An Interleaving Distance for Ordered Merge Trees

Thijs Beurskens[1], Tim Ophelders[1, 2], Bettina Speckmann[1], and Kevin Verbeek[1]

1 **Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands**
`{t.p.j.beurskens|k.a.b.verbeek|b.speckmann}@tue.nl`

2 **Department of Information and Computing Science, Utrecht University, The Netherlands**
`t.a.e.ophelders@uu.nl`

Abstract

Merge trees are a common topological descriptor for data with a hierarchical component. The interleaving distance, in turn, is a common distance measure for comparing merge trees. In this abstract, we introduce a form of ordered merge trees and extend the interleaving distance to a measure that preserves orders. Exploiting the additional structure of ordered merge trees, we then describe an $\mathcal{O}(n^2)$ time algorithm that computes a 2-approximation of this new distance with an additive term G that captures the maximum height differences of leaves of the input merge trees.

Related Version A full version of the paper is available at arxiv.org/abs/2312.11113.

1 Introduction

Merge trees are a common topological descriptor for data with a hierarchical component, such as terrains and scalar fields. However, standard merge trees focus solely on the hierarchy and do not represent other salient geometric features of the data. Specifically, our work is motivated by the study of braided rivers (see Figure 1). A braided river is a multi-channel river system, known to evolve rapidly [13, 17]. There exist methods to generate a *river network* from a snapshot of the terrain [7, 14, 16]. We model a river network as a hierarchy of *braids*. We use a merge tree to represent this hierarchy: each leaf represents a single channel in the network, and each internal vertex represents two braids merging (see Figure 1).

It is our goal to analyse the evolution of the channel network over time. The standard way to compare two merge trees is the interleaving distance [19]. However, the interleaving distance has two main drawbacks. Firstly, the standard interleaving distance is unable to capture any intrinsic order, e.g. from bank to bank in braided rivers, that might be present in

Figure 1 (Left) the Waimakariri River in New-Zealand is a braided river. Photo was taken by Greg O'Beirne [21]. (Right) representing a river network by a merge tree.

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

the data. Secondly, there is no known efficient algorithm to compute even an approximation of the interleaving distance.[1] To tackle both issues, we introduce the *monotone interleaving distance*: an order-preserving distance measure on ordered merge trees.

Contributions. We show that the monotone interleaving distance can be defined in terms of a single map between the input merge trees, two maps, or a labelling. Moreover, we give an algorithm that computes an approximation of this distance in $\mathcal{O}(n^2)$ time. Finding an efficient algorithm to compute the monotone interleaving distance exactly, or to prove NP-hardness of computing the distance, remains an open problem. We first review the relevant background in Section 2. In Section 3 we introduce a form of ordered merge trees and define monotone interleavings, monotone δ-good maps and monotone labellings. We give constructions to prove that all of these lead to the same distance. Finally, in Section 4 we describe an efficient algorithm to approximate the monotone interleaving distance. All omitted proofs can be found in the full version on arXiv.

Related work. The interleaving distance was first introduced as a measure for persistence modules [8]. It has since been well-studied from a categorical point of view [3, 4, 5, 6, 9, 10, 11, 15, 22], and has been transferred to numerous topological descriptors [2, 18, 20]. Morozov et al. [19] defined the interleaving distance for merge trees. Agarwal et al. [1] established a relation between the interleaving distance and the Gromov-Hausdorff distance.[1] The interleaving distance on merge trees was redefined by first Touli and Wang [23], and later Gasparovich et al. [12]. Touli and Wang also gave an FPT-algorithm to compute the interleaving distance. Recently, the result by Gasparovich et al. has been used to design algorithms for computing geometry aware labellings [24, 25].

2 Preliminaries

A *merge tree* is a pair (T, f), where T is a rooted tree and $f\colon T \to \mathbb{R} \cup \{\infty\}$ is a continuous height function that is increasing towards the root, with $f(v) = \infty$ if and only if v is the root. Here, f is defined not only on the vertices of T, but also on points of T interior to the edges. Specifically, f is linearly interpolated along the edges. For a point $x \in T$, we denote by T_x the subtree of T rooted at x. Furthermore, for a given value $\delta \geq 0$, we denote by x^δ the unique ancestor of x with $f(x^\delta) = f(x) + \delta \in T$.

Now consider two merge trees (T, f) and (T', f') and fix a value $\delta \geq 0$. Intuitively, a δ-*interleaving* describes a mapping α from T to T' that sends points exactly δ upwards, and a similar map β from T' to T, such that both compositions of α and β send any point to its unique ancestor 2δ higher. Figure 2 shows an example of a δ-interleaving.

▶ **Definition 1** (Morozov et al. [19]). Given two merge trees (T, f) and (T', f'), a pair of maps $\alpha\colon T \to T'$ and $\beta\colon T' \to T$ is called a δ-*interleaving* if for all $x \in T$ and $y \in T'$:

(C1) $f'(\alpha(x)) = f(x) + \delta$,
(C2) $\beta(\alpha(x)) = x^{2\delta}$,
(C3) $f(\beta(y)) = f'(y) + \delta$, and
(C4) $\alpha(\beta(y)) = y^{2\delta}$.

The *interleaving distance* d_{I} is defined as the smallest δ such that there exists a δ-interleaving.

[1] Agarwal et al. [1] actually prove that approximating the Gromov-Hausdorff distance with a factor better than 3 is NP-hard. As many have observed, this proof also applies to the interleaving distance.

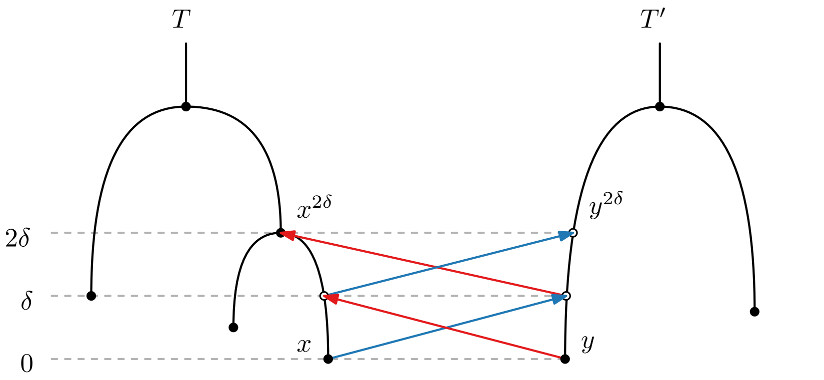

Figure 2 Two merge trees and part of a δ-interleaving. Mapping a point x from T to T' through α (in blue), and mapping it back to T via β (in red) gives the unique ancestor $x^{2\delta}$ of x.

The maps α and β are both *δ-shift maps*, i.e. continuous maps that send points exactly δ higher. Touli and Wang [23] give an alternative definition of the interleaving distance in terms of a single δ-shift map with additional requirements. They call this map a *δ-good* map. Gasparovich et al. [12] show that the interleaving distance can also be defined in terms of *labelled merge trees*: merge trees equipped with *label-maps*. Formally, let $n \geq 0$. We denote $[n] := \{1, \dots, n\}$. A map $\pi \colon [n] \to T$ is called a *label map* if each leaf in T is assigned at least one label. Note that π is not restricted to vertices, and may map different labels to the same point. The *induced matrix* $M = M(T, f, \pi)$ of a labelled merge tree is defined by $M_{i,j} = f(\text{lca}(\pi(i), \pi(j)))$, where $\text{lca}(\cdot, \cdot)$ is the *lowest common ancestor* of two points. See Figure 3 for an example of a labelled merge tree and its induced matrix.

For a matrix M, the ℓ^∞-norm is defined as $\|M\|_\infty = \max_{i,j} |M_{i,j}|$. For two *unlabelled merge trees* (T, f) and (T', f'), we refer to a pair of equally-sized label maps (π, π') as a *δ-labelling* if $\|M(T, f, \pi) - M(T', f', \pi')\|_\infty = \delta$. The δ-good interleaving distance $d_{\mathrm{I}}^{\mathrm{G}}$ and the label interleaving distance $d_{\mathrm{I}}^{\mathrm{L}}$ are defined as the smallest δ such that there exists a δ-good map or a δ-labelling. It has been shown that $d_{\mathrm{I}} = d_{\mathrm{I}}^{\mathrm{G}} = d_{\mathrm{I}}^{\mathrm{L}}$ [12, 23].

3 An Order-Preserving Interleaving Distance

We consider a new class of merge trees, which we call *ordered merge trees*. For a point $x \in T$ with $f(x) \leq h$, we denote by $x|^h$ the unique ancestor of x at height h. An ordered merge tree $(T, f, (\leq_h)_{h \geq 0})$ is a merge tree (T, f) equipped with a set of total orders on the level sets of T, such that these orders are *consistent* (see Figure 4). Formally, for two heights $h_1 \leq h_2$

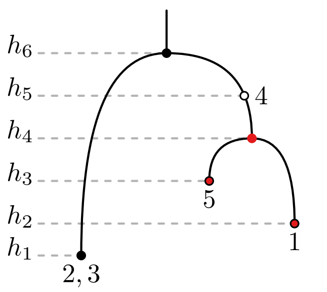

$$\begin{pmatrix} h_2 & h_6 & h_6 & h_5 & h_4 \\ . & h_1 & h_1 & h_6 & h_6 \\ . & . & h_1 & h_6 & h_6 \\ . & . & . & h_5 & h_5 \\ . & . & . & . & h_3 \end{pmatrix}$$

Figure 3 Example of a labelled merge tree and its induced matrix.

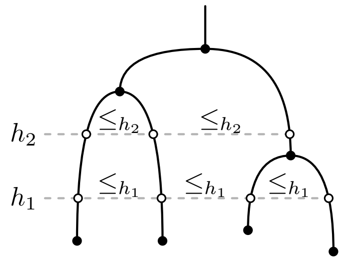

Figure 4 Two layers of an ordered merge tree.

and two points x_1, x_2 in $f^{-1}(h_1)$, we require that $x_1 \le_{h_1} x_2$ implies $x_1|^{h_2} \le_{h_2} x_2|^{h_2}$.

An ordered merge tree induces a binary relation $\sqsubseteq$ on the complete tree T: for $x_1, x_2 \in T$, we define $x_1 \sqsubseteq x_2$ if $x_1|^h \le_h x_2|^h$, where $h = \max(f(x_1), f(x_2))$. This *induced relation* is not antisymmetric, and thus not a total order. If we restrict $\sqsubseteq$ to the set of leaves of T, denoted $\sqsubseteq_L$, we do obtain a total order. We refer to $\sqsubseteq_L$ as the *induced leaf-order* of T.

Monotone interleaving distance. We now define an order-preserving distance measure for ordered merge trees. Specifically, we say a δ-shift map $\alpha\colon T \to T'$ is *monotone* if for all height values $h \ge 0$ and for any two points $x_1, x_2 \in f^{-1}(h)$ it holds that $x_1 \le_h x_2$ implies $\alpha(x_1) \le'_{h+\delta} \alpha(x_2)$. A *monotone δ-interleaving* is a δ-interleaving (α, β) such that the maps α and β are both monotone (see Figure 5). A *monotone δ-good map*, in turn, is a δ-good map α that is also monotone. Lastly, a δ-labelling (π, π') of size n is *monotone* if for all $\ell_1, \ell_2 \in [n]$ it holds that $\pi(\ell_1) \sqsubset \pi(\ell_2)$ implies $\pi'(\ell_1) \sqsubseteq' \pi'(\ell_2)$. The *monotone interleaving distance* d_{MI} is defined as the smallest δ such that there exists a monotone δ-interleaving. Similarly, we can define the monotone δ-good, and the monotone label interleaving distances, denoted $d^{\mathrm{G}}_{\mathrm{MI}}$ and $d^{\mathrm{L}}_{\mathrm{MI}}$. Our main result is the following.

▶ **Theorem 2.** *The distances d_{MI}, $d^{\mathrm{G}}_{\mathrm{MI}}$ and $d^{\mathrm{L}}_{\mathrm{MI}}$ are equal.*

To prove Theorem 2, we describe constructions between monotone δ-interleavings, monotone δ-good maps, and monotone δ-labellings. The first construction, from a monotone δ-interleaving to a monotone δ-good map, follows directly from the regular setting (by Touli and Wang [23]). The construction of a monotone δ-labelling from a monotone δ-good map follows from a refinement of the construction by Gasparovich et al. [12].[2] We use y^F to denote the lowest ancestor of $y \in T'$ in the image of α. Moreover, for two distinct points $x_1, x_2 \in T$, we say x_1 is *smaller* than x_2 if $x_1 \sqsubseteq x_2$. The existing construction is as follows.

[2] We remark that they use a slightly different (but equivalent) definition for a δ-good map.

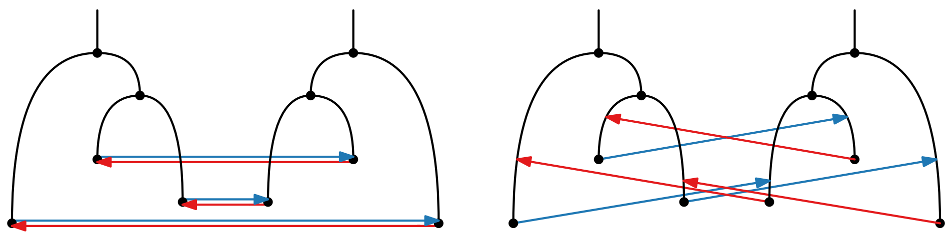

Figure 5 Parts of an optimal regular (left) and an optimal monotone (right) interleaving.

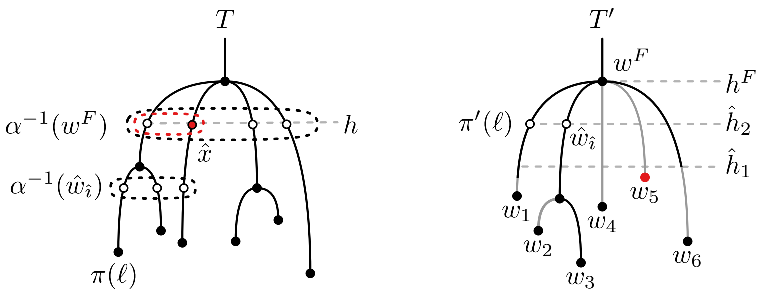

Figure 6 The refined step (S2). The grey parts of T' do not lie in the image of α. We add the pair $(\hat{x}, w_5)$ to Π. In this example, $w = w_5$, $S = \{1, 2, 3, 6\}$, $i = 5$, $S_i = \{1, 2, 3\}$ and $\hat{\imath} = 3$.

(S1) For every leaf $u \in L(T)$, add $(u, \alpha(u))$ to an initially empty set Π.
(S2) For every leaf $w \in L(T')$, take an arbitrary point $x \in \alpha^{-1}(w^F)$. Add (x, w) to Π.
(S3) Consider an arbitrary ordering $\Pi = \{(x_\ell, y_\ell) \mid \ell \in [n]\}$, and set $\pi(\ell) = x_\ell$, $\pi'(\ell) = y_\ell$.

We refine (S2), by choosing a specific $x \in \alpha^{-1}(w^F)$. Intuitively, we first identify all points $\bar{x}$ that lead to a violation of the monotonicity property if we choose x smaller than $\bar{x}$. Such a violation occurs if $\bar{x}$ is an ancestor of a labelled point whose corresponding labelled point in T' is smaller than w. We then take x to be the largest point among the points $\bar{x}$.

(S2) For every leaf $w \in L(T')$, sort the set of leaves in T'_{w^F} by induced leaf-order, denoted $W = \{w_1, \dots, w_m\}$. Define $S \subseteq [m]$ such that for $k \in S$, w_k^F is a strict descendant of w^F. Fix $i \in [m]$ such that $w_i = w$, and define $S_i = \{k \in S \mid k < i\}$. Now consider the set $X = \alpha^{-1}(w^F)$
- If S_i is empty, take x to be the smallest point in X and add (x, w) to Π.
- If S_i is not empty, consider the largest index $\hat{\imath} \in S_i$. Define Y as the set of strict descendants of w^F that were labelled in (S1). Consider the following height values:

$$\hat{h}_1 := \max\{f'(w_k^F) \mid k \in S\}, \quad \hat{h}_2 := \max\{f'(y) \mid y \in Y\}, \quad \hat{h} = \max(\hat{h}_1, \hat{h}_2) \qquad (1)$$

Consider the unique ancestor $\hat{w}_{\hat{\imath}}$ of $w_{\hat{\imath}}$ at height $\hat{h}$. Note that $\hat{w}_{\hat{\imath}}$ is a strict descendant of the point w^F and that it lies in the image of α. Let $X_{\hat{\imath}} \subset X$ be the set of ancestors of points in $\alpha^{-1}(\hat{w}_{\hat{\imath}})$. Take $\hat{x}$ to be the largest point in $X_{\hat{\imath}}$ and add $(\hat{x}, w)$ to Π.

See Figure 6 for an illustration. We can show that the resulting δ-labelling is monotone.

Lastly, we construct a monotone δ-interleaving from a monotone δ-labelling. To do so, we extend an existing construction of a δ-good map α from a δ-labelling by Gasparovich et al. [12]. Specifically, for all $x \in T$, they consider an arbitrary label ℓ from the subtree T_x and set $y_\ell = \pi'(\ell)|^{f(x)+\delta}$. Gasparovich et al. show that the point y_ℓ is well-defined, and argue that the resulting map α is a δ-good map. We can construct a δ-interleaving (α, β) by using this construction twice: first to build a map $\alpha \colon T \to T'$ and next to build a map $\beta \colon T' \to T$.

Monotone leaf-label interleaving distance. We now turn to a restriction of the (monotone) label interleaving distance. Specifically, if we restrict a label map to map *only* to the leaves of T, we obtain a *leaf-label map*. A *δ-leaf-labelling* is a pair of leaf-label maps that is also a δ-labelling. The *leaf-label interleaving distance* d_I^{LL}, in turn, is defined as the smallest δ for which there exists a δ-leaf labelling. We can show that this distance is an approximation of the interleaving distance, in both the regular and monotone setting. We define the *leaf-gap* G of two trees T and T' as the maximum height difference of any pair of leaves in T and T'.

▶ **Theorem 3.** *The monotone leaf-label interleaving distance between $(T, f, (\leq_h))$ and $(T', f', (\leq'_h))$ is bounded by $2\delta + G$, where $\delta = d_{\mathrm{MI}}(T, T')$ and G is the leaf-gap of T and T'.*

4 Approximating the Monotone Interleaving Distance

In this section we describe an algorithm to compute the monotone leaf-label interleaving distance between two ordered merge trees $(T, f, (\leq_h))$ and $(T', f', (\leq'_h))$. We denote the leaves of T and T', sorted by leaf-order, by $L(T) = \{u_1, \ldots, u_m\}$ and $L(T') = \{w_1, \ldots, w_{m'}\}$, respectively. We can show that we can permute the labels of a monotone labelling (ω, ω') on T and T', such that the resulting label maps *respect* the induced leaf-orders of their trees. That is, for any two labels $i, j \in [n]$ with $i < j$, we have $\omega(i) \sqsubseteq_L \omega(j)$ and $\omega'(i) \sqsubseteq'_L \omega'(j)$. Assume (ω, ω') is such a monotone leaf-labelling. Let $M = M(T, f, \omega)$ and $M' = M(T', f', \omega')$ be the corresponding induced matrices and set $\mathcal{M} = \mathcal{M}(T, T') := |M - M'|$. We refer to entries $\mathcal{M}_{i,i}$ as the *diagonal* of $\mathcal{M}$, and to entries $\mathcal{M}_{i,j}$ with $j - i = 1$ as the *upper-diagonal* of $\mathcal{M}$.

▶ **Lemma 4.** *The maximum of $\mathcal{M}$ lies on the diagonal or upper diagonal of $\mathcal{M}$.*

We now describe a dynamic program to compute the monotone leaf-label interleaving distance between T and T'. We denote by $T[i]$ the subtree of T consisting of only the first i leaves. For $i \in |L(T)|$ and $j \in |L(T')|$, we maintain a value $\Delta[i, j]$ that stores the monotone leaf-label interleaving distance between $T[i]$ and $T'[j]$. Consider an optimal monotone leaf-labelling (ω, ω') for $T[i]$ and $T'[j]$ with a minimum number of k labels, such that ω and ω' respect the leaf-orders of T and T', respectively. From Lemma 4 we know that to compute $d_{\mathrm{MI}}^{\mathrm{LL}}(T[i], T'[j])$, it suffices to compute the diagonal and upper-diagonal entries of $\mathcal{M} = \mathcal{M}(T[i], T'[j])$. As $\omega(k) = u_i$ and $\omega'(k) = w_j$, we have $\mathcal{M}_{k,k} = |f(u_i) - f'(w_j)| =: \varepsilon$. To compute the other relevant elements of $\mathcal{M}$, we consider the three options for label $k - 1$:

$$\begin{array}{llllll} (1) & \omega(k-1) = u_i, & (2) & \omega(k-1) = u_{i-1}, & (3) & \omega(k-1) = u_{i-1}, \\ & \omega'(k-1) = w_{j-1}, & & \omega'(k-1) = w_j, & & \omega'(k-1) = w_{j-1}. \end{array}$$

First assume case (1) applies. Then we know that $\mathcal{M}_{k-1,k} = |f(u_i) - f'(\mathrm{lca}(w_{j-1}, w_j))|$. Furthermore, $\Delta[i, j-1]$ captures the remaining relevant entries of $\mathcal{M}$. We set $\delta_1 = \max(\Delta[i, j-1], \mathcal{M}_{k-1,k})$. Similarly, we can set δ_2 and δ_3 for cases (2) and (3) respectively. Finally, at each iteration, we set $\Delta[i, j] = \max(\varepsilon, \min(\delta_1, \delta_2, \delta_3))$. We can show that $\Delta[i, j] = d_{\mathrm{MI}}^{\mathrm{LL}}(T[i], T'[j])$. The algorithm returns $\Delta[|L(T)|, |L(T')|]$.

We can use our algorithm to compute a monotone interleaving as follows. First, we compute the lowest common ancestors of all consecutive pairs of leaves in T or T'. This allows us to construct Δ in $\mathcal{O}(n^2)$ time, where $n = |L(T)| + |L(T')|$. Recovering an optimal leaf-labelling from the dynamic program can be done in a standard way. Next, we can construct two partial maps $\alpha_L \colon L(T) \to T'$ and $\beta_L \colon L(T') \to T$ using the construction from Section 3 in $\mathcal{O}(n)$ time. Lastly, one can recover a complete interleaving (α, β) from α_L and β_L using continuity and δ-shift map properties. Our final result follows from Theorem 3:

▶ **Theorem 5.** *Given two ordered merge trees T and T', there exists an $\mathcal{O}(n^2)$ algorithm that computes a monotone δ-interleaving between T and T', where $\delta \leq 2d_{\mathrm{MI}}(T, T') + G$.*

5 Acknowledgments

Thijs Beurskens and Tim Ophelders are supported by the Dutch Research Council (NWO) under project numbers OCENW.M20.089 and VI.Veni.212.260, respectively.

References

1. P.K. Agarwal, K. Fox, A. Nath, A. Sidiropoulos, and Y. Wang. Computing the Gromov-Hausdorff distance for metric trees. *ACM Transactions on Algorithms*, 14(2):1–20, 2018. doi:10.1145/3185466.
2. U. Bauer, X. Ge, and Y. Wang. Measuring distance between Reeb graphs. In *Proc. 30th Symposium on Computational Geometry (SoCG)*, pages 464–473, 2014. doi:10.1145/2582112.2582169.
3. H.B. Bjerkevik, M. B. Botnan, and M. Kerber. Computing the interleaving distance is np-hard. *Foundations of Computational Mathematics*, 20:1237–1271, 2020. doi:10.1007/s10208-019-09442-y.
4. H.B. Bjerkevik and M.B. Botnan. Computational Complexity of the Interleaving Distance. In *Proc. 34th Symposium on Computational Geometry (SoCG)*, pages 13:1–13:15, 2018. doi:10.4230/LIPIcs.SoCG.2018.13.
5. P. Bubenik, V. De Silva, and J.A. Scott. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, 15(6):1501–1531, 2015. doi:10.1007/s10208-014-9229-5.
6. P. Bubenik and J.A. Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51(3):600–627, 2014. doi:10.1007/s00454-014-9573-x.
7. A.J. Chadwick, E. Steel, R.A. Williams-Schaetzel, P. Passalacqua, and C. Paola. Channel migration in experimental river networks mapped by particle image velocimetry. *Journal of Geophysical Research: Earth Surface*, 127(1):e2021JF006300, 2022. doi:10.1029/2021JF006300.
8. F. Chazal, D. Cohen-Steiner, M. Glisse, L.J. Guibas, and S.Y. Oudot. Proximity of persistence modules and their diagrams. In *Proc. 25th Symposium on Computational Geometry (SoCG)*, pages 237–246, 2009. doi:10.1145/1542362.1542407.
9. J. Curry, H. Hang, W. Mio, T. Needham, and O.B. Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, 6(3):371–428, 2022. doi:10.1007/s41468-022-00089-3.
10. V. De Silva, E. Munch, and A. Patel. Categorified Reeb graphs. *Discrete & Computational Geometry*, 55(4):854–906, 2016. doi:10.1007/s00454-016-9763-9.
11. V. De Silva, E. Munch, and A. Stefanou. Theory of interleavings on a category with a flow. *Theory and Applications of Categories*, 33(1):583–607, 2018. URL: http://www.tac.mta.ca/tac/volumes/33/21/33-21a.pdf.
12. E. Gasparovich, E. Munch, S. Oudot, K. Turner, B. Wang, and Y. Wang. Intrinsic interleaving distance for merge trees. arXiv:1908.00063.
13. A.D. Howard, M.E. Keetch, and C.L. Vincent. Topological and geometrical properties of braided streams. *Water Resources Research*, 6(6):1674–1688, 1970. doi:10.1029/WR006i006p01674.
14. M.G. Kleinhans, M. van Kreveld, T. Ophelders, W. Sonke, B. Speckmann, and K. Verbeek. Computing representative networks for braided rivers. *Journal of Computational Geometry*, 10(1):423–443, 2019. doi:10.20382/jocg.v10i1a14.
15. M. Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015. doi:10.1007/s10208-015-9255-y.
16. Y. Liu, P.A. Carling, Y. Wang, E. Jiang, and P.M. Atkinson. An automatic graph-based method for characterizing multichannel networks. *Computers & Geoscience*, 166:105180, 2022. doi:10.1016/j.cageo.2022.105180.
17. W.A. Marra, M.G. Kleinhans, and E.A. Addink. Network concepts to describe channel importance and change in multichannel systems: test results for the jamuna river, bangladesh. *Earth Surface Processes and Landforms*, 39(6):766–778, 2014. doi:10.1002/esp.3482.

18 F. Mémoli and L. Zhou. Persistent homotopy groups of metric spaces. arXiv:1912.12399.
19 D. Morozov, K. Beketayev, and G. Weber. Interleaving distance between merge trees. Manuscript, 2013.
20 E. Munch and A. Stefanou. The ℓ^∞-Cophenetic metric for phylogenetic trees as an interleaving distance. *Research in Data Science*, 17:109–127, 2019. doi:10.1007/978-3-030-11566-1_5.
21 G. O'Beirne. Waimakariri River (New Zealand), 2017. Image is slightly cropped; accessed on 30/11/2023. Photo is licensed under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/). URL: https://commons.wikimedia.org/w/index.php?curid=61302377.
22 A. Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, 1(3-4):397–419, 2018. doi:10.1007/s41468-018-0012-6.
23 E.F. Touli and Y. Wang. FPT-algorithms for computing the Gromov-Hausdorff and interleaving distances between trees. *Journal of Computational Geometry*, 13:89–124, 2022. doi:10.20382/jocg.v13i1a4.
24 L. Yan, T. B. Masood, F. Rasheed, I. Hotz, and B. Wang. Geometry aware merge tree comparisons for time-varying data with interleaving distances. *IEEE Transactions on Visualization and Computer Graphics*, 29(8):3489–3506, 2022. doi:10.1109/TVCG.2022.3163349.
25 L. Yan, Y. Wang, E. Munch, E. Gasparovich, and B. Wang. A structural average of labeled merge trees for uncertainty visualisation. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 2020. doi:10.1109/TVCG.2019.2934242.