

Distributions of Matching Distances in Topological Data Analysis

So Mang Han, Taylor Okonek, Nikesh Yadav, Xiaojun Zheng
St. Olaf College, Northfield, Minnesota, USA

January 3, 2020

Abstract

Topological data analysis seeks to discern topological and geometric structure of data, and to understand whether or not certain features of data are significant as opposed to random noise. While progress has been made on statistical techniques for single-parameter persistence, the case of two-parameter persistence, which is highly desirable for real-world applications, has been less studied. This paper provides an accessible introduction to two-parameter persistent homology and presents results about matching distance between 2-parameter persistence modules obtained from families of simple point clouds. Results include observations of how differences in geometric structure of point clouds affect the matching distance between persistence modules. We offer these results as a starting point for the investigation of more complex data.

1 Introduction and Motivation

Topological data analysis (TDA) is a collection of methods used to discern the shape of data. TDA detects topological features, such as clusters, holes, and voids. Topological methods are especially useful for high-dimensional, noisy data. TDA has been applied in numerous settings, including image analysis [1], protein structure [3], texture representation in images [11], astronomical data [14], and neuroscience [15].

One of the main tools in TDA is **persistent homology**. Persistent homology associates to a dataset an algebraic object known as a persistence module, which

encodes topological features of the data. The study of persistence modules can then reveal insights about the data that underlies the modules.

One common problem is to compare two datasets via their persistence modules. In this setting, notions of distance between persistence modules are useful for quantifying the amount of difference between persistence modules. This paper examines one such distance, the matching distance, which is easily computed. Our goal is to understand how the matching distance quantifies similarity between datasets.

We computed matching distances between persistence modules arising from datasets of two types. The first type of dataset consists of three points in various configurations. The second type of dataset consists of two circles with radii r , and circles were separated by a distance d . We examined how changes to r or d affect the matching distance.

The organization of the paper is as follows: In Section 2, we provide mathematical background for persistent homology and the matching distance. In Section 3, we describe our data analysis and matching distance computations. Discussion and directions for future research are provided in Section 4.

2 Mathematical Background

Persistent homology is one of the main tools in TDA and can be applied to many types of data, including real-valued functions and sets of points in Euclidean space. It quantifies multi-scale topological features of data: connected components, holes, voids, and their higher-dimensional analogs.

Previous research has used persistent homology to discern topological structure in data from many fields [1, 3, 11, 14, 15]. Nearly all of this prior work has used one-parameter persistent homology, which produces easily-visualized descriptors called barcodes, but which is sensitive to outliers. This sensitivity can be avoided by using two-parameter persistence. We give here a brief introduction to persistent homology in both the one- and two-parameter settings; more detailed surveys of the subject are found in [6, 7].

2.1 One-parameter persistence

Given a set of point-cloud data, we first build a simplicial complex. Our building blocks are **simplices**: a point is a 0-simplex; an edge is a 1-simplex; a triangular face is a 2-simplex. More formally, an n -simplex is an n -dimensional

geometric object that is the convex hull of $n + 1$ points which are not contained in any $(n - 1)$ -dimensional plane. A **simplicial complex** X is a set of simplices such that if $v \in X$, then every face of v is also in X , and if $v, w \in X$, then $v \cap w$ is also in X . A common type of simplicial complex built from a point cloud is the Rips complex, which we now define.

Definition 1 Given a collection of points $\{x_\alpha\}$ in Euclidean space \mathbb{E}^n and $\epsilon > 0$, the **Rips complex** \mathcal{R}_ϵ is the simplicial complex whose k -simplices correspond to unordered $(k + 1)$ -tuples of points $\{x_\alpha\}_0^k$ whose pairwise distances are at most ϵ .

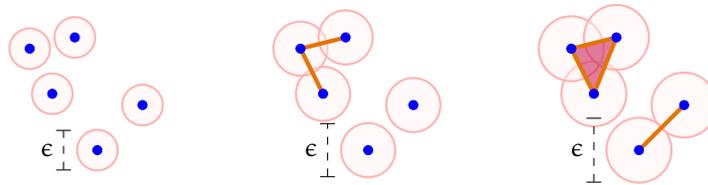


Figure 1: Three Rips complexes built from a dataset of five points, with different scale parameters ϵ .

In other words, the Rips complex depends on a scale parameter ϵ . The complex contains an edge between two points if and only if the distance between the points is at most ϵ . The complex contains a triangular face for any three points whose pairwise distances are at most ϵ . Figure 1 shows three Rips complexes built from the same underlying point cloud, but with different scale parameters. For illustration purposes, we have drawn a ball of diameter ϵ around each data point. The complex contains an edge for each pair of balls that intersect and a triangular face for each three balls that intersect pairwise.

A Rips complex is built with a fixed scale parameter ϵ , but usually no single choice of ϵ reveals all structure of the data. Instead, we consider many Rips complexes, one for *every* positive value ϵ . Imagine growing balls of diameter ϵ centered at each point: as ϵ increases from zero, an n -simplex appears whenever $n + 1$ balls pairwise intersect. Figure 1 shows three snapshots of this process, which leads us to the concept of a filtration.

A **filtration** is a sequence of simplicial complexes, each a subcomplex of the next:

$$X_1 \subset X_2 \subset X_3 \subset \cdots \subset X_n \subset \cdots . \quad (1)$$

Figure 2 illustrates a filtration. In the figure, $X_i \hookrightarrow X_{i+1}$ denotes a map that takes each simplex in X_i to its corresponding simplex in X_{i+1} ; this is possible because X_i is a subcomplex of X_{i+1} . Note that if a simplex appears in X_i , it must be present in X_j for all $j > i$.

If every complex in a filtration is a Rips complex, then we call the filtration a **Rips filtration**. Given a finite point set, simplices appear at only finitely many values of ϵ . Thus, a Rips filtration of a finite point set can be denoted

$$\mathcal{R}_{\epsilon_0} \subset \mathcal{R}_{\epsilon_1} \subset \mathcal{R}_{\epsilon_2} \subset \dots \subset \mathcal{R}_{\epsilon_n}$$

for some sequence $0 = \epsilon_0 < \epsilon_1 < \epsilon_2 < \dots < \epsilon_n$.

Figure 2 illustrates a Rips filtration; note that we have not drawn a complex for every ϵ at which a simplex appears. The shaded areas represent the triangular faces and, in X_6 , the boundary of a 3-simplex.

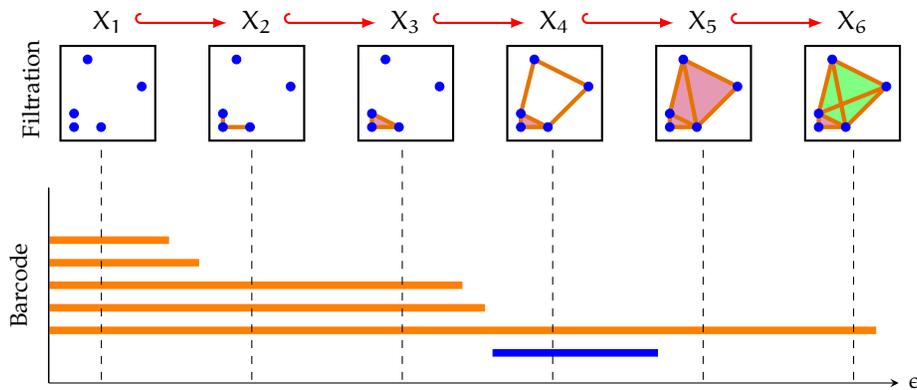


Figure 2: *Top:* Six complexes are shown from a Rips filtration built from a five-point dataset. *Bottom:* The corresponding barcode; orange bars are zeroth homology (which tracks connected components) and blue bars are first homology (which tracks holes).

Any topological feature (such as a component or a hole) appears in the filtration at some scale parameter ϵ_1 and disappears at some scale parameter ϵ_2 ; the pair (ϵ_1, ϵ_2) gives the persistence of the feature. Plotting each persistence pair as an interval along the scale axis produces a **barcode**, as seen at the bottom of Figure 2. The orange bars in Figure 2 represent components: each of the five points corresponds to an orange bar starting at $\epsilon = 0$. An orange bar ends at each ϵ value at which two components become connected. The blue bar represents the hole in the simplicial complex, which persists over a range of ϵ values.

The information in a barcode can also be visualized as a **persistence diagram**, which is a collection of points above the diagonal in the xy -plane. Bars in the barcode are in one-to-one correspondence with points in the persistence diagram. A bar from a to b is plotted as the point (a, b) in the persistence diagram.

In order to quantify the topological features of a simplicial complex, as illus-

trated in the barcode, we use the mathematics of **homology**. Homology associates a vector space to each simplicial complex and a linear map to each inclusion map in the filtration. The homology vector space $H_k(X)$ is generated by the k -dimensional holes of simplicial complex X . Making this precise requires some definitions, which we introduce briefly; for more details, see [16].

Let C_k be a vector space whose basis consists of all k -simplices in simplicial complex X . That is, C_k contains **k -chains**, which are sums of k -simplices with coefficients in a field \mathbb{F} . In TDA, \mathbb{F} is usually chosen to be the 2-element field, a choice we make in this paper as well. The **boundary operator** $\partial_k : C_k \rightarrow C_{k-1}$ maps a k -simplex to the sum of its $(k-1)$ -faces, extending by linearity to k -chains. Let $B_k \subseteq C_k$ be the subspace of **boundaries**, which are images of ∂_{k+1} . Let $Z_k \subseteq C_k$ be the subspace of **cycles**, defined by the property that $v \in Z_k$ if and only if $\partial_k(v) = 0$. Crucially, $B_k \subseteq Z_k$, since $\partial \circ \partial = 0$. We then define the **homology vector space** $H_k = Z_k/B_k$. Thus, H_k is a vector space consisting of all cycles that are not boundaries. The dimension of H_k is the number of equivalence classes of holes, in the sense that two holes are equivalent if they differ by a boundary.

The homology of a filtration is a one-parameter **persistence module**. The inclusion maps in the filtration induce linear maps between the homology vector spaces. Specifically, the degree i homology of the filtration in equation (1) is a persistence module consisting of the following vector spaces and linear maps:

$$H_i(X_1) \rightarrow H_i(X_2) \rightarrow H_i(X_3) \rightarrow \cdots \rightarrow H_i(X_n) \rightarrow \cdots$$

The structure theorem for persistence modules says each persistence module is the sum of interval modules; each interval gives the persistence of one topological feature in the filtration [5]. Thus, a barcode is a visualization of a persistence module, which each interval module shown as a bar.

In order to compare barcodes, we need a notion of distance between barcodes. We use the bottleneck distance, which is easily computable, though other options exist [2]. Before defining the bottleneck distance, we introduce the concept of a matching, which we explain in terms of persistence diagrams.

A **matching** η between persistence diagrams \mathcal{D}_1 and \mathcal{D}_2 pairs each point in \mathcal{D}_1 with a point in \mathcal{D}_2 or a point on the diagonal line, and pairs each point in \mathcal{D}_2 with a point in \mathcal{D}_1 or a point on the diagonal. For an illustration of a matching between two persistence diagrams, see Figure 3. By convention, we use the L_∞ metric to obtain the distance from a point $x = (x_1, x_2)$ to its matched point $\eta(x) = (y_1, y_2)$:

$$\|x - \eta(x)\|_\infty = \max(|x_1 - y_1|, |x_2 - y_2|).$$

Let the *size* of a matching refer to the supremum of the L_∞ distance between matched points. Among all possible matchings, we seek a matching with the

smallest size. The bottleneck distance between \mathcal{D}_1 and \mathcal{D}_2 is the size of this optimal matching, as defined below [6].

Definition 2 The *bottleneck distance* between persistence diagrams \mathcal{D}_1 and \mathcal{D}_2 is

$$d_B(\mathcal{D}_1, \mathcal{D}_2) = \inf_{\eta} \sup_x \|x - \eta(x)\|_{\infty},$$

where the supremum is taken over all matched points x and the infimum is taken over all matchings η .

Figure 3 shows the optimal matching between two persistence diagrams \mathcal{D}_1 and \mathcal{D}_2 . The size of this matching is given by the max L_{∞} distance between matched points, which is $\max(|a - c|, |b - d|)$. Since no other matching between these persistence diagrams has smaller size, the bottleneck distance $d_B(\mathcal{D}_1, \mathcal{D}_2)$ is equal to $\max(|a - c|, |b - d|)$.

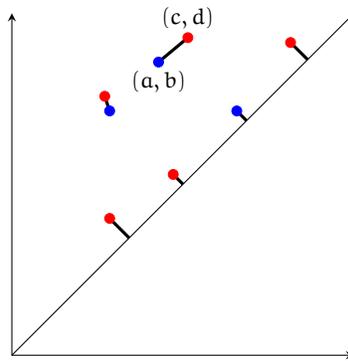


Figure 3: A matching between persistence diagrams \mathcal{D}_1 (plotted in blue) and \mathcal{D}_2 (plotted in red). Some points in each diagram are matched to the diagonal.

While one-parameter persistence is stable with respect to perturbation of the point cloud data, it is unstable with respect to the presence of outliers. A density estimator on the points (i.e., a function that indicates whether each point has many nearby neighbors) might be able to identify outliers, but this requires introducing a threshold. Instead, we prefer to use the density estimator as a second filtration parameter, which brings us into the realm of two-parameter persistence.

2.2 Two-parameter persistence

Two-parameter persistence arises from data that is simultaneously indexed by two parameters. For example, suppose we have a point cloud \mathcal{P} and a real-

valued function $f : \mathcal{P} \rightarrow \mathbb{R}$ on each point. In particular, f may arise from a density estimator. For any $r \in \mathbb{R}$, let

$$f^{-1}(-\infty, r] = \{p \in \mathcal{P} \mid f(p) \leq r\}.$$

We can then construct a Rips filtration from $f^{-1}(-\infty, r]$. Repeating this construction for an increasing sequence r_1, r_2, \dots, r_n , we obtain a sequence of Rips filtrations, which yields a bifiltration.

A **bifiltration** is a set of simplicial complexes, each indexed by two parameters, with inclusion maps in the direction of each increasing parameter. Specifically, the set of simplicial complexes $\{C_{i,j}\}_{i,j}$ forms a bifiltration if there exist commuting inclusion maps

$$C_{i,j} \hookrightarrow C_{i',j'}$$

whenever $i \leq i'$ and $j \leq j'$. Figure 4 (left) gives an example of a bifiltration.

The homology of a bifiltration is a 2-parameter persistence module, which is a set of vector spaces $H_p(C_{i,j})$ with commuting linear maps in the directions of increase of i and j , as illustrated in Figure 4 (right).

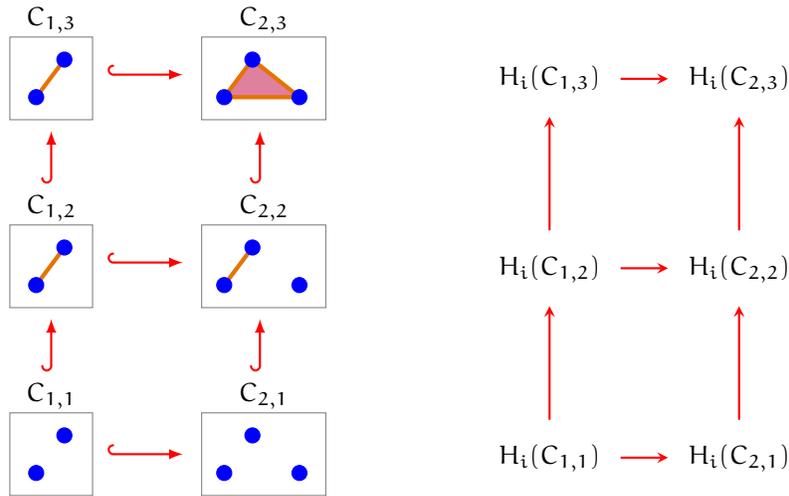


Figure 4: A sample bifiltration (left) and its corresponding 2-parameter persistence module in homology degree i (right).

Unfortunately, the algebraic structure of 2-parameter persistence modules is extremely complicated, and there is no reasonable “barcode” for such modules [9]. Instead, we can obtain a barcode along any line with nonnegative slope in the 2-parameter space by restricting the 2-parameter persistence module to such a line, as we now explain.

Let \mathcal{M} be a 2-parameter persistence module with parameter values in discrete indexing sets I and J . Denote the vector spaces in \mathcal{M} as $M_{i,j}$ for every $i \in I$ and $j \in J$, with linear maps $M_{i,j} \rightarrow M_{i',j'}$ whenever $i \leq i'$ and $j \leq j'$. We may then adopt a continuous perspective, assigning a vector space from \mathcal{M} to every point in $(x, y) \in \mathbb{R}^2$. If $x < \min(I)$ or $y < \min(J)$, then the point (x, y) is assigned the zero vector space; otherwise, the vector space assigned to (x, y) is $M_{a,b}$, where $a = \max\{i \in I \mid i \leq x\}$ and $b = \max\{j \in J \mid j \leq y\}$, as illustrated in Figure 5.

Let ℓ be a line in \mathbb{R}^2 with non-negative slope. Let \mathcal{M}_ℓ be the 1-parameter persistence module obtained by restricting \mathcal{M} to line ℓ : every point along ℓ is assigned the homology vector space of \mathcal{M} at that point in \mathbb{R}^2 , with linear maps induced from \mathcal{M} (as in Figure 5). Since \mathcal{M}_ℓ is a 1-parameter persistence module, it has a barcode, or equivalently, a persistence diagram.

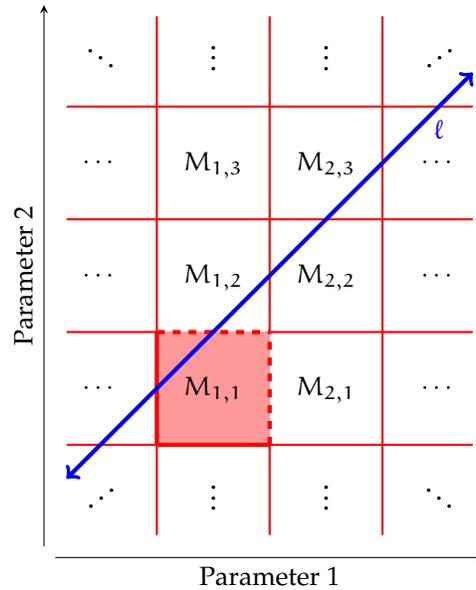


Figure 5: A 2-parameter persistence module with discrete parameter values may be viewed from a continuous perspective. Vector space $M_{1,1}$ is associated with all points inside the shaded region (with half-open boundary) in the 2-parameter plane, and similarly for the other vector spaces. The vector spaces along any line ℓ with nonnegative slope then form a 1-parameter persistence module, with linear maps induced from the 2-parameter persistence module.

Furthermore, we can define a distance between 2-parameter persistence modules by considering the bottleneck distances between persistence diagrams along all possible lines through the 2-parameter space. In the following definition,

$\mathcal{D}(\mathcal{M}_\ell)$ denotes the persistence diagram of the 1-parameter module \mathcal{M}_ℓ .

Definition 3 The **matching distance**, d_M , between two 2-parameter persistence modules \mathcal{M} and \mathcal{N} is the supremum of the bottleneck distances between the persistence diagrams on corresponding lines of non-negative slope in the two modules. Precisely,

$$d_M = \sup_{\ell} \{d_B(\mathcal{D}(\mathcal{M}_\ell), \mathcal{D}(\mathcal{N}_\ell)) \cdot \text{weight}(\text{slope}(\ell))\},$$

where the supremum is over all lines of nonnegative slope and $\text{weight}(m) = \frac{1}{\sqrt{1+q^2}}$, where $q = \max(m, \frac{1}{m})$.

In the definition of the matching distance, a weight is assigned to each line ℓ , which depends on the slope ℓ . A line with slope 1 gets the maximum weight, and the weight approaches zero as the slope approaches zero or infinity. The weight is chosen such that if the interleaving distance between two persistence modules is 1, then the weighted bottleneck distance is at most 1 [8, 9].¹

3 Computations and Analysis

Our datasets are point clouds with simple structure depending on a few parameters. Adjusting these parameters allows us to change the size of topological features (namely, components and holes) that are captured by the persistence modules. We expect the matching distance between these persistence modules to reflect the parameter differences in the underlying datasets.

From each dataset, we construct a **density-Rips bifiltration**. That is, our two parameters are a density estimator and Euclidean distance. A density estimator f is assigned to each point p such that $f(p)$ is small if p has many nearby neighbors and $f(p)$ is large if p has no nearby neighbors; this causes points with more neighbors to appear before points with few neighbors in the density filtration.² For any $r \in \mathbb{R}$, the Rips filtration is constructed on $f^{-1}(-\infty, r]$. This produces a 2-parameter family of simplicial complexes with inclusion maps in the increasing directions of both density and distance.

We begin with very simple point-cloud datasets, each consisting of three points in the xy -plane. We examine the matching distances while keeping two points fixed and moving the third point around the plane. Three points are enough

¹Prior to computing the matching distance, the persistence modules \mathcal{M} and \mathcal{N} are often *normalized*. That is, the parameter axes for each module are rescaled so that the parameter values for all generators and relations occur in specified intervals on each axis. For details, see [13].

²We used a k -nearest-neighbor density estimator, but many other options are available.

to produce a density-Rips bifiltration that is nontrivial in both density and distance. Thus, we regard these datasets as the simplest datasets that allow us to study the effect of moving a single point on the matching distance.

The second collection of datasets consists of points sampled from two circles with radii r and separated by distance d . These datasets have nontrivial structure in persistent homology of degree zero and one. We generated many datasets with various values of d and fixed r , investigating how varying d affects the matching distance computed from zero-degree persistent homology modules. We also fix d and vary r , investigating the matching distance between first-degree persistent homology modules.

We computed two-parameter persistent homology using RIVET, an interactive visualization software developed by Michael Lesnick and Matthew Wright [12]. A detailed description of RIVET and its algorithms appears in a comprehensive preprint by Lesnick and Wright [10]. Given point-cloud data, RIVET computes a 2-parameter persistence module, and then computes barcodes along linear slices of the persistence module. We approximated the matching distance from these barcodes using Python code written by Bryn Keller and Michael Lesnick [13], which uses a finite set of lines to approximate the supremum in Definition 3.³ We regard the matching distance as giving a measure of similarity between two point-cloud datasets.

3.1 Three-Point Datasets

For our first investigation, each data set consisted of two fixed points $A = (1, 1)$ and $B = (6.1, 1)$, and a third point C_i in the xy -plane. Figure 6 shows the locations of A and B (in red), as well as the locations of all C_i (blue). Since A and C_i are always the closest pair of points, we assigned to these points a density parameter of 1, and then we assigned point B a density parameter of 2.

For concreteness, let $X_{r,s}$ denote a 3-point dataset where C_i has coordinates (r, s) . That is, $X_{r,s} = \{A, B, (r, s)\}$. Let $X_{t,u} = \{A, B, (t, u)\}$ be another 3-point dataset. We compute the matching distance between the 2-parameter persistence modules constructed from $X_{r,s}$ and $X_{t,u}$.

First, we fix $s = u = 3$. Figure 7 displays the matching distance between the

³The approximation algorithm requires us to specify the number of lines used in the approximation. This is achieved by specifying a *grid-size* parameter, which determines the number of different slope and intercept values that the algorithm uses. For example, if grid-size is 20, then the algorithm uses 20 slope values and 20 intercept values to produce 400 lines, computing the bottleneck distance along each. In our experiments, we found we found little difference in the approximated matching distance when grid-size was set to 20 or to a large value, such as 50, though the computation time increases according to the square of the grid-size value.

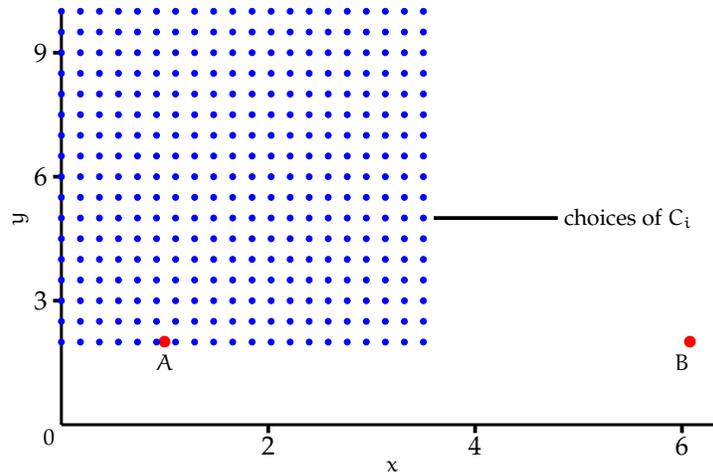


Figure 6: Data sets used in three point exploration. The red points A and B are two fixed points, and the blue points are the possibilities for the point C_i .

2-parameter persistence modules constructed from $X_{r,3}$ and $X_{t,3}$, for various choices of r and t . The horizontal axis gives r , and the color of each curve represents the value of t . Specifically, t ranges from 0, colored dark green, to 3.3, colored brown, and the step size is .184.

We observe that when s and u are small, such as $s = u = 3$, as shown in Figure 7, there is a relatively linear increase in matching distances as the t increases (above some threshold) with r fixed. Some of the curves display a nonlinear region for small values of t ; We note that these curves have t smaller than 1, which is the x -coordinate for A.

Next, we fix $s = u = 10$. Figure 7 displays the matching distance between the 2-parameter persistence modules constructed from $X_{r,10}$ and $X_{t,10}$, for various choices of r and t . Similarly, Figure 8 shows how the matching distance depends on r and t in this case. It is clear that no linear trend in matching distances is present in these cases, as the matching distance increases faster as r increases. We note a prominent feature in Figure 8 is that the matching distance attains the value 0; the following proposition explains why this occurs.

Proposition 1 *Suppose A and B are points on the xy -plane, such that the Euclidean distance between A and B is $d > 0$. Now suppose there are two points C_1 and C_2 , both distance $r < d$ away from A, with the distance between C_i and B as $h_i > d$ for $i \in \{1, 2\}$. Then the matching distance between the 2-parameter persistence modules constructed from the two point clouds $\{A, B, C_1\}$ and $\{A, B, C_2\}$ is 0. (See Figure 9)*

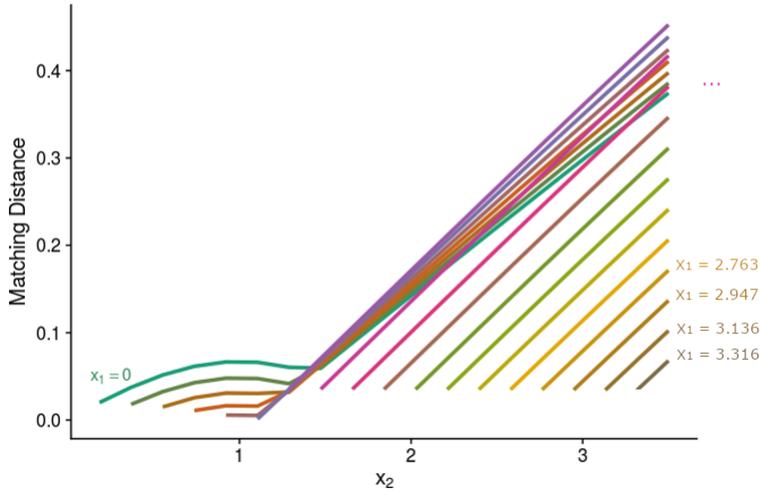


Figure 7: Distribution of matching distance between $X_{r,3}$ and $X_{t,3}$. The horizontal axis gives the value of t , the color of the plot gives the value of r , and the vertical axis gives the matching distance between two datasets.

PROOF: Since the distance between each C_i and A is smaller than the distance between each C_i and B , the density parameter 1 is assigned to C_1 , C_2 , and A , while point B is assigned density parameter 2. From such a point cloud $\{A, B, C_i\}$, we construct a bifiltration as shown on the left side of Figure 10, as we now explain, considering each density parameter value in turn.

Density 1: When the distance parameter $\epsilon = 0$, only C_1 and A appear, so we have two isolated points. When ϵ increases to r , C_1 and A will connect to form an edge. As ϵ increases from r , the simplicial complex remains unchanged, since no other points exist to produce edges at this density.

Density 2: When the scale parameter $\epsilon = 0$, all points A, B, C_1 in the point cloud appear, so we have three isolated points. When ϵ increases to r , C_1 and A will form an edge. This results in one connected component and one isolated point. When ϵ increases to d , an edge connects B to A . Since all points are connected at distance d , H_0 homology doesn't change as ϵ increases further.

Now consider the point cloud $\{A, B, C_2\}$. The bifiltration constructed from this point cloud is nearly the same as that described above; the only difference is the distance value h_2 at which the longest edge appears. However, this edge does not connect any components that were not already connected at distance d , so no new (zeroth) homology appears at distance h_2 . Thus, we obtain topologically equivalent bifiltrations for the two data sets. This implies that the two

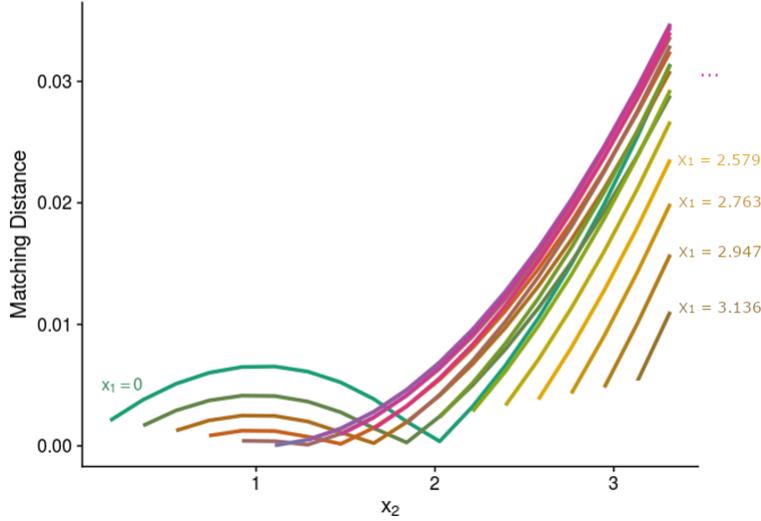


Figure 8: Distribution of matching distance between $X_{r,10}$ and $X_{t,10}$. The horizontal axis gives the value of t , the color of the plot gives the value of r , and the vertical axis gives the matching distance between two datasets.

2-parameter persistence modules are the *same*, and so barcodes along any linear slice of the two modules are also the same. Therefore, the matching distance between these modules is 0. ■

According to the proposition above, it would be easy to compute the probability of having matching distance of 0 if C_1 and C_2 are randomly selected on circle O_A and circle O_B in Figure 9. The following corollary gives such a probabilistic interpretation of the proposition; we leave the proof to the reader.

Corollary 1 *Suppose A and B are points on the xy -plane with distance between them $d > 0$. Let O_A be a circle of radius $r < d$ centered at A , and O_B a circle of radius $r < d$ centered at B . Now suppose there are two points C_1 and C_2 , which are randomly selected on the two circles. Then, the two point clouds $\{A, B, C_1\}$ and $\{A, B, C_2\}$ clouds have a matching distance of 0 with probability $(1 - \frac{1}{\pi} \arccos^{-1}(\frac{r}{2d}))^2$.*

The following theorem gives a n -point generalization of the 3-point proposition. In the theorem, a dataset consists of n vertices of a regular polygon in the xy -plane, as well as one additional point.

Theorem 1 *Suppose A_1, A_2, \dots, A_n are points on the xy -plane forming the vertices*

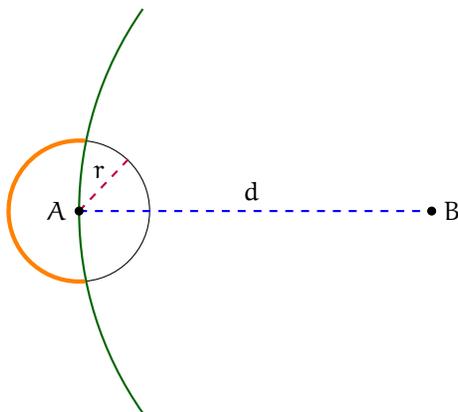


Figure 9: Diagram of the points in Proposition 1. Points C_1 and C_2 are located in the orange arc, at distance r from A and distance greater than d from B .

of a regular polygon with side length $d > 0$. Now suppose there are two points C_1 and C_2 , both distance $r > 0$ (also $r < d$) away from the A_i that they are closest to, with the distance away from all other A_i greater than d . Then, the matching distance between the 2-parameter persistence modules constructed from the two point clouds $\{A_1, A_2, \dots, A_n, C_1\}$, and $\{A_1, A_2, \dots, A_n, C_2\}$ will be 0.

PROOF: Similar to the proof for the proposition above, each point will be assigned a density parameter of 1 or 2. Since the distance between points C_1, C_2 and the respective A_i that they are closest to is smaller than the distance between the C_i and all other A_i , density 1 is assigned to C_1, C_2 , and the respective A_i that they are closest to (note that they could both be closest to the same A_i). All other A_i are assigned a density of 2. In the point clouds consisting of these points, we have a bifiltration for both density 1 and density 2.

In the bifiltration for density 1, only C_1 and the A_i that it is closest to appear, so the proof follows the same steps as the one for proposition.

In the bifiltration for density 2, all points in the point cloud appear (all A_i and C_1). At $\epsilon = r$, C_1 and the A_i it is closest to form an edge, but no other edges form between any points. As ϵ increases to d , all edges of the regular polygon appear.

Similarly, we have the exact same bifiltration for the point cloud consisting of C_2 and the A_i . Followed by the same reasonings in the proof for proposition, the matching distance between the two point clouds is 0. ■

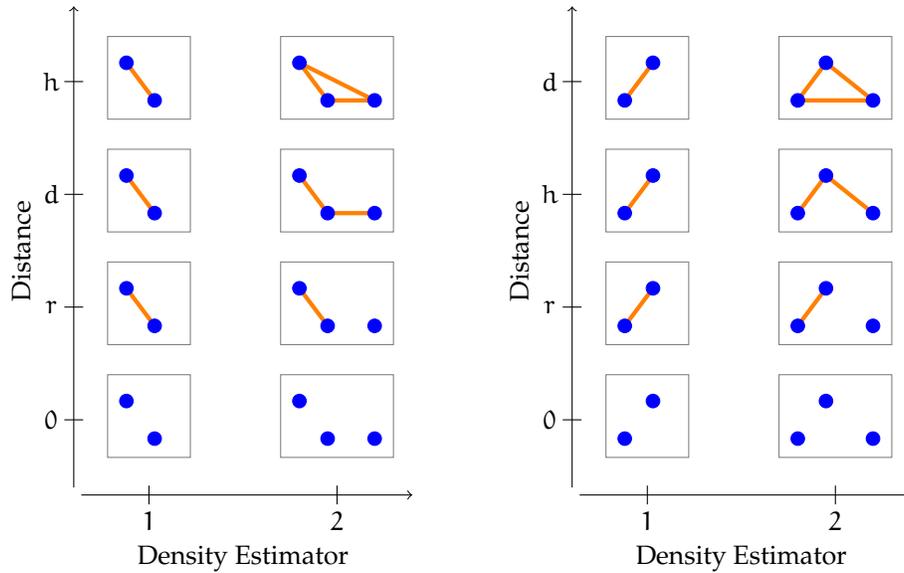


Figure 10: At left, the bifiltration constructed in the proof of Proposition 1. At right, a similar bifiltration that results in a nonzero matching distance.

3.2 Two-Circle Datasets

Our second investigation involved datasets consisting of points sampled from two circles with radii r , separated by distance d . For each dataset, two hundred points were selected randomly from a uniform distribution on each circle; Figure 11 displays an example. Each point was assigned a density estimator defined as the distance to the 20th nearest neighbor.

Suppose that we fix the radius r and vary the separation distance d . Intuitively, the larger the change in separation distance, the more different we regard the point clouds. Thus, we expect that a large change in d will result in a large matching distance between the 2-parameter persistence modules constructed from the point clouds.

We generated 60 datasets using $r = 3$ and $d \in \{0.5, 1, 1.5, \dots, 30\}$. We computed 2-parameter persistence modules for each dataset and computed the matching distance between each pair of modules.

Consider a pair of datasets, one with separation distance d_1 and the other with separation distance d_2 . Figure 12 displays the matching distance between the persistence modules plotted against the *difference* in separation distance $d_2 - d_1$. For the plot, we chose four representative values of d_2 from those listed above:

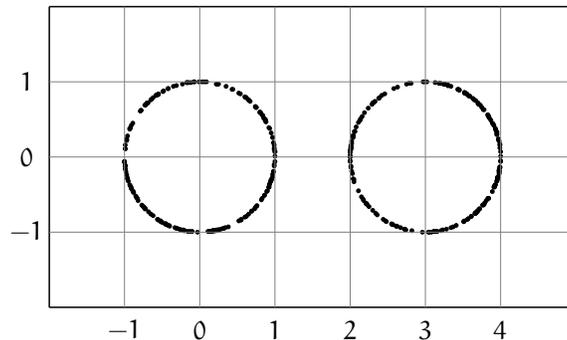


Figure 11: Example of circle dataset; the circles have radius $r = 1$ and are separated by distance $d = 1$.)

$d_2 \in \{10, 15, 22, 28\}$. For each of these four values of d_2 , we plot the matching distance for those datasets with separation distance d_1 such that the difference $d_2 - d_1$ in the range from 0.5 to $d_2 - 0.5$.

Figure 12 (left) shows the trend when d_2 is 10 or 15. As the separation distance $d_2 - d_1$ increases, the matching distance of the pairs of data sets increases initially, but then remains nearly constant. In the region of increasing matching distance, we note three linear segments, with a jump between each. Though the plots display nearly the same shape for $d_2 = 10$ and $d_2 = 15$, we note that the matching distance attains a higher value for $d_2 = 15$, indicating that the matching distance reveals a greater difference between the persistence modules when one dataset consists of circles that are farther apart.

When the separation distance d_2 is 22 or 28, as shown in Figure 12 (right), the matching distance of that pair follows almost the same trend as in Figure 12 (left). However, when $d_2 = 22$, we note some randomness, possibly due to sampling irregularities, when $d_2 - d_1$ is small. Also, we note only two linearly increasing segments, with a single jump, when $d_2 = 28$. Interestingly, the near-constant part of the plot is higher when d_2 is the smaller of the two values, which is contrary to what we observed in Figure 12 (left).

In order to understand the structure observed in Figure 12, we looked into the barcodes involved in the matching distance calculation. Recall that the matching distance between two 2-parameter persistence modules is the minimum weighted bottleneck distance between two barcodes, one from each persistence module. RIVET provides us access to these barcodes.

For example, consider the pair of datasets for $d_1 = 8$ and $d_2 = 10$. The matching distance between the persistence modules is plotted as one of the red dots in Figure 12 (left). The barcodes that realize this matching distance are shown

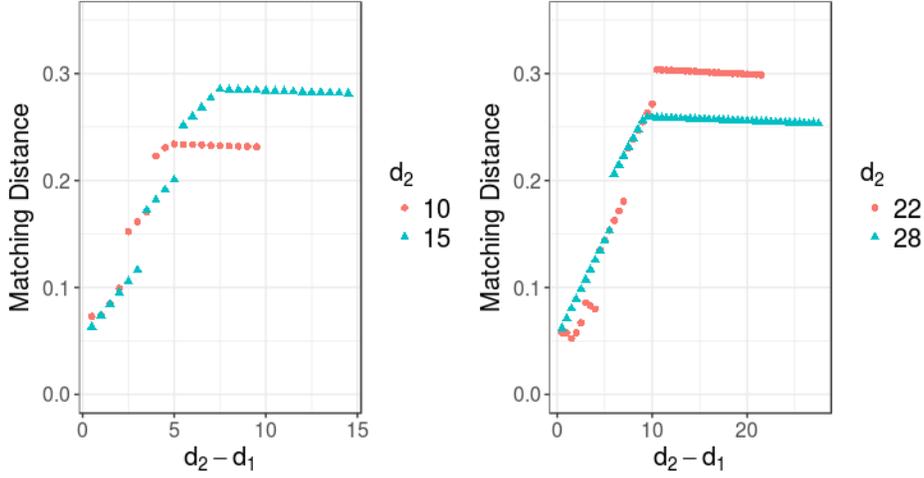


Figure 12: Distribution of matching distance for two-circle example

(as persistence diagrams) in Figure 13. We compared these persistence diagrams to understand the matching distance between the persistence modules. We observe that the persistence diagram for the dataset with separation distance $d_1 = 8$ has only three dots with finite c coordinates⁴ (i.e., the barcode has only three finite b ars). In comparison, the persistence diagram for the dataset with separation distance $d_2 = 10$ has five dots with finite coordinates (i.e., the barcode has only five finite bars). Given the small number of points, we manually matched the points and computed the matching distance. In this example, the two dots far from the diagonal are matched together in the optimal matching, and the distance between that pair is what gives the matching distance. By identifying how points are matched in persistence diagrams, we could better understand the structure in Figure 12.

To understand the near-constant part of the plot, we looked at the persistence diagram that realize the matching distances between pairs of datasets. Let \mathcal{M}_d be the 2-parameter persistence module computed from the dataset with separation distance d . In Figure 12 (left), the near-constant portion of dots for $d_2 = 10$ extends from $d_2 - d_1 = 5$ (that is, $d_1 = 5$) to $d_2 - d_1 = 9.5$ (that is, $d_1 = 0.5$). We examined the calculation of the matching distance between \mathcal{M}_5 and \mathcal{M}_{10} , and also the matching distance between $\mathcal{M}_{0.5}$ and \mathcal{M}_{10} . We found that in both matching distance calculations, the line ℓ that minimizes the bottleneck distance (thus realizing the matching distance) is the same. Let \mathcal{D}_d^ℓ be the persistence diagram obtained from \mathcal{M}_d along line ℓ . We found that the barcode \mathcal{D}_{10}^ℓ has only one finite b ar. When matching \mathcal{D}_{10}^ℓ to $\mathcal{D}_{0.5}^\ell$ or \mathcal{D}_5^ℓ , this

⁴Our persistence diagrams for zero-degree persistent homology always have a dot at $(0, \infty)$ since there is one connected component that persists at all distance scales.

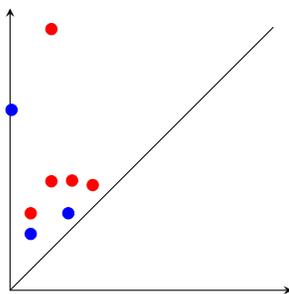


Figure 13: Persistence diagrams involved in the matching distance calculation for $d_1 = 8$ (blue) and $d_2 = 10$ (red).

finite bar is always matched to the diagonal, and this gives the maximum distance between matched points. Thus, for the persistence modules that produce the near-constant portion of the plot, the matching distance is determined by the distance of the finite bar in \mathcal{D}_{10}^ℓ from the diagonal.

Since circles have nontrivial first homology, we performed a second experiment to investigate the effect of varying the circle radius on the matching distance between first persistent homology modules. Specifically, we generated data sets with circle radius $r \in \{0.2, 0.4, 0.6, \dots, 6\}$ and fixed separation distance $d = 3$. We computed the 2-parameter persistence modules using first homology for each dataset, and compared them pointwise using the matching distance. In this experiment, we observed patterns very similar to those displayed in Figure 12: as the difference in radius increases, the matching distance first increases, but then becomes near-constant. Furthermore, the near-constant portion is slightly decreasing as the difference in radius increases.

Lastly, we wanted to determine whether our two-circle experiments are sensitive to the presence of outliers, given that an important advantage of 2-parameter persistent homology is robustness against noise. In our previous experiments, we sampled points precisely from the circles, but now we introduced some noise. We added a small error to 20% or 40% of the data points, and recomputed the matching distances. We found that the distribution of the matching distance for noisy data shares almost the same features as for the original data sets. We conclude that the matching distance still provides information about the change in separation distance or circle radius, even in the presence of noisy data. This confirms that two-parameter persistent homology is robust against outliers, which is one of the primary motivations for this study.

4 Discussion and Future Research

Our findings, though for very simple datasets, suggest that the matching distance can provide a notion of similarity for point-cloud data. This research provides a step towards a deeper understanding of what matching distance reveals about the similarity or difference between point-cloud datasets. Moreover, this leads to further questions regarding how to quantify the similarity between geometric data.

In order to better understand our results in Figure 12, we would like to study why the jumps appear in the increasing segments in the plot of matching distances. We would like to determine why the near-constant part of the plot is slightly decreasing as $d_2 - d_1$ increases. We are intrigued by the fact that the value of the matching distance along the near-constant part of the plot initially increases with d_2 , but then decreases when d_2 gets sufficiently large. In our experiments, the maximum matching distance occurs when $d_2 = 21$. We would like to study this further.

In this work, we obtained 1-parameter persistence modules by restricting 2-parameter persistence modules along lines of nonnegative slope. We would like to explore the structures that exist along lines of negative slope, but this is algebraically complicated and would likely involve zigzag persistence, as discussed in [4].

Furthermore, we would like to extend this research to the analysis of real-world data. To give one example, textual data such as Wikipedia articles can be converted to high-dimensional vectors — e.g., using a word2vec algorithm. We would like to use the matching distance to explore similarities between articles, and to compare collections of the article vectors with random point-cloud data.

Acknowledgements

We would like to thank Professor Matthew Wright and Professor Matthew Richey for their guidance in this project. We would also like to thank the anonymous reviewer for many suggestions that greatly improved this paper. In addition we are grateful for the Collaborative Undergraduate Research and Inquiry (CURI) program for the generous support of undergraduate research at St. Olaf College. This work was supported by NSF DMS-1606967 and NSF DMS-1045015.

References

- [1] A. ASAAD AND S. JASSIM, *Topological data analysis for image tampering detection*, in International Workshop on Digital Watermarking, Springer, 2017, pp. 136–146.
- [2] P. BUBENIK, V. DE SILVA, AND J. SCOTT, *Metrics for generalized persistence modules*, Foundations of Computational Mathematics, 15 (2015), pp. 1501–1531.
- [3] Z. CANG, L. MU, K. WU, K. OPRON, K. XIA, AND G.-W. WEI, *A topological approach for protein classification*, Computational and Mathematical Biophysics, 3 (2015).
- [4] G. CARLSSON AND V. DE SILVA, *Zigzag persistence*, Foundations of Computational Mathematics, 10 (2010), pp. 367–405.
- [5] F. CHAZAL, W. CRAWLEY-BOEVEY, AND V. DE SILVA, *The observable structure of persistence modules*, ArXiv e-prints, (2014).
- [6] H. EDELSBRUNNER AND J. HARER, *Persistent homology - a survey*, Contemporary Mathematics, 453 (2008), pp. 257–282.
- [7] R. GHRIST, *Barcodes: The persistent topology of data*, Bulletin of the American Mathematical Society, 45 (2008), pp. 61–75.
- [8] C. LANDI, *The rank invariant stability via interleavings*, ArXiv e-prints, (2014).
- [9] M. LESNICK, *The theory of the interleaving distance on multidimensional persistence modules*, Foundations of Computational Mathematics, 15 (2015), pp. 613–650.
- [10] M. LESNICK AND M. WRIGHT, *Interactive Visualization of 2-D Persistence Modules*, ArXiv e-prints, (2015).
- [11] J. A. PEREA AND G. CARLSSON, *A klein-bottle-based dictionary for texture representation*, International journal of computer vision, 107 (2014), pp. 75–97.
- [12] RIVET: *The rank invariant visualization and exploration tool*. <http://rivet.online>, 2015-2018.
- [13] *Rivet-Python: Python api for rivet*. <https://github.com/rivetTDA/rivet-python>, 2018.
- [14] P. ROSEN, B. WANG, A. SETH, B. MILLS, A. GINSBURG, J. KAMENETZKY, J. KERN, AND C. R. JOHNSON, *Using contour trees in the analysis and visualization of radio astronomy data cubes*, ArXiv e-prints, (2017).
- [15] A. E. SIZEMORE, J. E. PHILLIPS-CREMINS, R. GHRIST, AND D. S. BASSETT, *The importance of the whole: Topological data analysis for the network neuroscientist*, Network Neuroscience, 3 (2019), pp. 656–673.
- [16] L. WASSERMAN, *Topological data analysis*, Annual Review of Statistics and Its Application, 5 (2018), pp. 501–532.