Local Versus Global Distances for Zigzag and Multi-Parameter Persistence Modules

Ellen Gasparovic, Maria Gommel, Emilie Purvine, Radmila Sazdanovic, Bei Wang, Yusu Wang, and Lori Ziegelmeier

Abstract In this paper, we establish explicit and broadly applicable relationships between persistence-based distances computed locally and globally. In particular, we show that the bottleneck distance and the Wasserstein distance between two zigzag persistence modules restricted to an interval is always bounded above by the distance between the unrestricted versions. While this result is not surprising, it could have potential practical implications. We give two related applications for metric graph distances, as well as an extension for the matching distance between multi-parameter persistence modules. We also prove a similar restriction inequality for the interleaving distance between two multi-parameter persistence modules.

1 Introduction

Assessing similarity or dissimilarity of complex objects is an important problem that is ubiquitous in science, engineering, and mathematics. The complexity of the input objects can make this problem quite difficult; for instance, comparing graphs or even trees via the edit-distance can be NP-hard [6]. With the recent development of applied and computational topology, a new paradigm for comparing complex objects is to first map them to easier-to-compare yet still meaningful topological summaries, and then compare the resulting topological summaries as a proxy.

Indeed, one such topological summary is the persistence diagram obtained via persistent homology. By choosing appropriate filtrations, input objects are represented by so-called persistence modules. In this paper, we consider the more general *zigzag persistence modules* [7], which admit a unique decomposition up to isomorphism into certain elementary pieces (called indecomposable modules) that are intervals. The information encoded by these intervals can be combinatorially represented by its persistence diagram. Such a persistence diagram serves as a topological summary representation of the input object, enabling one to compare multiple input objects by comparing their associated persistence diagrams using metrics such as the bottleneck distance [12] or Wasserstein distance [14].

Ellen Gasparovic
Union College, Schenectady, NY, USA, e-mail: gasparoe@union.edu

Maria Gommel
North Central College, Naperville, IL, USA, e-mail: megommel@noctrl.edu

Emilie Purvine
Pacific Northwest National Laboratory, Seattle, WA, USA, e-mail: emilie.purvine@pnnl.gov

Radmila Sazdanovic
North Carolina State University, Raleigh, NC, USA, e-mail: rsazdanovic@math.ncsu.edu

Bei Wang
University of Utah, Salt Lake City, UT, USA, e-mail: beiwang@sci.utah.edu

Yusu Wang
University of California, San Diego, CA, USA, e-mail: yusuwang@ucsd.edu

Lori Ziegelmeier
Macalester College, Saint Paul, MN, USA, e-mail: lziegel1@macalester.edu

Motivation

In practice, there are many applications where one might be interested in only computing a local or restricted summary. For a first example, consider the question of determining or approximating graph motif counts. A graph motif is a subgraph on a small number of vertices contained within a larger, more complex graph. Graph motifs have proven useful for characterizing networks in domains such as biology [23] and cyber security [19]. The standard problem of counting the number of small motifs or patterns within a graph is equivalent to the subgraph isomorphism problem, which is NP-complete. Since restricted persistence modules reveal information about the local structure of a space, we posit that the restricted modules for a metric graph (see Section 4) can be utilized in a similar manner to the way in which graph motifs are currently used, e.g., as inputs to classification algorithms or anomaly detection algorithms in time-varying data [19, 21].

For a second example, consider persistent local homology, which studies a multi-scale notion of homology within a local neighborhood of the data relative to its boundary. It has applications in road network analysis [2], data visualization [24], graph reconstruction [11, 1], and clustering and stratification learning [5, 3]. Furthermore, persistent local homology extracts local geometric and topological information in data, which may then be used as input to machine learning algorithms [4]. In such applications, there is often a choice as to the size of local neighborhood that one should use, and a natural question to ask is how the local information relates to the global topological summaries. This paper aims to provide some answers to this question.

Our contributions

After a brief review of the relevant background information in Section 2, we show in Section 3 that the bottleneck distance and the Wasserstein distance between two zigzag persistence modules restricted over an interval of parameter values is always bounded from above by the distance between the unrestricted versions (Theorem 1) and state a corollary in the case of *level set zigzag persistence* (Corollary 2). As an example of how such a bound could be useful: imagine that one wishes to compare the persistence profiles of two very large data sets but finds that it is prohibitively computationally expensive. Instead, one could compute a restricted version of the bottleneck distance as an approximation to the global distance. As the interval size increases, the bottleneck (or Wasserstein) distance between the restricted versions approaches the distance for the global versions. In an opposite direction, our result means that the global distance between two modules does not “wash out" local information encoded in the same module. In Section 4, we discuss some further implications of our results, by establishing two results involving distance inequalities in the special case of metric graphs (Corollary 3 and Corollary 4), as well as by extending our result to the matching distance of multi-parameter persistence modules [20]. Finally, we close by providing a restriction inequality result (Theorem 2) on the interleaving distance between multi-parameter persistence modules. In this case, we are not operating on the level of persistence diagrams but rather on the modules themselves.

2 Background and Definitions

Our treatment of zigzag persistence is brief; for more details, see [7] and [8]. A *zigzag diagram* of topological spaces $\mathbb{X}_1, \mathbb{X}_2, \ldots, \mathbb{X}_n$ is a sequence

$$\mathbb{X}_1 \leftrightarrow \mathbb{X}_2 \leftrightarrow \cdots \leftrightarrow \mathbb{X}_n$$

where each bidirectional arrow between two topological spaces represents a continuous function mapping either forwards or backwards. Applying the p-th homology functor with coefficients in a field $\mathbb{K}$ yields a zigzag diagram of vector spaces

$$\mathsf{H}_p(\mathbb{X}_1;\mathbb{K}) \leftrightarrow \mathsf{H}_p(\mathbb{X}_2;\mathbb{K}) \leftrightarrow \cdots \leftrightarrow \mathsf{H}_p(\mathbb{X}_n;\mathbb{K}),$$

known as a *zigzag module*, denoted as $\mathbf{X}$, from which *zigzag persistence* may be computed. Since the maps in a zigzag diagram are allowed to go in either direction, the resulting zigzag module is the most general form of 1-dimensional persistence. A zigzag module decomposes into intervals $\mathbf{X} \cong \bigoplus_{j \in J} \mathbb{I}[b_j, d_j]$, where each $\mathbb{I}[b_j, d_j]$ is defined as

$$0 \longleftrightarrow \cdots \longleftrightarrow 0 \longleftrightarrow \mathbb{K} \longleftrightarrow \cdots \longleftrightarrow \mathbb{K} \longleftrightarrow 0 \cdots \longleftrightarrow 0$$

with nonzero values in the range $[b_j, d_j]$. We will use Dg$\mathbf{X}$ to denote the resulting *persistence diagram* of a fixed homology dimension p. Typically, a persistence diagram is considered to be a set of points $\{(b, d)\}$ for which $b < d$. Proposition 2.12 of [7] implies that restricting the module $\mathbf{X}$ to the range $[r_1, r_2]$ (denoted $\mathbf{X}[r_1, r_2]$) yields a decomposition as the direct sum of the intervals in $\mathbf{X}$ restricted to $[r_1, r_2]$; that is,

$$\mathbf{X}[r_1, r_2] \cong \bigoplus_{j \in J} \mathbb{I}([b_j, d_j] \cap [r_1, r_2]). \tag{1}$$

In particular, any summand $\mathbb{I}([b_j, d_j])$ of $\mathbf{X}$ with $[b_j, d_j] \cap [r_1, r_2] = \emptyset$ becomes a zero module in the direct sum for $\mathbf{X}[r_1, r_2]$. Moreover, the length of each summand, i.e., the number of spaces in the sequence, does not change in the restriction, meaning that the spaces of summands are 0 outside of the interval $[r_1, r_2]$.

The *bottleneck distance* between two persistence diagrams is equal to δ if there exists a matching between the points of the two diagrams (where points are allowed to be matched to diagonal elements) such that any pair of matched points are at distance at most δ. Formally, for a fixed homology dimension, the bottleneck distance is given by

$$d_B(\mathrm{Dg}\mathbf{X}, \mathrm{Dg}\mathbf{Y}) = \inf_{\mu:\mathrm{Dg}\mathbf{X}\to\mathrm{Dg}\mathbf{Y}} \sup_{x \in \mathrm{Dg}\mathbf{X}} ||x - \mu(x)||_\infty,$$

where μ ranges over all bijections between the two diagrams [17]. In order to compute the bottleneck distance, one adds countably many copies of the diagonal $\{(x, x) : x \in \mathbb{R}\}$, which may intuitively correspond to topological features that are born and simultaneously die (and thus, never really exist at all). This allows for a point in one persistence diagram to be matched to the diagonal if it is far away from any point in the other diagram, and also accounts for the fact that two persistence diagrams may have different numbers of off-diagonal points.

The *q-th Wasserstein distance* is defined as

$$d_q(\mathrm{Dg}\mathbf{X}, \mathrm{Dg}\mathbf{Y}) = \left[\inf_{\mu:\mathrm{Dg}\mathbf{X}\to\mathrm{Dg}\mathbf{Y}} \sum_{x \in \mathrm{Dg}\mathbf{X}} ||x - \mu(x)||_\infty^q \right]^{1/q}.$$

We conclude this section by defining a projection map that keeps track of the points in the global persistence diagram that disappear in the restricted version. We next define a restriction map for a persistence diagram and then show that this is the persistence diagram of a restricted zigzag module.

Definition 1 Given $I = [r_1, r_2] \subset \mathbb{R}$, we let Dg$\mathbf{X}^I$ denote the **restriction of the persistence diagram** Dg$\mathbf{X}$ **to the interval** I defined as the output of the following projection map (Types A-F reference Figure 1):

$$\Pi : \mathrm{Dg}\mathbf{X} \to \mathrm{Dg}\mathbf{X}^I$$
$$(b, d) \mapsto \begin{cases} (\max(b, r_1), \min(d, r_2)) \text{ (Types A, B, C, D)} \\ (b, b) \text{ if } r_2 \le b \text{ (Type E)} \\ (d, d) \text{ if } d \le r_1 \text{ (Type F)} \end{cases}$$

Notice that points like E and F in Figure 1 correspond to features that are born and die outside of the interval I (either completely before or completely after). The restriction result cited above from [7] would not include points $\Pi(E)$ or $\Pi(F)$ in its diagram. However, since both $\Pi(E)$ and $\Pi(F)$ are on the diagonal, including them in Dg$\mathbf{X}^I$ does not change the bottleneck distance between two restricted diagrams. We formalize this in the following Lemma. Note that we could choose to map to any point on the diagonal, such as a point in the support of $[r_1, r_2]$, but the choice of (b, b) and (d, d) simplifies the proof of Theorem 1 below.

Lemma 1 *Let $I = [r_1, r_2] \subset \mathbb{R}$. The restriction of the persistence diagram* Dg$\boldsymbol{X}$ *to the interval I is equal to the persistence diagram of the restricted module $\boldsymbol{X}[r_1, r_2]$, i.e.,* Dg$\boldsymbol{X}^I$ = Dg$\boldsymbol{X}[r_1, r_2]$.

Proof Let $\mathbf{X} \cong \bigoplus_{j \in J} \mathbb{I}[b_j, d_j]$. Then by (1) we have that Dg$\mathbf{X}[r_1, r_2]$ contains the following set of points:

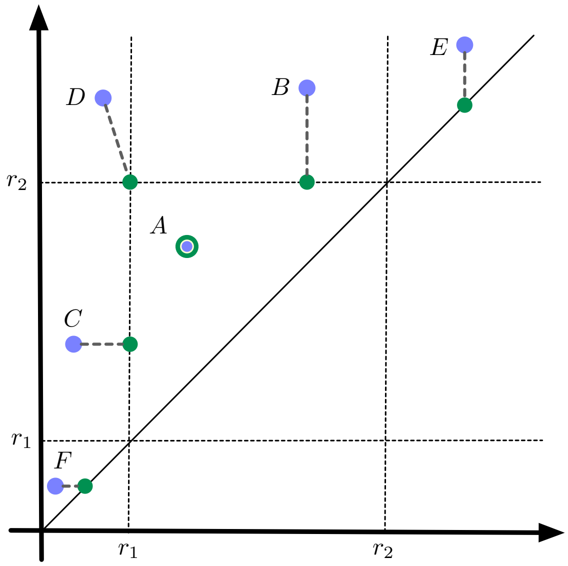

Fig. 1 An illustration of the six cases for the projection map Π. Green points are the images of the blue points under Π.

$$\mathrm{Dg}\mathbf{X}[r_1, r_2] = \left\{ \begin{matrix} (\max(b_j, r_1), \min(d_j, r_2)) \text{ if } r_1 \le b_j \le r_2 \text{ or } r_1 \le d_j \le r_2 \\ (r_1, r_2) \qquad\qquad \text{if } b_j < r_1 \le r_2 < d_j \end{matrix} : (b_j, d_j) \in \mathrm{Dg}\mathbf{X} \right\}.$$

The first case aligns with types A, B, and C in Figure 1, and the second case is type D.

Note that this does not take into account those intervals for which $[b_j, d_j] \cap [r_1, r_2] = \emptyset$. To create $\mathrm{Dg}\mathbf{X}^I$ we map them to (b_j, b_j) or (d_j, d_j) depending on whether they fall before the interval I or after. But these points are not represented directly in $\mathrm{Dg}\mathbf{X}[r_1, r_2]$. Rather, to complete the proof we rely on the fact that persistence diagrams by design contain infinitely many copies of the diagonal $\{(x, x) : x \in \mathbb{R}\}$. So, both diagrams contain the same mapped points of type A, B, C, and D, and infinitely many copies of the diagonal. □

3 Bottleneck and Wasserstein Distances in the Local vs. Global Settings

In this section, we prove our main result relating the bottleneck distance (resp. Wasserstein distance) between persistence diagrams with the bottleneck (resp. Wasserstein) distance between their interval-restricted versions.

Theorem 1 *Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be two zigzag modules and let $\mathrm{Dg}\boldsymbol{X}$ and $\mathrm{Dg}\boldsymbol{Y}$ be their corresponding zigzag persistence diagrams. Consider the interval $I = [r_1, r_2] \subset \mathbb{R}$ and let $\mathrm{Dg}\boldsymbol{X}^I$ and $\mathrm{Dg}\boldsymbol{Y}^I$ be the restrictions of these diagrams to I. Then*

$$d_B(\mathrm{Dg}\boldsymbol{X}^I, \mathrm{Dg}\boldsymbol{Y}^I) \le d_B(\mathrm{Dg}\boldsymbol{X}, \mathrm{Dg}\boldsymbol{Y})$$

and

$$d_q(\mathrm{Dg}\boldsymbol{X}^I, \mathrm{Dg}\boldsymbol{Y}^I) \le d_q(\mathrm{Dg}\boldsymbol{X}, \mathrm{Dg}\boldsymbol{Y}).$$

Proof Let $\mu \subseteq \mathrm{Dg}\mathbf{X} \times \mathrm{Dg}\mathbf{Y}$ be a partial matching, where any unpaired point in one of the persistence diagrams is matched to the nearest point (in the L_∞ norm) on the diagonal $\Delta = \{(x, x) : x \in \mathbb{R}\}$. Then μ can be modeled as a bipartite graph between finite sets of points from the two diagrams.

We define $\mu_I \subseteq \mathrm{Dg}\mathbf{X}^I \times \mathrm{Dg}\mathbf{Y}^I$ such that, for each $(p, q) \in \mu$, we have $(\Pi(p), \Pi(q)) \in \mu_I$. Since we are simply relabeling the coordinates of points in the bipartite graph of μ to define μ_I, μ_I remains a partial matching. What is left to show is that the maximal distance between the matched points in μ_I is less than that for μ. That is, we want to show that for the partial matching μ_I,

$$\sup_{(\Pi(p), \Pi(q)) \in \mu_I} ||\Pi(p) - \Pi(q)||_\infty \le \sup_{(p,q) \in \mu} ||p - q||_\infty. \tag{2}$$

We show that

$$\|\Pi(p) - \Pi(q)\|_\infty \leq \|p - q\|_\infty, \tag{3}$$

for all matched p and q. This is then sufficient to establish the claimed inequality.

Consider two points $p \in \mathrm{Dg}\mathbf{X}$ and $q \in \mathrm{Dg}\mathbf{Y}$ such that $(p, q) \in \mu$. Let $p = (b_p, d_p)$ and $q = (b_q, d_q)$. By definition of L_∞ distance, we have $\|p - q\|_\infty = \max(|b_p - b_q|, |d_p - d_q|)$. First, if both p and q are points of type **A, B, C,** or **D** as in Definition Definition 1 and Figure 1, then $\Pi(b_p, d_p) = (\max(b_p, r_1), \min(d_p, r_2))$ and $\Pi(b_q, d_q) = (\max(b_q, r_1), \min(d_q, r_2))$. Thus, we have

$$\|\Pi(p) - \Pi(q)\|_\infty = \max(|\max(b_p, r_1) - \max(b_q, r_1)|, |\min(d_p, r_2) - \min(d_q, r_2)|).$$

To establish inequality (3), we need only show that $|\max(b_p, r_1) - \max(b_q, r_1)| \leq |b_p - b_q|$, and likewise that $|\min(d_p, r_2) - \min(d_q, r_2)| \leq |d_p - d_q|$. Suppose without loss of generality that $b_p \leq b_q$, so that one of the following three cases holds: (i) $b_p \leq b_q \leq r_1$, (ii) $b_p \leq r_1 \leq b_q$, or (iii) $r_1 \leq b_p \leq b_q$. In all three cases, it is immediate that $|\max(b_p, r_1) - \max(b_q, r_1)| \leq |b_p - b_q|$. Similar reasoning shows that $|\min(d_p, r_2) - \min(d_q, r_2)| \leq |d_p - d_q|$.

Second, if either p or q is a point of type **E** or **F**, the non-trivial cases arise when such a point is paired with a point of type **B** or **C**. For instance, suppose $q = (b_q, d_q)$ is of type **B**, so that $\Pi(q) = (b_q, r_2)$. If $p = (b_p, d_p)$ is of type **E**, we have $\Pi(p) = (b_p, b_p)$ and $||\Pi(p) - \Pi(q)||_\infty = \max\{|b_p - b_q|, b_p - r_2\}$. Since $b_q \leq r_2 \leq b_p$, this implies that the horizontal distance between the projections must be larger than the vertical distance. Therefore, $||\Pi(p) - \Pi(q)||_\infty = b_p - b_q \leq \max\{b_p - b_q, |d_p - d_q|\} = ||p - q||_\infty$. If $p = (b_p, d_p)$ is of type **F**, then $\Pi(p) = (d_p, d_p)$ and $||\Pi(p) - \Pi(q)||_\infty = \max\{b_q - d_p, r_2 - d_p\}$. Since $b_p \leq d_p \leq r_1 \leq b_q \leq r_2 \leq d_q$, the horizontal distances satisfy $b_q - d_p \leq b_q - b_p$ and the vertical distances satisfy $r_2 - d_p \leq d_q - d_p$, yielding the desired inequality. The case analysis for the remaining pairings proceeds in a similar manner.

Therefore, if one takes μ to be the infimum over all matchings with respect to the bottleneck distance between $\mathrm{Dg}\mathbf{X}$ and $\mathrm{Dg}\mathbf{Y}$, and the cost of μ_I is smaller, then the bottleneck distance between $\mathrm{Dg}\mathbf{X}^I$ and $\mathrm{Dg}\mathbf{Y}^I$ will only be smaller still. This concludes the proof of the bottleneck distance portion of the theorem. The Wasserstein result follows similarly.□

Corollary 1 *Let $I = [\ell_I, r_I]$, and $J = [\ell_J, r_J]$ be two intervals, with $dist(I, J) = \max\{|\ell_I - \ell_J|, |r_I - r_J|\}$. Given zigzag persistence modules $\mathbf{X}$ and $\mathbf{Y}$, $d_B(\mathrm{Dg}\mathbf{X}^I, \mathrm{Dg}\mathbf{Y}^J) \leq d_B(\mathrm{Dg}\mathbf{X}, \mathrm{Dg}\mathbf{Y}) + dist(I, J)$.*

Proof The result follows from combining Theorem 1 with the inequality $d_B(\mathrm{Dg}\mathbf{X}^I, \mathrm{Dg}\mathbf{X}^J) \leq dist(I, J)$ and the triangle inequality. □

Given an $\mathbb{R}$-valued function, there is a natural construction of a level set zigzag persistence module [8] that sweeps its level sets from bottom to top [18] in a sense that we will describe in the next paragraph. Before going into those details, our motivation for introducing level set zigzag persistence is as follows. The level set zigzag persistence can be used to compute the ordinary persistent homology of an $\mathbb{R}$-valued function with good space efficiency. In particular, the level set zigzag module is related to the ordinary (extended) persistence module via the Mayer-Vietoris pyramid [8, Figure 3], where the zigzag sequence and the ordinary sequence are shown to contain the same information in their persistent homology. Therefore, one could use the algorithm for zigzag persistent homology to compute extended persistence while using space that depends only on the size of the largest level set instead of the entire domain [8, 22]. With this motivation in mind, we now discuss level set zigzag persistence a bit further.

Given a topological space $\mathbb{X}$ and a continuous function $f : \mathbb{X} \to \mathbb{R}$, $(\mathbb{X}, f)$ is said to be of *Morse type* if, for the finite set of critical values $a_1 < a_2 < \ldots < a_n$ of f, the open intervals $(-\infty, a_1), (a_1, a_2), \ldots, (a_{n-1}, a_n), (a_n, \infty)$ are such that for each interval I, $f^{-1}(I)$ is homeomorphic to $\mathbb{Y} \times I$ for some compact and locally connected space $\mathbb{Y}$ with f serving as the projection onto I [8]. The homeomorphisms should extend to continuous functions on $\mathbb{Y} \times \bar{I}$, where $\bar{I}$ is the closure of I in $\mathbb{R}$, and each $\mathbb{X}_t = f^{-1}(t)$ should also have finitely-generated homology, where $\mathbb{X}_t$ denotes the *level set* of f for any $t \in \mathbb{R}$. Let $\mathbb{X}_I = f^{-1}(I)$ denote the *slice* of $\mathbb{X}$ which f maps to the interval $I \subset \mathbb{R}$. If $I = [a, b]$, we may denote this as $\mathbb{X}_a^b$. Then, given $(\mathbb{X}, f)$ of Morse type with critical values a_i as above, we choose arbitrary s_i satisfying

$$-\infty < s_0 < a_1 < s_1 < a_2 < \cdots < s_{n-1} < a_n < s_n < \infty.$$

The *level set zigzag persistence* of $(\mathbb{X}, f)$ is defined to be the zigzag persistence for the sequence

$$\mathbb{X}_{s_0}^{s_0} \to \mathbb{X}_{s_0}^{s_1} \leftarrow \mathbb{X}_{s_1}^{s_1} \to \mathbb{X}_{s_1}^{s_2} \leftarrow \cdots \to \mathbb{X}_{s_{n-1}}^{s_n} \leftarrow \mathbb{X}_{s_n}^{s_n}.$$

We denote the persistence diagram by $\mathrm{Dg}f$.

We now state a straightforward corollary to Theorem 1 in the level set zigzag persistence setting, which we will refer to again in Section 4.

Corollary 2 *Let $f : \mathbb{X} \to \mathbb{R}$ and $g : \mathbb{Y} \to \mathbb{R}$ be Morse type functions defined on topological spaces $\mathbb{X}$ and $\mathbb{Y}$, and for an interval $I = [r_1, r_2]$, let $\mathrm{Dg}f^I$ and $\mathrm{Dg}g^I$ be the restrictions of the level set zigzag persistence diagrams $\mathrm{Dg}f$ and $\mathrm{Dg}g$ to the interval I. Then $d_B(\mathrm{Dg}f^I, \mathrm{Dg}g^I) \leq d_B(\mathrm{Dg}f, \mathrm{Dg}g)$.*

4 Applications to Metric Graphs and d-Parameter Persistence

4.1 Metric Graphs

For uses of Corollary 2, we turn to the *metric graph* setting. Metric graphs commonly arise when studying road networks as well as biological or chemical structure graphs. Given a graph G with a set of vertices and edges, a length function on the edges, and a geometric realization $|G|$ of the graph, one may specify a metric on G by taking the minimum length of any path between any pair of points (not necessarily vertices) in the geometric realization. Given a base point $v \in |G|$, the *geodesic distance function* $f_v : |G| \to \mathbb{R}$ is given by $f_v(x) = d_G(v, x)$. Then $\mathrm{Dg}f_v$ denotes the 0-dimensional level set zigzag persistence diagram induced by f_v. Equivalently, $\mathrm{Dg}f_v$ is the union of the 0- and 1-dimensional *extended persistence diagrams* for f_v (see [13] for the details of extended persistence). Corollary 2 can be used to compare local neighborhoods of two different metric graphs, G_1 and G_2, with base points $v \in G_1$ and $u \in G_2$. In particular, given $f_v : |G_1| \to \mathbb{R}$ and $g_u : |G_2| \to \mathbb{R}$, we have $d_B(\mathrm{Dg}{f_v}^I, \mathrm{Dg}{g_u}^I) \leq d_B(\mathrm{Dg}f_v, \mathrm{Dg}g_u)$ for any real interval I. Typically, for comparing local neighborhoods, $I = [0, r]$. The following corollary gives a stability-type result for comparing two local neighborhoods within a single metric graph.

Corollary 3 *Let G be a metric graph with geometric realization $|G|$. For a fixed interval I and points $u, v \in |G|$, we have $d_B(\mathrm{Dg}f_u^I, \mathrm{Dg}f_v^I) \leq d_G(u, v)$.*

Proof By Corollary 2, $d_B(\mathrm{Dg}{f_u}^I, \mathrm{Dg}{f_v}^I) \leq d_B(\mathrm{Dg}f_u, \mathrm{Dg}f_v)$. Since $f_u, f_v : |G| \to \mathbb{R}$ are two Morse type functions, $d_B(\mathrm{Dg}f_u, \mathrm{Dg}f_v) \leq ||f_u - f_v||_\infty$ by the level set zigzag stability theorem of [8]. Furthermore, by the triangle inequality, for any $x \in |G|$, $|d_G(x, u) - d_G(x, v)| \leq d_G(u, v)$, meaning that $||f_u - f_v||_\infty \leq d_G(u, v)$. Putting everything together proves the claim. □

Another application of Corollary 2 is as follows. Define $\Phi : |G| \to SpDg$, $\Phi(v) = \mathrm{Dg}f_v$, where $SpDg$ denotes the space of persistence diagrams. Given metric graphs (G_1, d_{G_1}) and (G_2, d_{G_2}), their *persistence distortion distance* [15] is

$$d_{PD}(G_1, G_2) := d_H(\Phi(|G_1|), \Phi(|G_2|)),$$

where d_H denotes the Hausdorff distance. In other words,

$$d_{PD}(G_1, G_2) = \max\left\{ \sup_{D_1 \in \Phi(|G_1|)} \inf_{D_2 \in \Phi(|G_2|)} d_B(D_1, D_2), \sup_{D_2 \in \Phi(|G_2|)} \inf_{D_1 \in \Phi(|G_1|)} d_B(D_1, D_2) \right\}.$$

Note that the diagram $\mathrm{Dg}f_v$ contains both 0- and 1-dimensional persistence points, but only points of the same dimension are matched under the bottleneck distance. A local version of the persistence distortion distance, which we will denote by d^r_{PD}, may be defined as follows: for each base point v, only consider the distance function to points within a fixed intrinsic radius r. Specifically, let $\Phi^r : |G| \to SpDg$, $\Phi^r(v) = \mathrm{Dg}f_v^{[0,r]}$. Then $d^r_{PD}(G_1, G_2) := d_H(\Phi^r(|G_1|), \Phi^r(|G_2|)$.

Corollary 4 *If $r \leq r'$, then $d^r_{PD}(G_1, G_2) \leq d^{r'}_{PD}(G_1, G_2)$.*

Proof Let D^r_1 be the persistence diagram for some base point $v \in |G_1|$, where the geodesic distance function is computed in the interval $[0, r]$. Let $D^{r'}_1$ be the persistence diagram for the same base point, but where the distance function is computed in the interval $[0, r']$. Define D^r_2 and $D^{r'}_2$ similarly for some base point in $|G_2|$. By viewing D^r_i as a restriction of $D^{r'}_i$ for $i = 1, 2$, we can apply Theorem 1 to show that $d_B(D^r_1, D^r_2) \leq d_B(D^{r'}_1, D^{r'}_2)$. Since our choice of base points was arbitrary, this inequality holds for persistence diagrams across all choices of base points in $|G_1|$ and $|G_2|$. Therefore, using the definition of the local version of the persistence distortion distance, we can conclude that $d^r_{PD}(G_1, G_2) \le d^{r'}_{PD}(G_1, G_2)$. □

4.2 Multi-Parameter Persistence

Theorem 1 can be applied to d-parameter persistence modules on any topological space (not restricted to the level set or metric graph settings). A d-parameter persistence module is indexed by a d-dimensional family of vector spaces, $\{\mathbf{X}_u\}_{u\in\mathbb{R}^d}$, together with a family of linear maps $\{\rho_{\mathbf{X}}(u,v) : \mathbf{X}_u \to \mathbf{X}_v\}_{u\le v}$ such that for $u \le v \le w \in \mathbb{R}^d$, we have $\rho_{\mathbf{X}}(u,u) = \mathrm{id}_{\mathbf{X}_u}$ and $\rho_{\mathbf{X}}(v,w) \circ \rho_{\mathbf{X}}(u,v) = \rho_{\mathbf{X}}(u,w)$ [9]. Here, $u \le v$ if and only if $u_i \le v_i$ for $i = 1,\ldots,d$, where u_i and v_i are the coordinates of u and v. Any line L in the set of all lines of $\mathbb{R}^d$ with direction $\mathbf{m} = (m_1,\ldots,m_d)$ such that $\min_i m_i$ is strictly positive gives a one-parameter slice of the d-parameter persistence module. Given two d-parameter persistence modules $\mathbf{X}$ and $\mathbf{Y}$, we define their *matching distance* [20] to be

$$d_{match}(\mathbf{X},\mathbf{Y}) := \sup_L \min_i m_i d_B(\mathrm{Dg}\mathbf{X}_L, \mathrm{Dg}\mathbf{Y}_L),$$

where $\mathrm{Dg}\mathbf{X}_L$ and $\mathrm{Dg}\mathbf{Y}_L$ are the persistence diagrams of the d-parameter persistence modules $\mathbf{X}$ and $\mathbf{Y}$ restricted along line L. Our result extends naturally to this linear relationship between these two parameters. Indeed, if we restrict both d-parameter persistence modules to a region $\mathbb{I} = I_1 \times \cdots \times I_d$, denoted $\mathbf{X}^{\mathbb{I}}$ and $\mathbf{Y}^{\mathbb{I}}$, where each I_i is an interval of the real line, then Theorem 1 implies the following corollary. Formally, $\mathbf{X}^{\mathbb{I}}$ is the module $\mathbf{X}$ restricted only to vector spaces in $\{\mathbf{X}_u\}_{u\in\mathbb{I}}$ (similarly for $\mathbf{Y}$).

Corollary 5

$$d_{match}(\boldsymbol{X}^{\mathbb{I}}, \boldsymbol{Y}^{\mathbb{I}}) \le d_{match}(\boldsymbol{X}, \boldsymbol{Y})$$

where $d_{match}(\boldsymbol{X}^{\mathbb{I}}, \boldsymbol{Y}^{\mathbb{I}})$ is computed by restricting $\mathrm{Dg}\boldsymbol{X}_L$ and $\mathrm{Dg}\boldsymbol{Y}_L$ to the subinterval of the line L passing through the region $\mathbb{I}$.

Proof For a fixed line L with direction $\mathbf{m}$, consider a region $\mathbb{I}$ restricted to L, denoted $I_L \subset \mathbb{I} \cap L$. Recall that $\mathrm{Dg}\mathbf{X}_L$ and $\mathrm{Dg}\mathbf{Y}_L$ are the persistence diagrams of the d-parameter persistence modules $\mathbf{X}$ and $\mathbf{Y}$ restricted along the line L. Based on Theorem 1,

$$d_B(\mathrm{Dg}\mathbf{X}_L^{I_L}, \mathrm{Dg}\mathbf{Y}_L^{I_L}) \le d_B(\mathrm{Dg}\mathbf{X}_L, \mathrm{Dg}\mathbf{Y}_L). \tag{4}$$

From the definition of supremum, we know that $\forall\epsilon > 0$, there is a line L_ϵ such that

$$d_{match}(\mathbf{X}^{\mathbb{I}}, \mathbf{Y}^{\mathbb{I}}) - \epsilon < \min_i m_i d_B(\mathrm{Dg}\mathbf{X}_{L_\epsilon}^{I_{L_\epsilon}}, \mathrm{Dg}\mathbf{Y}_{L_\epsilon}^{I_{L_\epsilon}}).$$

Using observation (4) above, we see that

$$d_{match}(\mathbf{X}^{\mathbb{I}}, \mathbf{Y}^{\mathbb{I}}) - \epsilon < \min_i m_i d_B(\mathrm{Dg}\mathbf{X}_{L_\epsilon}, \mathrm{Dg}\mathbf{Y}_{L_\epsilon}).$$

The right-hand side is, of course, less than the supremum over all lines L, the definition of $d_{match}(\mathbf{X},\mathbf{Y})$. Hence, for every $\epsilon > 0$, we have $d_{match}(\mathbf{X}^{\mathbb{I}}, \mathbf{Y}^{\mathbb{I}}) - \epsilon < d_{match}(\mathbf{X},\mathbf{Y})$; in other words, $d_{match}(\mathbf{X}^{\mathbb{I}}, \mathbf{Y}^{\mathbb{I}}) \le d_{match}(\mathbf{X},\mathbf{Y})$, as desired.□

The matching distance between two d-parameter persistence modules requires factoring through all possible 1-parameter persistence modules by restricting to lines. A more direct way to compare two d-parameter persistence modules is the *interleaving distance* which does not rely on persistence diagrams (which are only available for 1-parameter persistence modules). Our definition of the interleaving distance, $d_{\mathcal{I}}$, for d-parameter persistence modules follows the treatment in [16, Sect. 12.2]. We first need the notion of a δ-*interleaving* for two such d-parameter persistence modules $\mathbf{X}$ and $\mathbf{Y}$. For a given $\delta \ge 0$, let $\vec{\delta} = (\delta,\ldots,\delta) \in \mathbb{R}^d$; we use the notation $\mathbf{X}_{\vec{\delta}}$ to represent the module $\mathbf{X}$ shifted diagonally by $\vec{\delta}$. Also, $\rho^{\mathbf{X}}_{u\to v} := \mathbf{X}_{u\le v}$. A δ-*interleaving* between $\mathbf{X}$ and $\mathbf{Y}$ consists of two families of linear maps $\{\varphi_u : \mathbf{X}_u \to \mathbf{Y}_{u+\vec{\delta}}\}_{u\in\mathbb{R}^d}$ and $\{\psi_u : \mathbf{Y}_u \to \mathbf{X}_{u+\vec{\delta}}\}_{u\in\mathbb{R}^d}$ satisfying:

1. (Triangular commutativity) $\forall u \in \mathbb{R}^d$, $\rho^{\mathbf{X}}_{u \to u+2\vec{\delta}} = \psi_{u+\vec{\delta}} \circ \varphi_u$ and $\rho^{\mathbf{Y}}_{u \to u+2\vec{\delta}} = \varphi_{u+\vec{\delta}} \circ \psi_u$.
2. (Rectangular commutativity) $\forall u \leq v \in \mathbb{R}^d$, $\varphi_v \circ \rho^{\mathbf{X}}_{u \to v} = \rho^{\mathbf{Y}}_{u+\vec{\delta} \to v+\vec{\delta}} \circ \varphi_u$ and $\psi_v \circ \rho^{\mathbf{Y}}_{u \to v} = \rho^{\mathbf{X}}_{u+\vec{\delta} \to v+\vec{\delta}} \circ \psi_u$.

See Figure 2 for illustrations of both types of commutativity.

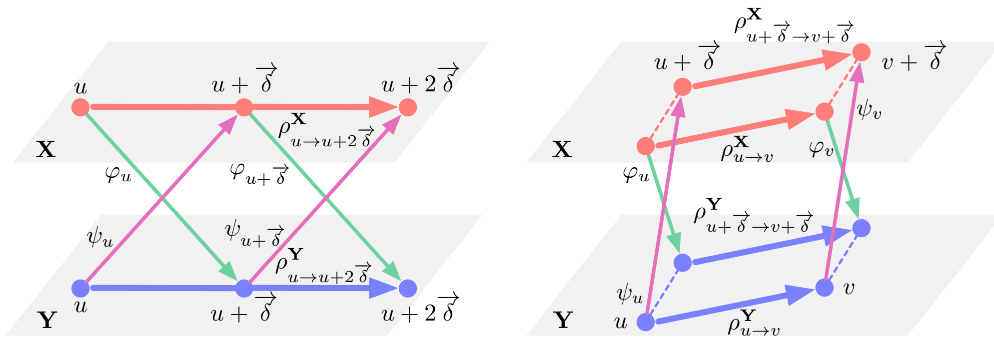

Fig. 2 Triangular commutativity (left) and rectangular commutativity (right) for a δ-interleaving.

Then the *interleaving distance* between $\mathbf{X}$ and $\mathbf{Y}$ is given by

$$d_{\mathcal{I}}(\mathbf{X}, \mathbf{Y}) = \inf_{\delta} \{\mathbf{X} \text{ and } \mathbf{Y} \text{ are } \delta\text{-interleaved}\}.$$

If no such $\delta \in \mathbb{R}^+$ exists, $\mathbf{X}$ and $\mathbf{Y}$ are said to be ∞-interleaved with $d_{\mathcal{I}}(\mathbf{X}, \mathbf{Y}) = \infty$.

Since the interleaving distance is based directly on the d-dimensional modules, and not on zigzag modules or their corresponding persistence diagrams, our Theorem 1 does not apply. However, a restricted distance inequality still holds through a different argument.

Theorem 2 $d_{\mathcal{I}}(\boldsymbol{X}^{\mathbb{I}}, \boldsymbol{Y}^{\mathbb{I}}) \leq d_{\mathcal{I}}(\boldsymbol{X}, \boldsymbol{Y})$.

Proof Let us say that $d_{\mathcal{I}}(\mathbf{X}, \mathbf{Y}) = \delta$ so that $\mathbf{X}$ and $\mathbf{Y}$ are δ-interleaved. Then, if we restrict the modules to the region $\mathbb{I}$ the same interleaving maps used for $\mathbf{X}$ and $\mathbf{Y}$ are still interleaving maps for $\mathbf{X}^{\mathbb{I}}$ and $\mathbf{Y}^{\mathbb{I}}$. But since there are fewer spaces in $\mathbf{X}^{\mathbb{I}}$ and $\mathbf{Y}^{\mathbb{I}}$ there may be additional interleaving maps in the restricted case. Therefore,

$$\{\delta : \mathbf{X} \text{ and } \mathbf{Y} \text{ are } \delta\text{-interleaved}\} \subseteq \{\delta : \mathbf{X}^{\mathbb{I}} \text{ and } \mathbf{Y}^{\mathbb{I}} \text{ are } \delta\text{-interleaved}\},$$

and so we have the desired result,

$$d_{\mathcal{I}}(\mathbf{X}, \mathbf{Y}) = \inf_{\delta} \{\mathbf{X} \text{ and } \mathbf{Y} \text{ are } \delta\text{-interleaved}\} \geq \inf_{\delta} \{\mathbf{X}^{\mathbb{I}} \text{ and } \mathbf{Y}^{\mathbb{I}} \text{ are } \delta\text{-interleaved}\} = d_{\mathcal{I}}(\mathbf{X}^{\mathbb{I}}, \mathbf{Y}^{\mathbb{I}}).$$

□

In the case where $\mathbf{X}$ and $\mathbf{Y}$ are 1-parameter persistence modules, Theorem 2 (together with the fact that $d_{\mathcal{I}}(\mathbf{X}, \mathbf{Y}) = d_B(\mathrm{Dg}\mathbf{X}, \mathrm{Dg}\mathbf{Y})$ [10]) is equivalent to Theorem 1. However, this is only in the case of traditional persistence modules where all linear maps go forwards, not in the more general case of zigzag modules where maps are allowed to go forwards and backwards. So, while special cases of Theorems 1 and 2 are equivalent neither theorem implies the other.

5 Discussion

Theorem 1 and its corollaries, as well as Theorem 2, provide explicit relationships between distances between persistence modules computed locally and globally, and the resulting inequalities are broadly applicable. For instance, the fact that the local bottleneck distance is bounded above by the global bottleneck distance allows for a single global computation to potentially rule out local differences if the global distance is low. If looking for local differences, starting with a global computation may save computational time if there are too many local comparisons to make. On the other hand, the global bottleneck distance being bounded below by the local version allows smaller computations to approach the global truth, while perhaps being more computationally tractable.

In future work, we would like to extend these ideas to generalized persistence, where instead of a linear sequence of topological spaces one considers topological spaces and transformations that form a poset. In contrast to zigzag persistence, this generalized persistence does not have the notion of a persistence diagram. Theorem 2 is the first step in this direction for the case of d-parameter persistence using interleaving distance. We expect similar techniques to be needed in the poset case. However, a notion of "local" would have to be defined in the poset setting.

Acknowledgements

We are grateful for the Women in Computational Topology (WinCompTop) workshop for initiating our research collaboration. In particular, participant travel support was made possible through the grant NSF-DMS-1619908. Our group was also supported by the American Institute of Mathematics (AIM) Structured Quartet Research Ensembles (SQuaRE) program. E.P. was supported by the High Performance Data Analytics (HPDA) program at Pacific Northwest National Laboratory. RS was partially supported by NSF grant DMS-1854705 and the SIMONS Collaboration grant 318086. B.W. was partially funded by NSF-IIS-1513616, NSF-DBI-1661375, and NIH-R01-1R01EB022876-01. Y.W. was supported by NSF grants CCF-1740761 and DMS-1547357. L.Z. was supported by NSF grant CDS&E-MSS-1854703.

References

1. Aanjaneya, M., Chazal, F., Chen, D., Glisse, M., Guibas, L., Morozov, D.: Metric graph reconstruction from noisy data. International Journal of Computational Geometry & Applications **22**(04), 305–325 (2012)
2. Ahmed, M., Fasy, B.T., Wenk, C.: Local persistent homology based distance between maps. Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems pp. 43–52 (2014)
3. Bendich, P., Cohen-Steiner, D., Edelsbrunner, H., Harer, J., Morozov, D.: Inferring local homology from sampled stratified spaces. In: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, pp. 536–546 (2007)
4. Bendich, P., Gasparovic, E., Harer, J., Izmailov, R., Ness, L.: Multi-scale local shape analysis and feature selection in machine learning applications. International Joint Conference on Neural Networks pp. 1–8 (2015)
5. Bendich, P., Wang, B., Mukherjee, S.: Local homology transfer and stratification learning. ACM-SIAM Symposium on Discrete Algorithms pp. 1355–1370 (2012)
6. Bille, P.: A survey on tree edit distance and related problems. Theoretical Computer Science **337**(1-3), 217–239 (2005)
7. Carlsson, G., de Silva, V.: Zigzag persistence. Foundations of Computational Mathematics **10**(4), 367–405 (2010)
8. Carlsson, G., de Silva, V., Morozov, D.: Zigzag persistent homology and real-valued functions. Proceedings of the 25th Annual Symposium on Computational Geometry pp. 247–256 (2009)
9. Carlsson, G., Zomorodian, A.: The theory of multidimensional persistence. Discrete & Computational Geometry **42**(1), 71–93 (2009)
10. Chazal, F., De Silva, V., Glisse, M., Oudot, S.: The structure and stability of persistence modules. Springer (2016)
11. Chernov, A., Kurlin, V.: Reconstructing persistent graph structures from noisy images. Image-a **3**(5), 19–22 (2013)
12. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. Discrete & Computational Geometry **37**, 103–120 (2007)
13. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Extending persistence using poincaré and lefschetz duality. Foundations of Computational Mathematics **9**(1), 79–103 (2009)
14. Cohen-Steiner, D., Edelsbrunner, H., Harer, J., Mileyko, Y.: Lipschitz functions have l_p-stable persistence. Foundations of Computational Mathematics **10**(2), 127–139 (2010)
15. Dey, T.K., Shi, D., Wang, Y.: Comparing graphs via persistence distortion. Proceedings of the 31st International Symposium on Computational Geometry **34**, 491–506 (2015)
16. Dey, T.K., Wang, Y.: Computational Topology for Data Analysis. Cambridge University Press (forthcoming)
17. Edelsbrunner, H., Harer, J.: Persistent homology - a survey. Contemporary Mathematics **453**, 257–282 (2008)
18. Edelsbrunner, H., Morozov, D.: Persistent homology. In: J.E. Goodman, J. O'Rourke, C.D. Tóth (eds.) Handbook of Discrete and Computational Geometry, chap. 24. CRC Press LLC (2017)
19. Harshaw, C.R., Bridges, R.A., Lannacone, M.D., Reed, J.W., Goodall, J.R.: Graphprints: Towards a graph analytic method for network anomaly detection. Proceedings of the 11th Annual Cyber and Information Security Research Conference p. 15 (2016)
20. Landi, C.: The rank invariant stability via interleavings. In: E.W. Chambers, B.T. Fasy, L. Ziegelmeier (eds.) Research in Computational Topology, pp. 1–10. Springer International Publishing, Cham (2018)

21. Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., Alon, U.: Network motifs: Simple building blocks of complex networks. Science **298**(5594), 824–827 (2002)
22. Milosavljević, N., Morozov, D., Skraba, P.: Zigzag persistent homology in matrix multiplication time. Proceedings of the 20th Annual Symposium on Computational Geometry pp. 216–225 (2011)
23. Shen-Orr, S.S., Milo, R., Mangan, S., Alon, U.: Network motifs in the transcriptional regulation network of Escherichia coli. Nature Genetics **31**(1), 64 (2002)
24. Wang, B., Summa, B., Pascucci, V., Vejdemo-Johansson, M.: Branching and circular features in high dimensional data. IEEE Transactions on Visualization and Computer Graphics **17**(12), 1902–1911 (2011)