Towards an Optimal Bound for the Interleaving Distance on Mapper Graphs

Erin Wolf Chambers[*1], Ishika Ghosh[†2], Elizabeth Munch[‡2], Sarah Percival[§3], and Bei Wang [¶4]

[1]University of Notre Dame
[2]Michigan State University
[3]University of New Mexico
[4]University of Utah

Abstract

Mapper graphs are a widely used tool in topological data analysis and visualization. They can be viewed as discrete approximations of *Reeb graphs*, offering insight into the shape and connectivity of complex data. Given a high-dimensional point cloud $\mathbb{X}$ equipped with a function $f : \mathbb{X} \to \mathbb{R}$, a *mapper graph* provides a summary of the topological structure of $\mathbb{X}$ induced by f, where each node represents a local neighborhood, and edges connect nodes whose corresponding neighborhoods overlap. Our focus is the interleaving distance for mapper graphs, which arises from a discretization of the interleaving distance for Reeb graphs. Notably, computing both distances is NP-hard in general. The interleaving distance for mapper graphs quantifies their similarity by measuring the extent to which they must be "stretched" to become comparable. Recent work introduced a loss function that provides an upper bound on the interleaving distance for mapper graphs. More specifically, this loss function evaluates how far a given assignment is from being a true interleaving and is computationally tractable, offering a practical way to estimate the distance.

In this paper, we employ a categorical formulation of mapper graphs and develop the first framework for computing the associated loss function. Since the quality of the bound depends on the chosen assignment, we optimize this loss function by formulating the problem of finding the best assignment as an integer linear programming (ILP) problem. To evaluate the effectiveness of our optimization, we apply it to small mapper graphs where the interleaving distance can be computed by hand, demonstrating that the optimized upper bound successfully matches the interleaving distance in these cases. Additionally, we conduct an experiment on the MPEG-7 dataset, computing pairwise optimal mapper loss on a collection of mapper graphs derived from images and leveraging the distance bound for image classification.

1 Introduction

In order to provide a graphical representation of data, mapper graphs [57] are an increasingly popular tool arising from the field of topological data analysis, with interest both in theory [2, 16, 20]

---

[*]echambe2@nd.edu

[†]ghois3@msu.edu

[‡]muncheli@msu.edu
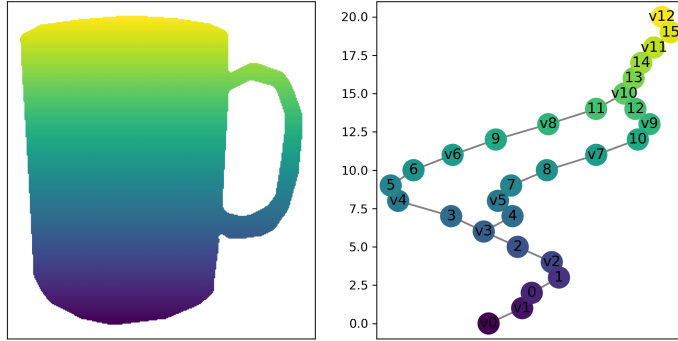
[§]spercival@unm.edu

[¶]beiwang@sci.utah.edu

Figure 1: An example of an input space in 2D (a *cup* from the MPEG-7 data set), and the mapper graph computed using filtration function given by pixel location on the y-axis is shown at right.

and practice [49, 52, 54, 61, 48]. The basic idea used in practice for defining the mapper graph is as follows. Given a dataset χ which has a distance measure $d : \chi \times \chi \to \mathbb{R}$ and a filter function $f : \chi \to \mathbb{R}$, fix a cover $\mathcal{U}$ of the image $f(\chi)$. Then the mapper graph has a vertex for each cluster in $f^{-1}(U)$ for all $U \in \mathcal{U}$, and an edge between vertices representing overlapping clusters. In the more continuous setting, this can be constructed by assuming an input topological space $\mathbb{X}$ with a function $f : \mathbb{X} \to \mathbb{R}$, and then the mapper graph is the nerve of the connected components of the inverse cover $\{f^{-1}(U) \mid U \in \mathcal{U}\}$; see Figure 1 for an example. With this latter viewpoint, we can view the mapper graph as a discretization of a Reeb graph [50], although that full generalization will not be used in this work.

A limitation of mapper graphs is that they have largely been restricted to qualitative evaluation and exploratory data analysis in practical settings. This is due in large part to the fact that quantitative analysis of the quality of a mapper graph has been restricted to theoretical guarantees only. Specifically, if we want to evaluate an output mapper graph relative to some ground truth mapper graph, we need access to a measurement of similarity. While distances between mapper graphs and Reeb graphs have been defined in theory (see Section 1.2), nearly all are severely limited in practical applications by computational complexity issues.

In this paper, we focus on the interleaving distance for mapper graphs. This construction arises from a discretization of the interleaving for Reeb graphs [28], which in turn came from a translation of this distance for algebraic topological data analysis representations [23, 17, 56] to graph objects. However, as with nearly all of the available metrics, it is NP-hard to compute in general [11, 17]. In this paper, we use the framework given in [21] to bound the interleaving distance. Formally, we encode our mapper inputs as cosheaves (see e.g. [26, 51]) of the form $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ for a given cover $\mathcal{U}$. We can define a thickening operation on the elements of $\mathbf{Open}(\mathcal{U})$, which amounts to a discretized version of thickening of open intervals in $\mathbb{R}$, and then define the smoothing of a functor F^n by post composing with this thickening operation. Then an n-interleaving is a pair of natural transformations $\varphi : F \to G^n$ and $\psi : G \to F^n$ which commute with each other up to this thickening operation. The interleaving distance $d_I(F, G)$ is given by the minimum n for which such an interleaving can be found.

2

1.1 Our contribution

This paper provides the first framework for providing an optimized bound for the interleaving distance between mapper graphs. In this section, we give a general description of our results, with details provided in the rest of the paper. We start with a simple discretization of the real line given by a simplicial complex K with vertices $\sigma_i = i\delta$ and edges τ_i between vertices σ_i and σ_{i+1}. This results in a cover of $\mathbb{R}$, with cover elements of the form

$$\{U_{\sigma_i} = ((i-1)\delta, (i+1)\delta) \mid i \in [-L+1, \cdots, L-1]\}$$

and intersections of the form

$$\{U_{\tau_i} = (i\delta, (i+1)\delta) \mid i \in [-L, \cdots, L-1]\}.$$

where we write $\mathcal{U} = \{U_{\sigma_i}\} \cup \{U_{\tau_i}\}$. We write $\rho \in K$ and similarly $U_\rho \in \mathcal{U}$ when we want to talk about cells of either dimension. Since $\mathcal{U}$ is a poset, we can give it the Alexandrov topology and then following [21], we write S^n for the n-thickening of a set $S \in \mathbf{Open}(\mathcal{U})$ representing a discretized thickening of intervals in $\mathbb{R}$, $(a, b) \mapsto (a-n, b+n)$. Details of this operation can be found in Section 2.2.

Definition 2.1 *Let $F, G : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ be cosheaves and $n \in \mathbb{Z}_{\geq 0}$. An n-interleaving is a pair of natural transformations $\varphi : F \Rightarrow G^n$ and $\psi : G \Rightarrow F^n$ such that the diagrams*

$$
\begin{array}{ccc}
F(S) & \xrightarrow{\;\;F[S \subseteq S^{2n}]\;\;} & F(S^{2n}) \\
\;\;\downarrow{\scriptstyle \varphi_S} & \nearrow{\scriptstyle \psi_{S^n}} & \\
G(S^n) & &
\end{array}
\qquad\qquad
\begin{array}{ccc}
& F(S^n) & \\
\nearrow{\scriptstyle \psi_S} & & \searrow{\scriptstyle \varphi_{S^n}} \\
G(S) & \xrightarrow[G[S \subseteq S^{2n}]]{} & G(S^{2n})
\end{array}
$$

commute for all $S \in \mathbf{Open}(\mathcal{U})$. The interleaving distance *is given by*

$$d_I(F, G) = \inf\{n \geq 0 \mid \text{there exists an } n\text{-interleaving}\},$$

and is set to be $d(F, G) = \infty$ if there is no interleaving for any n.

In this paper, we work with the idea of an *n-assignment*, (φ, ψ), which has the structure of the aforementioned pair of natural transformations but without the commutativity promises that come either from being natural transformations or from being an interleaving. The loss function given in [21], $L(\varphi, \psi)$, is a measure of how much additional shift is needed to turn the given assignment into an interleaving. The basic idea is that we can work with diagrams that look like those of the above definition, but that might not commute and define a distance (Equation (1)) between objects in the terminal set of the diagram. That is, we give a distance $d_S^F : F(S) \times F(S) \to \mathbb{R}$ for every open set S; defined similarly for G. This distance is used to define the amount of thickening needed before commutativity is promised for each of the diagrams required to ensure the pair gives an interleaving; these diagrams can also be seen in the Diagram column of Table 1. The losses for the two triangle-type diagrams for a given n-assignment are denoted $L_{\triangledown}^S(\varphi, \psi)$ and $L_{\triangle}^S(\varphi, \psi)$; while the losses for the parallelogram-type diagrams that ensure that the input assignment is actually a natural transformation are denoted $L_{\diamondsuit}^{S,T}(\varphi)$ and $L_{\diamondslash}^{S,T}(\psi)$. Then, these values are collected together to give a single maximum which defines the full loss function for an input n-assignment. Further specifics for these computations are given in the full version of the following definition.

	Loss Term	Diagram	Matrix Multiplication	Eval.
Edge-vertex Parallelogram	$L_{\lsem}^{S_\tau,S_\sigma}$	$F(S_\tau) \xrightarrow{F[\subseteq]} F(S_\sigma)$, φ_{S_τ}, φ_{S_σ}, $G^n(S_\tau) \xrightarrow{G[\subseteq]} G^n(S_\sigma)$	$D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right)$ $D_{G^n}^V \left(M_\varphi^V \cdot B_F^\downarrow - B_{G^n}^\downarrow \cdot M_\varphi^E \right)$	$\displaystyle\max_{x_{ij}\in A} x_{ij}$
	$L_{\square}^{S_\tau,S_\sigma}$	$F^n(S_\tau) \xrightarrow{F[\subseteq]} F^n(S_\sigma)$, ψ_{S_τ}, ψ_{S_σ}, $G(S_\tau) \xrightarrow{G[\subseteq]} G(S_\sigma)$	$D_{F^n}^V \left(M_\psi^V \cdot B_G^\uparrow - B_{F^n}^\uparrow \cdot M_\psi^E \right)$ $D_{F^n}^V \left(M_\psi^V \cdot B_G^\downarrow - B_{F^n}^\downarrow \cdot M_\psi^E \right)$	
Thickening Parallelogram	$L_{\lsem}^{S_\rho,S_\rho^n}$	$F(S_\rho) \xrightarrow{F[\subseteq]} F^n(S_\rho)$, φ_{S_ρ}, $\varphi_{S_\rho^n}$, $G^n(S_\rho) \xrightarrow{G[\subseteq]} G^{2n}(S_\rho)$	$D_{G^n}^V (M_{\varphi^n}^V \cdot I_F^V - I_{G^n}^V \cdot M_\varphi^V)$ $D_{G^n}^E (M_{\varphi^n}^E \cdot I_F^E - I_{G^n}^E \cdot M_\varphi^E)$	
	$L_{\square}^{S_\rho,S_\rho^n}$	$F^n(S_\rho) \xrightarrow{F[\subseteq]} F^{2n}(S_\rho)$, ψ_{S_ρ}, $\psi_{S_\rho^n}$, $G(S_\rho) \xrightarrow{G[\subseteq]} G(S_\rho^n)$	$D_{F^n}^V \left(M_{\psi^n}^V \cdot I_G^V - I_{F^n}^V \cdot M_\psi^V \right)$ $D_{F^n}^E \left(M_{\psi^n}^E \cdot I_G^E - I_{F^n}^E \cdot M_\psi^E \right)$	
Triangle	$L_{\triangledown}^{S_\rho}$	$F(S_\rho) \xrightarrow{F[\subseteq]} F^{2n}(S_\rho)$, φ_{S_ρ}, $\psi_{S_\rho^n}$, $G^n(S_\rho)$	$D_{F^{2n}}^V \left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V \right)$ $D_{F^{2n}}^E \left(I_{F^n}^E \cdot I_F^E - M_{\psi^n}^E \cdot M_\varphi^E \right)$	$\displaystyle\max_{x_{ij}\in A} \left\lceil \frac{x_{ij}}{2} \right\rceil$
	$L_{\triangle}^{S_\rho}$	$F^n(S_\rho)$, ψ_{S_ρ}, $\varphi_{S_\rho^n}$, $G(S_\rho) \xrightarrow{G[\subseteq]} G^{2n}(S_\rho)$	$D_{G^{2n}}^V \left(I_{E^n}^V \cdot I_G^V - M_{\varphi^n}^V \cdot M_\psi^V \right)$ $D_{G^{2n}}^E \left(I_{G^n}^E \cdot I_G^E - M_{\varphi^n}^E \cdot M_\psi^E \right)$	

Table 1: The list of loss terms, the relevant diagram, and the matrix multiplication representation. The final column gives the result taken from the final matrix: either the maximum, or the ceiling of half the maximum.
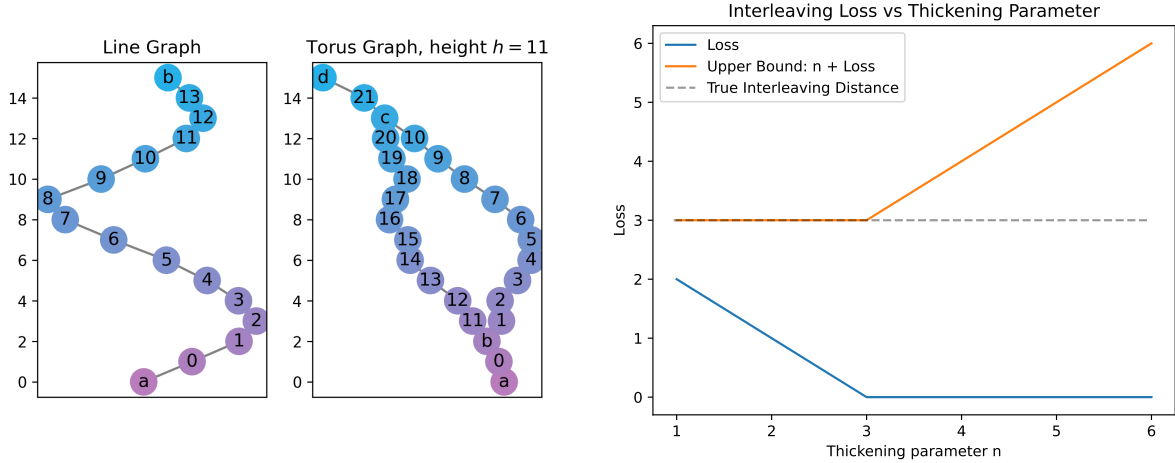
Figure 2: (Left) An example of a line and torus mapper graph, where the latter has a loop of height 11. (Right) Loss function and upper bound computed for this pair of graphs for varying n.

Definition 2.6 *Fix an n-assignment (φ, ψ). The* (extended basis) loss function *is defined to be*

$$L_B(\varphi, \psi) = \max_{\substack{\sigma < \tau \in K \\ \rho \in K}} \left\{ L_{\triangledown}^{S_\tau, S_\sigma}, L_{\square}^{S_\tau, S_\sigma}, L_{\triangledown}^{S_\rho, S_\rho^n}, L_{\square}^{S_\rho, S_\rho^n}, L_{\triangle}^{S_\rho}, L_{\triangledown}^{S_\rho} \right\} \tag{4}$$

where the max is taken over all cells $\rho \in K$ of any dimension; or of pairs of an edge τ adjacent to a vertex σ.

The results is that this loss function, combined with the n chosen for the assignment, can bound the interleaving distance as seen in the following theorem.

Theorem 2.7 (Cor. of [21, Thm 3.16]) *Given an (extended) basis n-assignment φ and ψ, we have*

$$d_I(F, G) \leq n + L_B(\varphi, \psi).$$

In Section 3.1, we show that the structure of the input mapper cosheaves can be stored as a graph, where the graph of $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ has a vertex for every element in every $F(S_\sigma)$ and an edge for every $F(S_\tau)$. Then, in Section 3.2, we show that we can represent all used natural transformations as well as the input assignment in matrices so that the loss function becomes the maximum entry in a collection of matrix multiplications. Specifically, the distance matrices are denoted by D, the boundary matrices for the graph representations by B, the inclusion natural transformations by I, and the assignment information as M_φ and M_ψ. Then each diagram shown in Table 1 can be checked for the k needed for commutativity using the multiplication shown in the Matrix Multiplication column. Then the loss contribution for each particular diagram comes from finding either a maximum entry or half the ceiling of the maximum as seen in the final column; then the theorem promises that the interleaving distance is bounded by $d_I(F, G) \leq n + k$ where k is the computed loss value by this procedure.

Note that this loss function is dependent on the input n-assignment (φ, ψ). With a poor choice of input comes a bad bound. For this reason, we utilize an integer linear programming (ILP) solver to find an optimal n-assignment, and thus the best available bound $n + L(\varphi, \psi)$ for a given n. As an example, see Figure 2. In this case we look at the interleaving distance between a line graph at
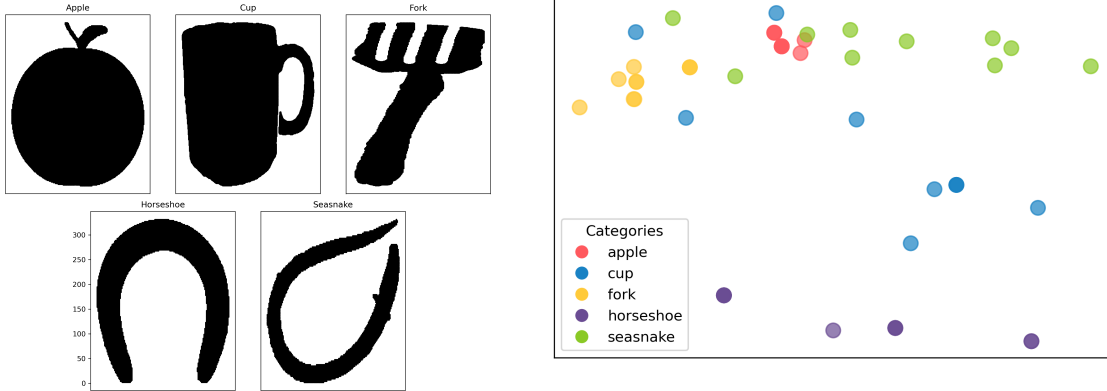
5

Figure 3: (Left) Example images for the used categories of images from the MPEG-7 dataset. (Right) MDS plot computed on the pairwise optimal upper bound for the interleaving distance.

left and the mapper graph of an upright torus in the middle. We can manually compute that the interleaving distance in this case is $\lceil \frac{h}{4} \rceil$ if the height of the torus loop is h. In the right of Figure 2, we show the bound found by the ILP for different choices of n for the input n-assignment structure. This graph includes both the computed loss function, as well as the upper bound which is $n + \text{loss}$. We see that up to $n = 3 = \lceil \frac{h}{4} \rceil$, the bound actually achieves the true interleaving distance. After finding the actual interleaving distance, the bound increases since the loss is always 0 so the bound is simply given by n.

From this example, it is tempting to think that we need only compute the bound for some n that is less than the true interleaving distance. Unfortunately, in Section 5.2 we give an example of a pair of graphs for which the bound cannot achieve the interleaving distance. For this reason, we perform an exponential search over possible n's to achieve an optimal bound for the interleaving distance between two given mapper graphs.

In addition to providing simple examples, we also run an experiment on the MPEG-7 data set, with details given in Section 5.3. The results can be seen in Figure 3. In this experiment, we used 5 classes of images: *apples, cups, forks, horseshoes, seasnakes.* We construct an MDS plot with the pairwise distance matrix constructed with our optimized mapper loss. We can see from this example that different shapes tend to cluster together, though some shapes have more difficulty clustering than the others. Often, this is due to what the interleaving distance is designed to measure; namely that the distance will be small for mapper graphs with the same graph *and* function, so seasnake examples in different orientations or cups that do or do not have connected handles have large interleaving distance despite coming from the same class. Section 6 gives additional details on potential improvements and extensions for the work presented here.

1.2 Related work

The mapper graph can be thought of as a discretization of the Reeb graph [50], which tracks the connected components of level sets of a topological space. Because a Reeb graph is indeed a graph with the added structure of a function, one can use any of the many existing graph distances (e.g. the graph edit distance [32]) to compare Reeb graphs by forgetting the function information. However, given the topologically-motivated construction of Reeb graphs, a notion of distance that is sensitive to the function is desirable.

There is now an extensive family of available distances between Reeb graphs, including the

interleaving distance [28, 22], the functional distortion distance [6], the universal distance [7, 5], the functional contortion distance [5], the bottleneck distance between the relevant persistence diagrams [19], and graph edit distances [4, 7]. See [13, 59] for surveys. There has been work using the interleaving distance in particular for evaluating mapper graphs, including showing convergence for the mapper graph to the Reeb graph [45, 16]. However, those works convert a mapper graph into a Reeb graph and use the continuous distance whereas in this paper, we directly construct an interleaving distance between the mapper graphs themselves. There is also extensive work understanding the effect of the parameter choices on the constructed mapper graph [20, 18, 1].

Further, the idea of the interleaving distance appears widely in the TDA literature for many kinds of input objects. This includes persistence modules [23], where the definition got its start. However, the family of distances has expanded to include multiparameter persistence modules [39, 8], filtered spaces [12], merge trees [43, 9], labeled Reeb graphs [38], labeled merge trees [44, 33, 60, 27], and more general categorical constructions [29, 41, 25, 15, 37, 55]. However, it has been shown that in a broad collection of these, computation is NP-hard [11, 10], thus rendering the interleaving distance to be a tool that is extensively used in theory but never used in practice.

Among these works, most do not include implementation in any form. There are a few exceptions which we will briefly note, although none of them apply to our setting of mapper/Reeb graphs. The interleaving distance for labeled merge trees [44] is in P, so implementations are easily available, such as in [60]. In [27], they use labelings found using Gromov-Wasserstein couplings to give bound on the interleaving distance between unlabeled merge trees. Finally, [47] uses an ILP on a formulation of the interleaving distance specific to merge trees [30] to provide bounds on the merge tree interleaving distance.

2 Background

We now provide the necessary background for this paper, largely following the terminology which was presented in [21].

2.1 Functors, cosheaves and mapper graphs

We work in a category-theoretic setting and provide a brief overview of the necessary definitions, with more details provided in [51, 26]. A *category* $\mathcal{C}$ is a collection of objects $X, Y, Z, \ldots$ along with morphisms $f, g, h, \ldots$ such that

1. each morphism f has a designated domain X and codomain Y,

2. every object has an identity morphism $\mathbb{1}_X : X \to X$,

3. any pair of morphisms $f : X \to Y$ and $g : Y \to Z$ has a composite morphism $g \circ f : X \to Z$ which satisfies an identity axiom, where $f : X \to Y = \mathbb{1}_Y \circ f = f \circ \mathbb{1}_X$, and an associativity axiom, where $h \circ (g \circ f) = (h \circ g) \circ f$.

Many settings in mathematics can be viewed as categories, for example: **Set** is the category whose objects are sets and morphisms are set maps; **Top**, whose objects are topological spaces and morphisms are continuous functions; and, for a given topological space X, **Open**(X) is the category whose objects are the objects are open sets and morphisms are given by inclusion. A category is *small* if the collections of objects and morphisms are both sets.

Given two categories $\mathcal{C}$ and $\mathcal{D}$, a *functor* $F : \mathcal{C} \to \mathcal{D}$ maps each object $x \in \mathcal{C}$ to an object $F(x) \in \mathcal{D}$ and each morphism $f : X \to Y$ to a morphism $F[f] : F(X) \to F(Y)$ that satisfies the

following properties: for any $X \in \mathcal{C}$, $F[\mathbb{1}_X] = \mathbb{1}_{F(X)}$; for any $f, g \in \mathcal{C}$ for which $gf[x]$ is defined, we have $F[gf] = F[g]F[f]$. Given two functors $F, G : \mathcal{C} \to \mathcal{D}$, a *natural transformation* $\eta : F \Rightarrow G$ is a collection of maps $\eta_X : F(X) \to G(X)$ such that for any morphism $f : X \to Y$ in $\mathcal{C}$, the diagram

$$
\begin{array}{ccc}
X & F(X) \xrightarrow{\eta_X} G(X) \\
\downarrow{f} & F[f]\downarrow \qquad \qquad \downarrow{G[f]} \\
Y & F(Y) \xrightarrow{\eta_Y} G(Y)
\end{array}
$$

commutes. One example of a functor is $\pi_0 : \mathbf{Top} \to \mathbf{Set}$, where $\pi_0(X)$ denotes the set of path-connected components of a topological space X, and the morphisms are set maps $\pi_0[f] : \pi_0(X) \to \pi_0(Y)$ sending a path-connected component A in X to the element $f(A)$ in Y.

In the special case where J is a small category, a diagram is simply a functor $F : J \to \mathcal{C}$. Letting $c \in \mathcal{C}$, we abuse notation to define the constant functor $c \colon J \to \mathcal{C}$, $c(j) \mapsto c$ for all $j \in J$. In this setting, a *cocone* is a natural transformation $\lambda : F \to c$. We denote $\lambda_j : F(j) \to c$ and note that for any $f : j \to k$ in J, $\lambda_k \circ F[f] = \lambda_j$. A cocone $\lambda : F \to c$ is a *colimit* if, for any other cocone $\lambda' : F \to c'$, there is a unique morphism $u : c \to c'$ such that the diagram

$$
\begin{array}{c}
F(j) \xrightarrow{\quad F[f] \quad} F(k)
\end{array}
$$

commutes for all $f : j \to k$ in J. We denote this by $\underrightarrow{\mathrm{colim}}_I(F)$.

For the rest of this paper, we turn our attention to functors of the form $F : \mathbf{Open}(X) \to \mathbf{Set}$, called *pre-cosheaves*. We are particularly interested in the case when a pre-cosheaf is a *cosheaf*. In short, a cosheaf is a presheaf that is entirely and uniquely defined by any open cover of $U \in \mathbf{Open}(X)$. More rigorously, given an object $U \in \mathbf{Open}(X)$ and a cover $\{U_\alpha\}_{\alpha \in A}$ of U, define a category $\mathcal{U} = \{U_\alpha \cap U_{\alpha'} \mid \alpha, \alpha' \in A\}$ with morphisms given by inclusion. As these open sets U_α are themselves objects in $\mathbf{Open}(X)$, the functor $F : \mathcal{U} \to \mathbf{Set}$ is well-defined, so we can consider its colimit $\lambda : F \to L$. A *cosheaf* is a functor F whose unique map $L \to F(U)$ given by the colimit definition is an isomorphism.

2.2 Mapper cosheaves and the interleaving distance

In this section, we show how we can represent the data of a mapper graph as a cosheaf. We assume that we are given data in the form of a compact topological space with a function $f : \mathbb{X} \to \mathbb{R}$. Fix a resolution $\delta > 0$. We assume $\mathrm{Im}(f) \subseteq [-\delta L, \delta L]$. This choice of resolution induces a discretization of $[-\delta L, \delta L]$ into a cell complex K where the 0-cells are given by the points $\{\sigma_i := i\delta \mid i \in [-L, \cdots, L]\}$, and edges given by $\{\tau_i := (i\delta, (i+1)\delta) \mid i \in [-L, L-1]\}$. Further, this cell complex K induces a cover of the image with cover elements of the form

$$\{U_{\sigma_i} = ((i-1)\delta, (i+1)\delta) \mid i \in [-L+1, \cdots, L-1]\}$$

and intersections of the form

$$\{U_{\tau_i} = (i\delta, (i+1)\delta) \mid i \in [-L, \cdots, L-1]\}.$$

Figure 4: The notation used for the cover of $\mathbb{R}$. At left, we emphasize the difference between the cover element U_σ, and the open set in the Alexandrov topology, S_σ, which is a discrete set of three objects. At right, we show the 1-thickenings S_{σ_i} and S_{τ_i} for the two types of basis open sets.

Together, we write $\mathcal{U} = \{U_{\sigma_i}\} \cup \{U_{\tau_i}\}$ and note that this is a poset $(\mathcal{U}, \subseteq)$ given by inclusions. We write $\rho \in K$ and similarly $U_\rho \in \mathcal{U}$ when we want to talk about cells of either dimension.

We define open sets in this poset using the Alexandrov topology, following [3] and [21]. Given this poset $(\mathcal{U}, \subseteq)$, for any set $S \subseteq \mathcal{U}$, the up-set is $S^\uparrow = \{U \in \mathcal{U} \mid V \subseteq U, V \in S\}$ and the down-set is $S^\downarrow = \{U \mid U \subseteq V, V \in S\}$. A set $S \subset \mathcal{U}$ is said to be open iff $S^\downarrow = S$; equivalently if for any $U \in S$, and any $V \subset U$, $V \in S$. It can be checked that this topology has a basis given by the set of downsets: write $S_\rho = \{U_\rho\}^\downarrow$, then the basis for $\mathbf{Open}(\mathcal{U})$ is given by $\{S_\rho^\downarrow \mid \rho \in K\}$ and we call S_ρ a basic open set. See the left of Figure 4 for a visualization of this notation.

We will often use these combinatorial subsets of $\mathcal{U}$ interchangeably with their geometric realization, given by $|S| := \bigcup_{U \in S} U \subset \mathbb{R}$. We can check that this notation is then reasonable on our set. First, for $\sigma_i = i\delta$ and $U_{\sigma_i} = ((i-1)\delta, (i+1)\delta) \subset \mathbb{R}$, then we have $\{U_{\sigma_i}\}^\downarrow = \{U_{\tau_i}, U_{\sigma_i}, U_{\tau_{i+1}}\}$ and $|\{U_{\sigma_i}\}^\downarrow| = U_{\sigma_i}$. Similarly, for $U_{\tau_i} = (i\delta, (i+1)\delta)$, $\{U_{\tau_i}\}^\downarrow = \{U_{\tau_i}\}$, and of course $|\{U_{\tau_i}\}^\downarrow| = U_{\tau_i}$. See Figure 4 for a visual of the notation.

We denote the functor returning the set of connected components of a given topological space by $\pi_0 : \mathbf{Top} \to \mathbf{Set}$. Then the mapper cosheaf of $(\mathbb{X}, f)$ at resolution δ is given by the functor

$$F : \begin{array}{ccc} \mathbf{Open}(\mathcal{U}) & \to & \mathbf{Set} \\ S & \mapsto & \pi_0(f^{-1}(|S|)) \end{array}$$

By functoriality of π_0, for open sets $S \subseteq T$, F induces a map

$$F[S \subseteq T] : \pi_0 f^{-1}(|S|) \to \pi_0 f^{-1}(|T|),$$

so that F satisfies the requirements of being a functor. Furthermore, this functor satisfies the definitions a cosheaf, which we leave as an exercise to the reader. Thus, we call any functor of the form $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ a *mapper cosheaf*. For notational simplicity, we write the induced map as $F[\subseteq] : F(S) \to F(T)$ when $S \subseteq T$ is clear from context.

We next show that we can define an interleaving on these mapper cosheaves. Given any open set $S \in \mathbf{Open}(\mathcal{U})$, we construct the *1-thickening* by taking the downset of the upset of S, denoted by $S^1 = (S^\uparrow)^\downarrow$. The *n-thickening* is defined to be n repetitions of this process, given recursively as

$$S^n = \begin{cases} S & n = 0 \\ (S^{n-1})^{\uparrow\downarrow} & n \geq 1. \end{cases}$$

9

Each S^n is an element of $\mathbf{Open}(\mathcal{U})$, if $S \subseteq T$, then $S^n \subseteq T^n$. Thus we can view the thickening operation as a functor $(-)^n : \mathbf{Open}(\mathcal{U}) \to \mathbf{Open}(\mathcal{U})$. It is easy to check that for a basic open S_{σ_i}, $|S_{\sigma_i}^n| = ((i-n)\delta, (i+n)\delta)$ and for S_{τ_i}, $|S_{\tau_i}^n| = ((i-n)\delta, (i+1+n)\delta)$. See the right of Fig. 4 for an example of 1-thickenings of basic opens.

As shown in previous work [21], the thickening can be used to build an interleaving distance which compares mapper cosheaf functors $F, G : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$. We use F^n to denote the composition of functors $F \circ (-)^n : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$, so that $F^n(S) = F(S^n)$, and similarly for for G^n. We define an n-interleaving to be a pair of natural transformations $\varphi : F \Rightarrow G^n$ and $\psi : G \Rightarrow F^n$ given by set maps $\varphi_S : F(S) \to G(S^n)$. We note the map at S^n, $\varphi_{S^n} : F(S^n) \to G(S^{2n})$, can also be viewed as constituent maps of a natural transformation $\varphi^n : F^n \Rightarrow G^{2n}$. Therefore, when φ is indeed a natural transformation, we use the notation φ_{S^n} and φ_S^n interchangeably.

Definition 2.1 *Let $F, G : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ be cosheaves and $n \in \mathbb{Z}_{\geq 0}$. An n-interleaving is a pair of natural transformations $\varphi : F \Rightarrow G^n$ and $\psi : G \Rightarrow F^n$ such that the diagrams*

$$
\begin{array}{ccc}
F(S) & \xrightarrow{\;F[S \subseteq S^{2n}]\;} & F(S^{2n}) \\
\varphi_S \searrow & & \nearrow \psi_{S^n} \\
& G(S^n) &
\end{array}
\qquad\qquad
\begin{array}{ccc}
& F(S^n) & \\
\psi_S \nearrow & & \searrow \varphi_{S^n} \\
G(S) & \xrightarrow[\;G[S \subseteq S^{2n}]\;]{} & G(S^{2n})
\end{array}
$$

commute for all $S \in \mathbf{Open}(\mathcal{U})$. The interleaving distance *is given by*

$$d_I(F, G) = \inf\{n \geq 0 \mid \text{there exists an } n\text{-interleaving}\},$$

and is set to be $d(F, G) = \infty$ if there is no interleaving for any n.

As shown in [21], this definition fits in the framework built by Bubenik et al. [17] and thus it is an extended pseudometric.

2.3 Loss Function Definition

In this section, we outline the loss function used in this paper, which will largely follow [21] with some necessary technical modifications noted as they are given. First, we have the idea of an assignment, inspired by [53] and [46], which gives names to collections of morphisms of the structure used to define an interleaving without guaranteeing these morphisms define a natural transformation.

Definition 2.2 (Modified from [21]) *Given functors $H, H' : \mathbf{Open}(\mathcal{U}) \to Set$, an* unnatural transformation[1] *$\eta : H \to H'$ is a collection of maps $\eta_S : H(S) \to H'(S)$ with no additional promise of commutativity. For a fixed $n \geq 0$ and cosheaves F and G, an* assignment, *or more specifically an n-assignment, is a pair of unnatural transformations $\varphi : F \Rightarrow G^n$ and $\psi : G \Rightarrow F^n$.*

An (extended) basis unnatural transformation *$\eta : H \rightsquigarrow H'$ for mapper cosheaves H and H' is a collection of maps*

$$\eta = \{\eta_{S_\rho} : H(S_\rho) \to H'(S_\rho^n) \mid \rho \in K\} \cup \{\eta_{S_\rho^n} : H(S_\rho^n) \to H'(S_\rho^{2n}) \mid \rho \in K\}$$

for all basis elements S_ρ from $\rho \in K$. A basis n-assignment *(or simply a basis assignment) is a pair of extended basis unnatural transformations*

$$\varphi : F \rightsquigarrow G^n \qquad and \qquad \psi : G \rightsquigarrow F^n$$

---

[1]A natural transformation is an unnatural transformation which obeys commutativity properties, so the two sets are not mutually exclusive.

The modification presented here from [21] is to include the maps $\eta_{S_\sigma^n}$ along with η_{S_σ} as part of the basis n-assignment. In the previous paper, we used the fact that the maps η_{S_σ} could be used to determine $\eta_{S_\sigma^n}$. However, in this paper, the matrix representations introduced in Section 3 do not allow for easy computation of this extension. Thus, we have modified the definition so that the maps $H(S_\sigma^n) \to H(S_\sigma^{2n})$ are also given as input. For the remainder of the paper, we use the term "basis unnatural transformation" to imply the extended definition given here.

We will be taking these basis unnatural transformations and extending them to natural transformations after a shift using a value k.

Definition 2.3 *Given a basis unnatural transformation $\varphi : F \Rightarrow G^n$ and a fixed k, define the basis unnatural transformation $\Phi : F \Rightarrow G^{n+k}$ by $\Phi_S = F[S^n \subseteq S^{n+k}] \circ \varphi_S$ for*

$$S \in \{S_\sigma \mid \sigma \in K\} \cup \{S_\sigma^n \mid \sigma \in K\}.$$

We adopt the following notation in this paper: when an assignment might not commute, we denote its maps by lower case $\varphi : F \rightsquigarrow G^n$ and $\psi : G \rightsquigarrow F^n$ with squiggly arrows; for assignments which are constructed to be natural transformations, we denote them by $\Phi : F \Rightarrow G^n$ and $\Psi : G \Rightarrow F^n$ with double arrows.

We measure the quality of an n-assignment φ, ψ using collections of distances for all open sets, $\{d_S^F \mid S \in \mathbf{Open}(\mathcal{U})\}$ and $\{d_S^G \mid S \in \mathbf{Open}(\mathcal{U})\}$, defined as follows.

Definition 2.4 *The distance $d_S^F(A, B)$ for $A, B \in F(S)$ is defined to be*

$$d_S^F(A, B) = \min\{k \geq 0 \mid F[S \subset S^k](A) = F[S \subset S^k](B)\}. \tag{1}$$

If no such k exists, then we set $d_S^F(A, B) = \infty$.

Note that geometrically, this can be thought of as the smallest integer n such that the connected components represented by A and B in $f^{-1}(|S|)$ are in the same connected component when included into $f^{-1}(|S^n|)$. Further, we can see that this distance is an ultrametric since if for $A, B, C \in F(S)$, $k = \max\{d_S^F(A, B), d_S^F(B, C)\}$, then A, B, and C all map to the same object in $F(S^k)$, and thus $d_S^F(A, C) \leq k$.

If φ defined a natural transformation, we would have that $G[\subseteq] \circ \varphi_S = \varphi_T \circ F[\subseteq]$ for all S. Since we do not have that assumption, we look instead to ensuring k is large enough that the compositions from $F(S)$ to $G(T^{n+k})$ in the diagram

$$
\begin{array}{ccc}
F(S) & \xrightarrow{\;\;F[\subseteq]\;\;} & F(T) \\
\;\;\Big\downarrow{\scriptstyle\varphi_S} & & \;\;\Big\downarrow{\scriptstyle\varphi_T} \\
G(S^n) & \xrightarrow[G[\subseteq]]{} G(T^n) & \longrightarrow G(T^{n+k})
\end{array}
\tag{2}
$$

are equal so that φ can be extended to an $n + k$ interleaving Φ as given in Defn. 2.3. With the distance definition, this is

$$k \geq \max_{\alpha \in F(S)} d_{T^n}^G(\varphi_S \circ F[\subseteq](\alpha), G[\subseteq] \circ \varphi_S(\alpha)).$$

If we also want to ensure that Φ and Ψ constitute an interleaving, we need to have that the two paths in

$$
\begin{array}{ccc}
F(S) & \xrightarrow{\;\;\;\;F[\subseteq]\;\;\;\;} & F(S^{2n}) \longrightarrow F(S^{2(n+k)}) \\
\;\;\searrow{\scriptstyle\varphi_S} & \nearrow{\scriptstyle\psi_{S^n}} & \\
& G(S^n) &
\end{array}
\tag{3}
$$

11

from $F(S)$ to $F(S^{2(n+k)})$ are the same, which is equivalent to

$$k \geq \max_{\alpha \in F(S)} \left\lceil \tfrac{1}{2} \cdot d_{S^{2n}}^F (F[\subseteq](\alpha), \psi_{S^n} \circ \varphi_S(\alpha)) \right\rceil.$$

We give the loss function for each diagram in the following setup.

Definition 2.5 *Fix an n-assignment* (φ, ψ). *We define four* diagram loss functions:

$$L_{\bigcirc}^{S,T}(\varphi) = \max_{\alpha \in F(S)} d_{T^n}^G (\varphi_S^n \circ F[S \subseteq T](\alpha), G[S^n \subseteq T^n] \circ \varphi_S(\alpha))$$

$$L_{\square}^{S,T}(\psi) = \max_{\alpha \in G(S)} d_{T^n}^F (\psi_S^n \circ G[S \subseteq T](\alpha), F[S^n \subseteq T^n] \circ \psi_S(\alpha))$$

$$L_{\bigtriangledown}^{S}(\varphi, \psi) = \max_{\alpha \in F(S)} \left\lceil \tfrac{1}{2} \cdot d_{S^{2n}}^F (F[S \subseteq S^{2n}](\alpha), \psi_{S^n} \circ \varphi_S(\alpha)) \right\rceil$$

$$L_{\bigtriangleup}^{S}(\varphi, \psi) = \max_{\alpha \in G(S)} \left\lceil \tfrac{1}{2} \cdot d_{S^{2n}}^G (G[S \subseteq S^{2n}](\alpha), \varphi_{S^n} \circ \psi_S(\alpha)) \right\rceil.$$

Putting these ideas together, we have the following definition for the loss function between given n-assignments.

Definition 2.6 *Fix an n-assignment* (φ, ψ). *The* (extended basis) loss function *is defined to be*

$$L_B(\varphi, \psi) = \max_{\substack{\sigma < \tau \in K \\ \rho \in K}} \left\{ L_{\bigcirc}^{S_\tau, S_\sigma}, L_{\square}^{S_\tau, S_\sigma}, L_{\bigcirc}^{S_\rho, S_\rho^n}, L_{\square}^{S_\rho, S_\rho^n}, L_{\bigtriangleup}^{S_\rho}, L_{\bigtriangledown}^{S_\rho} \right\} \tag{4}$$

where the max is taken over all cells $\rho \in K$ *of any dimension; or of pairs of an edge* τ *adjacent to a vertex* σ.

Again, note the slight modification from the definition given in [21] where two additional terms, $L_{\bigcirc}^{S_\rho, S_\rho^n}$ and $L_{\square}^{S_\rho, S_\rho^n}$, are included in this new loss function. Because the extended basis unnatural transformation provides the maps $\varphi_{S_\sigma^n} : F(S_\sigma^n) \to G(S_\sigma^{2n})$ as inputs rather than being determined from the maps φ_{S_σ}, these terms are used to ensure that the resulting φ will be a true natural transformation. Since we are taking a maximum over more inputs, it is an immediate corollary of Thm. 3.16 in [21] that this function gives a bound on the true interleaving distance between F and G.

Theorem 2.7 (Cor. of [21, Thm 3.16]) *Given an (extended) basis n-assignment* φ *and* ψ, *we have*

$$d_I(F, G) \leq n + L_B(\varphi, \psi).$$

3 Data Structures

In this section, we show how a given mapper cosheaf $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ can be stored as a graph, and then how given unnatural transformations φ, ψ can be stored as binary valued block matrices. Finally, we show that the entries of the basis loss function are represented by matrix multiplication.

Figure 5: Example portion of a mapper cosheaf graph representation. For example, the vertices from $F(S_{\sigma_i})$ are b and c; while the edges from $F(S_{\tau_i})$ are bd and cf.

3.1 Representing mapper cosheaves as graphs

Given a mapper cosheaf $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$, we will build a graph (V_F, E_F) as follows. Recall that K, the one-dimensional cubical complex discretizing $\mathbb{R}$, consists of vertices $\sigma_{-L}, \cdots, \sigma_L$ with heights contained in the bounding box $[-L\delta, L\delta]$, along with edges $\tau_j = (\sigma_j, \sigma_{j+1})$. The graph for the mapper cosheaf $F : \mathbf{Open}(\mathcal{U}) \to \mathbf{Set}$ is constructed by generating a vertex for every object in every $F(S_{\sigma_i})$, an edge for every object in every $F(S_{\tau_i})$ and connecting them using the constituent morphisms of F. This resulting vertex and edge sets are $V_F = \coprod_{i=1}^{B} F(S_{\sigma_i})$ and $E_F = \coprod_{i=1}^{B-1} F(S_{\tau_i})$, noting the difference in index sets is a result of having an edge for every object in every $F(S_{\tau_i})$, of which there is one fewer by construction. The endpoints of any edge $e \in F(S_{\tau_i}) \subseteq E_F$ can be found via the attaching maps:

$$F[S_{\tau_i} \subseteq S_{\sigma_i}](e) \in F(S_{\sigma_i}) \text{ and}$$
$$F[S_{\tau_i} \subseteq S_{\sigma_{i+1}}](e) \in F(S_{\sigma_{i+1}}).$$

For example, $(c, f) \in F(S_{\tau_i})$ in Figure 5 has endpoints $c \in F(U_{\sigma_i})$ and $f \in F(U_{\sigma_{i+1}})$. This data is stored in a standard adjacency list, noting that each vertex $v \in F(U_{\sigma_i})$ stores the associated value i as a representation of its height.

We next use F to construct graph representations for both the functors F^n and F^{2n}. See Figure 6 for a figure to follow along. Since building F^{2n} from F^n is the same as building F^n from F up to indexing, we only focus the construction of F^n. For any $\rho \in K$, $F^n(S_\rho) = F(S_\rho^n)$, and because of the cosheaf properties, this means

$$F^n(S_\rho) = \operatorname*{colim}_{S \in S_\rho^n} F(S) = \operatorname*{colim}_{S_{\rho'} \in S_\rho^n} F(S_{\rho'}). \tag{5}$$

If ρ is a vertex σ_i, then checking indices shows that the colimit is taken over the set

$$\{\rho' \in K \mid S_{\rho'} \in S_{\sigma_i}^n\} = \{\sigma_j \mid j \in [i-n, i+n]\} \cup \{\tau_j \mid j \in [i-n-1, i+n]\}. \tag{6}$$

The result is that $F^n(S_{\sigma_i})$ is the set of connected components of the graph induced by the vertex set

$$\coprod_{j \in [i-n, i+n]} F(S_{\sigma_j}) \subseteq V_F. \tag{7}$$

Each connected component results in a vertex in the graph representation of F^n. Note that this induced graph does not include any edges representing elements of $F(S_{\tau_{i-n-1}})$ or $F(S_{\tau_{i+n}})$ although
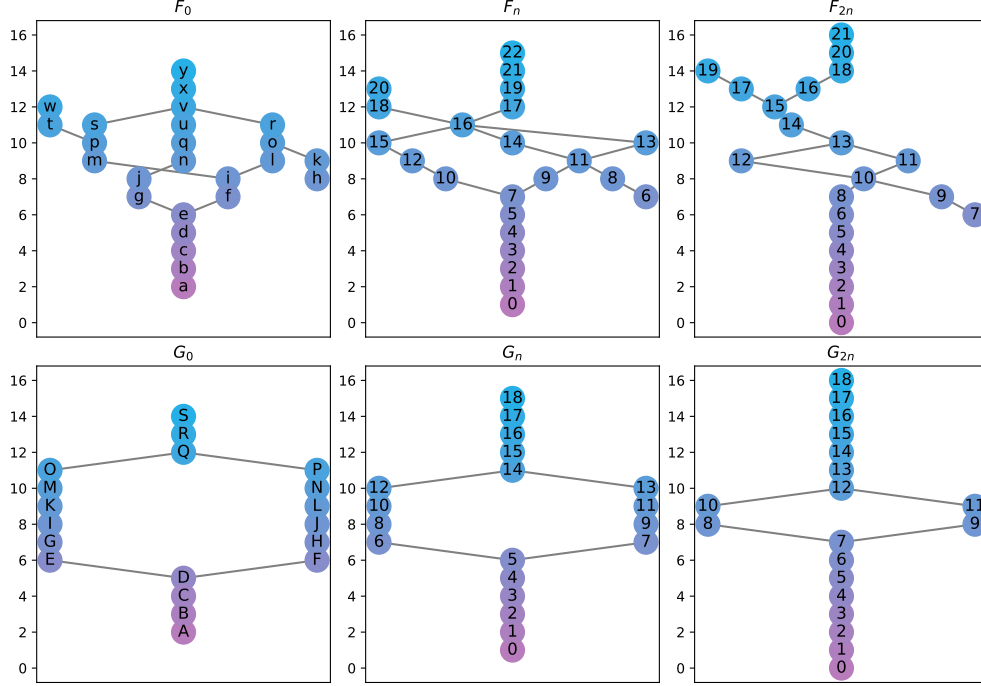
13

Figure 6: An example pair of input mapper cosheaf graphs, F and G, are shown in the left column. The middle and right columns give the n and $2n$ smoothing of these graphs, respectively. Note that vertex labels have no correspondence from one graph to the next; instead, they are an artifact of the construction methods for the later graphs. These examples correspond to the matrix examples shown in Figures 7 to 10.

these are both sets involved in the colimit of Equation (5). However, since these "half edge" elements have maps only to $F(S_{\sigma_{i-n}})$ and $F(S_{\sigma_{i+n}})$ respectively, dropping them does not change the resulting colimit.

Similarly, for an edge $\tau_i \in K$, the colimit is taken over the set

$$\{\rho' \in K \mid S_{\rho'} \in S_{\tau_i}^n\} = \{\sigma_j \mid j \in [i-n+1, i+n]\} \cup \{\tau_j \mid j \in [i-n, i+n]\}.$$

and thus is the connected components of the graph induced by vertex set

$$\coprod_{j \in [i-n+1, i+n]} F(S_{\sigma_j}) \subseteq V_F. \tag{8}$$

Each connected component of this graph gives an edge in the graph representation of F^n. The endpoints of the edges are determined by the induced map on colimits since $S_{\tau_i}^n \subseteq S_{\sigma_i}^n$ and $S_{\tau_i}^n \subseteq S_{\sigma_{i+1}}^n$.

Note that since connected components of a graph (V, E) can be determined in $O(|V| + |E|)$ time using breadth first search, F^n can be built from F in $O(L \cdot (|V_F| + |E_F|))$ time. Further, during this construction the natural transformation $F \to F^n$ can be stored at the same time; see Section 3.2.2 for more details on this. Similarly, F^{2n} can be built from F, so getting all four additional graphs from F and G can be done in $O(L \cdot (|V_F| + |E_F|))$ time.

Given a graph representation of one of the three functors F, F^n or F^{2n}, we need to determine the distance $d_{S_{\sigma_i}}(v, v')$ between any two vertices, as defined in Equation (1). Assume the graph representation for mapper cosheaf $H \in \{F, F^n, F^{2n}\}$ is given by vertex set V and edge set E, which

14

the stored information that each vertex $v \in V$ is associated to a σ_i and each edge $e \in E$ is associated to a τ_i. Let $v, v' \in H(S_{\sigma_i})$ be given. Ensuring that $H[S_{\sigma_i} \subseteq S_{\sigma_i}^k](v) = H[S_{\sigma_i} \subseteq S_{\sigma_i}^k](v')$ is equivalent to making sure that v and v' are in the same connected component of the graph induced by the vertices in $S_{\sigma_i}^k$. This amounts to finding the minimum height path between the two vertices, which can be done via a BFS traversal in $O(|V_F| + |E_F|)$ time.

3.2 Representing maps and relationships as block matrices

Our next job is to represent the information of the relevant (un)natural transformations. First, recall that an (un)natural transformation $\eta : H \to H'$ is a collection of set maps $\eta_S : H(S) \to H'(S)$ for $S \in \{S_\rho \mid \rho \in K\} \cup \{S_\rho^n \mid \rho \in K\}$. Denote the vertex and edge sets for the representations as (V_H, E_H) and $(V_{H'}, E_{H'})$. In particular, this means that each set map sends vertices (resp. edges) representing elements of $H(S_{\sigma_i})$ (resp. $H(S_{\tau_i})$) to vertices (resp. edges) with the same associated element of K, namely $H'(S_{\sigma_i})$ (resp. $H'(S_{\tau_i})$). Thus, we fix a total order on the vertices of each graph such that if v is associated to σ_i and v' is associated to σ_j with $i < j$, then $v < v'$. A similar procedure is done for the edges. Then we store η as a pair of block matrices, $M_V \in \{0,1\}^{|V_{H'}| \times |V_H|}$ and $M_E \in \{0,1\}^{|E_{H'}| \times |E_H|}$ with entry associated to pair (a, a') equal to 1 iff $\eta_{S_\rho}(a) = a'$ for the relevant choice of $\rho \in K$.

Note this means that all non-zero entries are in the blocks given by label sets $H'(S_\rho) \times H(S_\rho)$. Further, every column has exactly one entry of 1 and the rest are 0. We next go through constructing all the relevant matrices used in the loss function computation.

3.2.1 Boundary Matrices

The boundary matrices describe the relationship between the edges and vertices of a graph representation of a given mapper cosheaf. For a mapper cosheaf $H \in \{F, F^n, G, G^n\}$, this is a matrix $B_H \in \{0,1\}^{|E_H| \times |V_H|}$, with a 1 in entry $B_H[e, v]$ iff v is an endpoint of e. Unlike the other matrices used, this matrix has two non-zero entries per column. However, we can split it up into two single-entry column, block matrices as follows. Let $B_H^{\uparrow} \in \{0,1\}^{|E_H| \times |V_H|}$ have a 1 in entry $B[e, v]$ iff v is an endpoint of e for $v \in F(\sigma_i)$ and $e \in F(\tau_i)$; while $B_H^{\downarrow} \in \{0,1\}^{|E_H| \times |V_H|}$ have a 1 in entry $B[e, v]$ iff v is an endpoint of e for $v \in F(\sigma_i)$ and $e \in F(\tau_{i-1})$. Then $B_H = B_H^{\uparrow} + B_H^{\downarrow}$. See Figure 7 for an example. In this case, we will have eight matrices to work with, which we denote by $B_F^{\uparrow}$, $B_F^{\downarrow}$, $B_{F^n}^{\uparrow}$, $B_{F^n}^{\downarrow}$, $B_G^{\uparrow}$, and B_{G^n}. Note that we do not compute the boundary matrices for either F^{2n} or G^{2n} since, as can be seen in Table 1, these matrices are never used. Because these are simply matrix encodings of the adjacency information of the relevant graphs, they can be constructed in $O(V_{max} \cdot E_{max})$ time during construction, where $V_{max} = \max\{|V_H| \mid H \in F, F^n, G, G^n\}$ and $E_{max} = \max\{|E_H| \mid H \in F, F^n, G, G^n\}$.

3.2.2 Inclusion matrices

The inclusion matrices represent the natural transformations $F \Rightarrow F_n \Rightarrow F_{2n}$ and $G \Rightarrow G_n \Rightarrow G_{2n}$. Focusing on the first of these, $F \Rightarrow F^n$, the component morphisms take a vertex $v \in F(S_{\sigma_i}))$ to the vertex $v' \in F^n(S_{\sigma_i})$ representing the connected component containing v'. The edge maps are determined similarly. Thus the vertex inclusion matrix is written as $I_F^V \in \{0,1\}^{|V_{F^n}| \times |V_F|}$ with a 1 in entry $I_F[v', v]$ if v is in the connected component represented by v'. A similar matrix is built for the edges, $I_F^E \in \{0,1\}^{|E_{F^n}| \times |E_F|}$. The result is that each inclusion natural transformation is given by a pair of matrices, so from this construction we have the matrices $I_F^V, I_F^E, I_{F^n}^V, I_{F^n}^E, I_G^V, I_G^E, I_{G^n}^V$, $I_{G^n}^E$. See Figure 8 for an example. Note that these matrices all satisfy the block decomposition and
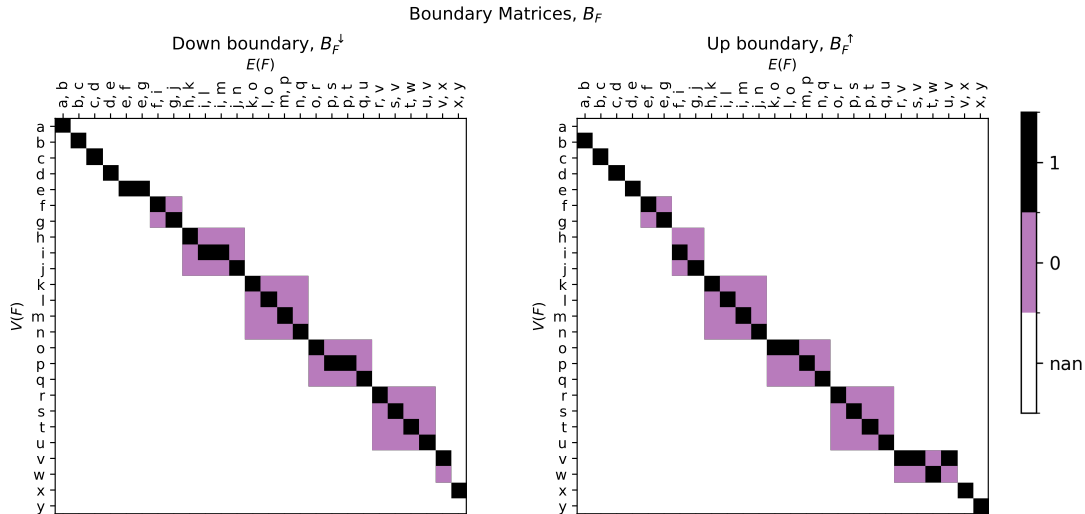
Figure 7: Example boundary matrices of the graph of F shown in Figure 6

have exactly one non-zero entry in each column. Also, because they track the representatives of objects in the n-smoothed connected components, they can likewise be built during the construction of the smoothed graphs, with only constant overhead.

3.2.3 Distance matrices

The next collection of matrices give the distance from Eqn. (1). In particular, for a pair of vertices $v, w \in F(S_\rho)$, we want to determine the minimum k such that they map to the same object in $F(S_\rho^k)$. Equivalently, this is the minimum k such that v and w are in the same connected component of the subgraph induced by the vertex set $\coprod_{S_{\sigma_i} \subseteq S_\rho^k} F(S_{\sigma_i})$. For example, in the graph of F from Figure 9, vertices 15 and 22, which are both at level σ_9, are not in the same connected component until $k = 3$ using the path that goes down to vertex 8 and back up; or alternatively the path that goes up to 1 and back down. Similarly, edges $(3, 15)$ and $(22, 23)$ have distance 3 since they are connected via the path through vertex 1.

This distance is stored as a pair of matrices $D_F^V \in \{0,1\}^{|V(F)|^2}$ and $D_F^E \in \{0,1\}^{|E(F)|^2}$. For the vertices, $D_F^V[u, v] = d_{S_{\sigma_i}}(u, v)$ if $u, v \in F(S_{\sigma_i}) \subset V_F$, and is 0 otherwise. The matrix D_F^E is defined similarly for the elements of $F(S_{\tau_i})$. As with previously mentioned matrices, these matrices have the block structure following the heights in the mapper cosheaf structure. See Figure 9 for examples corresponding to F from Figure 6. The result is that we have 12 distance matrices, two each for the 6 input matrices. However, as we can see from Table 1, only the eight for F^n, F^{2n}, G^n, and G^{2n} need to be calculated.

These distance matrices are calculated from the graph representations as follows. For each level $i \in [-L, L]$, we construct the block for σ_i in D_F^V and the block for τ_i in D_F^E. For the vertex version, we build a union find data structure by adding the edges in order of the thickening. Recall (and see Figure 4 for a reminder of the notation), that the elements of $F(S_{\sigma_i}^k)$ correspond to the subgraph induced by the vertices in $\coprod_{j \in [i-n, i+n]} F(S_{\sigma_j})$, which includes the edges in $\coprod_{j \in [i-k, i+k-1]} F(S_{\tau_j})$. So for $k = 1$, we add all the edges for τ_{i-1} and τ_i, and check to see if any of the vertices in $F(S_{\sigma_i}^k)$ are now in the same connected component. For any pair u and v that are, we set $D_F^V[u, v] = 1$. Then for $k = 2$, we add the edges for τ_{i-2} and τ_{i+2}, and repeat the same check for pairs that were not already in the same connected component at $k = 1$, setting these distances to 2 if they are now in the same
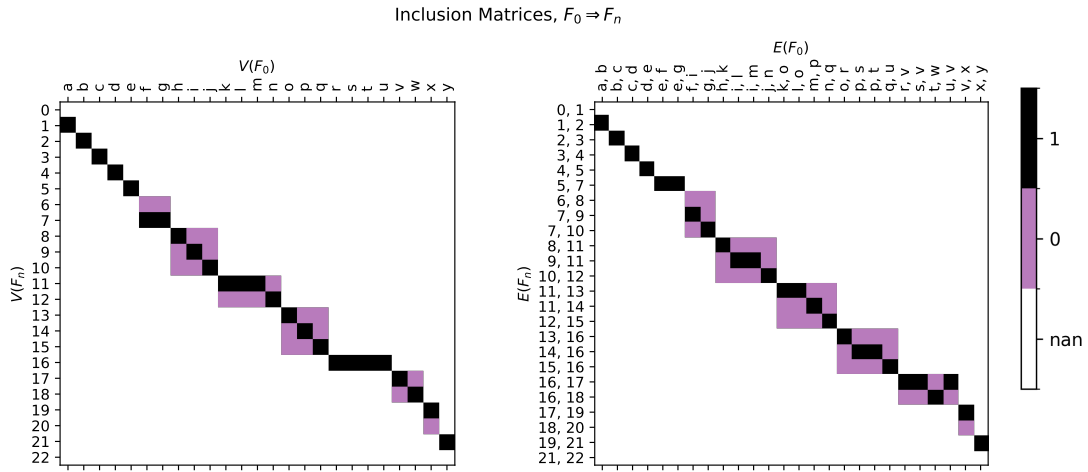
Figure 8: The inclusion matrices I_F^V (left) and I_F^E (right) giving the inclusion natural transformation $F \Rightarrow F_n$ for the example shown in Figure 6
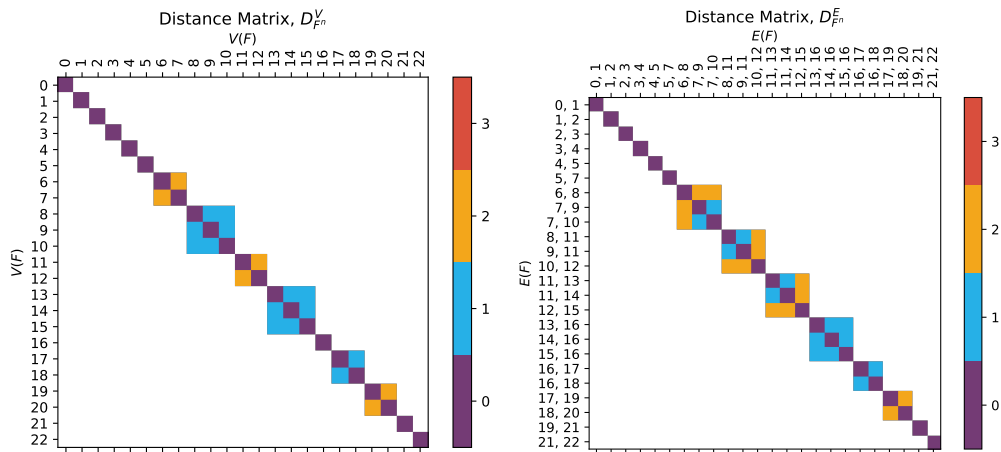

Figure 9: Distance matrix example for F^n from Figure 6

component. We continue until all pair distances are determined. A similar sweep for each level is used to construct D_F^E, with the difference being that the edges added at step k are those from $F(\tau_{i-k+1}, \tau_{i+k-1})$. We will have V_{max} entries in our union-find data structure, so amortized time per operation is $O(\alpha(V_{max}))$, with a total of $E_{\max}$ operations. Thus, since the union find structure is built at each level, the pair of matrices for a given graph takes time $O(L \cdot E_{max} \cdot \alpha(V_{max}))$, where $\alpha(\cdot)$ is the inverse Ackermann function.

3.2.4 Assignment matrices

The assignment matrices encode the given n-assignment φ and ψ. We split these representation into the portions defined on S_ρ and S_ρ^n; and then split these into the portions defined on vertices $\rho = \sigma_i$, and the portions defined on edges $\rho = \tau_i$. The four resulting matrices for φ are denoted

$$M_\varphi^V \in \{0,1\}^{|V(G^n)| \times |V(F)|}, \quad M_{\varphi_n}^V \in \{0,1\}^{|V(G^{2n})| \times |V(F^n)|},$$
$$M_\varphi^E \in \{0,1\}^{|E(G^n)| \times |E(F)|}, \quad M_{\varphi_n}^E \in \{0,1\}^{|E(G^{2n})| \times |E(F^n)|}.$$

the four matrices for ψ are denoted similarly.

As with other matrices, these are block matrices aligned with the σ_i and τ_i indices, and have a single one in every column. Matrix M_φ^V has $M_\varphi^V[u,v] = 1$ iff $\varphi_{S_{\sigma_i}}(u) = v$ for $u \in F(S_{\sigma_i})$ and $v \in G(S_{\sigma_i}^n)$. See Figure 10 for an example of the matrices for an input φ. In this particular case, the φ was randomly generated to place a 1 in each column. Of course, this is an unnatural transformation since the relevant parallelogram diagram does not commute for this example. In what follows, we will use this representation in the ILP formulation to find a good choice of unnatural transformations; thus we need only initialize empty matrices in setup.

3.3 Representing each term in L_B by matrix multiplication

To convert the problem as an ILP, we first show how the entries in the loss function

$$L_B(\varphi, \psi) = \max_{\substack{\sigma < \tau \in K \\ \rho \in K}} \left\{ L_{\seardiagdown}^{S_\tau, S_\sigma}, L_{\diagdown}^{S_\tau, S_\sigma}, L_{\seardiagdown}^{S_\rho, S_\rho^n}, L_{\diagdown}^{S_\rho, S_\rho^n}, L_{\bigtriangleup}^{S_\rho}, L_{\bigtriangledown}^{S_\rho} \right\}$$

correspond to matrix multiplication in this setting. Due to symmetry, we limit our discussion to $L_{\seardiagdown}^{S_\tau, S_\sigma}$, $L_{\seardiagdown}^{S_\rho, S_\rho^n}$, and $L_{\bigtriangledown}^{S_\sigma}$, however the full list can be found in Table 1.

First, consider the entry for $L_{\seardiagdown}^{S_\tau, S_\sigma}$. Further, assume for the moment that σ is the lower vertex of τ, so $\sigma = \sigma_i$ and $\tau = \tau_i$. See the top row of Figure 11 for examples of matrices used in the following proof.

Lemma 3.1 $\max_i L_{\seardiagdown}^{S_{\tau_i}, S_{\sigma_i}} = \max \left\{ x \mid x \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\downarrow - B_{G^n}^\downarrow \cdot M_\varphi^E \right) \right\}$

Proof 3.1 *To start, fix an i and for simplicity, throughout the proof we write $\sigma = \sigma_i$ and $\tau = \tau_i$. We are looking for the distance between images landing in the bottom right of the diagram*

$$
\begin{array}{ccc}
F(S_\tau) & \xrightarrow{F[\subseteq]} & F(S_\sigma) \\
& \searrow^{\varphi_{S_\tau}} & \downarrow^{\varphi_{S_\sigma}} \\
& G^n(S_\tau) & \xrightarrow[G[\subseteq]]{} G^n(S_\sigma).
\end{array}
$$

in order to determine $L_{\seardiagdown}^{S_\tau, S_\sigma}$. Fix any element $e \in F(S_\tau)$. If it has lower vertex $v \in F(S_\sigma)$ with $\varphi_{S_\sigma}(v) = v'$, then the composition along the top of the diagram sends $e \mapsto v'$. This means that the

Figure 10: All the matrices for a random input unnatural transformation φ for the example graphs shown in Figure 6.

column for e in $M_\varphi^V \cdot B_F^\downarrow$ has exactly one 1 in the row corresponding to v'. If $\varphi_{S_\tau}(e) = e' \in G^n(S_\tau)$ which has lower vertex v'', then the composition along the bottom of the diagram sends $e \mapsto v''$. Similarly, the column for e in $B_{G^n}^\downarrow \cdot M_\varphi^E$ has exactly one 1 in the row corresponding to v''. Writing $A := M_\varphi^V \cdot B_F^\downarrow - B_{G^n}^\downarrow \cdot M_\varphi^E$, the column for e in A is either entirely 0 if $v' = v''$; or its only non-zero entries are $A[v', e] = 1$ and $A[v'', e] = -1$.

We can then think about left-multiplying the matrix A by the distance matrix $D := D_{G^n}^V$. Clearly, if $v' = v''$ and thus the column in A for e is 0, the column for e in $D \cdot A$ is entirely 0. If $v' \neq v''$, the entries for the column e are

$$(D \cdot A)[u, e] = \begin{cases} -d(v', v'') & \text{if } u = v' \\ d(v', v'') & \text{if } u = v'' \\ d(u, v') - d(u, v'') & \text{else} \end{cases}$$

where d denotes the distance function $d_{S_\sigma^n}^G$. However, since $d(u, v') - d(u, v'') \leq d(v', v'')$ by the triangle inequality, the maximum entry in the column occurs at entry $[v'', e]$. This means that $L_{\diagdown}^{S_{\tau_i}, S_{\sigma_i}}$ is the maximum in the columns corresponding to $e \in F(S_{\tau_i})$. Further, the maximum of $L_{\diagdown}^{S_{\tau_i}, S_{\sigma_i}}$ over all i is the maximum over all columns, so the lemma follows.

A similar argument gives the following lemma.

Lemma 3.2 $\max_i L_{\diagdown}^{S_{\tau_i}, S_{\sigma_{i+1}}} = \max \left\{ x \mid x \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right) \right\}$

Next we turn our attention to the thickening matrix type, $L_{\diagdown}^{S_\rho, S_\rho^n}$, starting with the vertex case where $\rho = \sigma_i$.

Lemma 3.3 $\max_i L_{\diagdown}^{S_{\sigma_i}, S_{\sigma_i}^n} = \max \left\{ x \mid x \in D_{G^n}^V \left(M_{\varphi^n}^V \cdot I_F^V - I_{G^n}^V \cdot M_\varphi^V \right) \right\}.$

Proof 3.2 *We are looking for the distance between images in the bottom right of*

$$F(S_{\sigma_i}) \xrightarrow{F[\subseteq]} F^n(S_{\sigma_i})$$

with $\varphi_{S_{\sigma_i}}$ and $\varphi_{S_{\sigma_i}^n}$ mapping to

$$G^n(S_{\sigma_i}) \xrightarrow{G[\subseteq]} G^{2n}(S_{\sigma_i}).$$

A similar argument to the proof of Lemma 3.1 says that if we have elements in the diagram of the form

$$v \longmapsto v'$$
with maps down to
$$w \longmapsto w''$$

then column corresponding to v in $D_{G^n}^V \left(M_\varphi^V \cdot B_F^\downarrow - B_{G^n}^\downarrow \cdot M_\varphi^E \right)$ will be entirely 0 if $w' = w''$; and will be $d_{S_{\sigma_i}}^{G^{2n}}(w', w'')$ in the entry $[w'', v]$ otherwise. The lemma follows.

Again, a parallel argument shows the edge thickening diagram contribution.

Lemma 3.4 $\max_i L_{\diagdown}^{S_{\tau_i}, S_{\tau_i}^n} = \max \left\{ x \mid x \in D_{G^n}^E \left(M_{\varphi^n}^E \cdot I_F^V - I_{G^n}^E \cdot M_\varphi^E \right) \right\}.$
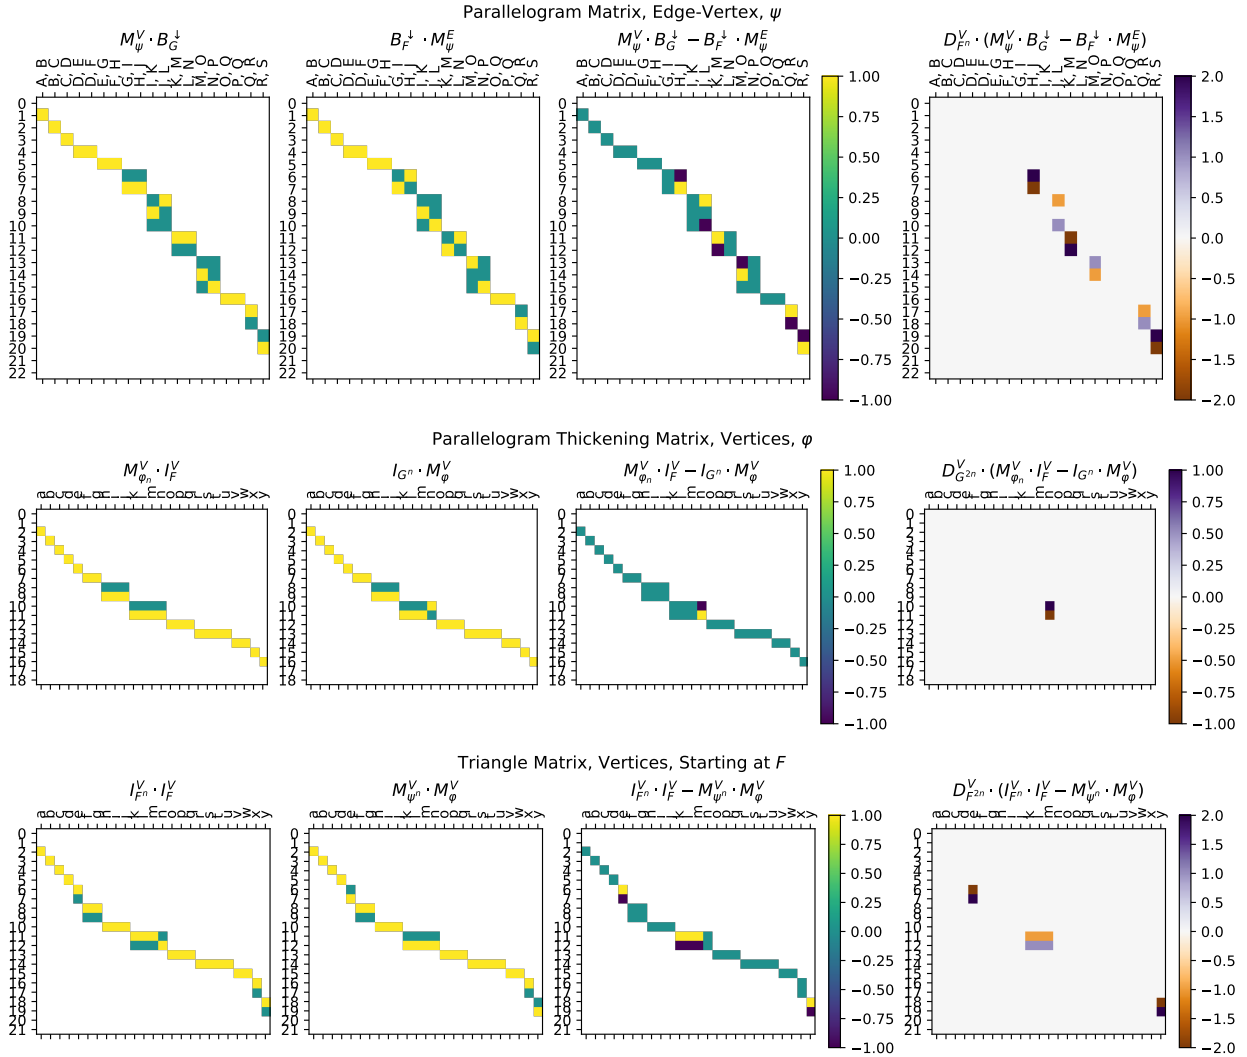
Figure 11: Examples of the matrix multiplications used to determine the value of the loss function. The top shows a case for the edge-vertex parallelogram; the middle shows a case for the parallelogram thickening; and the bottom shows a case for the triangle diagram. Note that the associated loss function is the maximum absolute value of the entries in the first two cases; and is the ceiling of half of the absolute maximum value in the last.

The last type of loss contribution is that of the triangle diagrams. The representation of the commutativity of the diagram is similar to the two previous examples. For example, for the vertex triangle loss $L_{\triangledown}^{S_{\sigma_i}}$, we are checking the diagram

$$F(S_{\sigma_i}) \xrightarrow{\quad F[\subseteq] \quad} F^{2n}(S_{\sigma_i})$$

with $\varphi_{S_{\sigma_i}}$ going down to $G^n(S_{\sigma_i})$ and $\psi_{S_{\sigma_i}^n}$ going up to $F^{2n}(S_{\sigma_i})$.

which amounts to looking for non-zero entries in $D_{F^{2n}}^V \left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V \right)$. However, the one difference in this case is that the resulting loss function is the ceiling of half of maximum, so in this case we have the following.

Lemma 3.5

$$\max_i L_{\triangledown}^{S_{\sigma_i},} = \max \left\{ \left\lceil \tfrac{x}{2} \right\rceil \mid x \in D_{F^{2n}}^V \left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V \right) \right\},$$

$$\max_i L_{\triangledown}^{S_{\tau_i},} = \max \left\{ \left\lceil \tfrac{x}{2} \right\rceil \mid x \in D_{F^{2n}}^E \left(I_{F^n}^E \cdot I_F^E - M_{\psi^n}^E \cdot M_\varphi^E \right) \right\}.$$

The result is that the loss function, Equation (4), can be computed by finding the maximum over the entries in the relevant matrices given in Table 1. Explicitly,

Theorem 3.6

$$L_B(\varphi, \psi) = \max\{$$
$$\left\{ x \mid x \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right) \right\}$$
$$\cup \left\{ x \mid x \in D_{G^n}^V \left(M_{\varphi^n}^V \cdot I_F^V - I_{G^n}^V \cdot M_\varphi^V \right) \right\}$$
$$\cup \left\{ x \mid x \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right) \right\}$$
$$\cup \left\{ x \mid x \in D_{G^n}^E \left(M_{\varphi^n}^E \cdot I_F^V - I_{G^n}^E \cdot M_\varphi^E \right) \right\}$$
$$\cup \left\{ \left\lceil \tfrac{x}{2} \right\rceil \mid x \in D_{F^{2n}}^V \left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V \right) \right\}$$
$$\cup \left\{ \left\lceil \tfrac{x}{2} \right\rceil \mid x \in D_{F^{2n}}^E \left(I_{F^n}^E \cdot I_F^E - M_{\psi^n}^E \cdot M_\varphi^E \right) \right\}$$
$$\}.$$

4 Implementation: Integer Linear Programming

Our matrix formulation above leads to the natural question: can we formulate loss function computation as a linear program? Indeed, we will show that both the objective function and constraints of the loss function optimization problem can be expressed as linear relationships. However, due to the discrete setting of our problem, our problem formulation will take the shape of an integer linear program. This comes with a potential trade-off in complexity, as integer linear programming is NP-complete, but is necessary since non-integer optima would be meaningless in our setting. Nevertheless, developing efficient LP solvers is an active area of research, resulting in many solvers that are extremely efficient in practice, which we will show work quite well on our problem in Section 4.2.

4.1 ILP model description

The goal of this optimization is to minimize the loss function. Hence, the objective function for the ILP is to minimize $L_B(\varphi, \psi)$. We denote the objective function as ℓ in the ILP. As the loss is calculated on assignments, the entries in the blocks of the eight assignment matrices M_η^A for $\eta \in \{\varphi, \varphi^n, \psi, \psi^n\}$ and $A \in \{V, E\}$ serve as decision variables. Following the discussion in Section 3.3, we know that the loss for each individual diagram can be computed by finding the maximum absolute value or the maximum of the ceiling of half the value among the entries of matrix multiplication terms listed in Table 1. Thus, we impose constraints in the ILP to connect the matrix representation of each term in $L_B(\varphi, \psi)$. We model the loss ℓ to be greater than or equal to the relevant entry in these matrices. Since it is a minimization problem, the ILP seeks the smallest ℓ that satisfies this condition, ensuring the equality. Below we describe how the three diagram types (see Table 1) contribute to the ILP formulation.

Reiterating above, our goal is

$$\text{Minimize} \quad \ell$$

with respect to the variables $\{x_{ij}^{\eta,A} \in M_\eta^A \mid \eta \in \{\varphi, \varphi^n, \psi, \psi^n\}, A \in \{V, E\}\}$ restricted to those entries in the blocks of the relevant matrices (we can implicitly assume $x_{ij}^{\eta,A} = 0$ outside of these blocks). To ensure that the resulting matrices have a single one in each column, we include the first collection of constraints,

$$\text{Subject to} \quad \sum_i x_{ij}^{\eta,A} = 1 \quad \forall x_{ij}^{\eta,A} \in M_\eta^A, \eta \in \{\varphi, \varphi^n, \psi, \psi^n\}, A \in \{V, E\},$$

which ensure that all columns sum to 1.

The next set of constraints are associated to the edge-vertex parallelograms, where the corresponding diagrams are $L_{\searrow}^{S_\tau, S_\sigma}$ and $L_{\swarrow}^{S_\tau, S_\sigma}$. For example, in the diagram $L_{\searrow}^{S_\tau, S_\sigma}$, the loss is given by

$$L_{\searrow}^{S_\tau, S_\sigma} = \max \left\{ x \mid x \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right) \right\}.$$

Therefore, we set the constraint that the objective function ℓ must be larger than any entry in the multiplied matrix. Thus, we include the constraints

$$
\begin{aligned}
\text{Subject to} \quad & \ell \geq x_{ij}^{F,\uparrow} \quad \forall \quad x_{ij}^{F,\uparrow} \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\uparrow - B_{G^n}^\uparrow \cdot M_\varphi^E \right) \\
& \ell \geq x_{ij}^{F,\downarrow} \quad \forall \quad x_{ij}^{F,\downarrow} \in D_{G^n}^V \left(M_\varphi^V \cdot B_F^\downarrow - B_{G^n}^\downarrow \cdot M_\varphi^E \right) \\
& \ell \geq x_{ij}^{G,\uparrow} \quad \forall \quad x_{ij}^{G,\uparrow} \in D_{F^n}^V \left(M_\psi^V \cdot B_G^\uparrow - B_{F^n}^\uparrow \cdot M_\psi^E \right) \\
& \ell \geq x_{ij}^{G,\downarrow} \quad \forall \quad x_{ij}^{G,\downarrow} \in D_{F^n}^V \left(M_\psi^V \cdot B_G^\downarrow - B_{F^n}^\downarrow \cdot M_\psi^E \right).
\end{aligned}
$$

We note that, the D and B matrices are constants in the ILP. Therefore, the matrix multiplication expression is linear and suitable for ILP.

For the thickening parallelograms, the corresponding losses are $L_{\searrow}^{S_\rho, S_\rho^n}$ and $L_{\swarrow}^{S_\rho, S_\rho^n}$. As in the vertex edge case, we need to ensure that all entries in the relevant multiplied matrices from Table 1 are bounded by ℓ. Thus we add the following constraints

$$
\begin{aligned}
\text{Subject to} \quad & \ell \geq y_{ij}^{\varphi,V} \quad \forall y_{ij}^{\varphi,V} \in D_{G^n}^V \left(M_{\varphi^n}^V \cdot I_F^V - I_{G^n}^V \cdot M_\varphi^V \right) \\
& \ell \geq y_{ij}^{\varphi,E} \quad \forall y_{ij}^{\varphi,E} \in D_{G^n}^E \left(M_{\varphi^n}^E \cdot I_F^E - I_{G^n}^E \cdot M_\varphi^E \right) \\
& \ell \geq y_{ij}^{\psi,V} \quad \forall y_{ij}^{\psi,V} \in D_{F^n}^V \left(M_{\psi^n}^V \cdot I_G^V - I_{F^n}^V \cdot M_\psi^V \right) \\
& \ell \geq y_{ij}^{\psi,E} \quad \forall y_{ij}^{\psi,E} \in D_{F^n}^E \left(M_{\psi^n}^E \cdot I_G^E - I_{F^n}^E \cdot M_\psi^E \right)
\end{aligned}
$$

Similar to the edge-vertex parallelograms, the D and I matrices are constants in ILP, which makes these constraints linear.

Finally, we include constraints to control the entries on the triangle diagram matrices with corresponding loss functions $L_{\bigtriangledown}^{S_\rho}$ and $L_{\bigtriangleup}^{S_\rho}$. We first discuss the diagram $L_{\bigtriangledown}^{S_{\sigma_i}}$ which operates on vertices. The loss contribution from this type of diagram can be obtained by finding

$$\max_i L_{\bigtriangledown}^{S_{\sigma_i}} = \max\left\{\left\lceil\tfrac{x}{2}\right\rceil \,\big|\, x \in D_{F^{2n}}^V\left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V\right)\right\}.$$

Unlike the previous two types of diagrams, the expression $D_{F^{2n}}^V\left(I_{F^n}^V \cdot I_F^V - M_{\psi^n}^V \cdot M_\varphi^V\right)$ is quadratic rather than linear, as it includes the term $M_{\psi^n}^V \cdot M_\varphi^V$, which is a multiplication of two variable matrices. However, we can linearize the problem by introducing a collection of integer and binary reparametrization variables and adding them to the constraints. Furthermore, the ceiling function that we need to compute here is also not a linear function either. This is taken care of another reparametrization integer variable. Specifically, we have the variables $\{z_{ijk}\}$, one to represent each entry $x_{ij}^{\psi_n,V} x_{jk}^{\varphi,V}$ so that if we construct the matrix M where $M_{ik} = \sum_j z_{ijk}$, then $M = (M_{\psi^n}^V)_{ij}(M_{\varphi^n}^V)_{jk}$. Then, we get the following constraints for $L_{\bigtriangledown}^{S_{\sigma_i}}$,

$$
\begin{aligned}
\text{Subject to} \quad & 2c_{ij} \geq k_{ij} && \forall \quad k_{ij} \in D_{F^{2n}}^V\left(I_{F^n}^V \cdot I_F^V - M\right) \\
& z_{ijk} \leq x_{ij}^{\psi^n,V} && \forall \quad x_{ij}^{\psi^n,V} \in M_{\psi^n}^V \\
& z_{ijk} \leq x_{jk}^{\varphi,V} && \forall \quad x_{jk}^{\varphi,V} \in M_\varphi^V \\
& z_{ijk} \geq x_{ij}^{\psi^n,V} + x_{jk}^{\varphi,V} - 1 \\
& \ell \geq c_{ij}.
\end{aligned}
$$

These constraints are only for one of the four relevant triangle diagrams, however the other setups are similar. Thus our full linear program is setup to minimize ℓ given all the relevant constraints listed above.

4.2   ILP optimization implementation approach

We implement the ILP using PuLP library in Python which is is an open-source software package in Python specializing in modeling and solving mixed-integer linear programming problems [42]. Specifically, PuLP is equipped to solve integer linear programming problems involving integer and binary variables, as in our case. PuLP also integrates seamlessly with popular ILP solvers, both commercial (e.g. Gurobi [34]) and open source (e.g. CBC [40], GLPK [31]), without having to alter the model formulation. It is important to note that PuLP is exclusively designed to model linear optimization problems, hence the extra work done in Section 4.1 to linearize our ILP. We choose PuLP over available non-linear modeling frameworks like Pyomo [35] because it is lightweight and has a simpler syntax, making it easier to interpret and update. Further, the design of PuLP is streamlined for ILP computation, reducing the computational overhead.

We test our ILP implementation with simple input examples to make a choice of solver for use in PuLP; the specifics of this particular example pair will be given in Section 5.1. In this example, increasing loop size increases the size of the input graphs, thus the size of the ILP increases. We test two open-source solvers, CBC (COIN-OR Branch and Cut) and GLPK (GNU Linear Programming Kit), as well as the commercial solver Gurobi. The results are shown in Figure 12. We see that for smaller problem size (up to a loop size of 20 in the torus mapper), all the solvers perform reasonably well. However, as the size of the loop increases, the GLPK solver becomes significantly
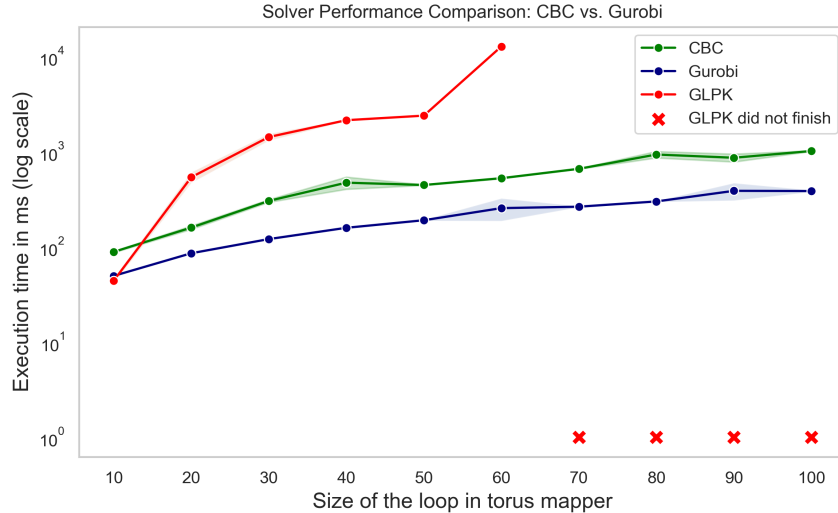
Figure 12: Time taken by different solvers to optimize the loss for the torus and line mappers

slower than CBC or Gurobi. In fact, for loop sizes of 70 and above, GLPK fails to complete computation even after 15 minutes of runtime. As expected, the commercial solver Gurobi is highly optimized, resulting in superior performance. Both CBC and Gurobi's computation times increased approximately linearly with the loop size. Among the open-source solvers we tested, CBC is slower than Gurobi but is still able to optimize large problems effectively. Since CBC is the default solver in PuLP, it is also more accessible from a user's perspective. Thus, for the rest of the experimental section, we exclusively use the CBC solver inside of PuLP. The code for our implementation is available in the Python package `ceREEBerus`[2].

4.3 Search for the optimal n

To this point, all discussion has been focused on the case where we are provided a fixed n, and then we can use the structure given to determine an optimal bound $L_B = k_n$ so that $n + k_n$ bounds the interleaving distance. However, as we will see in Section 5.2, there exists a pair of mapper graphs and a choice of n for which any n-assignment φ, ψ has $n + L_B(\varphi, \psi)$ strictly greater than the interleaving distance. For this reason, an additional step in the optimization must be taken to search over the possible choices of n to determine if this can also result in a better bound. To search for the most optimal upper bound efficiently, we implement a mixed exponential-binary search approach described as follows. Initially we find a plausible interval $[A, B]$ within which to search for n by doubling $n \in \{0, 1, 2, 4, \cdots\}$ iteratively until the computed loss found is $k_n = 0$. That is, $B = 2^j$ is the first n in the list that has $k_n = 0$, and then $A = 2^{j-1}$. Once the range is narrowed down, we set up a binary search scheme to find the n which provides an optimal bound $n + k_n$. Specifically, We select n to be the midpoint of the interval $[A, B]$ to be n and find the corresponding ILP optimized k_n, so that we know that $d_I \leq n + k_n$. If $k_n = 0$ (meaning an n-interleaving exists), choose the next n to be the midpoint of $[A, n]$. Otherwise, choose the next n to be the midpoint of $[n, B]$, and repeat the process.

---

[2]

25

Figure 13: (Left) Relationship between the computed upper bound and true interleaving distance between a line mapper graph, and a torus mapper graph of height given on the x-axis. (Right) Time taken to optimize the loss for different loop sizes ($n = 1$) including and excluding the initialization (to set up the B, D, I matrices).

5 Experiments

5.1 Experiments with Line and Torus Mappers

We validate our optimization of the loss function by testing it on two simple Mapper graphs, where the interleaving distance can be determined by hand. One is the *line mapper graph L*, which consists of a single vertex at each integral function value. The other is the *torus mapper T_h*, which features a loop in the middle with height h. Note that the justification for the name *torus* is that this is the mapper graph from the common example of a height function on an upright torus; not that the graph is isomorphic to a torus in some way. Both mappers are set to have the same range of function values; see the left of Figure 2 for an example with $h = 11$. We can check that the true interleaving distance between a *line mapper* and a *loop mapper* with loop height h is $d_I(L, T_h) = \lceil \frac{h}{4} \rceil$.

In the right of Figure 2, we see that the loss computed for varying choices of n for the n-assignments achieves the true interleaving distance bound for all n up to and including n equal to the true interleaving distance. Thus in this particular example, the search over n would not be necessary, but we do not know *a priori* whether this is the case; see Section 5.2 for an example where the search is necessary.

In Figure 13, we have the final results for the distance computed for varying height of the loop h. At left, we see that the computed bound in fact achieves the interleaving distance in these examples. At right, we have the time taken for the code to run, for both the entirety of the computation, as well as just the time for the PuLP to solve the ILP. We measure the time to optimize the loss function for these two mappers, with increasing torus heights for $n = 1$. We notice that the time increases linearly with the loop height h of *torus mappers*. We also run multiple iterations of loss optimization for a pair of these mappers and the ILP always outputs the optimal solution with the same optimized loss value, indicating that this optimization pipeline is stable.

5.2 An example requiring a binary search

As noted previously, a search over n is necessary. In this section, we provide an example of a pair of graphs and a choice of n for which any assignment has $d_I \lneq n + L_B(\varphi, \psi)$. Consider the two
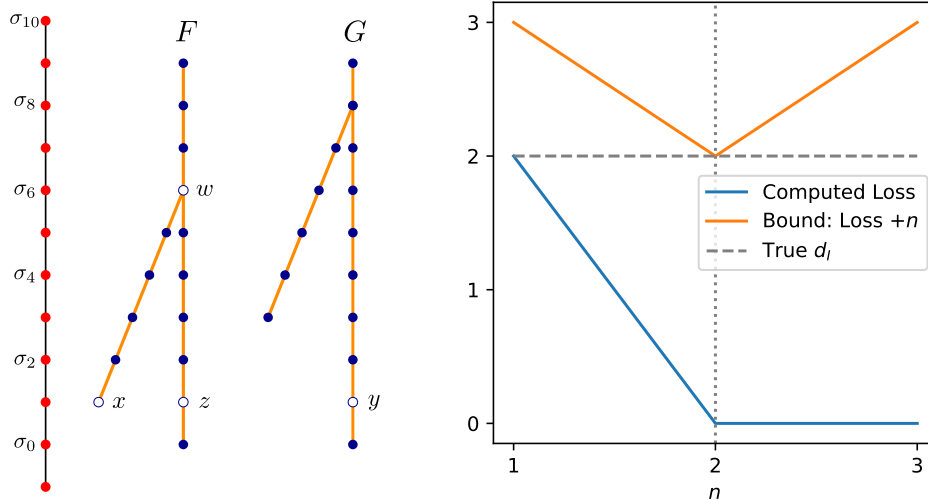
Figure 14: (Left) An example of a pair of mapper graphs where the binary search on n is required. (Right) The computed loss and interleaving bound for this example.

example mapper graphs in the left of Figure 14 and fix $n = 1$. Focusing on the vertices x and z, the only possible assignment for $\varphi(x) = \varphi(z) = [y]$ where we write $[y]$ for the connected component containing y in the 1-smoothing of G; i.e. $[y] \in G(S^1_{\sigma_1})$. In the other direction, there is a choice for y, so $\psi(y)$ is either $[x]$ or $[z]$ in $F(S^1_{\sigma_1})$. Then in the diagram

$$
\begin{array}{ccc}
F(S_{\sigma_1}) & \longrightarrow & F(S^2_{\sigma_1}) \\
& \searrow & \nearrow \\
& G(S^1_{\sigma_1}) &
\end{array}
\qquad
\begin{array}{ccc}
\genfrac{}{}{0pt}{}{x}{z} & \longrightarrow & \genfrac{}{}{0pt}{}{[x]}{[z]} \\
\searrow & \overset{?}{\nearrow} & \\
& [y] &
\end{array}
$$

either choice for y will result in the other object not commuting. Then the loss will be given by $k = \lceil d_F^{S^2_{\sigma_i}}([x], [y])/2 \rceil = \lceil 3/2 \rceil = 2$, so the interleaving bound is $n + k = 1 + 2 = 3$. However, we can build an interleaving for $n = 2$ by sending the left tail of F to the left tail of G, showing that the true interleaving $d_I(F, G) = 2$ is strictly less than any possible bound found by the solver.

5.3 Classification of binary images

To evaluate our approach, we conduct a simple classification experiment on MPEG-7 dataset [36], specifically MPEG-7 CE Shape-1 Part B, which is a widely used benchmark for shape analysis and classification with binary images. Our methodology involves computing mapper graphs from the dataset and analyzing their interleaving distance for classification tasks.

We choose five categories of images: *apple, cup, fork, horseshoe*, and *seasnake*; and select 10 images from each class. See Figure 3 (left) for examples from each category. We convert these images to Mapper graphs using *KeplerMapper* [58], where the lens function is given by the y-coordinate of each pixel. For each resulting mapper graph, we ensure that there is only one connected component in the graph, and that the node heights are in the range $[0, 20]$. This normalization ensures consistency in mapper size, simplifying loss computation. For every point cloud, we use a cover of 10 elements with a 0.3% overlap, and set the epsilon value for DBSCAN to 3. We choose these parameters to generate a mapper that best represents the point cloud.
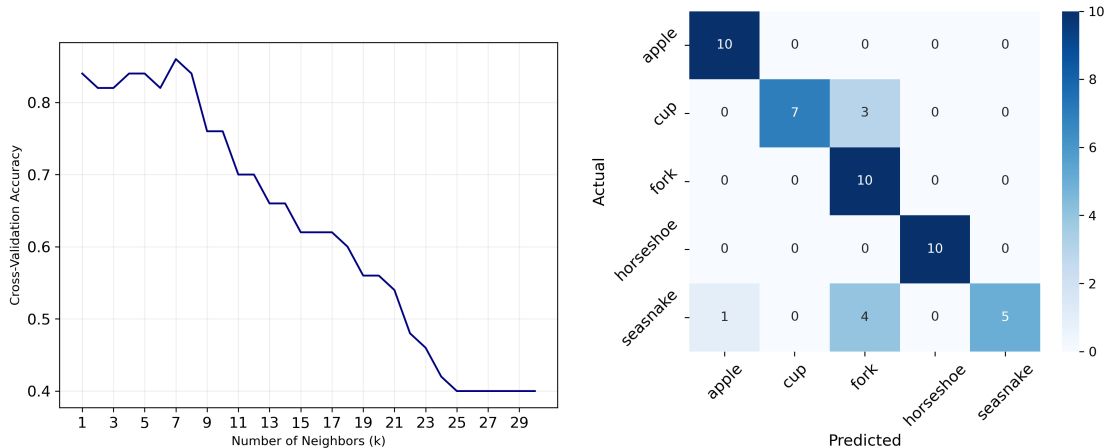
27

Figure 15: (Left) Accuracy of KNN classifier for different k values during 5-fold cross-validation. The optimal $k = 7$ achieved 84% accuracy.
(Right) Confusion matrix showing how the actual classes match with the predicted class. Each (i, j)-th cell indicate number of times an image from category i was classified as category j.

Once the mapper graphs are obtained, we compute the optimal upper bound of the interleaving distance between them using the optimization strategy discussed earlier, which serves as a measure of similarity between Mapper graphs. To visualize the distance performance, we construct a *multi-dimensional scaling* (MDS) [14] plot (Figure 3, right) using the computed distance matrix. MDS is chosen for visualization as it accepts an arbitrary distance matrix as input.

As we see in Figure 3 (right), the optimal upper bounds on distance successfully cluster similar shapes together. Among all categories, *apple* exhibits the most well-defined grouping. Additionally, some of the other shapes are well separated, such as *apple* and *fork* or *seasnake*. However, certain shapes are not differentiated by this distance. For instance, the *cup* category in the MPEG-7 dataset includes two types of cups: one with joined handles, forming a loop, and another without joined handles. This structural difference is reflected in the Mapper representation, causing several to appear far apart despite belonging to the same category. This is to be expected since the mapper interleaving distance is built to find similar topological structure in the resulting graph, and does not take the geometric similarity into account. Likewise, the horseshoes and sea snakes, despite being similar structurally have a large distance between them since the mapper graph lens function takes the embedding direction into account; thus the horseshoe with opening of the U-shape pointing down and the sea-snake with the opening point up are considered very different under this distance.

We also use the pairwise distances to classify the images into their correct classes using k-nearest neighbor (KNN) [24], based on the optimized loss computed on the mappers. To determine the most optimal number of neighbors (k) and avoid overfitting or underfitting, we first perform a 5-fold cross-validation on the data, testing values of k from 1 to 30. The optimal k was selected to be 7, which resulted in the highest and most stable accuracy of 84%, as shown in Figure 15 (Left). For a more robust evaluation, we also implement *Leave-one-out cross-validation* (LOO-CV), where data point is used as a test case while the rest are used for training the model. To visualize the classification and misclassification, we display a *confusion matrix*, which highlights the categories that are being misclassified, see Figure 15 (Right). The LOO-CV results confirm the model's accuracy of 84%, indicating that while the model is generally effective, there are still some misclassifications. It is important to note that this is a preliminary experiment with only

10 images per category (totaling 50 images), and expanding to the full MPEG-7 dataset might improve accuracy. As we see in the confusion matrix in Figure 15 (Right), the categories that suffer from misclassification are *cup* and *seasnake*. As discussed before, the some of the cups have joined handles and some do not, which contributes to the reclassification. Additionally, the sea snake images are highly varied in how they are oriented, leading to misclassification due to differences in their embeddings, which results in generation of very different mapper graphs.

6   Discussion

In this paper, we have given the first available method to compute a bound for the interleaving distance; and to our knowledge, the first code for understanding interleavings on mapper graphs. Note that we are being careful to not say that this result "computes" the interleaving distance, since the only theoretical guarantee we have is that the algorithm will return a *bound* on the interleaving distance. However, we have shown that in some small examples where the true interleaving distance is known, the bound achieved is indeed equal to the interleaving distance. We have also shown a basic proof of concept where our code to bound the interleaving distance can be used to compare input images from the MPEG-7 data set.

This work is the first step towards the truly practical use of the interleaving distance in applications. The computations in the implementation likely have many avenues for further optimization and speedups. The most clear place to look for this is in the computation of the distance matrix D (Section 3.2.3). Our current implementation uses a breadth first search with a union find data structure to compute the distances at each level individually; we suspect some variant of all-pairs shortest paths algorithms may be employable to somehow compute these distances in a single sweep rather than level-by-level.

In future work, we hope that this code and any future improvements can lead to new understanding and optimization of mapper graphs. For instance, the interleaving distance could be used to measure quality of parameters chosen in the mapper construction by either measuring how close a computed mapper graph is to either a ground truth. It would also be interesting to benchmark the classification done with the interleaving distance against existing graph distances. Once the computation time can be improved, a larger experiment comparing the MPEG-7 images as done in Section 5.3 would be very exciting.

In addition, the loss function definition is largely categorical; this means that it has the potential to be translated for us in the interleaving distance for other constructions from TDA such as merge trees and multiparameter persistence modules. We hope that this ILP viewpoint will make it possible to extend the use of computation for the interleaving distance to new settings where prior use was also limited to theoretical results.

7   Acknowledgments

References

[1] Enrique Alvarado, Robin Belton, Emily Fischer, Kang-Ju Lee, Sourabh Palande, Sarah Percival, and Emilie Purvine. G-mapper: Learning a cover in the mapper construction, 2025. URL:

`https://arxiv.org/abs/2309.06634`, `arXiv:2309.06634`.

[2] Enrique G Alvarado, Robin Belton, Kang-Ju Lee, Sourabh Palande, Sarah Percival, Emilie Purvine, and Sarah Tymochko. Any graph is a mapper graph, 2024. URL: `https://arxiv.org/abs/2408.11180`, `arXiv:2408.11180`.

[3] Jonathan A. Barmak. *Algebraic Topology of Finite Topological Spaces and Applications*. Springer, Berlin, Heidelberg, 2011. `doi:10.1007/978-3-642-22003-6`.

[4] U. Bauer, B. Di Fabio, and C. Landi. An edit distance for Reeb graphs. In *Proceedings of the Eurographics 2016 Workshop on 3D Object Retrieval*, 3DOR '16, page 27–34, Goslar, DEU, 2016. Eurographics Association.

[5] Ulrich Bauer, Håvard Bakke Bjerkevik, and Benedikt Fluhr. Quasi-universality of Reeb graph distances. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:18, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: `https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SoCG.2022.14`, `doi:10.4230/LIPIcs.SoCG.2022.14`.

[6] Ulrich Bauer, Xiaoyin Ge, and Yusu Wang. Measuring distance between Reeb graphs. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, SOCG'14, page 464–473, New York, NY, USA, 2014. Association for Computing Machinery. `doi:10.1145/2582112.2582169`.

[7] Ulrich Bauer, Claudia Landi, and Facundo Mémoli. The Reeb graph edit distance is universal. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. URL: `https://drops.dagstuhl.de/opus/volltexte/2020/12173`, `doi:10.4230/LIPIcs.SoCG.2020.15`.

[8] Nicolas Berkouk and Francois Petit. Ephemeral persistence modules and distance comparison. *Algebr. Geom. Topol. 21 (2021) 247-277*, 21(1):247–277, feb 2019. `arXiv:1902.09933`, `doi:10.2140/agt.2021.21.247`.

[9] Thijs Beurskens, Tim Ophelders, Bettina Speckmann, and Kevin Verbeek. Relating interleaving and Fréchet distances via ordered merge trees, 2025. URL: `https://arxiv.org/abs/2312.11113`, `arXiv:2312.11113`.

[10] Håvard Bakke Bjerkevik, Magnus Bakke Botnan, and Michael Kerber. Computing the interleaving distance is NP-hard. *Foundations of Computational Mathematics*, 20:1237–1271, 2020.

[11] Håvard Bakke Bjerkevik and Magnus Bakke Botnan. Computational complexity of the interleaving distance. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry (SoCG 2018)*, volume 99 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:15, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: `https://drops-dev.dagstuhl.de/entities/document/10.4230/LIPIcs.SoCG.2018.13`, `doi:10.4230/LIPIcs.SoCG.2018.13`.

[12] Andrew Blumberg and Michael Lesnick. Universality of the homotopy interleaving distance. *Transactions of the American Mathematical Society*, September 2023. URL: http://dx.doi.org/10.1090/tran/8738, doi:10.1090/tran/8738.

[13] Brian Bollen, Erin Chambers, Joshua A. Levine, and Elizabeth Munch. Reeb graph metrics from the ground up, 2022. URL: https://arxiv.org/abs/2110.05631, arXiv:2110.05631.

[14] Ingwer Borg and Patrick JF Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2007.

[15] Magnus Bakke Botnan, Justin Curry, and Elizabeth Munch. A relative theory of interleavings, 2020. URL: https://arxiv.org/abs/2004.14286, arXiv:2004.14286.

[16] Adam Brown, Omer Bobrowski, Elizabeth Munch, and Bei Wang. Probabilistic convergence and stability of random mapper graphs. *Journal of Applied and Computational Topology*, 5(1):99–140, December 2020. URL: http://dx.doi.org/10.1007/s41468-020-00063-x, doi:10.1007/s41468-020-00063-x.

[17] Peter Bubenik, Vin De Silva, and Jonathan Scott. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, 15:1501–1531, 2015.

[18] Mathieu Carrière, Bertrand Michel, and Steve Oudot. Statistical analysis and parameter selection for mapper. *Journal of Machine Learning Research*, 19(12):1–39, 2018. URL: http://jmlr.org/papers/v19/17-291.html.

[19] Mathieu Carrière and Steve Oudot. Local equivalence and intrinsic metrics between Reeb graphs. In Boris Aronov and Matthew J. Katz, editors, *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:15, Dagstuhl, Germany, March 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: http://drops.dagstuhl.de/opus/volltexte/2017/7179, doi:10.4230/LIPIcs.SoCG.2017.25.

[20] Mathieu Carrière and Steve Oudot. Structure and stability of the one-dimensional mapper. *Foundations of Computational Mathematics*, 18(6):1333–1396, October 2017. URL: http://dx.doi.org/10.1007/s10208-017-9370-z, doi:10.1007/s10208-017-9370-z.

[21] Erin W Chambers, Elizabeth Munch, Sarah Percival, and Bei Wang. Bounding the interleaving distance for geometric graphs with a loss function. *arXiv preprint arXiv:2307.15130*, 2023.

[22] Erin Wolf Chambers, Elizabeth Munch, and Tim Ophelders. A family of metrics from the truncated smoothing of Reeb graphs. In Kevin Buchin and Éric Colin de Verdière, editors, *37th International Symposium on Computational Geometry (SoCG 2021)*, volume 189 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:17, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SoCG.2021.22, doi:10.4230/LIPIcs.SoCG.2021.22.

[23] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th annual symposium on Computational geometry*, SCG '09, pages 237–246, New York, NY, USA, 2009. ACM. URL: http://doi.acm.org/10.1145/1542362.1542407, doi:10.1145/1542362.1542407.

[24] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.

[25] Joshua Cruz. Metric limits in categories with a flow, 2019. URL: https://arxiv.org/abs/1901.04828, arXiv:1901.04828.

[26] Justin Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania, 2014. arXiv:1303.3255.

[27] Justin Curry, Haibin Hang, Washington Mio, Tom Needham, and Osman Berat Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, 6(3):371–428, March 2022. doi:10.1007/s41468-022-00089-3.

[28] Vin de Silva, Elizabeth Munch, and Amit Patel. Categorified Reeb graphs. *Discrete & Computational Geometry*, pages 1–53, 2016. URL: http://dx.doi.org/10.1007/s00454-016-9763-9, doi:10.1007/s00454-016-9763-9.

[29] Vin de Silva, Elizabeth Munch, and Anastasios Stefanou. Theory of interleavings on categories with a flow. *Theory and Applications of Categories*, 33(21):583–607, 2018. URL: http://www.tac.mta.ca/tac/volumes/33/21/33-21.pdf.

[30] Elena Farahbakhsh Touli and Yusu Wang. FPT-algorithms for computing Gromov-Hausdorff and interleaving distances between trees. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms (ESA 2019)*, volume 144 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 83:1–83:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: http://drops.dagstuhl.de/opus/volltexte/2019/11204, doi:10.4230/LIPIcs.ESA.2019.83.

[31] GLPK GNU Project Free Software Foundation (FSF). URL: https://www.gnu.org/software/glpk/.

[32] Xinbo Gao, Bing Xiao, Dacheng Tao, and Xuelong Li. A survey of graph edit distance. *Pattern Analysis and applications*, 13(1):113–129, 2010.

[33] Ellen Gasparovic, Elizabeth Munch, Steve Oudot, Katharine Turner, Bei Wang, and Yusu Wang. Intrinsic interleaving distance for merge trees. *La Matematica*, December 2024. URL: http://dx.doi.org/10.1007/s44007-024-00143-9, doi:10.1007/s44007-024-00143-9.

[34] Gurobi Optimization, LLC. Gurobi optimizer reference manual, 2024. URL: https://www.gurobi.com.

[35] William E Hart. Python optimization modeling objects (Pyomo). In *Operations Research and Cyber-Infrastructure*, pages 3–19. Springer, 2009.

[36] Sylvie Jeannin and Miroslaw Bober. Description of core experiments for MPEG-7 motion/shape. *MPEG-7, ISO/IEC/JTC1/SC29/WG11/MPEG99 N*, 2690, 1999.

[37] Woojin Kim, Facundo Mémoli, and Anastasios Stefanou. Interleaving by parts: Join decompositions of interleavings and join-assemblage of geodesics. *Order*, 41(2):497–537, September 2023. URL: http://dx.doi.org/10.1007/s11083-023-09643-9, doi:10.1007/s11083-023-09643-9.

[38] Fangfei Lan, Salman Parsa, and Bei Wang. Labeled interleaving distance for Reeb graphs. *Journal of Applied and Computational Topology*, 8(8):2367–2399, September 2024. doi:10.1007/s41468-024-00193-6.

[39] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015. URL: http://dx.doi.org/10.1007/s10208-015-9255-y, doi:10.1007/s10208-015-9255-y.

[40] R. Lougee-Heimer. The common optimization interface for operations research: Promoting open-source software in the operations research community. *IBM Journal of Research and Development*, 47(1):57–66, 2003. doi:10.1147/rd.471.0057.

[41] Killian Meehan and David Meyer. Interleaving distance as a limit, 2017. URL: https://arxiv.org/abs/1710.11489, arXiv:1710.11489.

[42] Stuart Mitchell, Michael OSullivan, and Iain Dunning. Pulp: a linear programming toolkit for python. *The University of Auckland, Auckland, New Zealand*, 65:25, 2011.

[43] Dmitriy Morozov, Kenes Beketayev, and Gunther Weber. Interleaving distance between merge trees. In *Proceedings of TopoInVis*, 2013.

[44] Elizabeth Munch and Anastasios Stefanou. The ℓ_∞-cophenetic metric for phylogenetic trees as an interleaving distance. In *Association for Women in Mathematics Series*, pages 109–127. Springer International Publishing, 2019. doi:10.1007/978-3-030-11566-1_5.

[45] Elizabeth Munch and Bei Wang. Convergence between categorical representations of Reeb space and Mapper. In Sándor Fekete and Anna Lubiw, editors, *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 53:1–53:16, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: http://drops.dagstuhl.de/opus/volltexte/2016/5945, doi:http://dx.doi.org/10.4230/LIPIcs.SoCG.2016.53.

[46] nLab authors. Unnatural transformation. https://ncatlab.org/nlab/show/unnatural+transformation, November 2023. Revision 1.

[47] Matteo Pegoraro. A graph-matching formulation of the interleaving distance between merge trees, 2024. URL: https://arxiv.org/abs/2111.15531, arXiv:2111.15531.

[48] Emilie Purvine, Davis Brown, Brett Jefferson, Cliff Joslyn, Brenda Praggastis, Archit Rathore, Madelyn Shapiro, Bei Wang, and Youjia Zhou. Experimental observations of the topology of convolutional neural network activations. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023. doi:10.1609/aaai.v37i8.26134.

[49] Archit Rathore, Yichu Zhou, Vivek Srikumar, and Bei Wang. TopoBERT: Exploring the topology of fine-tuned word representations. *Information Visualization*, 22(3):186–208, 2023. arXiv:https://doi.org/10.1177/14738716231168671, doi:10.1177/14738716231168671.

[50] Georges Reeb. Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique. *Comptes Rendus de l'Académie des Sciences*, 222:847–849, 1946.

[51] Emily Riehl. *Category theory in context*. Courier Dover Publications, 2017.

[52] Abbas H Rizvi, Pablo G Camara, Elena K Kandror, Thomas J Roberts, Ira Schieren, Tom Maniatis, and Raul Rabadan. Single-cell topological RNA-seq analysis reveals insights into

cellular differentiation and development. *Nature Biotechnology*, 35(6):551–560, May 2017. URL: http://dx.doi.org/10.1038/nbt.3854, doi:10.1038/nbt.3854.

[53] Michael Robinson. Assignments to sheaves of pseudometric spaces. *Compositionality*, 2, June 2020. doi:10.32408/compositionality-2-2.

[54] Manish Saggar, Olaf Sporns, Javier Gonzalez-Castillo, Peter A. Bandettini, Gunnar Carlsson, Gary Glover, and Allan L. Reiss. Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature Communications*, 9(1), April 2018. URL: http://dx.doi.org/10.1038/s41467-018-03664-4, doi:10.1038/s41467-018-03664-4.

[55] Luis Scoccola. *Locally persistent categories andmetric properties of interleaving distances*. PhD thesis, Western University, 2020.

[56] 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. URL: http://www.tac.mta.ca/tac/volumes/33/21/33-21.pdf, arXiv:1706.04095.

[57] Gurjeet Singh, Facundo Mémoli, and Gunnar Carlsson. Topological methods for the analysis of high dimensional data sets and 3D object recognition. In *Eurographics Symposium on Point-Based Graphics*, 2007. doi:10.2312/SPBG/SPBG07/091-100.

[58] Hendrik Jacob Van Veen, Nathaniel Saul, David Eargle, and Sam W Mangham. Kepler mapper: A flexible python implementation of the mapper algorithm. *Journal of Open Source Software*, 4(42):1315, 2019.

[59] Lin Yan, Talha Bin Masood, Raghavendra Sridharamurthy, Farhan Rasheed, Vijay Natarajan, Ingrid Hotz, and Bei Wang. Scalar field comparison with topological descriptors: Properties and applications for scientific visualization. *Computer Graphics Forum*, 40(3):599–633, jun 2021. doi:10.1111/cgf.14331.

[60] Lin Yan, Yusu Wang, Elizabeth Munch, Ellen Gasparovic, and Bei Wang. A structural average of labeled merge trees for uncertainty visualization. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019. arXiv:1908.00113, doi:10.1109/tvcg.2019.2934242.

[61] Youjia Zhou, Helen Jenne, Davis Brown, Madelyn Shapiro, Brett Jefferson, Cliff Joslyn, Gregory Henselman-Petrusek, Brenda Praggastis, Emilie Purvine, and Bei Wang. Comparing mapper graphs of artificial neuron activations. In *2023 Topological Data Analysis and Visualization (TopoInVis)*, pages 41–50, 2023. doi:10.1109/TopoInVis60193.2023.00011.