.$$

We will construct the map α and show that it is δ-good. For any $x \in |T|$, let $S_x \subseteq [n]$ be the labels in the subtree of x. Let y_i be the unique ancestor of $\pi'(i) \in |T'|$ for $i \in S_x$ with $f'(y_i) = f(x) + \delta$. First, we note that $y_i = y_j$ for all $i, j \in S_x$. Indeed, let $M = \mathcal{M}(T,f,\pi)$ and $M' = \mathcal{M}(T',f',\pi')$. Then we know $M'_{ij} \leq \delta + M_{ij}$ and so

$$f'(y_i) = f(x) + \delta \geq f(\mathrm{LCA}(\pi(S_x))) = M_{ij} + \delta \geq M'_{ij} = f'(\mathrm{LCA}(\pi'(S_x))).$$

Because every y_i has function value greater than the lowest common ancestor of $\pi'(S_x)$, the tree property implies that all y_i are equal. Thus, we can set $\alpha(x) = y_i$ for any $i \in S_x$ and it is well-defined.

We need to ensure that the α constructed is δ-good as given in Section 2.11. The map satisfies property (i) by construction, so we move on to (ii). Let $w \in |T'| \cap \mathrm{Im}(\alpha)$ and set $x' = \mathrm{LCA}(\alpha^{-1}(w)) \in |T|$. Fix any $u \in \alpha^{-1}(w)$, and clearly $f(u) \leq f(x')$. Now x' must be $\mathrm{LCA}(u, u')$ for some other $u' \in \alpha^{-1}(w)$. Let i be a label in the subtree of u, and let j be a label in the subtree of u'. This further implies that $x' = \mathrm{LCA}(\pi(i), \pi(j))$. Set $w' = \mathrm{LCA}(\pi'(i), \pi'(j))$ and note that as $\pi'(i) \preceq w$ and $\pi'(j) \preceq w$, this implies that $w' \preceq w$. In particular, this means $f'(w') \leq f'(w)$. Further, by assumption $|f(x') - f'(w')| = |M_{ij} - M'_{ij}| \leq \delta$. Thus,

$$f(x') - f(u) \leq (f'(w) - f(u)) + (f(x') - f'(w')) + (f'(w') - f'(w)) \leq 2\delta$$

as the first part of the middle term is exactly δ, the second is $\leq \delta$, and the last is negative, showing that α satisfies property (ii).

Finally, we ensure property (iii). Let $w \in |T'| \setminus \mathrm{Im}(\alpha)$. Let i be the label of any leaf in the subtree of w, and set $y = \alpha(\pi(i))$ to be the image of the vertex labeled i in T. Then the tree property implies that $\pi'(i) \preceq w \preceq y$ and thus $f'(\pi'(i)) \leq f'(w) \leq f'(y)$. So,

$$|f'(w) - f'(\pi'(i))| \leq |f'(\pi'(i)) - f'(y)| \leq |f'(\pi'(i)) - f(\pi(i))| + \delta = |M_{ii} - M'_{ii}| + \delta \leq 2\delta.$$

As this is true for every leaf in the subtree of w, $\mathrm{depth}(w) \leq 2\delta$ and so α satisfies property (iii).

Thus, we have that $d_I((T,f),(T',f')) \leq d_I^L((T,f,\pi),(T',f',\pi'))$ for any given Π, completing the proof of the proposition. □

We can use the construction from the proof to state something stronger. Recall that we work with finite labeled and unlabeled merge tree throughout the paper.

Corollary 4.4. *There exist an n and a pair of labelings π, π' so that*

$$d_I((T,f),(T',f')) = d_I^L((T,f,\pi),(T',f',\pi')).$$

Thus, the interleaving distance for finite merge trees is always achieved by a map α.

Proof. The right side of equation (1) is taken over labelings using at most $N = |L| + |L'|$ labels, which is finite. Up to reordering, we can use the first $|L|$ numbers to label the leaves in T and the last $|L'|$ to label the leaves in T'. All that remains to show is that there are finitely many possible locations to place the remaining labels in each tree. Indeed, if $d_I(T,T') = \delta$, then for each $i \in \{1, \cdots, |L|\}$, one has the option of placing i at any point in $(f')^{-1}(f(\pi(i)) + \delta) \subset T'$. Note that $|(f')^{-1}(f(\pi(i)) + \delta)|$ is finite. Similarly, there are $|f^{-1}(f'(\pi'(i)) + \delta)|$ possible locations available for $i \in \{|L|+1, N\}$ to be placed in T. For any fixed choice from this set for every i, let M and M' be the associated matrices for T and T', respectively.

The options are set up so that any choice of location for label i in the opposite tree will automatically satisfy $|M_{ii} - M'_{ii}| = \delta$, so we need only ensure that some choice in each tree of these locations for every i promises $|M_{ij} - M'_{ij}| \leq \delta$. For every choice of remaining labels, say there is some i, j for which $|M_{ij} - M'_{ij}| > \delta$. As we have finitely many options, there is an ϵ so that $|M_{ij} - M'_{ij}| > \delta + \epsilon$. However, there is certainly a $(\delta + \epsilon/2)$-good map α that does not take the labels into consideration, and we could then build the labeling as discussed in Section 4.1, giving a contradiction. Thus, one of finitely many options achieves the left infimum of (1), and thus there is a δ-good map α that also achieves the unlabeled distance. □

We conclude this section by showing that the interleaving distance is intrinsic on the space of finite (unlabeled) merge trees.

Theorem 4.5. *For the space of finite (unlabeled) merge trees, $d_I = \hat{d}_I$.*

Proof. Let T and T' be two merge trees, and set $\delta = d_I((T,f),(T',f'))$. Let π, π' be optimal labelings such that $d_I((T,f),(T',f')) = d_I^L((T,f,\pi),(T',f',\pi')) = \delta$, as established by Section 4.4.

Now consider the space of labeled merge trees LMT. By Section 3.3, there exists a geodesic $\gamma : (T,f,\pi) \rightsquigarrow (T',f',\pi')$ in LMT such that the length $L_{d_I^L}(\gamma) = \delta$.

Note that γ can be projected to a path γ' from T to T' in the space of (unlabeled) merge trees MT by simply forgetting the labeling. As $d_I((T,f),(T',f')) \leq d_I^L((T,f,\pi_1),(T',f',\pi_2))$ for any labelings π_1, π_2 between any two trees T and T', we have

$$\hat{d}_I(T,T') \leq L_{d_I}(\gamma') \leq L_{d_I^L}(\gamma) = \delta. \tag{2}$$

On the other hand, by definition of the intrinsic metric $\hat{d}_I$ induced by d_I,

$$\hat{d}_I(T,T') \geq d_I(T,T') = \delta. \tag{3}$$

Combining equations (2) and (3), we conclude that $\hat{d}_I(T,T') = d_I(T,T')$ for any two merge trees T and T'. □

5 Concluding Remarks and Discussion

In this paper, we investigated whether interleaving-type distances for (finite) labeled or unlabeled merge trees are intrinsic or not, and presented positive answers in both cases. In the case of labeled trees, the geodesic between two labeled merge trees can be characterized and computed easily, and we also showed how to compute the 1-center of a set of labeled merge trees. For unlabeled merge trees, however, computing the geodesic (even if just numerically estimating it) between two merge trees appears to be significantly harder, part of the reason being that it is NP-hard to approximate the interleaving distance between two merge trees.

On the other hand, a simpler and easier to compute object is the bottleneck distance $d_B(T_1, T_2)$ between two (unlabeled) merge trees. We conjecture that the intrinsic distance $\hat{d}_B$ induced by d_B is in fact equivalent to $\hat{d}_I (= d_I)$.

Another natural question is whether (some of the) results for merge trees in this paper can be extended to contour trees. As a first question, can we characterize and compute the midpoint (i.e., the contour tree representing the 1-center) for two labeled contour trees under either $\hat{d}_I$, $\hat{d}_B$, or $\hat{d}_{FD}$? One idea is to compute the join and split trees of input contour trees, and compute the midpoint of the pair of join trees (resp., the pair of split trees). Note that each join or split tree can be viewed as a merge tree. Next we need to use the common ancestor information in both trees to construct a midpoint for the two contour trees. This step could be subtle: in particular, it is known [52] that in general, given a descending (join) tree T_J and an ascending (split) tree T_S with consistent functions associated to them, there may not exist a contour tree (or even a graph) whose join and split trees are equal to T_J and T_S, respectively. If such a contour tree exists, then it is unique, and the algorithm by Carr et al. [13] will compute this tree in near linear time.

Acknowledgments. Our initial research collaboration began during the Dagstuhl Seminar 17292: Topology, Computation and Data Analysis in July 2017. We thank all members of the breakout session on Reeb graphs for stimulating discussions. We are grateful to the Institute for Computational and Experimental Research in Mathematics (ICERM) for supporting us through the Collaborate@ICERM program in August 2018. EM was partially supported by National Science Foundation (NSF) through grants CMMI-1800466 and DMS-1800446. YW was partially supported by NSF through grants CCF-1740761 and DMS-1547357, as well as by National Institute of Health (NIH) under grant R01EB022899. BW was partially supported by NSF IIS-1513616, NSF DBI-1661375 and NIH R01EB022876.

References

[1] U. Bauer, B. Di Fabio, and C. Landi. An edit distance for Reeb graphs. *Proceedings of the Eurographics Workshop on 3D Object Retrieval*, pages 27–34, 2016.

[2] U. Bauer, X. Ge, and Y. Wang. Measuring distance between Reeb graphs. *Symposium on Computational Geometry*, 2014. Full version available at arXiv:1307.2839.

[3] U. Bauer, C. Landi, and F. Memoli. The Reeb graph edit distance is universal. *Foundations of Computational Mathematics*, 2017.

[4] U. Bauer, E. Munch, and Y. Wang. Strong equivalence of the interleaving and functional distortion metrics for Reeb graphs. In L. Arge and J. Pach, editors, *31st International Symposium on Computational Geometry*, volume 34 of *Leibniz International Proceedings in Informatics*

(LIPIcs), pages 461–475, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[5] K. Beketayev, D. Yeliussizov, D. Morozov, G. H. Weber, and B. Hamann. Measuring the distance between merge trees. In *Mathematics and Visualization*, pages 151–165. Springer International Publishing, 2014.

[6] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4):733–767, 2001.

[7] D. Bogdanowicz and K. Giaro. Matching split distance for unrooted binary phylogenetic trees. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 9(1):150–160, 2012.

[8] D. Bogdanowicz and K. Giaro. On a matching distance between rooted phylogenetic trees. *International Journal of Applied Mathematics and Computer Science*, 23(3):669–684, 2013.

[9] P. Bubenik, V. de Silva, and J. Scott. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, 15(6):1501–1531, 2014.

[10] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry.* American Mathematical Society, 2001.

[11] G. Cardona, M. Llabrés, F. Rosselló, and G. Valiente. Nodal distances for rooted phylogenetic trees. *Journal of Mathematical Biology*, 61(2):253–276, 2009.

[12] G. Cardona, A. Mir, F. Rosselló, L. Rotger, and D. Sánchez. Cophenetic metrics for phylogenetic trees, after sokal and rohlf. *BMC Bioinformatics*, 14(1):3, 2013.

[13] H. Carr, J. Snoeyink, and U. Axen. Computing contour trees in all dimensions. *Computational Geometry*, 24(2):75–94, 2003.

[14] M. Carrière, B. Michel, and S. Oudot. Statistical analysis and parameter selection for mapper. *Journal of Machine Learning Research*, 19:1–39, 2018.

[15] M. Carrière and S. Oudot. Local equivalence and intrinsic metrics between Reeb graphs. In B. Aronov and M. J. Katz, editors, *33rd International Symposium on Computational Geometry*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:15, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[16] M. Carrière and S. Oudot. Structure and stability of the one-dimensional mapper. *Foundations of Computational Mathematics*, 18(6):1333–1396, 2018.

[17] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th Annual Symposium on Computational Geometry*, pages 237–246, New York, NY, USA, 2009. ACM.

[18] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. *The Structure and Stability of Persistence Modules.* Springer International Publishing, 2016.

[19] K. Choi and S. M. Gomez. Comparison of phylogenetic trees through alignment of embedded evolutionary distances. *BMC Bioinformatics*, 10(1):423, 2009.

[20] J. Curry. *Sheaves, Cosheaves and Applications.* PhD thesis, University of Pennsylvania, 2014.

[21] B. DasGupta, X. He, T. Jiang, M. Li, and J. Tromp. On the linear-cost subtree-transfer distance between phylogenetic trees. *Algorithmica*, 25(2-3):176–195, 1999.

[22] B. DasGupta, X. He, T. Jiang, M. Li, J. Tromp, and L. Zhang. On distances between phylogenetic trees. In *ACM-SIAM Symposium on Discrete Algorithms*, volume 97, pages 427–436, 1997.

[23] V. de Silva, E. Munch, and A. Patel. Categorified reeb graphs. *Discrete & Computational Geometry*, pages 1–53, 2016.

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

[25] B. Di Fabio and C. Landi. The edit distance for Reeb graphs of surfaces. *Discrete & Computational Geometry*, 55(2):423–461, 2016.

[26] P. W. Diaconis and S. P. Holmes. Matchings and phylogenetic trees. *Proceedings of the National Academy of Sciences*, 95(25):14600–14602, 1998.

[27] O. Dovgoshey and E. Petrov. From isomorphic rooted trees to isometric ultrametric spaces. *p-Adic Numbers, Ultrametric Analysis and Applications*, 10(4):287–298, 2018.

[28] O. Dovgoshey, E. Petrov, and H.-M. Teichert. How rigid the finite ultrametric spaces can be? *Journal of Fixed Point Theory and Applications*, 19(2):1083–1102, 2016.

[29] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28:511–533, 2002.

[30] G. F. Estabrook, F. McMorris, and C. A. Meacham. Comparison of undirected phylogenetic trees based on subtrees of four evolutionary units. *Systematic Biology*, 34(2):193–200, 1985.

[31] C. Flamm, I. L. Hofacker, P. F. Stadler, and M. T. Wolfinger. Barrier trees of degenerate landscapes. *Zeitschrift für Physikalische Chemie*, 216(2), 2002.

[32] A. Gavryushkin and A. J. Drummond. The space of ultrametric phylogenetic trees. *Journal of Theoretical Biology*, 403:197–208, 2016.

[33] R. Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.

[34] V. Gurvich and M. Vyalyi. Characterizing (quasi-)ultrametric finite spaces in terms of (directed) graphs. *Discrete Applied Mathematics*, 160(12):1742–1756, 2012.

[35] B. Hughes. Trees and ultrametric spaces: A categorical equivalence. *Advances in Mathematics*, 189(1):148–191, 2004.

[36] A. J. Lemin. The category of ultrametric spaces is isomorphic to the category of complete, atomic, tree-like, and real graduated lattices LAT^*. *Algebra Universalis*, 50(1):35–49, 2003.

[37] A. Markin and O. Eulenstein. Cophenetic median trees under the manhattan distance. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology,and Health Informatics*. ACM Press, 2017.

[38] E. Miller, M. Owen, and J. S. Provan. Polyhedral computational geometry for averaging metric phylogenetic trees. *Advances in Applied Mathematics*, 68:51–91, 2015.

[39] J. Milnor. *Morse Theory.* Princeton University Press, New Jersey, NY, USA, 1963.

[40] D. Morozov, K. Beketayev, and G. Weber. Interleaving distance between merge trees. *Proceedings of Topology-Based Methods in Visualization*, 2013.

[41] E. Munch and A. Stefanou. The ℓ^∞-cophenetic metric for phylogenetic trees as an interleaving distance. ArXiv:1803.07609, 2018.

[42] E. Munch and B. Wang. Convergence between categorical representations of Reeb space and Mapper. In S. Fekete and A. Lubiw, editors, *32nd International Symposium on Computational Geometry*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 53:1–53:16, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[43] P. Oesterling, C. Heine, H. Jaenicke, G. Scheuermann, and G. Heyer. Visualization of high-dimensional point clouds using their density distribution's topology. *IEEE Transactions on Visualization and Computer Graphics*, 17(11):1547–1559, 2011.

[44] G. Reeb. Sur les points singuliers d'une forme de pfaff complèment intégrable ou d'une fonction numérique. *Comptes Rendus de L'Académie ses Séances*, 222:847–849, 1946.

[45] D. F. Robinson and L. R. Foulds. Comparison of weighted labelled trees. In *Combinatorial mathematics VI*, pages 119–126. Springer, 1979.

[46] D. F. Robinson and L. R. Foulds. Comparison of phylogenetic trees. *Mathematical Biosciences*, 53(1):131–147, 1981.

[47] C. Semple, M. Steel, and R. A. Caplan. *Phylogenetics.* Oxford University Press, 2003.

[48] G. Singh, F. Mémoli, and G. Carlsson. Topological methods for the analysis of high dimensional data sets and 3D object recognition. In *Eurographics Symposium on Point-Based Graphics*, pages 91–100, 2007.

[49] R. R. Sokal and F. J. Rohlf. The comparison of dendrograms by objective methods. *Taxon*, 11(2):33, 1962.

[50] R. Sridharamurthy, T. B. Masood, A. Kamakshidasan, and V. Natarajan. Edit distance between merge trees. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2018.

[51] E. F. Touli and Y. Wang. FPT-algorithms for computing Gromov-Hausdorff and interleaving distances between trees. ArXiv:1811.02425, 2018.

[52] S. Wang, Y. Wang, and R. Wenger. The JS-graph of join and split trees. In *Proceedings of the 30th Annual Symposium on Computational Geometry*, pages 539–548, 2014.

[53] G. Weber, P.-T. Bremer, and V. Pascucci. Topological landscapes: A terrain metaphor for scientific data. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1416–1423, 2007.

[54] W. Widanagamaachchi, A. Jacques, B. Wang, E. Crosman, P.-T. Bremer, V. Pascucci, and J. Horel. Exploring the evolution of pressure-perturbations to understand atmospheric phenomena. *IEEE Pacific Visualization Symposium*, 2017.