



Approximating Continuous Functions on Persistence Diagrams Using Template Functions

Jose A. Perea¹ · Elizabeth Munch² · Firas A. Khasawneh³

Received: 16 December 2020 / Revised: 8 October 2021 / Accepted: 28 March 2022 © SFoCM 2022

Abstract

The persistence diagram is an increasingly useful tool from Topological Data Analysis, but its use alongside typical machine learning techniques requires mathematical finesse. The most success to date has come from methods that map persistence diagrams into vector spaces, in a way which maximizes the structure preserved. This process is commonly referred to as featurization. In this paper, we describe a mathematical framework for featurization called *template functions*, and we show that it addresses the problem of approximating continuous functions on compact subsets of the space of persistence diagrams. Specifically, we begin by characterizing relative compactness with respect to the bottleneck distance, and then provide explicit theoretical methods for constructing compact-open dense subsets of continuous functions on persistence diagrams. These dense subsets—obtained via template functions—are leveraged for supervised learning tasks with persistence diagrams. Specifically, we test the method for classification and regression algorithms on several examples including shape data and dynamical systems.

Communicated by Shmuel Weinberger.

Elizabeth Munch muncheli@msu.edu

Firas A. Khasawneh khasawn3@egr.msu.edu

Published online: 01 June 2022

- Department of Mathematics and Khoury College of Computer Sciences, Northeastern University, 43 Leon St, Boston, MA 02115, USA
- Department of Computational Mathematics, Science and Engineering and Department of Mathematics, Michigan State University, 428 S Shaw Ln, East Lansing, MI 48824, USA
- Department of Mechanical Engineering, Michigan State University, 474 S Shaw Ln, East Lansing, MI 48824, USA





Keywords Topological data analysis \cdot Persistent homology \cdot Machine learning \cdot Featurization \cdot Bottleneck distance

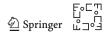
Mathematics Subject Classification $55N31 \cdot 30L05 \cdot 68T05$

1 Introduction

Many machine learning tasks can be reduced to the following problem: Approximate a continuous function defined on a topological space, the "ground truth," given the function values (or approximations thereof) on some subset of the points. This task has been studied extensively for data in Euclidean space, but more work is necessary to extend these ideas to arbitrary topological spaces. In this paper, we focus on the task of learning continuous functions on the space \mathcal{D} of persistence diagrams endowed with the bottleneck distance d_B , as input to supervised learning (e.g., regression and classification) on \mathcal{D} .

Persistence diagrams are mathematical objects arising in the field of Topological Data Analysis (TDA). They are signatures giving insight into the underlying structure of data sets, and typically arise in the following pipeline. Given a sequence of simplicial complexes $K_1 \subseteq K_2 \subseteq \cdots \subseteq K_n$ obtained, for example, by connecting the points in a data set at increasing proximity scales (e.g., as in the Vietoris-Rips filtration), then their homology with coefficients in a field and the maps induced by the inclusions $K_i \hookrightarrow K_{i+1}$ yield a sequence $H_p(K_1) \to H_p(K_2) \to \cdots \to H_p(K_n)$ of vector spaces and linear transformations. One can then interrogate this sequence to determine when homological features appear (are born) and disappear (die), and encode each such feature as a point (birth, death) in a so-called persistence diagram. It is possible to weaken the implicit finiteness assumptions of the input data and instead start with a family $\{K_a \mid a \in \mathbb{R}\}$ of spaces with $K_a \subseteq K_b$ for $a \leq b$, giving rise to a so-called persistence module $\{H_p(K_a) \mid a \in \mathbb{R}\}$ with linear maps $\phi_a^b: H_p(K_a) \to H_p(K_b)$ such that ϕ_a^a is the identity map of $H_p(K_a)$ and $\phi_b^c \phi_a^b = \phi_a^c$ for $a \le b \le c$. To simplify definitions, we will assume that our input diagrams are defined over positive indices $a \ge 0$. There has been extensive study on the various niceness restrictions that can be placed on a persistence module [12]. With enough assumptions, a persistence module can be represented up to isomorphism as a persistence diagram, which is simply a collection of points with multiplicity in the "wedge" $\mathbb{W} = \{(x, y) \in \mathbb{R}^2 \mid 0 \le x < y\}$. Persistence diagrams are particularly useful for data analysis due to the availability of metrics, and their stability [45].

The downside of all this mathematical structure is that the geometry of the space of persistence diagrams (\mathcal{D}, d_B) is not directly amenable to the application of existing machine learning methodologies. Thus, methods for utilizing persistence diagrams in statistics and machine learning contexts have taken two basic forms. The first are attempts at working with the persistence diagrams directly; however, the issues with the geometry (particularly the lack of unique means) mean that the work in this direction is rather limited. More recently, a great deal of success has been found in *featurization*; that is, transforming each persistence diagram into a point in a vector space, in a way that preserves as much of the structure as possible. The work in this paper is part of



the latter category, and provides a new method for featurization which sits on a solid mathematical foundation with respect to the structure of (\mathcal{D}, d_B) .

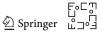
Mathematically, we are working with the following framework. Suppose one has a set $S \subset \mathcal{D}$, typically compact, and a continuous function $F: S \longrightarrow \mathbb{R}$ encoding the ground truth of the phenomenon under study. **Our goal** is to devise provably-correct and computationally feasible approaches to approximating F, given a finite sample $D_1, \ldots, D_n \in S$ and their values $F(D_1), \ldots, F(D_n) \in \mathbb{R}$. This encompasses, for instance, supervised learning tasks such as regression and classification. The problem at hand is thus (1) to characterize compactness in (\mathcal{D}, d_B) , for these are the sets where sequential approximations are guaranteed to converge, (2) to construct dense subsets of the space of continuous functions from \mathcal{D} to \mathbb{R} , for these will be the search space for approximation (i.e., supervised learning) tasks, and (3) to devise algorithms using said families to approximate real valued functions on compact subsets of \mathcal{D} .

1.1 Our Contribution

The first contribution of this paper is a characterization of (relative) compactness for subsets of \mathcal{D} with respect to d_B (see Fig. 2 and Theorem 3.7). These results can be viewed in parallel to the characterization in [48] of relative compactness with respect to the Wasserstein distance, though our results are shown to capture different phenomena (see Sect. 3.6). Our characterization also comes with some unexpected consequences for the topology of (\mathcal{D}, d_B) : (1) Every compact subset of \mathcal{D} has empty interior (hence \mathcal{D} is not locally compact); (2) \mathcal{D} cannot be written as a countable union of compact subsets; and if $C(\mathcal{D}, \mathbb{R})$ denotes the set of continuous functions from \mathcal{D} to \mathbb{R} , then (3) the compact-open topology on $C(\mathcal{D}, \mathbb{R})$ —which captures approximations on compact subsets of \mathcal{D} —is not metrizable. The main consequence of this last point is the impossibility of purely metric-based objective functions for approximations (e.g., supervised learning) in $C(\mathcal{D}, \mathbb{R})$ with respect to the topology of convergence on compact sets.

To circumvent this, we turn our attention to the problem of finding compact-open dense subsets of $C(\mathcal{D}, \mathbb{R})$. Ideally, the elements of these sets should be succinctly represented (e.g., with a few parameters) and efficiently searched (e.g., via appropriate optimization routines), in order to devise general computational schemes. Our second contribution is a methodology for constructing infinitely many examples of said families. The strategy goes as follows: First, we continuously embed \mathcal{D} in an appropriate topological vector space V—consistent with the monoidal structure of \mathcal{D} given by disjoint union \square of multisets—and then restrict the continuous \mathbb{R} -linear maps on V to yield elements in $C(\mathcal{D}, \mathbb{R})$. Specifically, let $C_c(\mathbb{W})$ denote the set of compactly supported continuous functions from \mathbb{W} to \mathbb{R} , endowed with the *strict inductive limit topology* (see Sect. 4.3). Let $C_c(\mathbb{W})'$ be its topological dual, endowed with the corresponding *weak-* topology*, and let

$$\nu: \mathcal{D} \longrightarrow C_c(\mathbb{W})'$$



be the function which assigns to each $D \in \mathcal{D}$ the Radon measure on \mathbb{W} consisting of a Dirac delta mass at each $\mathbf{x} \in D$. We show in Theorem 4.9 that ν is continuous, injective and satisfies $\nu(D \sqcup D') = \nu(D) + \nu(D')$ for all $D, D' \in \mathcal{D}$, providing the aforementioned linear embedding. Hence, each bounded linear operator $T: C_c(\mathbb{W})' \longrightarrow \mathbb{R}$ yields a map $T \circ \nu \in C(\mathcal{D}, \mathbb{R})$ —a feature, in machine learning parlance—which respects the monoidal structure \square of \mathcal{D} .

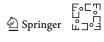
We show (see Theorem 5.1 and discussion thereafter) that each such T uniquely determines an $f \in C_c(\mathbb{W})$, and viceversa, so that $T \circ \nu(D)$ is exactly the same as integrating f against the Radon measure $\nu(D)$. We then show how to construct compact-open dense subsets of $C(\mathcal{D}, \mathbb{R})$ by taking countably many dilations and translations of any nonzero $f \in C_c(\mathbb{W})$ (Theorems 5.3 and 5.4). This is why we refer to the elements of $C_c(\mathbb{W})$ as *template functions*.

As the final contribution of this paper, we provide two explicit families of template functions—called respectively tent functions and interpolating polynomials (see Sect. 6)—so that the algebras they generate in $C(\mathcal{D}, \mathbb{R})$ are compact-open dense. We then provide algorithms to perform regularized regression and classification using template functions (in Sect. 7), and finally, we compare tent functions and interpolating polynomials in several tasks including shape classification and inference in dynamical systems (Sect. 8).

1.2 Related Work

Existing methods for applying statistics and machine learning methods to persistence diagrams can be loosely divided into two categories. The first attempts to work in the space of persistence diagrams directly. This can be done by studying the Fréchet mean for collections of diagrams [48, 49, 62], confidence sets for persistence diagrams [33], or simply passing the Wasserstein or bottleneck distance matrix to metric learning methods for classification tasks (e.g. [46]). The main issue with this viewpoint is that the geometry of the space of persistence diagrams is ill-behaved [64], so directly working with these objects is often not advisable.

The second collection of methods maps the space of persistence diagrams into another, more well-behaved space where available mathematical machinery can be readily applied. Our work on approximating continuous functions of persistence diagrams with templates fits into this category. These approaches take inspiration from different viewpoints including algebraic geometry [2, 14, 28, 39], functional representations [8, 10, 20, 51], image encodings [1, 21, 31, 34, 56], path representations [22], kernel methods [3, 15–17, 26, 40–44, 55, 67, 68], and other more ad hoc methodologies [5, 18, 23, 50, 59, 69, 70]. Our work is most closely related to the persistence images of [1], where each point in a persistence diagram contributes to a Gaussian bump, and the sum of these Gaussians provides a function on the upper half plane. Our work is, in some sense, dual to this idea, where we start with bump functions and evaluate this function at the points of the persistence diagram.



1.3 Outline

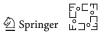
We go over the background needed for understanding persistence diagrams in Sect. 2. In Sect. 3, we give a full characterization of compact sets in persistence diagram space with the bottleneck distance (Sect. 3.7); the reader more interested in the featurization method than in its mathematical justification may safely skip this section. We provide the mathematical justification for the template functions in Sect. 4, and fit this into a function approximation scheme in Sect. 5. In Sect. 6 we give two options for template functions: tent functions and (Chebyshev) interpolating polynomials. In Sect. 7 we fit these into a regression framework. We give results of our experiments in Sect. 8 and discuss implications and future directions in Sect. 9.

2 Basics

Traditionally, persistence diagrams arise in the course of the following procedure. Given a function $f: \mathbb{X} \to \mathbb{R}$ on a topological space \mathbb{X} , denote the sublevel set by $\mathbb{X}_a = f^{-1}(-\infty, a]$. For example, given a point cloud $X \subseteq \mathbb{R}^d$, one can define $f: \mathbb{R}^d \to \mathbb{R}$ as $f(y) = \inf_{x \in X} \|x - y\|$. Such a function induces a filtration $\mathbb{X}_a \subseteq \mathbb{X}_b$, $a \leq b$, on \mathbb{X} and applying k-dimensional homology yields the persistence module $(H_k(\mathbb{X}_a), \phi_a^b)$. Namely, the collection of vector spaces $H_k(\mathbb{X}_a)$ with induced maps $\phi_a^b: H_k(\mathbb{X}_a) \to H_k(\mathbb{X}_b)$ for all $a \leq b$. In full generality, persistence modules can simply be viewed as a collection of vector spaces and linear maps $\mathcal{V} = (V_a, \phi_a^b)$ where $\phi_a^b: V_a \to V_b, \phi_a^a = \mathbb{1}_{V_a}$, and $\phi_b^c \phi_a^b = \phi_a^c$.

Under appropriate tameness conditions, a persistence module can be decomposed uniquely. The pieces of the decomposition are called *interval modules*; these are persistence modules $\mathcal{I}_U = (I_a, i_a^b)$ where $U \subset \mathbb{R}$ is an (open, closed or half-open) interval with end-points $-\infty \leq r \leq s \leq \infty$, $I_a = \mathbf{k}$ if $a \in U$ and 0 otherwise. The maps i_a^b are identities whenever possible. A persistence module $\mathcal{V} = (V_a, \phi_a^b)$ is called pointwise-finite if V_a is finite dimensional for every a. Every pointwise-finite persistence module decomposes uniquely as a direct sum of interval modules, $\mathcal{V} = \bigoplus_{U \in \mathcal{A}} \mathcal{I}_U$ [27]. This decomposition is often visualized as a persistence diagram as seen in the center panel of Fig. 1. The diagram consists of a point at $(r,s) \in (\mathbb{R} \cup \{\pm \infty\})^2$ for the endpoints of each $U \in \mathcal{A}$. We remark that diagonal points (r,r) will be discarded as they encode non-persistent features (these are called ephemeral modules), and are not detected by our template functions $f \in C_c(\mathbb{W})$ for in this case $\operatorname{supp}(f) \subset \mathbb{W}$ does not intersect the diagonal. For the same reason, points (r,s) with $\{r,s\} \cap \{\pm \infty\} \neq \emptyset$ (i.e., with infinite persistence) will also be discarded.

For simplicity of notation and definitions, we will be working with positive diagrams: those with $r \ge 0$ for each point (r, s) in the persistence diagram. Note that we can always convert a finite set of persistence diagrams to have this assumption by shifting the index by adding the minimum birth value. For infinite sets, we can reindex using any positive, monotone increasing function of the original index parameter to achieve the same results.



¹ Homology is computed with coefficients in a field **k**.



Fig. 1 An example point cloud is shown at left, with its persistence diagram shown in the middle. At right, we show the conversion of the persistence diagram into the birth-lifetime plane which is used throughout this paper

Certain visualizations will be on birth-lifetime coordinates, consisting of a point at $(r, s - r) \in \mathbb{R}^2$ for each interval $U \in \mathcal{A}$ with endpoints (either closed or open) $r \leq s$. See the right of Fig. 1 for an example. We note that the methods here developed also apply to zigzag persistence modules. These consist of vector spaces V_a and linear maps between them that can point in either direction: i.e., $V_a \to V_b$ or $V_a \leftarrow V_b$ for $a \leq b$. It turns out that Zigzag modules can also be decomposed as sums of (zigzag) intervals modules [13], which in turn can be represented as persistence diagrams.

2.1 The Space of Persistence Diagrams

A persistence diagram D, then, can be thought of as a collection of points

$$S \subset \{(x, y) \in \mathbb{R}^2 \mid 0 \le x < y\}$$

with a notion of multiplicity, which we write as a function

$$\mu: S \to \mathbb{N} = \{1, 2, \ldots\}.$$

We will often write $D = (S, \mu)$. In order to make statements about the structure of the space of persistence diagrams we will need a few notions.

Definition 2.1 Given $D = (S, \mu)$ and $U \subset \mathbb{R}^2$, the **multiplicity of** D **in** U is

$$\mathsf{Mult}(D,U) = \begin{cases} \sum\limits_{\mathbf{x} \in S \cap U} \mu(\mathbf{x}) & \text{if this is finite, or} \\ \infty & \text{else.} \end{cases}$$

Other common notions we will use repeatedly are as follows. The diagonal is denoted $\Delta := \{(x, x) \in \mathbb{R}^2 \mid x \ge 0\}$, and the wedge is

$$\mathbb{W} := \left\{ (x, y) \in \mathbb{R}^2 \mid 0 \le x < y \right\}.$$

Note that the boundary of \mathbb{W} is Δ , but it is not included in \mathbb{W} . The persistence of a point $\mathbf{x} = (x, y) \in \mathbb{W}$ is $\mathsf{pers}(\mathbf{x}) = y - x$, and the portion of \mathbb{W} where persistence is greater than ε is

$$\mathbb{W}^{\varepsilon} := \{ \mathbf{x} \in \mathbb{W} \mid \mathsf{pers}(\mathbf{x}) > \varepsilon \}.$$

Note that the (ε -diagonal) lower boundary is not included; in order to do so we write

$$\overline{\mathbb{W}^{\varepsilon}} = \{\mathbf{x} \in \mathbb{W} \mid \mathsf{pers}(\mathbf{x}) \ge \varepsilon\}.$$

If we want to work with the portion of points in $D = (S, \mu)$ in a region $U \subseteq \mathbb{R}^2$, we write

$$D \cap U := (S \cap U, \mu|_{S \cap U}).$$

We further abuse notation by writing $D \subset U$ if $S \subset U$. If $S = \emptyset$, we follow the settheoretic convention $\mu = \emptyset$ and denote by $\emptyset = (\emptyset, \emptyset)$ the resulting (empty) persistence diagram. For the sake of figures, we sometimes plot persistence diagrams in the birthlifetime plane. That is, we plot $\mathbf{x} = (x, y)$ at the point $(x, y - x) = (x, \mathsf{pers}(\mathbf{x}))$. In this representation, Δ gets mapped to the x-axis. It should be noted that this transformation is different from the rotation used by [10] and [1]. With these notions we define the space of persistence diagrams as follows.

Definition 2.2 The **space of persistence diagrams**, denoted \mathcal{D} , is the collection of pairs $D = (S, \mu)$ where:

- 1. $S \subset \mathbb{W}$ is the **underlying set of** D, and $\mu : S \to \mathbb{N} = \{1, 2, ...\}$ encodes the **multiplicity** $\mu(\mathbf{x}) \in \mathbb{N}$ of each $\mathbf{x} \in S$.
- 2. $\mathsf{Mult}(D, \mathbb{W}^{\varepsilon}) < \infty \text{ for any } \epsilon > 0.$

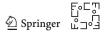
The space of finite persistence diagrams is $\mathcal{D}_0 := \{(S, \mu) \in \mathcal{D} \mid S \text{ is finite}\}.$

Finite persistence diagrams were the first to appear in the literature, and many current papers implicitly assume finiteness. We do not do so here, as diagrams with infinitely many points can be used to encode fractal behavior—e.g., in attractors from dynamical systems. Since we will be interested in studying subsets of \mathcal{D} , we will extend Definition 2.1 as follows.

Definition 2.3 Given $S \subset \mathcal{D}$ and $U \subset \mathbb{R}^2$, then the total multiplicity of S in U is

$$\mathsf{Mult}(\mathcal{S},U) = \begin{cases} \sum_{D \in \mathcal{S}} \mathsf{Mult}(D,U) & \text{if this is finite, or} \\ \infty & \text{else.} \end{cases}$$

In particular, for each $D \in \mathcal{D}$ we write Mult(D, U) instead of $Mult(\{D\}, U)$.



2.2 The Bottleneck Distance

The space of persistence diagrams can be endowed with a distance, which we now define. To each persistence diagram $D = (S, \mu)$ one can associate a set

$$S_{\mu} := \left\{ (\mathbf{x}, k) \mid \mathbf{x} \in S \text{ and } 1 \le k \le \mu(\mathbf{x}) \right\}$$
 (1)

obtained by replicating the elements of S and decorating them with integer labels according to their multiplicity. A **partial matching** between two persistence diagrams (S, μ) , (T, α) is a bijection (with notation as in Eq. 1)

$$\begin{array}{ccc} \mathsf{M} : A & \longrightarrow & B \\ & & & & \\ & & & & \\ S_{\mu} & & & T_{\alpha} \end{array}$$

between a subset A of S_{μ} and a subset B of T_{α} . If $(\mathbf{y}, n) = \mathsf{M}(\mathbf{x}, k)$ we say that (\mathbf{x}, k) is matched with (\mathbf{y}, n) and, conversely, that (\mathbf{y}, n) is matched with (\mathbf{x}, k) . If (\mathbf{z}, m) is in either $S_{\mu} \setminus A$ or $T_{\alpha} \setminus B$, then we call it unmatched. Given $\delta > 0$, a partial matching M between (S, μ) and (T, α) is a δ -matching if two things happen:

- 1. If $(\mathbf{x}, k) \in S_{\mu}$ and $(\mathbf{y}, n) = \mathsf{M}(\mathbf{x}, k)$ are matched, then $\|\mathbf{x} \mathbf{y}\|_{\infty} < \delta$, where $\|(x_1, x_2)\|_{\infty} = \max\{|x_1|, |x_2|\}$ denotes the L^{∞} norm on \mathbb{R}^2 .
- 2. If $(\mathbf{z}, m) \in S_{\mu} \cup T_{\alpha}$ is unmatched, then pers $(\mathbf{z}) < 2\delta$.

Definition 2.4 The **bottleneck distance**, $d_B : \mathcal{D} \times \mathcal{D} \longrightarrow [0, \infty)$, is given by

$$d_B(D, D') := \inf \{ \delta > 0 \mid \text{there is a } \delta\text{-matching between } D \text{ and } D' \}$$

It has been shown that d_B defines a metric on \mathcal{D} [24], and that \mathcal{D} is the metric completion of \mathcal{D}_0 [9].

For simplicity (though without loss of generality) we assume in this paper that all persistence points are finite; i.e. for each point (b,d) in the diagram, the lifetime d-b is finite. That is, we discount homological features with infinite lifetime which occur when $d=\infty$. The assumptions on $\mathcal D$ make it so that the bottleneck distance is still finite between diagrams in $\mathcal D$. In particular, this comes from the triangular inequality and the fact that

$$d_B(D,\varnothing) = \frac{1}{2} \max{\{\mathsf{pers}(\mathbf{x}) \mid \mathbf{x} \in S\}}, \quad \text{for all } D = (S,\mu) \in \mathcal{D}.$$

3 Compactness in \mathcal{D}

Our first contribution is Thm. 3.7, which gives a criterion for characterizing (relatively) compact sets in (\mathcal{D}, d_B) . This work can be viewed in parallel to Theorem 21 of [48], which does the same using the related Wasserstein distance d_{W_p} for persistence diagrams. For other structural properties of families of persistence modules see [12].

Definition 3.1 A subspace of a topological space is **relatively compact** if its closure is compact.

In what follows we will provide a criterion to check whether a subset $S \subseteq \mathcal{D}$ is relatively compact. Specifically, we will show that a set is relatively compact if it satisfies the following three properties.

3.1 Bounded

The first property of interest is boundedness. A subset of a metric space is said to be bounded if it is contained in an open ball of finite radius. Let

$$B_C(D) := \{ D' \in \mathcal{D} \mid d_B(D, D') < C \}$$

denote the ball of radius C > 0 about the diagram D. In particular, it can be seen from the definition that $S \subset D$ is bounded if and only if there exists C > 0 so that $S \subset B_C(\emptyset)$.

Proposition 3.2 *Relatively compact subsets of* (\mathcal{D}, d_B) *are bounded.*

Proof Let $S \subseteq \mathcal{D}$ be relatively compact. To see that S is in fact bounded, consider the cover $\{B_1(D) \mid D \in \overline{S}\}$ by open balls of radius 1 and let $\{B_1(D_i)\}_{i=1}^N$ be a finite subcover. If

$$C > 1 + \max\{d_B(D_i, \varnothing) \mid 1 \le j \le N\}$$

then it follows that $S \subset \overline{S} \subset B_C(\emptyset)$, as claimed.

Note that this proof works for a general metric space, but we work in \mathcal{D} for clarity.

3.2 Off-Diagonally Birth Bounded

The second property of interest controls persistence diagrams with unbounded birth.

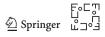
Definition 3.3 A set $S \subset D$ is said to be **off-diagonally birth bounded (ODBB)** if for every $\varepsilon > 0$ there exists a constant $C_{\varepsilon} \geq 0$ so that if $\mathbf{x} \in S \cap \overline{\mathbb{W}^{\varepsilon}}$ (i.e., $\mathsf{pers}(\mathbf{x}) \geq \varepsilon$) for $(S, \mu) \in S$, then $\mathsf{birth}(\mathbf{x}) \leq C_{\varepsilon}$.

See Fig. 2 for a visualization of the notation.

Proposition 3.4 *Relatively compact subsets of* (\mathcal{D}, d_B) *are ODBB.*

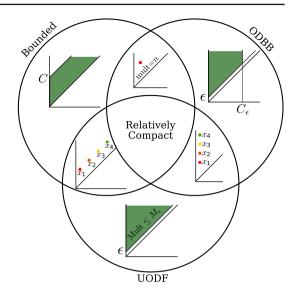
Proof By way of contradiction, assume $S \subseteq \mathcal{D}$ is relatively compact but not ODBB. Then there exist (i) $\varepsilon > 0$; (ii) a sequence $\{D_n\}_{n \in \mathbb{N}} \subseteq S$ with $D_n = (S_n, \mu_n)$; (iii) a fixed diagram $D = (S, \mu) \in \mathcal{D}$ such that

$$\lim_{n\to\infty} D_n = D,$$



П

Fig. 2 The three criterion for compact sets with examples given on intersections



which exists because $\overline{\mathcal{S}} \subset \mathcal{D}$ is compact; and (iv) a chosen point in each diagram D_n , namely $\mathbf{x}_n \in S_n$, with $\mathsf{pers}(\mathbf{x}_n) \geq \varepsilon$ for all $n \in \mathbb{N}$, and $\lim_{n \to \infty} \mathsf{birth}(\mathbf{x}_n) = \infty$.

Let $\delta < \frac{\varepsilon}{2}$ and let $N \in \mathbb{N}$ be large enough so that $d_B(D_n, D) < \delta$ for all $n \ge N$. For each $n \ge N$ fix a δ -matching

Since $\mathbf{x}_n \in S_n$ has $\mathsf{pers}(\mathbf{x}_n) \geq \varepsilon$ and $\delta < \varepsilon/2$, then $(\mathbf{x}_n, 1) \in A_n$. Let $\mathbf{y}_n \in S$ be such that $\gamma_n(\mathbf{x}_n, 1) = (\mathbf{y}_n, k_n)$. As γ_n is a δ -matching, then $\|\mathbf{x}_n - \mathbf{y}_n\|_{\infty} < \delta$ which means, in particular, that $\mathsf{pers}(\mathbf{y}_n) \geq \delta$. Hence $\{\mathbf{y}_n\}_{n \in \mathbb{N}} \subseteq S$ is an infinite set in \mathbb{W}^{δ} , contradicting $D \in \mathcal{D}$.

3.3 Uniformly Off-Diagonally Finite

The final property of interest controls the multiplicity of points across all diagrams:

Definition 3.5 A set $S \subset D$ is said to be **uniformly off-diagonally finite (UODF)** if for every $\varepsilon > 0$ there exists $M_{\varepsilon} \in \mathbb{N}$ so that

$$\mathsf{Mult}\left(D,\overline{\mathbb{W}^{arepsilon}}
ight) \leq \ \mathit{M}_{arepsilon}$$

for all $D \in \mathcal{S}$.

Again, see Fig. 2 for a visualization of the notation.

Proposition 3.6 *Relatively compact subsets of* \mathcal{D} *are uniformly off-diagonally finite.*

Proof By the contrapositive, if $S \subset D$ is not UODF, then there exist $\varepsilon > 0$ and a sequence $\{D_n\}_{n \in \mathbb{N}} \subset S$ so that

$$\mathsf{Mult}\left(D_n,\overline{\mathbb{W}^arepsilon}
ight)<\mathsf{Mult}\left(D_{n+1},\overline{\mathbb{W}^arepsilon}
ight)$$

for all $n \in \mathbb{N}$. In particular, using the pigeonhole principle, any partial matching between D_n and D_{n+1} must have at least one point in D_{n+1} unmatched. As this point has persistence greater than ε , it follows that $d_B(D_n, D_{n+1}) \geq \varepsilon$ for every $n \in \mathbb{N}$ and therefore $\{D_n\}_{n \in \mathbb{N}}$ cannot have a convergent subsequence. This shows that $\overline{\mathcal{S}}$ is not compact.

3.4 Helpful Counterexamples

The three conditions bounded, off-diagonally birth bounded and uniformly off-diagonally finite are independent. Indeed, here are three examples of sets $S = \{(S_n, \mu_n) \mid n \in \mathbb{N}\}$ which satisfy only two out of the three conditions. See Fig. 2.

- 1. Bounded and ODBB, but not UODF. $S_n = \{(0, 1)\}$ with $\mu_n(0, 1) = n$.
- 2. Bounded and UODF, but not ODBB. $S_n = \{(n, n+1)\}$ with $\mu_n(n, n+1) = 1$.
- 3. UODF and ODBB, but not bounded. $S_n = \{(0, n)\}$ with $\mu_n(0, n) = 1$.

3.5 Characterizing Compactness in (\mathcal{D}, d_B)

With these definitions, we can now state our main compactness theorem.

Theorem 3.7 (Characterization of compactness in (\mathcal{D}, d_B)) A set $\mathcal{S} \subset \mathcal{D}$ is relatively compact in (\mathcal{D}, d_B) if and only if it is bounded, off-diagonally birth bounded (ODBB) and uniformly off-diagonally finite (UODF).

Note that one direction is already provided by Propositions 3.2, 3.4 and 3.6, so our main job is to show that a set which satisfies the three conditions is relatively compact. Before we prove this, however, we will need to build a bit of machinery. First, notice that if $S \subset \mathcal{D}$ is bounded ($S \subset B_C(\emptyset)$) and ODBB, then there exist a collection of finite "boxes" \mathbb{B}_k in \mathbb{W} whose union contain all points in the diagrams. Specifically, if S is bounded and ODBB, then there is a C > 0 and $\{C_k\}_{k \in \mathbb{N}} \subset \mathbb{R}_{>0}$ non-decreasing, so that if

$$\mathbb{B}_k = \left\{ \mathbf{x} \in \mathbb{W} \mid 0 \le \mathsf{birth}(\mathbf{x}) \le C_k \quad \text{and} \quad \frac{C}{k+1} < \mathsf{pers}(\mathbf{x}) \le \frac{C}{k} \right\}$$
 (2)

then for all $D \in \overline{\mathcal{S}}$, $D \subset \bigcup_{k \in \mathbb{N}} \mathbb{B}_k$. While these are parallelograms in the birth-death plane, they become rectangles in the birth-lifetime plane, hence the moniker "box", see Fig. 3.

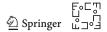
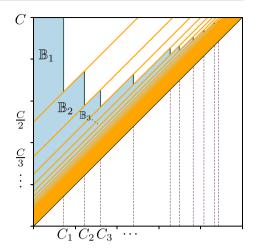


Fig. 3 An example to show the notation used in Eq. 2



Second, we can control the multiplicity of the diagrams in these boxes. Indeed, if $S \subset \mathcal{D}$ is simultaneously bounded, ODBB, and UODF, then there exists a sequence $\{M_k\}_{k\in\mathbb{N}}\subset\mathbb{N}$ so that for every $D\in\overline{\mathcal{S}}$ and every $k\in\mathbb{N}$ we have that

$$Mult(D, \mathbb{B}_k) < M_k$$
 (3)

We can use this to prove the following useful, technical lemma:

Lemma 3.8 Let $R \subset \mathbb{W}$ be a relatively compact subset of \mathbb{R}^2 . If $\{D_n\}_{n\in\mathbb{N}} \subset \mathcal{D}$ is so that

$$\sup_{n\in\mathbb{N}} \; \mathit{Mult}(D_n,R) < \infty$$

then the restricted sequence $\{D_n \cap R\}_{n \in \mathbb{N}}$ has a convergent subsequence. Specifically, there exists a diagram $D \in \mathcal{D}$ with $D \subset \overline{R}$ and a strictly increasing function $\varphi : \mathbb{N} \to \mathbb{N}$ so that

$$\lim_{n\to\infty} \left(D_{\varphi(n)} \cap R \right) = D.$$

Proof Let $M \in \mathbb{N}$ be so that

$$Mult(D_n, R) < M$$

for every $n \in \mathbb{N}$, which exists by hypothesis. For each $n \in \mathbb{N}$, let $\mathbf{m}_n \in \{0, \dots, M\}^M$ be the vector having as entries the integers $\mu_n(\mathbf{x})$ for $\mathbf{x} \in D_n \cap R$, sorted in descending order and padded with zeros at the end as necessary. The pigeonhole principle implies that there is a vector $\mathbf{m} \in \{0, \dots, M\}^M$ which repeats infinitely often. That is, there is a strictly increasing function $\phi : \mathbb{N} \longrightarrow \mathbb{N}$ such that $\mathbf{m} = \mathbf{m}_{\phi(n)}$ for all $n \in \mathbb{N}$.

If **m** is the zero vector, then $D_{\phi(n)} \cap R = \emptyset$ for all $n \in \mathbb{N}$, and thus we can let $\varphi = \phi$, $D = \emptyset$, which trivially satisfy

$$\lim_{n\to\infty} (D_{\varphi(n)}\cap R) = \varnothing.$$

If on the other hand **m** is nonzero, then let $J \ge 1$ be so that

$$\mathbf{m} = (m_1, \dots, m_J, 0, \dots, 0)$$

and $M \geq m_1 \geq m_2 \geq \cdots \geq m_J > 0$. For $n \in \mathbb{N}$, order the underlying set of $D_{\phi(n)} \cap R$ as $\{\mathbf{x}_1^n, \cdots, \mathbf{x}_J^n\}$ in such a way that $\mu_{\phi(n)}(\mathbf{x}_j^n) = m_j$ for all $j = 1, \ldots, J$. Then the collection $\{(\mathbf{x}_1^n, \ldots, \mathbf{x}_J^n)\}_{n \in \mathbb{N}}$ is an infinite sequence in R^J which, by compactness of \overline{R} , has an accumulation point $(\mathbf{x}_1, \ldots, \mathbf{x}_J) \in \overline{R}^J$. Thus, let $\psi : \mathbb{N} \longrightarrow \mathbb{N}$ be strictly increasing with the property that

$$\lim_{n\to\infty} \left(\mathbf{x}_1^{\psi(n)},\ldots,\mathbf{x}_J^{\psi(n)}\right) = (\mathbf{x}_1,\ldots,\mathbf{x}_J).$$

Define D as the disjoint union

$$D = \bigsqcup_{j=1}^{J} (\{\mathbf{x}_j\}, m_j)$$

where $(\{\mathbf{x}_j\}, m_j)$ is the persistence diagram having one off-diagonal point at $\mathbf{x}_j \in \overline{R}$ with multiplicity m_j .

We contend that the subsequence of $\{D_n \cap R\}_{n \in \mathbb{N}}$ defined by $\varphi = \phi \circ \psi$ converges to D. Indeed, given $\varepsilon > 0$, let $N \in \mathbb{N}$ be so that $n \geq N$ implies

$$\max_{1 \le j \le J} \left\| \mathbf{x}_{j}^{\psi(n)} - \mathbf{x}_{j} \right\|_{\infty} < \varepsilon. \tag{4}$$

Since the diagram $D_{\varphi(n)} \cap R$ has the collection $\left\{\mathbf{x}_1^{\psi(n)}, \dots, \mathbf{x}_J^{\psi(n)}\right\}$ as underlying set, then

$$\gamma_n: D_{\varphi(n)} \cap R \longrightarrow D$$

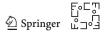
$$\mathbf{x}_j^{\psi(n)} \mapsto \mathbf{x}_j$$

defines a bijection (of multisets). As no points are unmatched, Eq. 4 implies that this is an ε -matching. Thus $d_B(D_{\varphi(n)} \cap R, D) \leq \varepsilon$ for all $n \geq N$, and convergence follows.

With this lemma in place, we can return to the proof of the main theorem.

Thm. 3.7 (\Rightarrow) If $S \subset D$ is relatively compact, then it being bounded, ODBB and UODF follow from Propositions 3.2, 3.4 and 3.6, respectively.

(⇐) Let $S \subset D$ be bounded, ODBB, and UODF. Fix $\{M_k\}_{k \in \mathbb{N}} \subset \mathbb{N}$ and $\mathbb{B}_k \subset \mathbb{W}$, $k \in \mathbb{N}$, as in Eqs. 2 and 3, and let $\{D_n\}_{n \in \mathbb{N}} \subset \overline{S}$ be arbitrary. We will use Lemma 3.8



inductively to construct a sequence $\varphi_1, \varphi_2, \dots, \varphi_k, \dots$ of strictly increasing functions $\varphi_k : \mathbb{N} \longrightarrow \mathbb{N}$ so that if $\Phi_k = \varphi_1 \circ \dots \circ \varphi_k$, then the subsequence of restricted diagrams

$$\left\{D_{\Phi_k(n)} \cap \mathbb{B}_k\right\}_{n \in \mathbb{N}}$$

converges to a diagram $D^k \subset \overline{\mathbb{B}}_k$, for each $k \geq 1$. Once we have built this, we will let

$$\varphi: \mathbb{N} \longrightarrow \mathbb{N}$$

$$n \mapsto \Phi_n(n)$$

and the main task for the proof is to show that $\{D_{\varphi(n)}\}_{n\in\mathbb{N}}$ converges to the diagram

$$D = \bigsqcup_{k=1}^{\infty} D^k.$$

We now proceed inductively in k. The base case follows from applying Lemma 3.8 to the sequence $\{D_n\}_{n\in\mathbb{N}}$ and the relatively compact set \mathbb{B}_1 . This results in a strictly increasing function $\varphi_1: \mathbb{N} \to \mathbb{N}$ and a diagram $D^1 \subset \overline{\mathbb{B}}_1$, so that

$$\lim_{n\to\infty} D_{\varphi_1(n)} \cap \mathbb{B}_1 = D^1.$$

Now the inductive step. Let $k \geq 1$ and assume that $\varphi_j : \mathbb{N} \to \mathbb{N}$ and $D^j \subset \overline{\mathbb{B}}_j$, $1 \leq j \leq k$, have been constructed in such a way that if $\Phi_j = \varphi_1 \circ \cdots \circ \varphi_j$, then

$$\lim_{n\to\infty} D_{\Phi_j(n)} \cap \mathbb{B}_j = D^j.$$

The sequence $\{D_{\Phi_k(n)}\}_{n\in\mathbb{N}}$ and the set \mathbb{B}_{k+1} satisfy the hypotheses of Lemma 3.8. Thus, there exist a strictly increasing function $\varphi_{k+1}:\mathbb{N}\longrightarrow\mathbb{N}$ and a diagram $D^{k+1}\subset\overline{\mathbb{B}}_{k+1}$ so that if $\Phi_{k+1}=\Phi_k\circ\varphi_{k+1}$, then

$$\lim_{n\to\infty} D_{\Phi_{k+1}(n)} \cap \mathbb{B}_{k+1} = D^{k+1}.$$

Now, let $\varphi(n) = \Phi_n(n)$, and let us show that $\{D_{\varphi(n)}\}_{n \in \mathbb{N}}$ converges to $D = \bigsqcup_{k=1}^{\infty} D^k$.

To this end, fix $\varepsilon > 0$ and let $K \in \mathbb{N}$ be large enough so that $\frac{C}{K} < \frac{\varepsilon}{2}$. For each $1 \le k \le K$, let $N_k \in \mathbb{N}$ be so that $n \ge N_k$ implies

$$d_B\left(D_{\Phi_k(n)}\cap\mathbb{B}_k,D^k\right)<\frac{\varepsilon}{2}$$

and let $N = \max\{K, N_1, \dots, N_K\}$. Note that if n > N and $1 \le k \le K$, then

$$\varphi(n) = \Phi_k(\varphi_{k+1} \circ \cdots \circ \varphi_n(n)).$$

Moreover, since $\varphi_{k+1} \circ \cdots \circ \varphi_n(n) \ge n > N_k$, then $d_B\left(D_{\varphi(n)} \cap \mathbb{B}_k, D^k\right) < \frac{\varepsilon}{2}$. Thus, we can assume we have an $\varepsilon/2$ -matching

$$\gamma_n^k: D_{\varphi(n)} \cap \mathbb{B}_k \longrightarrow D^k.$$

As the \mathbb{B}_k 's are disjoint, then the union of the γ_n^k 's yields a bijection (of multisets)

$$\Gamma_n^K: D_{\varphi(n)} \cap \bigcup_{k < K} \mathbb{B}_k \longrightarrow \bigcup_{k < K} D^k$$

Moreover, since all points in $D_{\varphi(n)} \cap \bigcup_{k>K} \mathbb{B}_k$ have persistence at most $\varepsilon/2$, then Γ_n^K defines an $\varepsilon/2$ -matching between $D_{\varphi(n)}$ and D, and hence $d_B(D_{\varphi(n)}, D) < \varepsilon$ for all n > N.

3.6 Bottleneck versus Wasserstein Compactness

The Wasserstein distance d_{W_p} , for $p \in \mathbb{N}$, is another common measure of similarity between persistence diagrams. It is given by

$$d_{W_p}(D, D') := \inf_{\mathsf{M}} \left(\sum_{\substack{\mathbf{x}, \mathbf{x}' \\ \text{matched}}} \|\mathbf{x} - \mathbf{x}'\|_{\infty}^p + \sum_{\substack{\mathbf{z} \\ \text{unmatched}}} \left(\frac{\mathsf{pers}(\mathbf{z})}{2} \right)^p \right)^{1/p}$$
 (5)

where the infimum ranges over all partial matchings M between D and D'. One can show that d_{W_p} defines a metric on the set

$$\mathcal{D}_p \ := \ \left\{D \in \mathcal{D} \mid d_{W_p}(D, \varnothing) < \infty \right\}$$

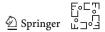
and that (\mathcal{D}_p, d_{W_p}) is a complete separable metric space [48]. Moreover,

$$\mathcal{D}_1 \subset \mathcal{D}_2 \subset \dots \subset \mathcal{D}_p \subset \dots \subset \mathcal{D} \tag{6}$$

and all inclusions can be shown to be continuous with respect to the appropriate metrics. In particular, if $\iota_p:\mathcal{D}_p\hookrightarrow\mathcal{D}$ is the inclusion map and $\mathcal{S}\subset\mathcal{D}_p$ is relatively-compact with respect to d_{W_p} , then $\iota_p\left(\overline{\mathcal{S}}\right)$ is compact (hence closed) with respect to d_B , and the equality

$$\iota_p\left(\overline{\mathcal{S}}\right) = \overline{\iota_p\left(\overline{\mathcal{S}}\right)} = \overline{\iota_p(\mathcal{S})}$$

shows that relatively-compact subsets of (\mathcal{D}_p, d_{W_p}) are also relatively compact in (\mathcal{D}, d_B) . This implies that relatively-compact subsets of (\mathcal{D}_p, d_{W_p}) also satisfy the three conditions of Theorem 3.7. The characterization of relative-compactness in (\mathcal{D}_p, d_{W_p}) is work of [48], and we will describe it next.



Definition 3.9 A set $S \subset \mathcal{D}_p$ is **uniform** if for every $\epsilon > 0$ there exists $\alpha > 0$ so that

$$\sum_{\substack{\mathbf{x} \in S \\ \mathsf{pers}(\mathbf{x}) < \alpha}} \mu(\mathbf{x}) \mathsf{pers}(\mathbf{x})^p \leq \epsilon$$

for all $(S, \mu) \in \mathcal{S}$.

The relevant characterization is as follows:

Theorem 3.10 ([48]) A set $S \subset \mathcal{D}_p$ is relatively compact in (\mathcal{D}_p, d_{W_p}) if and only if it is bounded (with respect to d_{W_p}), off-diagonally birth bounded (in the sense of Definition 3.3) and uniform.

The next two examples illustrate how relative compactness in (\mathcal{D}, d_B) and (\mathcal{D}_p, d_{W_p}) can exhibit very different behaviors, even if Theorems 3.7 and 3.10 seem similar at first glance.

Examples:

1. Let S be the set of persistence diagrams $D_{jk} = (S_{jk}, \mu_{jk})$, for $j, k \in \mathbb{N}$, with

$$S_{jk} = \left\{ \left(0, \frac{1}{kn^{1/k}} \right) \middle| 1 \le n \le j \right\}$$

and multiplicity function $\mu_{jk}\left(0,\frac{1}{kn^{1/k}}\right)=1$. Notice that each S_{jk} is finite, and hence $\mathcal{S}\subset\mathcal{D}_p$ for all $p\in\mathbb{N}$. However, for k=p the set $\{D_{jp}\}_{j\in\mathbb{N}}\subset\mathcal{S}$ is not bounded with respect to d_{W_p} , and thus \mathcal{S} is not relatively compact in (\mathcal{D}_p,d_{W_p}) for any $p\in\mathbb{N}$.

On the other hand, $d_B(D_{jk},\varnothing) \leq 1$ for every $j,k \in \mathbb{N}$, so \mathcal{S} is bounded with respect to d_B , and it is clearly birth-bounded (hence ODBB). In order to see that \mathcal{S} is UODF, fix $\epsilon > 0$ and let $k_0 \in \mathbb{N}$ be so that $\frac{1}{\epsilon} < k_0$. It follows that $D_{jk} \cap \overline{\mathbb{W}^\epsilon} = \varnothing$ for $k \geq k_0$ and every $j \in \mathbb{N}$, so assume $1 \leq k < k_0$, $j \in \mathbb{N}$, and let $\left(0, \frac{1}{kn^{1/k}}\right) \in S_{jk} \cap \overline{\mathbb{W}^\epsilon}$. Hence

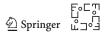
$$n \leq \frac{1}{(k\epsilon)^k} \leq \frac{1}{\epsilon^k} < k_0^{k_0}$$

and therefore Mult $\left(D,\overline{\mathbb{W}^{\epsilon}}\right) < k_0^{k_0}$ for every $D \in \mathcal{S}$. This shows that \mathcal{S} is UODF and thus relatively compact in $(\mathcal{D}_{\cdot}d_B)$.

2. Let S be the set of persistence diagrams $D_k = (S_k, \mu_k)$, for $k \in \mathbb{N}$, with

$$S_k = \left\{ (0, 1), \left(0, n^{-1/k} \right) \mid n \ge 2 \right\}$$

and multiplicity function $\mu_k(0,1) = k$, $\mu_k\left(0,n^{-1/k}\right) = 1$. The first thing to note is that $\mathcal{S} \subset \mathcal{D}$, and that $\text{Mult}\left(D_k,\overline{\mathbb{W}^{0.5}}\right) > k$ for every $k \in \mathbb{N}$. Therefore, \mathcal{S} is not UODF, and thus—by Theorem 3.7—it is not relatively compact in (\mathcal{D},d_B) .



On the other hand, for each $p \in \mathbb{N}$ we have that $S \cap \mathcal{D}_p = \{D_k \mid 1 \leq k < p\}$, and thus

$$\mathcal{S} = \bigcup_{p \in \mathbb{N}} \mathcal{S} \cap \mathcal{D}_p.$$

Since $S \cap D_p$ is finite, then it is compact in (D_p, d_{W_p}) for every $p \in \mathbb{N}$.

In summary, in order to understand relative compactness in (\mathcal{D}, d_B) it is not enough to do so at each (\mathcal{D}_p, d_{W_p}) ; the class of relatively compact subsets of (\mathcal{D}, d_B) is in fact larger than the union over $p \in \mathbb{N}$ of those in (\mathcal{D}_p, d_{W_p}) .

3.7 Consequences of Thm. 3.7

We note a few consequences of our characterization of relative compactness in (\mathcal{D}, d_B) .

Theorem 3.11 *Relatively compact subsets of* (\mathcal{D}, d_B) *have empty interior.*

Proof Let $S \subset \mathcal{D}$ be relatively compact, and let

$$\mathbb{B}_k = \left\{ \mathbf{x} \in \mathbb{W} \;\middle|\; 0 \leq \mathsf{birth}(\mathbf{x}) < C_k \; \text{ and } \; \frac{C}{k+1} \leq \mathsf{pers}(\mathbf{x}) < \frac{C}{k} \right\}, \;\; k \in \mathbb{N}$$

be a sequence of boxes (as defined in Eq. 2) so that $D \subset \bigcup_k \mathbb{B}_k$ for every $D \in \mathcal{S}$. Fix $D \in \mathcal{S}$. We will show that any open ball around D contains a persistence diagram D' whose underlying set is not in the union of these boxes. Indeed, given $\epsilon > 0$, there exists $k_0 \in \mathbb{N}$ so that $\frac{C}{k_0} < \frac{\epsilon}{2}$, and if D' is the persistence diagram obtained from D by adding the point $\left(C_{k_0} + 1, C_{k_0} + 1 + \frac{C}{k_0}\right)$ with multiplicity one, then $d_B(D, D') < \epsilon$, but $D' \not\subset \bigcup_k \mathbb{B}_k$.

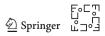
Recall that a topological space is *locally compact* if every point has an open neighborhood contained in a compact set; said open set is called a compact neighborhood of the point. The following corollary is a direct consequence of Theorem 3.11.

Corollary 3.12 The space of persistence diagrams (\mathcal{D}, d_B) is not locally compact. Moreover, no diagram $D \in \mathcal{D}$ has a compact neighborhood.

Proof Let $D \in \mathcal{D}$ and suppose, by way of contradiction, that D has a compact neighborhood. That is, that there exist an open set $\mathcal{U} \subset \mathcal{D}$ and a compact set $\mathcal{S} \subset \mathcal{D}$ so that $D \in \mathcal{U} \subset \mathcal{S}$. Since taking interiors preserves the order of inclusions, then

$$\mathcal{U} = \operatorname{int}(\mathcal{U}) \subset \operatorname{int}(\mathcal{S}) = \emptyset$$

which contradicts $D \in \mathcal{U}$.



As we will see next, the lack of enough compact regions in the space of persistence diagrams also has global implications for learning tasks. Recall that a subset of a topological space is called *nowhere dense* if its closure has empty interior. It follows from Corollary 3.12 that in the space of persistence diagrams all compact sets are nowhere dense. Moreover, since (\mathcal{D}, d_B) is complete [9], then the Baire Category Theorem [4]—which contends that no complete metric space can be written as the countable union of nowhere dense subsets—implies the following.

Corollary 3.13 \mathcal{D} *cannot be written as the countable union of compact subsets.*

Many optimization tasks (e.g., gradient descent) leverage the compactness of a space to argue that solutions can be found as limits to sequential processes. The fact that (\mathcal{D}, d_B) cannot be written as the countable union of compact subsets, implies that non-convex global optimization problems in \mathcal{D} cannot be guaranteed to be solvable via this type of local decompositions. As we will see below, this also implies that it is not possible to find a metric in $C(\mathcal{D}, \mathbb{R})$ that characterizes approximations with respect to the topology of uniform convergence on compact sets.

Definition 3.14 Let X, Y be topological spaces and let C(X, Y) denote the set of continuous functions from X to Y. Given $K \subset X$ compact and $V \subset Y$ open, let

$$U(K, V) = \{ f \in C(X, Y) \mid f(K) \subset V \}.$$

The collection

$$\{U(K, V) \mid K \subset X \text{ compact}, V \subset Y \text{ open}\}$$

forms a subbase for a topology on C(X,Y), called the *compact-open topology*. When Y is a metric space, a sequence of continuous functions $f_n:X\longrightarrow Y,\,n\in\mathbb{N}$, converges to f in the compact-open topology, if and only if $\{f_n|_K\}_{n\in\mathbb{N}}$ converges uniformly to $f|_K$ for each compact set $K\subset X$.

Since \mathcal{D} cannot be written as a countable union of compact sets, then we have the following.

Corollary 3.15 *The compact-open topology on* $C(\mathcal{D}, \mathbb{R})$ *is not metrizable.*

Sequential optimization tasks in machine learning typically proceed by updating the current solution until no further progress is made. Solutions are often functions (e.g., a classifier) and progress in the optimization is measured through loss functionals (e.g., like mean-squared error) and/or the distance between the current and prior state. The latter being particularly useful in convex optimization. Since compact-open sequential convergence in $C(\mathcal{D}, \mathbb{R})$ cannot be measured with a metric (Corollary 3.15), then other alternatives are needed. We describe such methodologies in the next sections.



4 Linearizing \mathcal{D}

The fact that the compact-open topology on $C(\mathcal{D}, \mathbb{R})$ is not metrizable, implies that optimization with respect to compact convergence needs to be handled with care. The goal of this section is to provide methods for doing this. We begin with a definition.

Definition 4.1 A **coordinate system** for \mathcal{D} is a collection $\mathcal{F} \subset C(\mathcal{D}, \mathbb{R})$ which separates points. That is, if $D, D' \in \mathcal{D}$ are distinct then there exists $F \in \mathcal{F}$ for which $F(D) \neq F(D')$.

Of course one could take \mathcal{F} to be the space of all real-valued continuous functions on \mathcal{D} , but this is an extreme case; the quality of a coordinate system is determined by its size—the smaller the better. The metaphor to keep in mind is Euclidean space, \mathbb{R}^n . In this case, an oblique coordinate system (e.g. Cartesian coordinates) is uniquely determined by a linear basis for the space $\mathcal{L}(\mathbb{R}^n)$ of (continuous) linear functions from \mathbb{R}^n to \mathbb{R} .

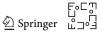
Our goal is to coordinatize the space of persistence diagrams by finding a continuous embedding of \mathcal{D} into an appropriate topological vector space V, and taking the restriction to \mathcal{D} of elements from $\mathcal{L}(V)$. The problem of embedding persistence diagrams into topological vector spaces has received considerable recent attention. For instance, it is known that (\mathcal{D}_p, d_{W_p}) does not admit an inner product compatible with d_{W_p} for any p > 2, and that (\mathcal{D}_p, d_{W_p}) does not admit a coarse embedding (this is weaker than being bi-Lipschitz) into any Hilbert space for any p > 2 [64].

In order to choose V we will use two principles. First, that persistence diagrams can be interpreted as Radon measures on \mathbb{W} [19, 30], which suggests embedding \mathcal{D} into the dual space of some set of continuous real-valued functions on \mathbb{W} . Second, that \mathcal{D} is a topological monoid: the sum of two persistence diagrams D, $D' \in \mathcal{D}$ is their disjoint union $D \sqcup D'$ as multisets, the empty diagram \varnothing is the identity, i.e. $D \sqcup \varnothing = D$, and the operation $\sqcup : \mathcal{D} \times \mathcal{D} \longrightarrow \mathcal{D}$ is associative and continuous (see also [11]). In what follows we will construct an embedding $v : \mathcal{D} \hookrightarrow V$ which recovers the measure-theoretic interpretation of persistence diagrams, and preserves the monoidal structure of \mathcal{D} (see Theorem 4.9). In addition, we will show that appropriate subsets of $\mathcal{L}(V)$ will yield coordinate systems for \mathcal{D} (see Theorem 5.4), and these in turn will generate dense subsets of $\mathcal{L}(\mathcal{D}, \mathbb{R})$ with respect to the compact-open topology (see Theorem 5.3).

4.1 Topological Vector Spaces, Duals and Their Topologies

We will first review some basics of topological vector spaces, following [25]. Let V be a topological vector space; that is, a vector space endowed with a topology so that addition and scalar multiplication are continuous functions. Its (topological) dual is the vector space

$$V' = \{T : V \longrightarrow \mathbb{R} \text{, so that } T \text{ is linear and continuous} \}$$



In particular, if the topology on V comes from a norm $\|\cdot\|_V$, then we write V^* instead of V'. If V^* is endowed with the operator norm

$$||T||_* = \sup_{\|\mathbf{v}\|_V = 1} |T(\mathbf{v})|,$$

then V^* is in fact a Banach space. There are three standard topologies on V^* :

Strong: The strong topology is the one generated by the operator norm $\|\cdot\|_*$. A basis for open neighborhoods of a point $T \in V^*$ is given by sets of the form

$$B_{\epsilon}(T) = \left\{ T' \in V^* \mid \sup_{\|\mathbf{v}\|_V = 1} |T(\mathbf{v}) - T'(\mathbf{v})| < \epsilon \right\}$$

where $\epsilon > 0$. In particular, a sequence $\{T_n\}_{n \in \mathbb{N}} \subset V^*$ converges to $T \in V^*$ in the strong topology if and only if $\{T_n(\mathbf{v})\}_{n \in \mathbb{N}}$ converges to $T(\mathbf{v})$ uniformly in $\mathbf{v} \in V$.

Weak: If V^{**} denotes the dual of the normed space $(V^*, \|\cdot\|_*)$, then the weak topology on V^* is the smallest topology so that every $T \in V^{**}$ is continuous. A basis for open neighborhoods of a point $T \in V^*$ is given by sets of the form

$$N(\mathcal{T}_1, \dots, \mathcal{T}_I; \epsilon)(T) = \left\{ T' \in V^* \mid \max_{1 \le i \le I} |\mathcal{T}_i \left(T' \right) - \mathcal{T}_i \left(T \right)| < \epsilon \right\}$$

where $\mathcal{T}_1, \ldots, \mathcal{T}_I \in V^{**}$ and $\epsilon > 0$. In particular, $\{T_n\}_{n \in \mathbb{N}} \subset V^*$ converges to $T \in V^*$ in the weak topology if and only if $\{\mathcal{T}(T_n)\}_{n \in \mathbb{N}}$ converges to $\mathcal{T}(T)$ for all $T \in V^{**}$.

Weak-*: The weak-* topology is the smallest topology so that for each $\mathbf{v} \in V$, the resulting evaluation function

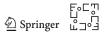
$$e_{\mathbf{v}}: V^* \longrightarrow \mathbb{R}$$
 $T \mapsto T(\mathbf{v})$

is continuous. A basis for open neighborhoods of $T \in V^*$ is given by sets of the form

$$N(\mathbf{v}_1,\ldots,\mathbf{v}_I;\epsilon)(T) = \left\{ T' \in V^* \mid \max_{1 \le i \le I} |T'(\mathbf{v}_i) - T(\mathbf{v}_i)| < \epsilon \right\}$$

where $\mathbf{v}_1, \ldots, \mathbf{v}_I \in V$ and $\epsilon > 0$. A sequence $\{T_n\}_{n \in \mathbb{N}} \subset V^*$ converges to $T \in V^*$ in the weak-* topology if and only if $\{T_n(\mathbf{v})\}_{n \in \mathbb{N}}$ converges to $T(\mathbf{v})$ for each $\mathbf{v} \in V$. The convergence, however, need not be uniform in \mathbf{v} .

One can check that the weak-* topology is weaker than the weak topology, which in turn is weaker than the strong topology.



4.2 Linearizing the Set of Finite Diagrams

It is useful to first illustrate some of the difficulties associated to finding embeddings for the set of finite diagrams \mathcal{D}_0 . In what follows we will prove several negative results which will inform the choices in embedding \mathcal{D} . Indeed, the first thing to notice is that the set of compactly supported continuous functions from \mathbb{W} to \mathbb{R} , denoted $C_c(\mathbb{W})$, is a normed vector space if endowed with the sup norm $\|\cdot\|_{\infty}$.

The Dirac mass centered at $\mathbf{x} \in \mathbb{W}$ is the linear function

$$\delta_{\mathbf{x}}: C_c(\mathbb{W}) \longrightarrow \mathbb{R}$$
$$f \mapsto f(\mathbf{x})$$

and since $|\delta_{\mathbf{x}}(f)| \leq ||f||_{\infty}$ for each $f \in C_c(\mathbb{W})$, it follows that $\delta_{\mathbf{x}} \in C_c(\mathbb{W})^*$. Let

$$\begin{array}{ccc}
\nu_0: & \mathcal{D}_0 & \longrightarrow & C_c(\mathbb{W})^* \\
\varnothing & \mapsto & 0 \\
\varnothing \neq (S, \mu) & \mapsto & \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) \delta_{\mathbf{x}}.
\end{array}$$

It is not hard to see that

Proposition 4.2 $\nu_0 : \mathcal{D}_0 \longrightarrow C_c(\mathbb{W})^*$ is injective, and satisfies $\nu_0(D \sqcup D') = \nu_0(D) + \nu_0(D')$ for every $D, D' \in \mathcal{D}_0$.

Whether v_0 is continuous or not depends on the topology with which $C_c(\mathbb{W})^*$ is endowed. We start with the coarser topologies, but immediately have the following negative results.

Proposition 4.3 *If* $C_c(\mathbb{W})^*$ *is endowed with the weak topology, then* v_0 *is discontinuous at every point.*

Proof Fix $D \in \mathcal{D}_0$, and let $D_n \in \mathcal{D}_0$ be the diagram obtained from D by adding the point (1, 1+1/n) with multiplicity n. It follows that $\{D_n\}_n$ converges to D with respect to the bottleneck distance. We contend that $\nu_0(D_n)$ does not converge to $\nu_0(D)$ with respect to the weak topology; in other words, we will show that there exist $\epsilon_0 > 0$ and a linear operator $\mathcal{T}: C_c(\mathbb{W})^* \longrightarrow \mathbb{R}$, continuous with respect to the strong topology, so that $|\mathcal{T}(\nu_0(D_n)) - \mathcal{T}(\nu_0(D))| \ge \epsilon_0$ for infinitely many values of n.

Indeed, the first thing to notice is that since \mathbb{W} is not compact, $C_c(\mathbb{W})$ is not complete. Its completion is the space $C_0(\mathbb{W})$ of continuous functions on \mathbb{W} which vanish at the diagonal Δ and at infinity. More explicitly, $f \in C_0(\mathbb{W})$ if and only if it is continuous, and $f(x,y) \to 0$ whenever $y-x \to 0$ or $x+y \to \infty$. Let $\iota: C_c(\mathbb{W}) \hookrightarrow C_0(\mathbb{W})$ be the inclusion and let $\iota^*: C_0(\mathbb{W})^* \to C_c(\mathbb{W})^*$ be the induced homomorphism. Since $C_c(\mathbb{W})$ is dense in $C_0(\mathbb{W})$, then ι^* is an isometric isomorphism; its inverse $J^*: C_c(\mathbb{W})^* \to C_0(\mathbb{W})^*$ sends a continuous linear map $T: C_c(\mathbb{W}) \to \mathbb{R}$ to its unique continuous linear extension $J^*(T): C_0(\mathbb{W}) \to \mathbb{R}$.

Now, let $\varphi : \mathbb{R}^2 \longrightarrow [0, 1]$ be a continuous (bump) function so that

$$\varphi(x, y) = \begin{cases} 1 & \text{if } \max\{x, y - x\} < 2\\ 0 & \text{if } \max\{x, y - x\} \ge 3. \end{cases}$$

It follows that $f(x, y) = (y-x) \cdot \varphi(x, y) \in C_0(\mathbb{W})$, and hence the evaluation function $e_f : C_0(\mathbb{W})^* \longrightarrow \mathbb{R}$ is a bounded linear operator. Let $\mathcal{T} : C_c(\mathbb{W})^* \longrightarrow \mathbb{R}$ be the composition $e_f \circ J^*$. Then, for each n > 1 we have that

$$\mathcal{T}(\nu_0(D_n)) = \mathcal{T}(\nu_0(D)) + 1,$$

and letting $\epsilon_0 = 1$ completes the proof.

Corollary 4.4 If $C_c(\mathbb{W})^*$ is endowed with the strong topology, then v_0 is discontinuous at every point.

Proof Indeed, the strong topology contains the weak topology.

It is not until we pass to the weakest of the three standard topologies that we approach a useful result.

Proposition 4.5 If $C_c(\mathbb{W})^*$ is endowed with the weak-* topology, then $v_0: \mathcal{D}_0 \longrightarrow C_c(\mathbb{W})^*$ is continuous.

Before presenting the proof, we have the following useful lemma.

Lemma 4.6 For each $f \in C_c(\mathbb{W})$, the function

$$v_f: \mathcal{D} \longrightarrow \mathbb{R}$$

$$(S, \mu) \mapsto \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x})$$
(7)

is continuous.

Proof The first observation is that the sum defining v_f is always finite, since the support of $f \in C_c(\mathbb{W})$ intersects the underlying set of any persistence diagram at only finitely many points. Let $D = (S, \mu) \in \mathcal{D}$ and fix $\epsilon > 0$. Since $\text{supp}(f) \subset \mathbb{W}$ is compact, then f is uniformly continuous. Further, there exists $\delta > 0$ for which $\text{supp}(f) \subset \mathbb{W}^{2\delta}$, and

$$|f(\mathbf{x}) - f(\mathbf{y})| < \frac{\epsilon}{\sum_{\mathbf{z} \in S \cap W\delta} \mu(\mathbf{z})}$$
(8)

whenever $\|\mathbf{x} - \mathbf{y}\|_{\infty} < \delta$.

Let $D' = (T, \alpha) \in \mathcal{D}$ be given with $d_B(D, D') < \delta$. We will show that $|\nu_f(D) - \nu_f(D')| < \epsilon$. Fix a δ -matching

$$\begin{array}{ccc} \mathsf{M}: \ A & \xrightarrow{\cong} & B \\ & & \mathsf{I} \cap & & \mathsf{I} \cap \\ & S_{\mu} & & T_{\alpha}. \end{array}$$

This means that if $(\mathbf{x}, k) \in A$ and $(\mathbf{y}, n) = \mathsf{M}(\mathbf{x}, k)$ are matched, then $\|\mathbf{x} - \mathbf{y}\|_{\infty} < \delta$; and if $(\mathbf{z}, m) \in (S_{\mu} \setminus A) \cup (T_{\alpha} \setminus B)$ is unmatched, then $\mathsf{pers}(\mathbf{z}) < 2\delta$. In this case $\mathbf{z} \notin \mathsf{supp}(f)$, which implies $f(\mathbf{z}) = 0$. Hence

$$\nu_f(D) = \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x})$$
$$= \sum_{(\mathbf{x}, k) \in A} f(\mathbf{x})$$

and similarly,

$$\nu_f(D') = \sum_{(\mathbf{y}, n) \in B} f(\mathbf{y}).$$

Therefore,

$$|v_f(D) - v_f(D')| = \left| \sum_{\substack{(\mathbf{x}, k) \in A}} f(\mathbf{x}) - \sum_{\substack{(\mathbf{y}, n) \in B \\ (\mathbf{y}, k) \in A}} f(\mathbf{y}) \right|$$

$$\leq \sum_{\substack{(\mathbf{y}, n) = \mathsf{M}(\mathbf{x}, k) \\ (\mathbf{x}, k) \in A}} |f(\mathbf{x}) - f(\mathbf{y})|$$

where each term $|f(\mathbf{x}) - f(\mathbf{y})|$ is potentially nonzero only when \mathbf{x} or \mathbf{y} are in $\operatorname{supp}(f) \subset \mathbb{W}^{2\delta}$. Since in this case $\|\mathbf{x} - \mathbf{y}\|_{\infty} < \delta$, we would get $\mathbf{x}, \mathbf{y} \in \mathbb{W}^{\delta}$. Combining this observation with Eq. (8) completes the proof.

Proposition 4.5 Let $D \in \mathcal{D}_0$, and fix a weak-* basic neighborhood $N(f_1, \ldots, f_I; \epsilon)$ for $\nu_0(D)$. Notice that for each $i = 1, \ldots, I$ we have $\nu_0(D)(f_i) = \nu_{f_i}(D)$. Since ν_{f_i} is continuous at D, then given $\epsilon > 0$ there exists $\delta_i > 0$ so that $d_B(D, D') < \delta_i$ implies $|\nu_{f_i}(D) - \nu_{f_i}(D')| < \epsilon$. If we let $\delta = \min\{\delta_1, \ldots, \delta_I\}$, it follows that whenever $d_B(D, D') < \delta$ then for all $i = 1, \ldots, I$

$$|v_0(D)(f_i) - v_0(D')(f_i)| = |v_{f_i}(D) - v_{f_i}(D')| < \epsilon.$$

This shows that $\nu_0(D') \in N(f_1, ..., f_I; \epsilon)$ and hence ν_0 is continuous.

These results imply that out of the three standard topologies on $C_c(\mathbb{W})^*$, the weak* topology is the only one for which ν_0 yields a continuous embedding of \mathcal{D}_0 into $C_c(\mathbb{W})^*$. The question now is whether this embedding can be extended to \mathcal{D} . The answer, as it turns out, is negative.

Proposition 4.7 If $C_c(\mathbb{W})^*$ is endowed with the weak-* topology, then $v_0 : \mathcal{D}_0 \longrightarrow C_c(\mathbb{W})^*$ cannot be continuously extended to any $D \in \mathcal{D} \setminus \mathcal{D}_0$.

Proof Assume, by way of contradiction, that ν_0 extends continuously to some $D = (S, \mu) \in \mathcal{D} \setminus \mathcal{D}_0$. If for each $n \in \mathbb{N}$ we let D_n be the restriction of D to $\mathbb{W}^{1/n}$, then $D_n \in \mathcal{D}_0$ for all $n \in \mathbb{N}$, and the sequence $\{D_n\}_{n \in \mathbb{N}}$ converges to D with respect to the bottleneck distance. By the continuity assumption of ν_0 at D, we have that

$$\nu_0(D) = \lim_{n \to \infty} \nu_0(D_n)$$

where convergence is with respect to the weak-* topology. In other words,

$$\nu_0(D)(f) = \lim_{n \to \infty} \nu_0(D_n)(f)$$

for every $f \in C_c(\mathbb{W})$. It follows that, given $f \in C_c(\mathbb{W})$, there exists $N_f \in \mathbb{N}$ so that $\operatorname{supp}(f) \subset \mathbb{W}^{1/n}$ for all $n \geq N_f$, and therefore the sequence $\nu_0(D_n)(f)$ becomes constant and equal to

$$\sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x}). \tag{9}$$

We claim that if $C_c(\mathbb{W})$ is endowed with the sup norm $\|\cdot\|_{\infty}$, then the linear function

$$\nu_0(D): C_c(\mathbb{W}) \longrightarrow \mathbb{R}$$

$$f \mapsto \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x})$$

is discontinuous at every point. To this end, we will show that given $f \in C_c(\mathbb{W})$ there exists a sequence $\{f_n\}_{n\in\mathbb{N}} \subset C_c(\mathbb{W})$ which converges to f with respect to $\|\cdot\|_{\infty}$, but for which $\{\nu_0(D)(f_n)\}_{n\in\mathbb{N}}$ does not converge to $\nu_0(D)(f)$. This would contradict $\nu_0(D) \in C_c(\mathbb{W})^*$.

Indeed, since $(S, \mu) = D \notin \mathcal{D}_0$, then there exists a sequence $\{\mathbf{x}_n\}_{n \in \mathbb{N}} \subset S \setminus \text{supp}(f)$ so that $\text{pers}(\mathbf{x}_n)$ is strictly decreasing as n goes to infinity. Therefore, it is possible to construct a sequence $\{r_n\}_{n \in \mathbb{N}}$ of positive real numbers, so that the balls $B_{r_n}(\mathbf{x}_n) \subset \mathbb{W}$ are all disjoint and disjoint with the support of f. Let $\phi_n : \mathbb{R}^2 \longrightarrow [0, \infty)$ be the bump function

$$\phi_n(\mathbf{x}) = \frac{\max\{0, r_n - 2\|\mathbf{x} - \mathbf{x}_n\|\}}{r_n}$$

supported on the closure of $B_{\frac{r_n}{2}}(\mathbf{x}_n)$, and let

$$f_n = f + \frac{\phi_1 + \dots + \phi_n}{n}.$$

It follows that $\{f_n\}_{n\in\mathbb{N}}$ is a sequence of continuous and compactly supported functions on \mathbb{W} , so that $\|f_n - f\|_{\infty} < \frac{1}{n}$ for all $n \in \mathbb{N}$, and for which

$$\nu_0(D)(f_n) - \nu_0(D)(f) = \frac{1}{n}\nu_0(D)(\phi_1 + \dots + \phi_n)$$

$$\geq \frac{1}{n} \left(\mu(\mathbf{x}_1)\phi_1(\mathbf{x}_1) + \dots + \mu(\mathbf{x}_n)\phi_n(\mathbf{x}_n)\right)$$

$$\geq 1.$$

Hence $\{\nu_0(D)(f_n)\}_{n\in\mathbb{N}}$ does not converge to $\nu_0(D)(f)$, and $\nu_0(D)$ is discontinuous at f.

4.3 Linearizing Infinite Diagrams

There are two main lessons to draw from the previous results: First, that even though there is a candidate for extending v_0 to infinite diagrams, namely Eq. 9, the topology on $C_c(\mathbb{W})$ induced by the sup norm is inadequate as it does not have enough open sets. The second lesson is that a weak-* topology on the dual of $C_c(\mathbb{W})$ is the most likely to ensure continuity when embedding \mathcal{D} . In what follows we will describe a (locally convex) topology on $C_c(\mathbb{W})$, and a corresponding weak-* topology on the topological dual $C_c(\mathbb{W})'$ with the required properties. We will utilize the theory of *locally convex topological vector spaces*, which generalize Banach spaces, and provide a rich framework in which to study weak topologies. For a more detailed account we direct the interested reader to Chapters IV and V of [25].

Let $\{K_n\}_{n\in\mathbb{N}}$ be a sequence of compact subsets of \mathbb{W} so that $K_n\subset K_{n+1}$ for all $n\in\mathbb{N}$, and for which

$$\mathbb{W}=\bigcup_{n\in\mathbb{N}}K_n.$$

It follows that each vector space

$$C_c(K_n) = \{ f \in C(\mathbb{W}) \mid \mathsf{supp}(f) \subset K_n \}$$

is a Banach space if endowed with the sup norm $\|\cdot\|_{\infty}$; in particular it is a *locally convex space*.

Definition 4.8 A locally convex space is a topological vector space V, whose topology is generated by a family $\mathscr P$ of seminorms on V which separate points. More specifically, $\mathscr P$ is a collection $\{\rho_\alpha\}_{\alpha\in\Gamma}$ of continuous functions $\rho_\alpha:V\longrightarrow [0,\infty)$ so that

- 1. $\rho_{\alpha}(\mathbf{u} + \mathbf{v}) \leq \rho_{\alpha}(\mathbf{u}) + \rho_{\alpha}(\mathbf{v})$ for all $\mathbf{u}, \mathbf{v} \in V$,
- 2. $\rho_{\alpha}(\lambda \mathbf{u}) = |\lambda| \rho_{\alpha}(\mathbf{u})$ for all scalars λ ,
- 3. $\rho_{\alpha}(\mathbf{u}) = 0$ for all $\alpha \in \Gamma$ implies $\mathbf{u} = \mathbf{0}$

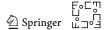
and the topology of V is the weakest for which all the ρ_{α} 's are continuous.

In particular, all normed spaces are locally convex: any norm is a seminorm, and the norm topology is the smallest for which the norm is a continuous function. Notice also that each inclusion

$$C_c(K_n) \subset C_c(K_{n+1})$$
 $n \in \mathbb{N}$

is continuous and that

$$C_c(\mathbb{W}) = \bigcup_{n \in \mathbb{N}} C_c(K_n).$$



The **strict inductive limit topology** on $C_c(\mathbb{W})$ is the finest locally convex topology so that each inclusion $C_c(K_n) \hookrightarrow C_c(\mathbb{W})$ is continuous. In this topology, a linear map $T: C_c(\mathbb{W}) \longrightarrow Y$ to a locally convex space Y is continuous if and only if the restriction of T to each $C_c(K_n)$ is continuous. Moreover, this topology is independent of the choice of compact sets $\{K_n\}_{n\in\mathbb{N}}$ filtering \mathbb{W} .

Let $C_c(\mathbb{W})'$ denote the topological dual of $C_c(\mathbb{W})$ with respect to the strict inductive limit topology, and endow $C_c(\mathbb{W})'$ with the weakest topology so that for each $f \in C_c(\mathbb{W})$ the resulting evaluation function

$$e_f: C_c(\mathbb{W})' \longrightarrow \mathbb{R}$$
 $T \mapsto T(f)$

is continuous. This is the corresponding weak-* topology. It follows that a basis for neighborhoods of a point $T \in C_c(\mathbb{W})'$ is given by sets of the form

$$N(f_1,\ldots,f_I;\epsilon)(T) = \left\{ \tilde{T} \in C_c(\mathbb{W})' : |(T-\tilde{T})(f_i)| < \epsilon, i = 1,\ldots, I \right\}$$

where $f_1, \ldots, f_I \in C_c(\mathbb{W})$ and $\varepsilon > 0$. Here is the main theorem of this section.

Theorem 4.9 Given a persistence diagram $D = (S, \mu) \in \mathcal{D}$ and a function $f \in C_c(\mathbb{W})$, define

$$\nu_D(f) := \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x}). \tag{10}$$

If $C_c(\mathbb{W})$ is endowed with the strict inductive limit topology, and $C_c(\mathbb{W})'$ is its topological dual endowed with the corresponding weak-* topology, then

$$\nu: \mathcal{D} \longrightarrow C_c(\mathbb{W})'$$
$$D \mapsto \nu_D$$

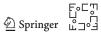
is continuous, injective and satisfies $v(D \sqcup D') = v(D) + v(D')$ for all $D, D' \in \mathcal{D}$.

Proof First, we ensure that ν_D is well defined. Fix $D = (S, \mu) \in \mathcal{D}$ and $n \in \mathbb{N}$. Then $S \cap K_n$ is a finite set and hence for each $f \in C_c(K_n)$ it follows that

$$v_D(f) = \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) f(\mathbf{x}) < \infty.$$

As for continuity of ν_D , fix $f_0 \in C_c(K_n)$, let $\varepsilon > 0$, and let

$$\delta < \frac{\varepsilon}{\sum_{\mathbf{z} \in S \cap K} \mu(\mathbf{z})}.$$



If $f \in C_c(K_n)$ is so that $||f - f_0||_{\infty} < \delta$, then

$$|\nu_D(f) - \nu_D(f_0)| = \left| \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) \left(f(\mathbf{x}) - f_0(\mathbf{x}) \right) \right|$$

$$\leq \sum_{\mathbf{z} \in S \cap K_n} \mu(\mathbf{z}) ||f - f_0||_{\infty}$$

$$< \varepsilon.$$

Therefore, ν_D is a real-valued continuous linear function on $C_c(K_n)$ for each n. This shows that $\nu(D) = \nu_D \in C_c(\mathbb{W})'$ for all $D \in \mathcal{D}$.

To see that $\nu: \mathcal{D} \longrightarrow C_c(\mathbb{W})'$ is continuous, we proceed exactly as in the proof of Proposition 4.5. Indeed, let $D \in \mathcal{D}$, and fix a basic neighborhood $N(f_1, \ldots, f_I; \epsilon)$ for $\nu(D)$. For each $i = 1, \ldots, I$ the function $\nu_{f_i}: \mathcal{D} \longrightarrow \mathbb{R}$, $\nu_{f_i}(D) = \nu_D(f_i)$, is continuous by Lemma 4.6, and hence there exists $\delta > 0$ such that

$$|v(D)(f_i) - v(D')(f_i)| = |v_{f_i}(D) - v_{f_i}(D')| < \epsilon$$

for all i = 1, ..., I, whenever $d_B(D, D') < \delta$.

Injectivity of ν is deduced from the following observation. If $(S, \mu), (T, \alpha) \in \mathcal{D}$ are distinct, then we can assume without loss of generality that there exists $\mathbf{x} \in S$ such that either: $\mathbf{x} \notin T$, or $\mathbf{x} \in T$ and $\mu(\mathbf{x}) \neq \alpha(\mathbf{x})$. Let $f \in C_c(\mathbb{W})$ be such that $\operatorname{supp}(f) \cap (S \cup T) = \{\mathbf{x}\}$, and for which $f(\mathbf{x}) = 1$. If $\mathbf{x} \notin T$, then

$$\nu(S, \mu)(f) = \mu(\mathbf{x}) \neq 0 = \nu(T, \alpha)(f).$$

Similarly, for the case where $\mathbf{x} \in T$ we have

$$\nu(S, \mu)(f) = \mu(\mathbf{x}) \neq \alpha(\mathbf{x}) = \nu(T, \alpha)(f),$$

which completes the proof.

The Riesz-Markov representation theorem—see for instance Theorem 2.14 and Theorem 2.17 of [57]—contends that if $T:C_c(\mathbb{W})\longrightarrow \mathbb{R}$ is linear and satisfies $T(f)\geq 0$ whenever $f(\mathbf{x})\geq 0$ for all $\mathbf{x}\in \mathbb{W}$, then there exists a unique positive Radon measure η on \mathbb{W} so that

$$\int_{\mathbb{W}} f d\eta = T(f)$$

for all $f \in C_c(\mathbb{W})$. Specifically, η is Borel regular and $\eta(K) < \infty$ for every compact set $K \subset \mathbb{W}$. Applying this theorem to elements in the image of $\nu : \mathcal{D} \longrightarrow C_c(\mathbb{W})'$ implies that $\nu(D)$ is a Radon measure on \mathbb{W} for each $D \in \mathcal{D}$. This, of course, can be derived directly from the definition of $\nu(D)$ by writing it (see Eq. 10) as a sum of Dirac delta masses



$$\nu(S, \mu) = \sum_{\mathbf{x} \in S} \mu(\mathbf{x}) \delta_{\mathbf{x}}$$
 (11)

The representation-theoretic view, however, has the following advantages. The first is that it recovers the interpretation of persistence diagrams as rectangular measures introduced by Chazal et. al. in [19]. Indeed, Eq. 11 yields exactly the counting measure (Theorem 3.19, Eq. 3.6) from [19]. The second advantage is that it provides a natural framework in which to generalize persistence diagrams: from those in Eq. 11, to general Radon measures on W. This viewpoint has been recently studied by Divol and Lacombe in [30]. As they show, it allows one to apply the mature theoretical and computational tools from (partial) optimal transport, to problems in the space of persistence diagrams endowed with the Wasserstein distance. It would be interesting to see—though outside the scope of this work—if the approximation methods presented here, in particularly those in the next section, apply in the greater generality of persistence diagrams as Radon measures.

5 Approximating Continuous Functions on Persistence Diagrams

As we saw in Theorem 4.9, the function $\nu: \mathcal{D} \longrightarrow C_c(\mathbb{W})'$ provides a continuous embedding so that $\nu(D \sqcup D') = \nu(D) + \nu(D')$ for all $D, D' \in \mathcal{D}$. We can now proceed to the task of finding coordinate systems for \mathcal{D} (see Definition 4.1). The first thing to note is that composing ν with elements from $C_c(\mathbb{W})''$, the topological dual of $C_c(\mathbb{W})'$, yields continuous real-valued functions on \mathcal{D} . By construction, these functions also respect the monoidal structure \sqcup of \mathcal{D} . The elements of $C_c(\mathbb{W})''$ can be characterized as follows.

Theorem 5.1 Let V be a locally convex space, and endow its topological dual V' with the associated weak-* topology. That is, the smallest topology such that all the evaluations

$$\begin{array}{ccc} e_{\mathbf{v}} \ V' \longrightarrow & \mathbb{R} \\ T & \mapsto & T(\mathbf{v}) \end{array}$$

for $v \in V$, are continuous. Then the function

$$e: V \longrightarrow V''$$

 $\mathbf{v} \mapsto e_{\mathbf{v}}$

is an isomorphism of locally convex spaces.

Applying this theorem to the locally convex space $C_c(\mathbb{W})$, topologized with the strict inductive limit topology, implies that the elements of $C_c(\mathbb{W})''$ are evaluations e_f , with $f \in C_c(\mathbb{W})$ uniquely determined. Composing e_f with ν yields a continuous function $e_f \circ \nu : \mathcal{D} \longrightarrow \mathbb{R}$ which preserves the monoidal structure \sqcup of \mathcal{D} . Moreover, given $D \in \mathcal{D}$ we have that

$$e_f \circ v(D) = v_D(f) = v_f(D)$$

where $v_f : \mathcal{D} \longrightarrow \mathbb{R}$ is defined by Eq. 7. We saw in Lem. 4.6 that these types of functions are indeed continuous, but now we have the full picture: they arise exactly as the continuous linear functions on a linearization of \mathcal{D} . The goal now is to construct coordinate systems for \mathcal{D} by selecting appropriate subsets of $C_c(\mathbb{W})$. The sets of interest are defined next.

Definition 5.2 A **template system** for \mathcal{D} is a collection $\mathcal{T} \subset C_c(\mathbb{W})$ so that

$$\mathcal{F}_{\mathcal{T}} = \{ v_f \mid f \in \mathcal{T} \}$$

is a coordinate system (see Defn. 4.1) for \mathcal{D} . The elements of \mathcal{T} are called template functions.

The point of working with these template systems is that they can be used to approximate continuous functions on persistence diagrams, as given by the following theorem.

Theorem 5.3 Let $T \subset C_c(\mathbb{W})$ be a template system for \mathcal{D} , let $C \subset \mathcal{D}$ be compact, and let $F : C \longrightarrow \mathbb{R}$ be continuous. Then for every $\varepsilon > 0$ there exist $N \in \mathbb{N}$, a polynomial $p \in \mathbb{R}[x_1, \dots, x_N]$ and template functions $f_1, \dots, f_N \in \mathcal{T}$ so that

$$|p(v_{f_1}(D),\ldots,v_{f_N}(D))-F(D)|<\varepsilon$$

for every $D \in C$. That is, the collection of functions of the form

$$D \mapsto p(\nu_{f_1}(D), \dots, \nu_{f_N}(D)) \tag{12}$$

is dense in $C(\mathcal{D}, \mathbb{R})$ with respect to the compact-open topology.

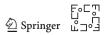
Proof Let $\mathcal{T} \subset C_c(\mathbb{W})$ be a template system for \mathcal{D} and let

$$\mathcal{F} = \{ v_f \mid f \in \mathcal{T} \} \subset C(\mathcal{D}, \mathbb{R})$$

be the corresponding coordinate system. Let $\mathcal{A} \subset C(\mathcal{D}, \mathbb{R})$ denote the algebra generated by $\mathcal{F} \cup \{1\}$. In other words, \mathcal{A} is the set of finite linear combinations of finite products of elements from $\mathcal{F} \cup \{1\}$. It follows that every element of \mathcal{A} can be written as

$$p(v_{f_1},\ldots,v_{f_N})$$

for some collection of templates $f_1, \ldots, f_N \in \mathcal{T}$ and some polynomial $p \in \mathbb{R}[x_1, \ldots, x_N]$. Let $\iota : \mathcal{C} \hookrightarrow \mathcal{D}$ be the inclusion and $\iota^* : C(\mathcal{D}, \mathbb{R}) \longrightarrow C(\mathcal{C}, \mathbb{R})$ the corresponding restriction homomorphism. Now, since \mathcal{F} separates points in \mathcal{D} and $\mathcal{F} \subset \mathcal{A}$, then $\iota^*(\mathcal{A})$ is an algebra which separates points in $C(\mathcal{C}, \mathbb{R})$ and contains the nonzero constant functions. The result follows from the Stone-Weierstrass theorem, which contends that any such algebra is dense with respect to the sup norm.



The main question now is how to go about constructing template systems in practice. The next theorem elucidates a method for producing countable template systems for \mathcal{D} by translating and re-scaling the support of any nonzero $f \in C_c(\mathbb{W})$. This shows, in particular, that there are plenty of coordinate systems for the space of persistence diagrams, and helps explain why we refer to nonzero elements in $C_c(\mathbb{W})$ as templates.

Theorem 5.4 Let $f \in C_c(\mathbb{W})$, $n \in \mathbb{N}$, $\mathbf{m} \in \mathbb{Z}^2$ and define the re-scales and translates

$$f_{n,\mathbf{m}}(\mathbf{x}) = f\left(n\mathbf{x} + \frac{\mathbf{m}}{n}\right).$$

If f is nonzero, then

$$\mathcal{T} = \left\{ f_{n,\mathbf{m}} \mid n \in \mathbb{N}, \mathbf{m} \in \mathbb{Z}^2 \right\} \cap C_c(\mathbb{W})$$

is a template system for \mathcal{D} . Moreover, if f is Lipschitz, then the elements of the associated coordinate system

$$\{v_{f_{n,\mathbf{m}}} = f_{n,\mathbf{m}} \circ v \mid f_{n,\mathbf{m}} \in \mathcal{T}\} = \mathcal{F}_{\mathcal{T}}$$

are Lipschitz on any relatively compact set $S \subset D$. That is, the coordinate system associated to a nonzero Lipschitz template function is stable on relatively compact subsets of D.

Proof In order to show that $\mathcal{F}_{\mathcal{T}}$ separates points in \mathcal{D} , let (S, μ) , $(T, \alpha) \in \mathcal{D}$ be distinct diagrams, and assume without loss of generality that there exists $\mathbf{y} = (y_1, y_2) \in S$ so that either: $\mathbf{y} \notin T$; or $\mathbf{y} \in T$ and $\mu(\mathbf{y}) \neq \alpha(\mathbf{y})$. Our strategy will be to find an element $f_{n,\mathbf{m}} \in \mathcal{T}$ so that $f_{n,\mathbf{m}}(\mathbf{y}) \neq 0$, and $f_{n,\mathbf{m}}(\mathbf{x}) = 0$ for all other $\mathbf{x} \in S \cup T$.

To begin, let $\mathbf{z} = (z_1, z_2) \in \mathbb{W}$ be so that $f(\mathbf{z}) \neq 0$. By continuity of f with respect to the Euclidean norm $\|\cdot\|$, which is equivalent to the sup norm $\|\cdot\|_{\infty}$, there exists r > 0 so that

$$B_r^{\infty}(\mathbf{z}) := \left\{ \mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x} - \mathbf{z}\|_{\infty} < r \right\} \subset \operatorname{supp}(f).$$

Moreover, since $\operatorname{supp}(f) \subset \mathbb{W}$ is compact, then there exists s > r so that $\operatorname{supp}(f) \subset B_s^{\infty}(\mathbf{z})$. Putting this together, we have r < s so that

$$B_r^{\infty}(\mathbf{z}) \subset \operatorname{supp}(f) \subset B_s^{\infty}(\mathbf{z}).$$

Fix $\varepsilon > 0$ small enough so that $B_{\varepsilon}^{\infty}(\mathbf{y}) \subset \mathbb{W}$ and $B_{\varepsilon}^{\infty}(\mathbf{y}) \cap (T \cup S) = \{\mathbf{y}\}$. Also, let $n \in \mathbb{N}$ be large enough so that $n \geq \max\left\{\frac{1}{r}, \frac{2s}{\epsilon}\right\}$. What we will show now is that it is possible to find $\mathbf{m} \in \mathbb{Z}^2$ for which $n\mathbf{y} + \frac{\mathbf{m}}{n} \in B_r^{\infty}(\mathbf{z})$, and so that $\mathbf{x} \notin B_{\varepsilon}^{\infty}(\mathbf{y})$ implies $n\mathbf{x} + \frac{\mathbf{m}}{n} \notin B_s^{\infty}(\mathbf{z})$. Indeed, define

$$L_j(t) = nt + (z_j - ny_j)$$

for j = 1, 2 with $\mathbf{z} = (z_1, z_2)$ and $\mathbf{y} = (y_1, y_2)$. This function has the property that $L_j(y_j) = z_j$. Further, if $|t - y_j| > \varepsilon$, then

$$|L_j(t)-z_j| = |(nt+z_j)-(ny_j+z_j)| = n|t-y_j| > \frac{2s}{\varepsilon}\varepsilon = 2s.$$

Let $k_j \in \mathbb{Z}$ be the unique integer so that

$$k_i \le z_i - ny_i < k_i + 1.$$

By dividing the interval $[k_j, k_j + 1)$ into n subintervals of length $\frac{1}{n}$, we have that there exists a unique integer $0 \le \ell_j < n$ so that

$$k_j + \frac{\ell_j}{n} \le z_j - ny_j < k_j + \frac{\ell_j + 1}{n}.$$

Let $m_i = nk_i + \ell_i$, and $\mathbf{m} = (m_1, m_2)$. It follows that

$$\left\|\mathbf{z} - \left(n\mathbf{y} + \frac{\mathbf{m}}{n}\right)\right\|_{\infty} < \frac{1}{n} \le r$$

and therefore $f_{n,\mathbf{m}}(\mathbf{y}) = f\left(n\mathbf{y} + \frac{\mathbf{m}}{n}\right) \neq 0$. Moreover, if $\mathbf{x} \notin B_{\epsilon}^{\infty}(\mathbf{y})$ and $j \in \{1,2\}$ is so that $|x_j - y_j| \geq \epsilon$, then

$$|nx_j - ny_j| = |L_j(x_j) - z_j| > 2s,$$

and therefore

$$\left\| \left(n\mathbf{x} + \frac{\mathbf{m}}{n} \right) - \mathbf{z} \right\|_{\infty} \ge \left| \left(nx_j + \frac{m_j}{n} \right) - z_j \right|$$

$$= \left| nx_j - ny_j - \left(z_j - \left(ny_j + \frac{m_j}{n} \right) \right) \right|$$

$$\ge \left| nx_j - ny_j \right| - \left| z_j - \left(ny_j + \frac{m_j}{n} \right) \right|$$

$$> 2s - r$$

$$\ge s,$$

showing that $n\mathbf{x} + \frac{\mathbf{m}}{n} \notin B_s^{\infty}(\mathbf{z})$, which in turn implies $f_{n,\mathbf{m}}(\mathbf{x}) = 0$.

Let us see that the support of $f_{n,\mathbf{m}}$ is a bounded subset of \mathbb{W} . To this end, let $\mathbf{x} \in \operatorname{supp}(f_{n,\mathbf{m}})$. Hence $n\mathbf{x} + \frac{\mathbf{m}}{n} \in \operatorname{supp}(f)$, and $\|(n\mathbf{x} + \frac{\mathbf{m}}{n}) - \mathbf{z}\|_{\infty} < s$. Then for j = 1, 2,

$$|nx_{j} - ny_{j}| \leq \left| (nx_{j} - ny_{j}) + z_{j} - \left(ny_{j} + \frac{m_{j}}{n} \right) \right| + \left| z_{j} - \left(ny_{j} + \frac{m_{j}}{n} \right) \right|$$

$$= \left| nx_{j} + z_{j} - \frac{m_{j}}{n} \right| + \left| z_{j} - \left(ny_{j} + \frac{m_{j}}{n} \right) \right|$$

$$\leq s + r < 2s,$$

and thus $|x_j - y_j| < \frac{2s}{n} < \varepsilon$. Therefore, $\mathbf{x} \in B_{\varepsilon}^{\infty}(\mathbf{y}) \subset \mathbb{W}$, and so $f_{n,\mathbf{m}} \in C_{\varepsilon}(\mathbb{W})$.

Thus far we have that $f_{n,\mathbf{m}}(\mathbf{y}) \neq 0$, and that if $\mathbf{x} \notin B_{\epsilon}^{\infty}(\mathbf{y})$ then $f_{n,\mathbf{m}}(\mathbf{x}) = 0$. This observation, paired with $B_{\epsilon}^{\infty}(\mathbf{y}) \cap (S \cup T) = \{\mathbf{y}\}$, implies that

$$v_{f_{n,\mathbf{m}}}(S,\mu) = \mu(\mathbf{y}) f_{n,\mathbf{m}}(\mathbf{y}) \neq 0.$$

If $\mathbf{y} \notin T$ then we have that $v_{f_{n,\mathbf{m}}}(T,\alpha) = 0$; and if $\mathbf{y} \in T$ then

$$\nu_{f_{n,\mathbf{m}}}(T,\alpha) = \alpha(\mathbf{y}) f_{n,\mathbf{m}}(\mathbf{y}) \neq \mu(\mathbf{y}) f_{n,\mathbf{m}}(\mathbf{y}) = \nu_{f_{n,\mathbf{m}}}(S,\mu),$$

showing that $\mathcal{F}_{\mathcal{T}}$ separates points in \mathcal{D} .

Let us now show that if f is Lipschitz and $\mathcal{S} \subset \mathcal{D}$ is relatively compact, then the elements of $\mathcal{F}_{\mathcal{T}}$ are Lipschitz on \mathcal{S} . Indeed, let $D, D' \in \mathcal{S}$, and let $\delta > 0$ be so that $d_B(D, D') < \delta$. Moreover, fix a δ -matching between D and D'. Recall that this means that if \mathbf{x}, \mathbf{x}' are matched, then $\|\mathbf{x} - \mathbf{x}'\|_{\infty} < \delta$, and that if for $\mathbf{z} = (z_1, z_2) \in \mathbb{W}$ unmatched we let

$$\bar{\mathbf{z}} = \left(\frac{z_1 + z_2}{2}, \frac{z_1 + z_2}{2}\right)$$

then $\|\mathbf{z} - \bar{\mathbf{z}}\|_{\infty} < \delta$.

Since S is relatively compact, then it is uniformly off-diagonally finite (see Def. 3.5), and hence there exists a uniform upper bound $\beta>0$ for the multiplicity in $\operatorname{supp}(f)$ of any diagram in S—see Definition 2.1. Now, if L>0 is the Lipschitz constant of f with respect to the $\sup \operatorname{norm} \|\cdot\|_{\infty}$ on \mathbb{W} , and $n\in\mathbb{N}$, $\mathbf{m}\in\mathbb{Z}^2$ are so that $f_{n,\mathbf{m}}\in\mathcal{T}$, then $f_{n,\mathbf{m}}\in C_c(\mathbb{W})$ is also Lipschitz with constant nL. Moreover, if $\mathbf{z}\in\mathbb{W}$ is unmatched, then $f_{n,\mathbf{m}}(\bar{\mathbf{z}})=0$ and

$$\begin{split} |\nu_{f_{n,\mathbf{m}}}(D) - \nu_{f_{n,\mathbf{m}}}(D')| &\leq \sum_{\mathbf{x},\mathbf{x}'} |f_{n,\mathbf{m}}(\mathbf{x}) - f_{n,\mathbf{m}}(\mathbf{x}')| + \sum_{\mathbf{z} \text{ unmatched}} |f_{n,\mathbf{m}}(\mathbf{z})| \\ &\leq \sum_{\mathbf{x} \text{ or } \mathbf{x}' \in \text{ Supp}(f)} nL\delta + \sum_{\mathbf{z} \in \text{ Supp}(f)} |f_{n,\mathbf{m}}(\mathbf{z}) - f_{n,\mathbf{m}}(\bar{\mathbf{z}})| \\ &\leq 2n\beta L\delta. \end{split}$$

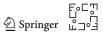
Since this inequality holds for any $\delta > d_B(D, D')$, it readily follows that

$$|\nu_{f_{n,\mathbf{m}}}(D) - \nu_{f_{n,\mathbf{m}}}(D')| \leq 2n\beta L d_B(D,D'),$$

and hence $v_{f_{n,\mathbf{m}}}$ is Lipschitz on S.

Remark 1 Recall that if $S \subset \mathcal{D}$ is relatively-compact, then there exist boxes $\mathbb{B}_k \subset \mathbb{W}$, $k \in \mathbb{N}$, so that $D \subset \bigcup_{k \in \mathbb{N}} \mathbb{B}_k$ for every $D \in \overline{S}$ (see Eq. 2 and Fig. 3). This implies that in order to find approximations to a continuous function $F : \overline{S} \longrightarrow \mathbb{R}$, it suffices to start with a nonzero $f \in C_c(\mathbb{W})$ and take only the re-scaled translates $f_{n,\mathbf{m}}$ for which

$$\operatorname{supp}(f_{n,\mathbf{m}})\cap\bigcup_{k\in\mathbb{N}}\mathbb{B}_k\neq\emptyset.$$



Remark 2 Let us say a few words on aligning the approximation methodologies outlined in Theorems 5.3 and 5.4—where one has access to infinitely many template functions—with practical algorithmic implementations that are inherently finite. A useful point of reference is the implementation of generalized linear models of a real variable x, in order to approximate a continuous function $F:I\subset\mathbb{R}\longrightarrow\mathbb{R}$ on a compact interval I. The Stone-Weierstrass theorem implies that F can be uniformly approximated in I via polynomials, or equivalently, via linear combinations of the monomials $1, x, x^2, \dots, x^n$ for $n \in \mathbb{N}$ arbitrary. In theory, one would need infinitely many monomials for arbitrary approximations, but in reality with finite training data, only finitely many monomials are relevant and even advisable. Indeed, this is the common bias-variance tradeoff where tools like cross validation can be used. The same view applies to learning with template functions on persistence diagrams. Given finite training data $D_1, \ldots, D_N \in \mathcal{S} \subset \mathcal{D}$ and $F(D_1), \ldots, F(D_N)$, then only finitely many re-scaled translates $f_{n,m}$ of $f \in C_c(\mathbb{W})$ are relevant to the problem at hand. These can be interpreted as a user-provided hyperparameter for the model—like the maximum degree of monomials in linear regression—or they can be derived from adaptive methods as described in [63] or [54]. The main idea being that one can identify those compact regions in W which are most relevant to the learning task at hand.

5.1 The Need for Compactness

In Sect. 3.4 we provided examples of non-relatively-compact sets $\mathcal{S} \subset \mathcal{D}$ which satisfied only two out of the three conditions from Theorem 3.7. We will revisit these examples next to see how the approximation strategy described in Theorem 5.3 can fail in the absence of compactness. Specifically, we will construct continuous functions $F: \overline{\mathcal{S}} \longrightarrow \mathbb{R}$ on the closure of $\mathcal{S} \subset \mathcal{D}$, which cannot be uniformly approximated by functions of the form (12). Indeed, let $\mathcal{S} = \{D_n = (S_n, \mu_n) \mid n \in \mathbb{N}\}$ be:

1. Not Uniformly Off-Diagonally Finite: $S_n = \{(0, 1)\}$ with $\mu_n(0, 1) = n$. Let $f \in C_c(\mathbb{W})$ be so that f(0, 1) = 1, and let $F : \mathcal{D} \longrightarrow \mathbb{R}$ be defined as $F(D) = e^{\nu_f(D)}$. Note that F is continuous on \mathcal{D} , and that $F(D_n) = e^n$ for every $n \in \mathbb{N}$. Let $f_1, \ldots, f_N \in C_c(\mathbb{W})$ and let $p \in \mathbb{R}[x_1, \ldots, x_N]$ be a polynomial of degree $k \geq 0$. Since

$$|v_{f_j}(D_n)|=n|f_j(0,1)|\leq n\|f_j\|_\infty$$

for every $1 \le j \le N$ and every $n \in \mathbb{N}$, then

$$\lim_{n\to\infty}\frac{\left|p\left(\nu_{f_1}(D_n),\ldots,\nu_{f_N}(D_n)\right)\right|}{n^{k+1}}=0.$$

If we had the uniform bound

for some $0 < \epsilon < \infty$ and every $n \in \mathbb{N}$, then dividing both sides of the inequality by n^{k+1} and taking the limit as $n \to \infty$ would imply that

$$\lim_{n\to\infty} \frac{e^n}{n^{k+1}} = 0.$$

This is a contradiction since the above limit is ∞ , and therefore the restriction of F to \overline{S} cannot be uniformly approximated by functions of the form (12).

2. Not Off-Diagonally Birth Bounded: $S_n = \{(n, n+1)\}$ with $\mu_n(n, n+1) = 1$. The first thing to note is that $d_B(D_n, D_m) = \frac{1}{2}$ for every $n \neq m \in \mathbb{N}$, and thus S has no accumulation points. This implies that $\overline{S} = S$, and that any real-valued function on \overline{S} is continuous. Let $F : \overline{S} \longrightarrow \mathbb{R}$ be defined as

$$F(D_n) = n$$
 , $n \in \mathbb{N}$.

If $f_1, \ldots, f_N \in C_c(\mathbb{W})$ and $p \in \mathbb{R}[x_1, \ldots, x_N]$, then

$$v_{f_1}(D_n) = \cdots = v_{f_N}(D_n) = 0$$

for all *n* large enough, and thus there exists $N_0 \in \mathbb{N}$ so that

$$p\left(\nu_{f_1}(D_n),\ldots,\nu_{f_N}(D_n)\right)=p(0,\ldots,0)$$

for every $n > N_0$. It follows that

$$\lim_{n\to\infty} \left| F(D_n) - p\left(\nu_{f_1}(D_n), \dots, \nu_{f_N}(D_n)\right) \right| = \infty$$

and thus F cannot be uniformly approximated in \overline{S} by functions of the form (12).

3. Not Bounded: $S_n = \{(0, n)\}$ with $\mu_n(0, n) = 1$. Let $F : \mathcal{D} \longrightarrow \mathbb{R}$ be

$$F(D) = 2d_R(D, \varnothing).$$

It follows that F is continuous and that $F(D_n) = n$ for every $n \in \mathbb{N}$. The argument now proceeds exactly as in (2) above.

5.2 Compact Approximations with the Wasserstein Distance

Thus far we have established that $\nu: \mathcal{D} \longrightarrow C_c(\mathbb{W})'$ is a linearization of (\mathcal{D}, d_B) (see Theorem 4.9), and that the rescales/translates of any nonzero $f \in C_c(\mathbb{W})$ can be used to construct compact-open dense subsets of $C(\mathcal{D}, \mathbb{R})$ (see Theorems 5.4 and 5.3). A natural question is whether similar results hold for (\mathcal{D}_p, d_{W_p}) —i.e., for the p-Wasserstein distance, $p \geq 1$. We will see next that this is indeed the case.

Recall that $\mathcal{D}_p \subset \mathcal{D}$ for every $p \in \mathbb{N}$, and that this inclusion is (uniformly) continuous. It follows that,



Corollary 5.5 ν restricts to a d_{W_p} -continuous function $\nu: \mathcal{D}_p \longrightarrow C_c(\mathbb{W})'$ with the same properties as in Theorem 4.9

This implies that $v_f : \mathcal{D}_p \longrightarrow \mathbb{R}$ is continuous for every $f \in C_c(\mathbb{W})$, and thus any template system $\mathcal{T} \subset C_c(\mathbb{W})$ for \mathcal{D} is also a template system for \mathcal{D}_p . Using the first half of Theorem 5.4, and following the same proof as in Theorem 5.3, we have that

Corollary 5.6 Let $f \in C_c(\mathbb{W})$ be nonzero, and let

$$\mathcal{T} = \left\{ f_{n,\mathbf{m}} \mid n \in \mathbb{N}, \mathbf{m} \in \mathbb{Z}^2 \right\} \cap C_c(\mathbb{W}).$$

Then the set of functions of the form

$$D \mapsto P(\nu_{f_1}(D), \ldots, \nu_{f_N}(D))$$

for $N \in \mathbb{N}$, $P \in \mathbb{R}[x_1, ..., x_N]$ and $f_1, ..., f_N \in \mathcal{T}$, is compact-open dense in $C(\mathcal{D}_p, \mathbb{R})$.

6 Example Template Functions

At this point in the story, we shift our view from theory to practice, as the mathematical framework built to this point leaves open the choice of template system. In our experiments, we use two collections of functions, but we have no reason to suspect that these are the only or even the best available options. The first, which we call tent functions, are described in Sect. 6.1. The second are interpolating polynomials, traditionally used for approximating functions, which are described in Sect. 6.2.

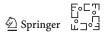
For the entirety of this section, we will define functions on the birth-lifetime plane as this simplifies notation substantially. We use the tilde to denote the portions that are defined in this plane to emphasize the change from the birth-lifetime plane. So, let $\widetilde{\mathbb{W}} = \{(x,y) \mid x \geq 0, y > 0\}$; that is, the conversion of \mathbb{W} to the birth-lifetime plane. Likewise, let $\widetilde{\mathbb{W}}^{\varepsilon} = \{(x,y) \in \widetilde{\mathbb{W}} \mid y > \varepsilon\}$ so that it is the conversion of \mathbb{W}^{ε} to the birth-lifetime plane. Given $\mathbf{x} = (a,b) \in \mathbb{W}$, we write $\widetilde{\mathbf{x}} = (a,b-a) \in \widetilde{\mathbb{W}}$ for the converted point. Given a diagram $D = (S,\mu)$, we write $\widetilde{D} = (\widetilde{S},\widetilde{\mu})$ where $\widetilde{S} = \{\widetilde{\mathbf{x}} \mid \mathbf{x} \in S\}$ and $\widetilde{\mu}(\widetilde{\mathbf{x}}) = \mu(\mathbf{x})$.

6.1 Tent Functions

We first define a template system in the birth-lifetime plane which we call tent functions. Given $\mathbf{a} = (a,b) \in \widetilde{\mathbb{W}}$ and a radius $0 < \delta < b$, define the tent function on $\widetilde{\mathbb{W}}$ to be

$$g_{\mathbf{a},\delta}(x, y) = \left| 1 - \frac{1}{\delta} \max\{|x - a|, |y - b|\} \right|_{+}$$

where $|r|_+ = \max\{r, 0\}$ for $r \in \mathbb{R}$. As $\delta < b$, this function has support in the compact box $[a - \delta, a + \delta] \times [b - \delta, b + \delta] \subset \widetilde{\mathbb{W}}$.



Given a persistence diagram $D = (S, \mu)$, the tent function is defined to be the sum over the evaluation on the points in the diagram, namely

$$G_{\mathbf{a},\delta}(D) = \widetilde{G}_{\mathbf{a},\delta}(\tilde{D}) = \sum_{\tilde{\mathbf{x}} \in \tilde{\mathbf{x}}} \widetilde{\mu}(\mathbf{x}) \cdot g_{\mathbf{a},\delta}(\tilde{\mathbf{x}}).$$

We use G or \widetilde{G} depending on whether we want our input to be a diagram in the birth-death or birth-lifetime plane, respectively, but all subscript notation is written in the birth-lifetime plane for ease of notation.

We then have the following theorem.

Theorem 6.1 The collection of tent functions

$$\{G_{\mathbf{a},\delta} \mid \mathbf{a} = (a,b) \in \widetilde{\mathbb{W}}, 0 < \delta < b\}$$

separates points in \mathcal{D} .

Proof Let $D_1 = (S_1, \mu_1)$ and $D_2 = (S_2, \mu_2) \in \mathcal{D}$ be distinct. WLOG there is an $\mathbf{x} \in S_1$ for which either (i) $\mathbf{x} \notin S_2$ or (ii) $\mathbf{x} \in S_2$ but $\mu_1(\mathbf{x}) > \mu_2(\mathbf{x})$. For ease of notation, assume in case (ii) that $\mathbf{x} \in S_2$ and $\mu_2(\mathbf{x}) = 0$. Then we always have $\mathbf{x} \in S_2$ and $\mu_1(\mathbf{x}) > \mu_2(\mathbf{x})$.

Let $\widetilde{\mathbf{x}} = (a, b)$. For any δ , define $\widetilde{B_{\delta}} = [a - \delta, a + \delta] \times [b - \delta, b + \delta]$ and note that this is the support of $g_{\mathbf{x},\delta}$. As D_1 and D_2 are in \mathcal{D} , both diagrams have finite multiplicity in $\widetilde{\mathbb{W}}^{b/2}$. So, there exists a $\delta < b/2$ so that $\widetilde{S_1} \cap \widetilde{B_{\delta}} = \{\mathbf{x}\} = \widetilde{S_2} \cap B$. As \mathbf{x} is the only point in either diagram in the support of $g_{\mathbf{x},\delta}$,

$$G_{\mathbf{x} \delta}(D_1) = \mu_1(\mathbf{x}) > \mu_2(\mathbf{x}) = G_{\mathbf{a} \delta}(D_2).$$

Thus, the collection of tent functions separates points.

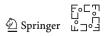
For practical purposes, we pick a subset of these tent functions. Let $\delta > 0$ be the partition scale, let d the number of subdivisions along the diagonal (resp. y axis), and let $\varepsilon > 0$ be the upward shift. In our experiments described in Sect. 8, we use the collection of tent functions given by

$$\left\{G_{(\delta i,\delta j+\varepsilon),\delta}\mid 0\leq i\leq d, 1\leq j\leq d\right\}. \tag{13}$$

That is, these are the tent functions centered at a regular grid shifted up by ε to ensure that g is supported on a compact set in $\widetilde{\mathbb{W}}$. See Fig. 4 for an example.

6.2 Interpolating Polynomials

Say we are given a nonempty, finite set of distinct mesh values $\mathcal{A} = \{a_i\}_{i=0}^m \subset \mathbb{R}$ and a collection of evaluation values $\{c_i \in \mathbb{R}\}$, the first goal is to build a polynomial such that $f(a_i) = c_i$ for all i. The Lagrange polynomial $\ell_j^{\mathcal{A}}(x)$ corresponding to node a_j is defined as



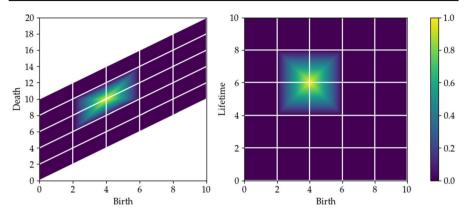


Fig. 4 Tent function $g_{(4,6),2}$ drawn in the birth-death plane (left) and in the birth-lifetime plane (right) for the grid determined by d=5, $\delta=2$, and $\varepsilon=0$. Note that the template system for tent functions require $\varepsilon>0$, which would shift this function up by ε

$$\ell_j^{\mathcal{A}}(x) = \prod_{i \neq j} \frac{x - a_i}{a_j - a_i}.$$
 (14)

Note that this function satisfies

$$\ell_j^{\mathcal{A}}(a_k) = \begin{cases} 1, & j = k, \\ 0, & \text{otherwise,} \end{cases}$$
 and $\sum_{j=0}^m \ell_j^{\mathcal{A}}(x) = 1.$

The Lagrange interpolation polynomial is then simply $f(x) = \sum_{j=0}^{m} c_j \ell_j^{\mathcal{A}}(x)$. Note that for numerical stability, one must work with the barycentric form of Lagrange interpolation formula described by [7] and shown in Appendix A.

We will now use these polynomials to create functions on $\widetilde{\mathbb{W}}$. Assume we have two collections of mesh points $\mathcal{A} = \{a_i\}_{i=0}^m \subset \mathbb{R} \text{ and } \mathcal{B} = \{b_i\}_{i=0}^n \subset \mathbb{R}_{>0} \text{ so that } (a_i, b_j) \in \widetilde{\mathbb{W}} \text{ for all } i, j.$ Then, given a collection of evaluation points $\mathcal{C} = \{c_{i,j}\}_{i,j} \subset \mathbb{R}$, we want to build a polynomial for which $f(a_i, b_j) = c_{i,j}$. Note that in general the evaluation points \mathcal{C} are not part of the persistence diagrams to be evaluated; nevertheless, these values are not needed in our construction but we do keep track of their coefficients. We define the 2D interpolating polynomial for the collection \mathcal{A} , \mathcal{B} , \mathcal{C} to be

$$f(x,y) = \sum_{i=0}^{m} \sum_{j=0}^{n} c_{i,j} \cdot g\left(\ell_i^{\mathcal{A}}(x) \cdot \ell_j^{\mathcal{B}}(y)\right). \tag{15}$$

where $g(\cdot)$ is either the identity function or $g(\cdot) = |\cdot|$; in our experiments we used the latter for the simple reason that it seemed to give better results. We now evaluate f at each of the N query points which are the points of a persistence diagram in $\widetilde{\mathbb{W}}$ to obtain N equations that we can write as

$$\mathbf{f} = \mathbf{L} \, \mathbf{c},\tag{16}$$



where **L** is an $((m + 1) \times (n + 1)) \times N$ matrix, and **f** is an $(m + 1) \times (n + 1)$ vector obtained by concatenating a 2D mesh, similar to the one shown in Fig. 6, rowwise. Note that this representation implies an ordering of the points of the persistence diagram, but after the next step, the order will not matter.

Each column of matrix **L** in Eq. (16) represents a vector that describes the contributions of all the query points to the corresponding entry in **f**. Renumbering the entries in vector **f** according to $(i, j) \mapsto i(n+1)+j+1=r$ where $r \in \{0, \ldots, (m+1)(n+1)\}$, we can now assign a score for each point in the mesh using the map $S_r : \mathbb{R}^N \to \mathbb{R}$, i.e., by operating on the rows of matrix **L** according to

$$S_r = \sum_{j=0}^{N-1} L_{j,r}.$$
 (17)

Choosing a larger base mesh implies using a higher degree polynomial in the interpolation. Therefore, the role of increasing the degree of the polynomial is similar to the role of increasing the number of tent functions. A larger mesh leads to more features which gives a tool for either increasing or reducing the number of features. The former improves the fit to the training set, while the latter reduces the number of features which allows mitigating overfitting effects. While any class of interpolating polynomials can be used, in this study we chose Chebyshev interpolating polynomials due to their excellent approximation properties, see [61]. Appendix A describes how to use the interpolation matrices separately obtained for each of the birth times and lifetimes of a given persistence diagram to construct $L_{j,r}$.

Fix a compact region $K \subset \widetilde{\mathbb{W}}$, an $\varepsilon > 0$ such that

$$\overline{K^{\varepsilon}} = \{ \mathbf{x} \in \mathbb{R}^2 \mid d_K(\mathbf{x}) \le \varepsilon \} \subset \widetilde{\mathbb{W}} \qquad \text{where} \qquad d_K(\mathbf{x}) = \inf_{\mathbf{y} \in K} \|\mathbf{x} - \mathbf{y}\|$$

and a collection of mesh points $\{(a_i, b_j)\}\subset K$ given by $\mathcal A$ and $\mathcal B$ as above. Define $h_{K,\varepsilon}$ to be a continuous function on $\widetilde{\mathbb W}$ such that

$$h_{K,\varepsilon}(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in K \\ 0 & d_K(\mathbf{x}) \ge \varepsilon \end{cases}$$

For instance, one can let $h_{K,\varepsilon}(\mathbf{x}) = \max\left\{0, 1 - \frac{1}{\epsilon}d_K(\mathbf{x})\right\}$. Note that the support of this function is contained in the compact set $\overline{K^\varepsilon} \subset \widetilde{\mathbb{W}}$.

Let $C_{i,j}$ be the collection of evaluation values which are entirely 0 except for $c_{i,j}=1$. Define $f_{i,j}^{\mathcal{A},\mathcal{B}}=f_{i,j}$ to be the interpolating polynomial (Eq. 15) for this setup. Then the function on diagrams is defined to be

$$F_{i,j}^{\mathcal{A},\mathcal{B},K,\varepsilon}(D) = \widetilde{F}_{i,j}^{\mathcal{A},\mathcal{B},K,\varepsilon}(\widetilde{D}) := \sum_{\tilde{\mathbf{x}} \in \widetilde{S}} \mu(\tilde{\mathbf{x}}) \cdot f_{i,j}^{\mathcal{A},\mathcal{B}}(\tilde{\mathbf{x}}) \cdot h_{K,\varepsilon}(\tilde{\mathbf{x}}).$$



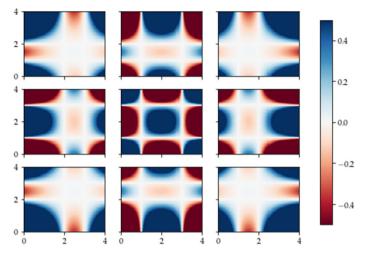


Fig. 5 An example of interpolating polynomials for the mesh $\mathcal{A} = \mathcal{B} = \{1, 2, 3\}$ for evaluation values $\mathcal{C}_{i,j} = \{c_{i,j} = 1, c = 0 \text{ else}\}$. Viewing the grid of images, $f_{i,j}^{\mathcal{A},\mathcal{B}}$ is drawn at location (i,j) in the figure, with (1,1) at the bottom left

See Fig. 5 for examples of these functions for different parameter choices. We have the following theorem to show that these interpolating polynomials can be used as template functions.

Theorem 6.2 *The collection of interpolating polynomials*

$$\left\{F_{i,j}^{\mathcal{A},\mathcal{B},K,\varepsilon}\right\}$$

separates points, where the collection varies over all choices of compact $K \subset \widetilde{\mathbb{W}}$, $\delta \in \mathbb{R}_{>0}$, and of mesh A, B as specified above.

Proof We are given two persistence diagrams $D_1 = (S_1, \mu_1)$ and $D_2 = (S_2, \mu_2) \in \mathcal{D}$, with $D_1 \neq D_2$. WLOG there is an $\mathbf{x} = (a, b) \in S_1$ for which either (i) $\mathbf{x} \notin S_2$ or (ii) $\mathbf{x} \in S_2$ but $\mu_1(\mathbf{x}) > \mu_2(\mathbf{x})$. To avoid case checking, we assume as before that $\mathbf{x} \in S_2$ with $\mu_2(\mathbf{x}) = 0$ in the later case. This way, in both cases have $\mathbf{x} \in S_2$ and $\mu_1(\mathbf{x}) > \mu_2(\mathbf{x})$.

Choose a compact set $K \ni \mathbf{x}$ and ε both small enough so that $\overline{K^{\varepsilon}} \cap S_1 = \mathbf{x}$ and $\overline{K^{\varepsilon}} \cap S_2 = \mathbf{x}$. If $\widetilde{\mathbf{x}} = (a, b)$, set $\mathcal{A} = \{a\}$ and $\mathcal{B} = \{b\}$. Note that in this overly simplistic setup, $\ell^{\mathcal{A}}(x) = x/a$ and $\ell^{\mathcal{B}}(y) = y/b$, so the only interpolating polynomial is f(x, y) = g((xy)/(ab)). Whether g is the identity or the absolute value function, f evaluates to 1 at $\widetilde{\mathbf{x}}$. Because the only point in either diagram inside $\overline{K^{\varepsilon}}$ is \mathbf{x} , $h_{K,\varepsilon}(\widetilde{\mathbf{x}}) = 1$ and $h_{K,\varepsilon}(\widetilde{\mathbf{y}}) = 0$ for every other $\mathbf{y} \in S_1 \cup S_2$, so

$$F(D_i) = \sum_{\tilde{\mathbf{x}} \in \tilde{\mathbf{x}}} \mu(\tilde{\mathbf{x}}) \cdot f(\tilde{\mathbf{x}}) \cdot h_{K,\varepsilon}(\tilde{\mathbf{x}}) = \mu_i(\mathbf{x})$$

for i = 1, 2. Thus F separates the two diagrams.

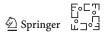
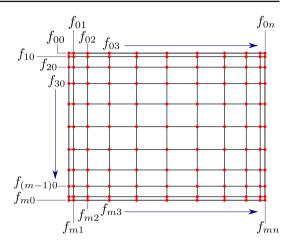


Fig. 6 An example 11×11 2D grid with m = n = 10 defined using the Chebyshev points of the second kind



In our experiments, we set K to be a box $[A, A'] \times [B, B']$ with B > 0, ε to be either machine precision or B/2, and use the non-uniform Chebyshev mesh as seen in Fig. 6 ([61]). Our naieve implementation of this featurization was quite slow; see Sect. A for an explanation of the vectorization used to implement and speed up the code.

7 Implementing Regularized Regression/Classification with Templates

The results thus far imply that template systems on \mathcal{D} can be used to featurize persistence diagrams; these vectorizations, in turn, can be used as inputs to machine learning algorithms for classification and regression tasks. We describe next one avenue for implementing these ideas in practice. Indeed, given a finite collection of labeled persistence diagrams

$$\{(D_m, \ell_m)\}_{m=1}^M \subset \mathcal{D} \times L$$

with $L \subset \mathbb{R}$, and a template system $\mathcal{T} \subset C_c(\mathbb{W})$, the goal is to find $N \in \mathbb{N}$, template functions $f_1, \ldots, f_N \in \mathcal{T}$, and a polynomial $p \in \mathbb{R}[x_1, \ldots, x_N]$, such that the function

$$P: \mathcal{D} \longrightarrow \mathbb{R}$$

$$D \mapsto p(\nu_{f_1}(D), \dots, \nu_{f_N}(D))$$

satisfies $P(D_m) \approx \ell_m$ for m = 1, ..., M. It follows from Thm. 5.3 that this process results in arbitrarily accurate approximations on compact subsets of \mathcal{D} , provided the labels ℓ_m vary continuously. In practice the template functions $f_1, ..., f_N$ can be either provided by the user from a specific class, like tents or interpolating polynomials (Sect. 6), or can be derived from adaptive and data-driven strategies as described



in related work of the authors; please refer to [63] or [54] for several adaptive strategies and comparisons. Given template functions $f_1, \ldots, f_N \in C_c(\mathbb{W})$, the optimal polynomial $p \in \mathbb{R}[x_1, \ldots, x_N]$ is uniquely determined by its vector of coefficients, $\mathbf{a} \in \mathbb{R}^k$. We will make this explicit with the notations $p_{\mathbf{a}}$ and $P_{\mathbf{a}}$, and an optimization will be set up in order to determine $\mathbf{a} \in \mathbb{R}^k$ from the available labeled data.

The error of fit $P_{\mathbf{a}}(D_m) \approx \ell_m$ is measured in the usual way via a loss function

$$\mathcal{E}: \mathbb{R} \times L \longrightarrow \mathbb{R}$$

where common choices include:

Square is given by $\mathcal{E}_{sq}(t,\ell)=(t-\ell)^2, L=\mathbb{R}$, and yields a least-squares regression. Can handle multi-class classification.

Hinge is given by $\mathcal{E}_{hg}(t,\ell) = \max\{0, 1-\ell \cdot t\}$, with $L = \{-1, 1\}$, and appears in the soft-margin classifier of support vector machine.

Logistic is given by the log-loss $\mathcal{E}_{log}(t,\ell) = \ln(1 + e^{-\ell \cdot t})$, with $L = \{-1,1\}$, and yields logistic regression.

Meanwhile, the complexity of the model can be measured, for instance, via a regularization function

$$\Omega: \mathbb{R}^k \longrightarrow [0, \infty).$$

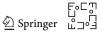
The regularized optimization scheme looking to minimize the regularized mean loss is

$$\mathbf{a} = \underset{\mathbf{v} \in \mathbb{R}^k}{\operatorname{argmin}} \ \frac{1}{M} \sum_{m=1}^{M} \mathcal{E}\left(P_{\mathbf{v}}(D_m), \ell_m\right) + \alpha \Omega(\mathbf{v})$$

where $\alpha > 0$ is the regularization parameter, often chosen from the set $\{10^n\}_{n \in \mathbb{Z}}$.

Visualization of Coefficients.

Our collections of template functions have a uniquely 2d geometric flavor. In particular, for both tent functions and Lagrange polynomials on our formulation, we have a function for each a_i , b_j location on a mesh. This means that we can pull the coefficients \mathbf{v} determined in the optimization back to the grid which built them for visualization by drawing a heat map with $v_{i,j}$ drawn at (i,j) to more fully understand the model. Examples of this are shown in Figs. 10 and 12. We can use these heat maps can be used to help localize the important features in the learning task. However, since we did not put a sparsity penalization term (e.g., L_1 regularization or lasso) in the regression problem, the highlighted pixels may be more than what is needed to localize the problem.



8 Experiments

8.1 Code

Code for doing classification and regression using tent and interpolating polynomial functions is available in the python teaspoon package. Classification and regression were done using ridge regression, where the loss and regularization functions are both square. Computation is done using RidgeClassifierCV and RidgeCV functions from the sklearn package. These functions are the counterparts of Ridge and RidgeCV functions with built-in cross-validation of the regularization parameter, α . Unless otherwise noted, we have used the default in sklearn, which chooses α from the set $\{10^n \mid n \in \{-1,0,1\}\}$.

The main hyperparameters to be chosen prior to running experiments are those controlling the number of template functions used. The number of tent functions is controlled by d in Eq. 13; the default is d=10 unless otherwise noted. We then choose δ and ε which control the support of the functions. If they are not specified in advance, we use the following procedure. First, determine a bounding box is for the points in all diagrams in the test set. This is done by setting ε to be half the minimum persistence of all points in the diagrams, and then setting δ to be the smallest value so that the box $[0, \delta d] \times [0\delta d + \varepsilon]$ contains all points in the birth-lifetime plane. Then in the notation of Eq. 13, the support of all tent functions is contained in the box $[-\delta, (d+1)\delta] \times [\varepsilon, (d+1)\delta + \varepsilon]$.

For interpolating polynomials, the number and the type of the mesh points must be specified. We used the roots of the Chebyshev polynomials of the first kind. The number of the used functions is controlled by the length of the meshes \mathcal{A} and \mathcal{B} in Eq. 15, i.e., the number of the chosen Chebyshev points in each direction. The default length is m = n = 10 unless otherwise noted.

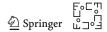
All experiments were run using seed = 48824 and 33% of the data reserved for testing. Scores for classification experiments are reported using the percent that were correctly classified. Scores for regression experiments are reported using the coefficient of determination, R^2 . Note that this latter score can potentially take negative values; perfect regression would score 1, and a method which returns the constant prediction of the expected value is given a score of 0.

8.2 Off-Diagonal, Normally Distributed Points

We generate diagrams from the following procedure. Given μ and σ , draw n points from the gaussian $N(\mu, \sigma)$ on \mathbb{R}^2 . Retain all points which are are in \mathbb{W} . For our simulations, we fixed $\sigma=1$ and varied μ . Examples of two overlaid example diagrams are shown in Fig. 7.

Classification. We tested our classification accuracy with the following experiment. We chose two collections, A and B, of 500 persistence diagrams each generated by drawing n=20 points (note that this means there are at most n points). The means μ were different: $\mu_A=(1,3)$ and μ_B was varied along the line (1,3)t+(2,5)(1-t) for

² https://github.com/lizliz/teaspoon.



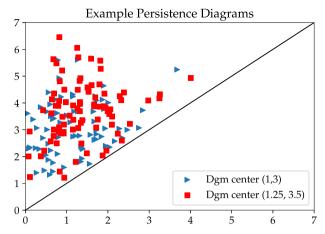


Fig. 7 Example diagrams randomly generated by the procedure described in Sect. 8.2. Two diagrams drawn from a distributions with different choices of μ are shown

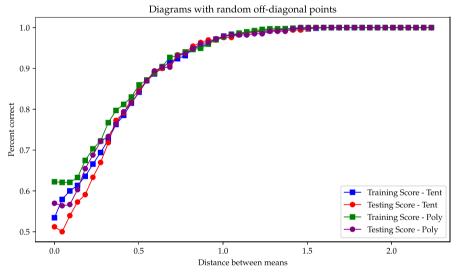
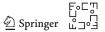


Fig. 8 Results from classification test for pairs of choices of μ with $\mu_A = (1,3)$ and μ_B chosen on the line from (1,3) to (2,5)

 $t \in [0, 1]$. Classification accuracy using tent and interpolating polynomial functions is shown in Fig. 8. As expected, the correct classification percentage for the test set is around 50% when $\mu_A \sim \mu_B$, and improves as they move farther apart. In particular, by the time the means are at distance apart equal to the standard deviation used for the normal distribution ($\sigma = 1$), classification is well above 90%. For this particular experiment, we do not see any difference between the choice of template function used.



Regression. We further ran two regression versions of the experiment as follows. In the first test, we drew 500 diagrams from the above procedure with a choice of center μ drawn uniformly on the line segment $t(1,3)+(1-t)(6,8), t\in [0,1]$. Then, we predicted the distance from μ to (1,3). We call this the "line" experiment. Second, we drew μ from the normal distribution N((1,3),1) and again predicted the distance of μ from the point (1,3). We call this the "ball" experiment.

Each of these experiments was run 10 times, and the results can be seen in Table 1. Example predictions for single runs can be seen in Fig. 9, and the coefficients for these examples are in Fig. 10. Note that the coefficients are drawn at the location of their index. In particular, the interpolating polynomials are determined using a non-uniform mesh, so the heatmap for these coefficients does not align with the location of the associated point.

Table 1 The R^2 results of the regression tests described in Sect. 8.2

Tents-train		Tents-test	Polynomials-train	Polynomials-test	
Line	0.977 ± 0.002	0.970 ± 0.004	0.979 ± 0.001	0.970 ± 0.002	
Ball	0.823 ± 0.023	0.782 ± 0.023	0.832 ± 0.021	0.786 ± 0.021	

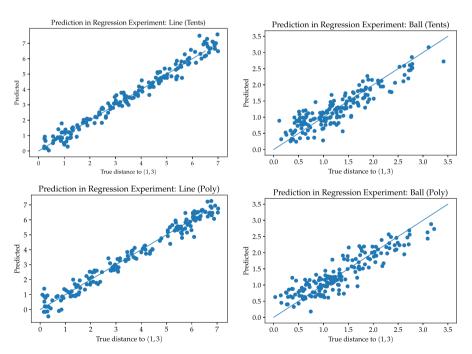
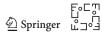


Fig. 9 True versus predicted of distance to starting mean for the ball and line experiments described in Sect. 8.2



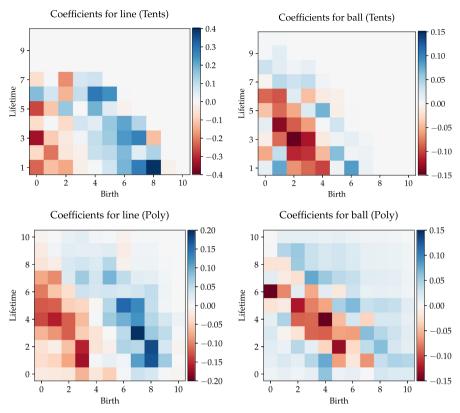


Fig. 10 Coefficients for regression experiment with mean drawn from a line (left) and from a ball (right). The top row uses tent functions, the bottom uses interpolating polynomials. The *x* and *y* coordinates correspond to the index of the test function used

8.3 Manifold Experiment

Following an experiment run in [1], we generated collections of point clouds drawn from different manifolds embedded in \mathbb{R}^2 or \mathbb{R}^3 . Each point cloud has N=200 points. The categories are as follows:

Annulus.	Points drawn uniformly from an annulus with inner radius 1 and
	outer radius 2.

3 clusters. The 200 points are drawn from one of three different normal distributions, with means (0, 0), (0, 2) and (2, 0) respectively, and all with standard deviation 0.05.

3 clusters of The points are drawn from normal distributions with standard devi-3 clusters. The points are drawn from normal distributions with standard deviation 0.05 centered at the points (0, 0), (0, 1.5), (1.5, 0), (0, 4),

(1, 3), (1, 5), (3, 4), (3, 5.5), (4.5, 4).

Cube. Points drawn uniformly from $[0, 1]^2 \subseteq \mathbb{R}^2$.



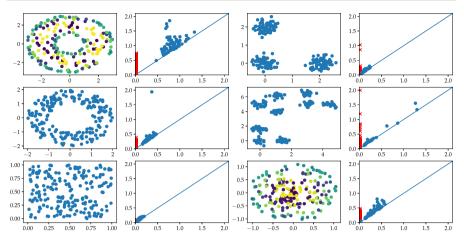


Fig. 11 Example point clouds from the experiment described in Sect. 8.3. From top left reading across rows, the point clouds are a torus (in \mathbb{R}^3 but drawn with color as the third coordinate), 3 clusters, annulus, three clusters of three clusters, uniform box, and a sphere (again drawn with third coordinate as color). The associated persistence diagrams are shown to the right of the point clouds; the 1-dimensional diagram is given by blue dots and the 0-dimensional diagram is shown by red x's. Diagrams are drawn with the same axis values

Torus. Points drawn uniformly from a torus thought of as rotating a circle

of radius 1 in the xz-plane centered at (2,0) around the z-axis. The generation of the points is done using the method from [29].

Sphere. Points drawn from a sphere of radius 1 in \mathbb{R}^3 with uniform noise in [-0.05, 0.05] added to the radius.

Examples of each of these can be seen in Fig. 11. Code for generation of these point clouds as well as the full dataset can be found in the teaspoon package at teaspoon.MakeData.PointCloud.testSetManifolds.

The choice of tent function parameters was done as follows. We determined the bounding box necessary to enclose the training set diagrams in the (birth, lifetime) plane and added padding of 0.05. We fixed d=10; ϵ was chosen to be half the minimum lifetime over all training set diagrams; then δ was chosen to ensure coverage of the bounding box.

We reserved 33% of the data for testing and trained a regression model on the remaining data. The results of this experiment averaged over 10 runs can be seen in Table 2. In this experiment, particularly when we have 50 or more diagrams per class, we see excellent (\geq 99%) classification.

8.4 Shape Data

We compared our results to the kernel method results reported in [55] by applying feature functions to the same data set from that paper. In particular, the synthetic SHREC 2014 data set ([53]) consists of 3D meshes of humans in different poses. The people are labeled as male, female, and child (five each); and each person assumes one of 20 poses. Reininghaus et al. defined a function on each mesh using the heat kernel



No. Dgms	Tents		Polynomials		
	Train	Test	Train	Test	
10	$99.8\% \pm 0.9$	$96.5\% \pm 3.2$	$99.8\% \pm 0.9$	$95.0\% \pm 3.9$	
25	$99.9\% \pm 0.3$	$99.0\% \pm 1.0$	$99.7\% \pm 0.5$	$97.6\% \pm 1.5$	
50	$99.9\% \pm 0.2$	$99.9\% \pm 0.3$	$100\% \pm 0$	$99.2\% \pm 0.9$	
100	$99.8\% \pm 0.1$	$99.7\% \pm 0.4$	$99.6\% \pm 0.2$	$99.3\% \pm 0.5$	
200	$99.5\% \pm 0.1$	$99.5\% \pm 0.3$	$99.2\% \pm 0.2$	$98.9\% \pm 0.5$	

Table 2 Results from the manifold test described in Sect. 8.3 for different numbers of examples drawn for each type of manifold

The reported results are averaged over 10 experiments each

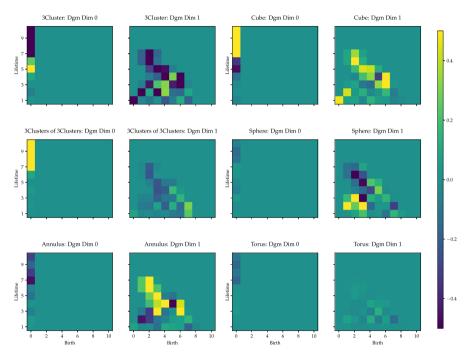


Fig. 12 Coefficients for the manifold experiment using tent functions run with 100 diagrams each

signature ([60]) for 10 parameters and computed the 0- and 1-dimensional diagrams of each.

We start with this data set of 300 pairs of persistence diagrams (0- and 1-dimensional) for each of the 10 parameter values, and predicted the human model; i.e. which of the 15 people were represented by each mesh. A comparison of the results reported in [55] with our method using polynomial functions is shown in Table 3. Additional results using tent functions are provided in the appendix (Table 4).

For this experiment, polynomial features (Table 3) were considerably more successful than the tent functions. Further, using the 0- and 1-dimensional persistence diagrams together was largely better than the 1-dimensional diagram alone, and



freq	MSK	Dim 0		Dim 1		Dim 0 & Dim 1	
		Train	Test	Train	Test	Train	Test
1	94.7% ± 5.1	$94.3\% \pm 0.5$	$67.1\% \pm 4.7$	$99.1\% \pm 0.3$	$85.4\% \pm 3.0$	$99.8\% \pm 0.3$	90.4% ± 5.3
2	$\textbf{99.3\%} \pm \textbf{0.9}$	$92.1\% \pm 1.4$	$60.8\% \pm 6.3$	$99.9\% \pm 0.3$	$89.9\% \pm 1.5$	$100.\% \pm 0.0$	$95.1\% \pm 2.4$
3	$96.3\% \pm 2.2$	$83.4\% \pm 2.4$	$45.1\% \pm 2.9$	$99.6\% \pm 0.5$	$88.9\% \pm 3.0$	$99.7\% \pm 0.5$	$90.0\% \pm 2.0$
4	$\textbf{97.3\%} \pm \textbf{1.9}$	$74.7\% \pm 2.0$	$37.4\% \pm 4.7$	$99.1\% \pm 0.7$	$85.2\% \pm 2.5$	$98.6\% \pm 0.9$	$84.8\% \pm 3.9$
5	$96.3\%\pm2.5$	$65.3\% \pm 2.9$	$27.8\% \pm 5.0$	$99.2\% \pm 0.7$	$93.0\%\pm2.2$	$99.7\% \pm 0.4$	$93.3\%\pm2.2$
6	$93.7\% \pm 3.2$	$67.2\% \pm 2.5$	$36.5\% \pm 3.6$	$99.2\% \pm 0.5$	$93.4\% \pm 2.8$	$98.8\% \pm 0.5$	$92.9\%\pm1.8$
7	$88.0\% \pm 4.5$	$71.5\% \pm 2.8$	$40.9\% \pm 4.1$	$98.3\% \pm 0.7$	$96.6\%\pm0.7$	$99.0\% \pm 0.4$	$95.6\%\pm1.4$
8	$\textbf{88.3\%} \pm \textbf{6.0}$	$84.2\% \pm 3.3$	$63.0\% \pm 4.5$	$99.0\% \pm 0.5$	$93.0\%\pm1.8$	$99.6\% \pm 0.4$	$94.0\%\pm2.2$
9	$88.0\%\pm5.8$	$83.5\% \pm 2.7$	$62.4\% \pm 5.0$	$98.4\% \pm 1.2$	$92.9\%\pm1.5$	$98.5\% \pm 1.3$	$92.6\%\pm2.1$
10	$91.0\% \pm 4.0$	$79.8\% \pm 2.7$	$59.0\% \pm 4.6$	$96.9\% \pm 0.6$	$\textbf{92.1\%} \pm \textbf{1.7}$	$97.7\% \pm 1.1$	$89.5\% \pm 4.6$

Table 3 Results of classification of shape data discussed in Sect. 8.4

The function used was the Chebyshev polynomial of the second kind. The MSK column gives the original results from [55]; the subsequent columns use the 0-dimensional diagrams only, the 1-dimensional diagrams only, and both, respectively. Scores highlighted in bold give best average score MSK versus template functions and overlapping intervals with the best score. Compare this to the results with tent functions, Table 4

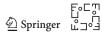
considerably better than the 0-dimensional diagram alone. The average classification rates were improved in four out of the ten parameter choices; results with intersecting confidence intervals occurred in an additional four out of ten parameter choices.

8.5 Rossler Periodicity

We tested our machine learning approach on time series simulated from the Rossler system ([47])

$$\dot{x} = -y - z,
\dot{y} = x + \alpha y,
\dot{z} = \beta + z, (x - \gamma),$$
(18)

where the overdot denotes a derivative with respect to time. We used an explicit Runge-Kutta (4,5) formula to solve the Rossler system for $\beta=2$, $\gamma=4$, and 1201 evenly spaced values of the bifurcation parameter α where $0.37 \le \alpha \le 0.43$. For each value of α a set of initial conditions was sampled from uniformly distributed values in [0, 1]. We simulated 2×10^4 points using a time step of 0.2 seconds. Half of the simulated points were discarded, and only the second half of the x variable data was used in the current analysis. The left and right columns of Fig. 13 show two examples of the resulting time series: one periodic with $\alpha=0.37$, and one chaotic with $\alpha=0.42$, respectively. The first row of the figure shows the time series after dropping the first half of the simulated data, and the bottom row shows the corresponding phase space. The black dots in Fig. 13 represent the extrema of x which were accurately computed using a modified version of Henon's algorithm [38, 52]. These dots are basically the



Poincaré points obtained by finding all the intersections of the x trajectory with the surface $\dot{x} = 0$.

The two examples in Fig. 13 show how the bifurcation parameter α can influence the system behavior. This dependence on α is further illustrated in the top graph of Fig. 15 which depicts the maximum Lyapunov exponent computed using the algorithm described by [6, 32, 58]. The bottom graph of Fig. 15 shows the score of the zero-one test for chaos [35–37]: a binary test that yields a score of 0 for regular dynamics, and 1 for chaotic dynamics. In order to avoid the failure of the 0–1 test due to oversampling, the test was applied to the subsampled data which was obtained by retaining every sixth point from the original signal. The periodic windows shown in Figs. 14 and 15 were identified by examining the plot of the bifurcation diagram in Fig. 14 as well as plots of the maximum Lyapunov exponent and the 0-1 test scores in Fig. 15.

We applied the feature function method on the resulting data set using the tent functions. In this experiment, we set d=10, $\delta=0.4$, and ε to be machine precision. For this test, we got a score of 98.9% on the training set, and 97.2% on the test set. The misclassified time series are superimposed on the bifurcation diagram in Fig. 14. In this figure, green circles show the values of α that were tagged chaotic but that the algorithm identified as periodic. Similarly, purple diamonds indicate the α values that were identified as chaotic even though they were tagged as periodic. It can be seen, unsurprisingly, that misclassification occurs near the transitions of the system behavior from chaotic to periodic or vice versa.

We note that tagging of the data used for both training and testing was performed by inspecting the bifurcation diagram, the maximum Lyapunov exponent plot, and the 0-1 test. Therefore, for the very few misclassified α values we would actually conjecture that our approach can provide a check for the correctness of tagging in the testing and training sets especially for the boundary cases.

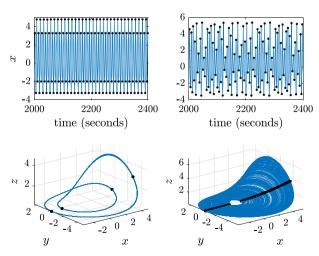


Fig. 13 Time series for the x value in a Rossler system (top row) and the corresponding phase portrait in the (x, y, z) space (bottom row) for a periodic case with $\alpha = 0.37$ (left column), and a chaotic case with $\alpha = 0.42$ (right column). The superimposed black dots correspond to the extrema of x

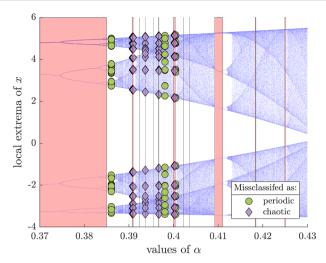


Fig. 14 The bifurcation diagram for the Rossler system with α as the bifurcation parameter and the extrema of x as the response parameter. The shaded windows indicate the regions that were tagged as periodic. The misclassified points are superimposed with diamonds indicating the points that were incorrectly identified as chaotic, while dots indicate that the algorithm incorrectly identified chaotic points as periodic

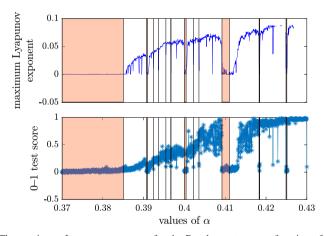


Fig. 15 Top: The maximum Lyapunov exponent for the Rossler system as a function of the bifurcation parameter α . Bottom: The score of the 0–1 test for chaos where 0 indicates periodicity, while 1 indicates chaos. The shaded windows denote the regions that were tagged as periodic

9 Discussion

In this paper, we have provided a new method for the featurization of persistence diagrams through the use of template functions; that is, collections of functions compactly supported on the upper half plane away from the diagonal whose induced functions on diagrams separate points. To do this, we further gave a complete description of compact sets in persistence diagram space endowed with the bottleneck distance.



This method of featurization allows for a great deal of flexibility for the end user. In particular, we have provided two options for template functions, tent functions and interpolating polynomials, but surely there are many other collections of functions which could be tested for optimizing classification and regression tasks. We showed these two functions worked quite well on standard experiments, as well as in comparison to other methods available in the literature.

We find the particular results of the SHREC data set (Sect. 8.4) to be quite fascinating due to the vast improvement seen from tent functions to interpolating polynomials. The usual knee-jerk reaction to setting up these featurization methods for persistence diagrams is that localization is key. This was the impetus for creation of the tent functions as they have support contained in a small box, so each tent function truly only sees a small window of the diagram. Meanwhile, the interpolating polynomials are nonzero well away from their chosen "basepoint" so the fact that these functions work at all is surprising to say the least.

Template function featurization has also found successful applications in the literature. Specifically, in [65] the authors show a comparison between several TDA-based featurization methods such as persistence landscapes, persistence images, and Carlsson coordinates for classifying chatter in metal cutting from vibration signals. That paper shows that template functions yield higher accuracy in comparison to some of the existing featurization methods. Further, template functions (and Carlsson coordinates) were shown to be computationally the fastest methods. Another application of template functions can be found in [66], where the authors show that template functions can be used to detect chatter from simulated, noisy vibration signals in milling—another metal cutting process. In addition to its speed, template functions were shown to yield high classification accuracy and did not require the user to manually identify the classification features, which is the case in traditional chatter classification methods.

Future work is certainly needed to explore the available options that can be utilized within the provided mathematical framework. Because the template system definition is rather broad, we suspect there are many possible collections which could be utilized to improve the results seen here. We have begun to see this in follow-up work which takes a data driven approach to adaptively chose the relevant template functions [54, 63]. We hope to better understand this behavior in future work.

Acknowledgements JAP acknowledges the support of the National Science Foundation (NSF) under grants DMS-1622301, CCF-2006661, CAREER award DMS-1943758, and DARPA under grant HR0011-16-2-003. EM was supported by the NSF through grants CMMI-1800466, DMS-1800446, CCF-1907591, and CCF-2106578. FAK was supported by the NSF through grants CMMI-1759823 and DMS-1759824.

Appendix A: Implementation of the Interpolating Polynomials Algorithm

In this appendix, we give more details on the implementation of the interpolating polynomials described in Sect. 6.2. The barycentric formula for Lagrange interpolation described by [7] is given by



$$f(x) := \sum_{j=0}^{m} \ell_{j}^{\mathcal{A}}(x)c_{j} = \frac{\sum_{j=0}^{m} \frac{w_{j}}{x - \tilde{x}_{j}} c_{j}}{\sum_{j=0}^{k} \frac{w_{j}}{x - \tilde{x}_{j}}}; \text{ where}$$

$$w_{j} = \frac{1}{\ell'(a_{j})}; \quad \ell'(a_{j}) = \prod_{i=0, i \neq j}^{m} (a_{j} - a_{i}), \tag{A1}$$

while $\mathcal{A} = \{a_i\}_{i=0}^m \subset \mathbb{R}$ is a finite set of distinct mesh values, and $\{c_i \in \mathbb{R}\}$ is a collection of evaluation values. The function in Eq. (A1) has the property that $f(a_i) = c_i$ for all i, and it also satisfies the partition of unity condition $\sum_{j=0}^m f(x) = 1$, $\forall x$.

Barycentric Lagrange interpolation is often used for approximating \mathbb{R} -valued functions and there are efficient algorithms for obtaining the weights associated with it. However, in our formulation we need to an interpolating polynomial over an \mathbb{R}^2 -valued function. Therefore, we next describe how to expand the algorithm for interpolating a scalar valued function to interpolating a function on the plane. Note that the notation used here is self-contained from Sect. 6.2.

We assume that our planar mesh is the outer product of m+1 mesh points along the birth time x-axis, and n+1 points along the lifetime y-axis. We also assume that the persistence diagram has N pairs of (birth, lifetime) points.

- 1. Get $\tilde{\gamma}$ and ϕ which correspond to the interpolation matrices along the *x*-mesh and the *y*-mesh, respectively. These are the matrices that describe the linear transformation from the m+1 mesh points of birth times (n+1 mesh of lifetimes) to the corresponding interpolated values of the N query birth times (N query lifetimes) for a given diagram. This step is equivalent to separately obtaining the interpolation matrices for the birth times and the lifetimes.
- 2. Set $\gamma = \tilde{\gamma}^T$.
- 3. (a) Replicate each column in γ n+1 times to obtain Γ whose dimensions are $(m+1)\times (N\times (n+1))$.
 - (b) Unravel ϕ row-wise into a row vector, then replicate each row m+1 times to obtain Φ whose dimensions are $(m+1)\times (N\times (n+1))$.
- 4. Use element-wise multiplication to obtain $\tilde{\Psi} = \Gamma \cdot \Phi$, where \cdot means element-wise multiplication, and $\tilde{\Psi}$ has dimension $(m+1) \times (N \times (n+1))$.
- 5. (a) Split $\tilde{\Psi}$ into N chunks of $(m+1) \times (n+1)$ matrices along the columns axis.
 - (b) Concatenate the split pieces row-wise to obtain an $(N \times (m+1)) \times (n+1)$ matrix Ψ .
- 6. Reshape Ψ by concatenating each $(m+1) \times (n+1)$ piece row-wise to obtain an $N \times ((m+1) \times (n+1))$ matrix Ξ .
- 7. Let the 2D base mesh be given as

$$\begin{bmatrix} f_{00} & f_{01} & \dots & f_{0n} \\ f_{10} & f_{11} & \dots & f_{1n} \\ \vdots & & & \vdots \\ f_{m0} & f_{m1} & \dots & f_{mn} \end{bmatrix},$$

where $f_{ij} = f(x_i, y_j)$ and (x_i, y_j) is a unique point in the 2D mesh. Define the vector $[f_{00} f_{01} \dots f_{mn}]$ which is obtained by unraveling the 2D mesh row-wise.

8. We can interpolate the query points (x_q, y_q) using

$$p(x_q, y_q) = \begin{bmatrix} \ell_0(x_0)\ell_0(y_0) & \dots & \ell_m(x_0)\ell_n(y_0) \\ \ell_0(x_1)\ell_0(y_1) & \dots & \ell_m(x_1)\ell_n(y_1) \\ \vdots & & \vdots \\ \ell_0(x_{N-1})\ell_0(y_{N-1}) & \dots & \ell_m(x_{N-1})\ell_n(y_{N-1}) \end{bmatrix} \begin{bmatrix} f_{00} \\ f_{01} \\ \vdots \\ f_{mn} \end{bmatrix}.$$

Here is a sketch of the resulting matrices:

$$\tilde{\gamma} = \begin{bmatrix} \ell_0(x_0) & \ell_1(x_0) & \dots & \ell_m(x_0) \\ \vdots & & \vdots & & \vdots \\ \ell_0(x_{N-1}) & \ell_1(x_{N-1}) & \dots & \ell_m(x_{N-1}) \end{bmatrix}_{N \times (m+1)},$$

$$\phi = \begin{bmatrix} \ell_0(y_0) & \ell_1(y_0) & \dots & \ell_n(y_0) \\ \vdots & & \vdots & & \vdots \\ \ell_0(y_{N-1}) & \ell_1(y_{N-1}) & \dots & \ell_n(y_{N-1}) \end{bmatrix}_{N \times (n+1)},$$

$$\gamma = \tilde{\gamma}^T = \begin{bmatrix} \ell_0(x_0) & \ell_0(x_1) & \dots & \ell_0(x_{N-1}) \\ \ell_1(x_0) & \ell_1(x_1) & \dots & \ell_1(x_{N-1}) \\ \vdots & & & \vdots \\ \ell_m(x_0) & \ell_m(x_1) & \dots & \ell_m(x_{N-1}) \end{bmatrix}_{(m+1) \times N},$$

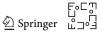
$$\Gamma = \begin{bmatrix} \ell_0(x_0) & \ell_0(x_0) & \dots & \ell_0(x_0) & \dots & \ell_0(x_{N-1}) & \ell_0(x_{N-1}) & \dots & \ell_0(x_{N-1}) \\ \ell_1(x_0) & \ell_1(x_0) & \dots & \ell_1(x_0) & \dots & \ell_1(x_{N-1}) & \ell_1(x_{N-1}) & \dots & \ell_1(x_{N-1}) \\ \vdots & & \vdots & & \vdots & & \vdots \\ \ell_m(x_0) & \ell_m(x_0) & \dots & \ell_m(x_0) & \dots & \ell_m(x_{N-1}) & \ell_m(x_{N-1}) & \dots & \ell_m(x_{N-1}) \end{bmatrix}$$

where Γ has dimension $(m+1) \times (N \times (n+1))$.

$$\Phi = \begin{bmatrix} \ell_0(y_0) \ \ell_1(y_0) \dots \ell_n(y_0) \dots \ell_0(y_{N-1}) \ \ell_1(y_{N-1}) \dots \ell_n(y_{N-1}) \\ \ell_0(y_0) \ \ell_1(y_0) \dots \ell_n(y_0) \dots \ell_0(y_{N-1}) \ \ell_1(y_{N-1}) \dots \ell_n(y_{N-1}) \\ \vdots & \vdots & \vdots & \vdots \\ \ell_0(y_0) \ \ell_1(y_0) \dots \ell_n(y_0) \dots \ell_0(y_{N-1}) \ \ell_1(y_{N-1}) \dots \ell_n(y_{N-1}) \end{bmatrix}$$

where Φ has dimension $(m+1) \times (N \times (n+1))$.

We can now compute the elementwise product $\Psi = \Gamma \cdot \Phi$, which has the dimension $(m+1) \times (N \times (n+1))$.



We then need to apply the following operations: (i) reshaping Ψ to obtain $\hat{\Psi}_1$ given by

$$\hat{\Psi}_{1} = \begin{bmatrix} \ell_{0}(x_{0})\ell_{0}(y_{0}) & \ell_{0}(x_{0})\ell_{1}(y_{0}) & \dots & \ell_{0}(x_{0})\ell_{n}(y_{0}) \\ \ell_{1}(x_{0})\ell_{0}(y_{0}) & \ell_{1}(x_{0})\ell_{1}(y_{0}) & \dots & \ell_{1}(x_{0})\ell_{n}(y_{0}) \\ \vdots & & \vdots & & \vdots \\ \ell_{m}(x_{0})\ell_{0}(y_{0}) & \ell_{m}(x_{0})\ell_{1}(y_{0}) & \dots & \ell_{m}(x_{0})\ell_{n}(y_{0}) \\ \vdots & & \vdots & & \vdots \\ \ell_{0}(x_{N-1})\ell_{0}(y_{N-1}) & \ell_{0}(x_{N-1})\ell_{1}(y_{N-1}) & \dots & \ell_{0}(x_{N-1})\ell_{n}(y_{N-1}) \\ \ell_{1}(x_{N-1})\ell_{0}(y_{N-1}) & \ell_{1}(x_{N-1})\ell_{1}(y_{N-1}) & \dots & \ell_{1}(x_{N-1})\ell_{n}(y_{N-1}) \\ \vdots & & & \vdots & & \vdots \\ \ell_{m}(x_{N-1})\ell_{0}(y_{N-1}) & \ell_{m}(x_{N-1})\ell_{1}(y_{N-1}) & \dots & \ell_{m}(x_{N-1})\ell_{n}(y_{N-1}) \end{bmatrix}.$$

(ii) unraveling $\hat{\Psi}_1$ into an $N \times ((m+1) \times (n+1))$ matrix $\hat{\Psi}_2$ given by

$$\hat{\Psi}_{2} = \begin{bmatrix}
\ell_{0}(x_{0})\ell_{0}(y_{0}) & \dots & \ell_{0}(x_{0})\ell_{n}(y_{0}) & \dots & \ell_{m}(x_{0})\ell_{n}(y_{0}) \\
\vdots & & \vdots & & \vdots \\
\ell_{0}(x_{k})\ell_{0}(y_{k}) & \dots & \ell_{0}(x_{k})\ell_{n}(y_{k}) & \dots & \ell_{m}(x_{k})\ell_{n}(y_{k}) \\
\vdots & & \vdots & & \vdots \\
\ell_{0}(x_{N-1})\ell_{0}(y_{N-1}) & \dots & \ell_{0}(x_{N-1})\ell_{n}(y_{N-1}) & \dots & \ell_{m}(x_{N-1})\ell_{n}(y_{N-1})
\end{bmatrix} .$$
(A2)

The collection of all the scores constitutes the feature vector corresponding to the chosen base mesh point and to the query points where the latter are the persistence diagram points. In this study we summed the rows of $\hat{\Psi}_2$ after taking the absolute value of each entry. The resulting number represents the score at each base mesh point. If the persistence diagram contains the mesh points and we want to find the interpolated values at query points p_{interp} , then we would compute $p_{\text{interp}} = \hat{\Psi}_2 f$.

The implementation of this algorithm can be found in the teaspoon package at teaspoon.ML.feature_functions.interp_polynomial.

Appendix B: Additional Shape Data Results

This appendix gives additional results for the SHREC data set described in Sect. 8.4 using tent functions instead of interpolating polynomials. Table 4 should be compared to the results of Table 3.

freq	MSK	Dim 0		Dim 1		Dim 0 & Dim 1	
		Train	Test	Train	Test	Train	Test
1	94.7% ± 5.1	$8.3\% \pm 0.5$	$3.4\% \pm 1.1$	$8.1\% \pm 0.2$	$3.7\% \pm 0.5$	$8.2\% \pm 0.3$	$3.5\% \pm 0.5$
2	$99.3\%\pm0.9$	$8.3\% \pm 0.3$	$3.4\% \pm 0.7$	$8.2\% \pm 0.5$	$3.5\% \pm 1.1$	$8.56\% \pm 0.4$	$3.0\% \pm 1.0$
3	$96.3\%\pm2.2$	$66.5\% \pm 2.7$	$31.8\% \pm 4.8$	$50.6\% \pm 2.1$	$31.1\% \pm 4.0$	$80.5\% \pm 1.3$	$44.4\% \pm 4.3$
4	$\textbf{97.3\%} \pm \textbf{1.9}$	$46.2\%\pm2.5$	$27.0\% \pm 3.8$	$83.1\% \pm 1.6$	$63.5\% \pm 4.6$	$89.1\% \pm 1.5$	$69.0\% \pm 4.9$
5	$96.3\% \pm 2.5$	$28.5\% \pm 1.4$	$18.9\% \pm 4.0$	$75.2\% \pm 2.6$	$58.3\% \pm 4.6$	$76.8\%\pm2.7$	$58.4\% \pm 7.9$
6	$\textbf{93.7\%} \pm \textbf{3.2}$	$25.4\% \pm 1.8$	$19.0\% \pm 2.4$	$96.5\% \pm 1.1$	$\textbf{88.7\%} \pm \textbf{2.4}$	$96.8\% \pm 0.67$	$89.9\% \pm 1.7$
7	$88.0\% \pm 4.5$	$19.4\% \pm 2.6$	$10.0\% \pm 3.4$	$98.2\% \pm 0.5$	$93.6\%\pm1.9$	$98.3\% \pm 0.6$	$94.1\%\pm2.5$
8	$\textbf{88.3\%} \pm \textbf{6.0}$	$10.8\% \pm 2.6$	$3.6\% \pm 2.4$	$91.9\% \pm 0.9$	$\textbf{88.8\%} \pm \textbf{2.7}$	$91.9\% \pm 1.2$	$89.7\% \pm 3.3$
9	$88.0\% \pm 5.8$	$10.6\% \pm 2.7$	$4.3\% \pm 2.2$	$63.8\% \pm 2.7$	$53.3\% \pm 5.9$	$64.9\% \pm 2.3$	$53.7\% \pm 3.8$
10	$91.0\% \pm 4.0$	$9.2\% \pm 2.3$	$3.6\% \pm 1.7$	$27.0\% \pm 3.9$	$16.2\% \pm 3.2$	$27.3\% \pm 3.4$	$18.6\% \pm 5.6$

Table 4 Results of classification of shape data discussed in Sect. 8.4

The functions used are the tent functions with d=20, and a ridge regression classifier. The MSK column gives the original results from [55]; the subsequent columns use the 0-dimensional diagrams only, the 1-dimensional diagrams only, and both, respectively. Scores highlighted in bold give best average score MSK versus template functions and overlapping intervals with the best score

References

- H. Adams, T. Emerson, M. Kirby, R. Neville, C.Peterson, P.Shipman, S. Chepushtanova, E. Hanson, F. Motta, and L. Ziegelmeier, *Persistence images: A stable vector representation of persistent homology*, Journal of Machine Learning Research 18 (2017), no. 8, 1–35.
- 2. A. Adcock, E. Carlsson, and G. Carlsson, *The ring of algebraic functions on persistence bar codes*, Homology, Homotopy and Applications **18** (2016), no. 1, 381–402.
- 3. R. Anirudh, V. Venkataraman, K. N. Ramamurthy, and P. Turaga, *A Riemannian framework for statistical analysis of topological persistence diagrams*, The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops.
- 4. R. Baire, *Sur les fonctions de variables réelles*, Annali di Matematica Pura ed Applicata (1898-1922) **3** (1899), no. 1, 1–123.
- P. Bendich, J. S. Marron, E. Miller, A. Pieloch, and S. Skwerer, Persistent homology analysis of brain artery trees, The Annals of Applied Statistics 10 (2016), no. 1, 198–218.
- G. Benettin, L. Galgani, A. Giorgilli, and J. M. Strelcyn, Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. part 2: Numerical application, Meccanica 15 (1980), no. 1, 21–30.
- J. Berrut and L. N. Trefethen, Barycentric Lagrange interpolation, SIAM Review 46 (2004), no. 3, 501–517.
- 8. E. Berry, Y. C. Chen, J. Cisewski-Kehe, and B. T. Fasy, *Functional summaries of persistence diagrams*, Journal of Applied and Computational Topology **4** (2020), no. 2, 211–262.
- 9. A. J. Blumberg, I. Gal, M. A. Mandell, and M. Pancia, *Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces*, Foundations of Computational Mathematics **14** (2014), no. 4, 745–789.
- P. Bubenik, Statistical topological data analysis using persistence landscapes, Journal of Machine Learning Research 16 (2015), 77–102.
- 11. P. Bubenik and Alex Elchesen, *Universality of persistence diagrams and the bottleneck and wasserstein distances*, arXiv preprint arXiv:1912.02563 (2019).
- 12. P. Bubenik and T. Vergili, *Topological spaces of persistence modules and their properties*, Journal of Applied and Computational Topology (2018).
- 13. G. Carlsson and V. De Silva, *Zigzag persistence*, Foundations of computational mathematics **10** (2010), no. 4, 367–405.



- 14. G. Carlsson and S. Kalisnik Verovsek, *Symmetric and r-symmetric tropical polynomials and rational functions*, Journal of Pure and Applied Algebra (2016), 3610–3627.
- 15. M. Carrière and U. Bauer, On the metric distortion of embedding persistence diagrams into reproducing kernel hilbert spaces, arXiv:1806.06924 (2018).
- M. Carrière, F. Chazal, Y.i Ike, T. Lacombe, M. Royer, and Y. Umeda, *Perslay: A neural network layer* for persistence diagrams and new graph topological signatures, Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (Palermo, Sicily, Italy) (S. Chiappa and R. Calandra, eds.), Proceedings of Machine Learning Research, vol. 108, PMLR, 2020, pp. 2786– 2796.
- M. Carrière, M. Cuturi, and S. Oudot, *Sliced Wasserstein kernel for persistence diagrams*, Proceedings of the 34th International Conference on Machine Learning (Sydney NSW Australia) (D. Precup and Y. W. Teh, eds.), Proceedings of Machine Learning Research, vol. 70, PMLR, 06–11 Aug 2017, pp. 664–673.
- 18. M. Carrière, S. Oudot, and M. Ovsjanikov, *Stable topological signatures for points on 3d shapes*, Computer Graphics Forum **34** (2015), no. 5, 1–12.
- F. Chazal, V. de Silva, M. Glisse, and S. Oudot, The structure and stability of persistence modules, Springer International Publishing, Switzerland, 2016.
- F. Chazal, B. T. Fasy, F. Lecci, A. Rinaldo, and L. Wasserman, Stochastic convergence of persistence landscapes and silhouettes, Proceedings of the Thirtieth Annual Symposium on Computational Geometry (New York, NY, USA), SOCG'14, ACM, 2014, pp. 474:474

 474:483.
- Y. C. Chen, D. Wang, A. Rinaldo, and L. Wasserman, Statistical analysis of persistence intensity functions, arXiv preprint arXiv:1510.02502 (2015).
- 22. I. Chevyrev, V. Nanda, and H. Oberhauser, *Persistence paths and signature features in topological data analysis*, IEEE transactions on pattern analysis and machine intelligence **42** (2018), no. 1, 192–202.
- M. K. Chung, P. Bubenik, and P. T. Kim, Persistence diagrams of cortical surface data, Information Processing in Medical Imaging (J. L. Prince, D. L. Pham, and K. J. Myers, eds.), Lecture Notes in Computer Science, vol. 5636, Springer Berlin Heidelberg, Williamsburg, VA, USA, 2009, pp. 386–397.
- 24. D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, *Stability of persistence diagrams*, Discrete Comput. Geom. **37** (2007), no. 1, 103–120.
- 25. J. B. Conway, A course in functional analysis, vol. 96, Springer, New York, NY, USA, 2013.
- 26. R. Corbet, U. Fugacci, M. Kerber, C. Landi, and B. Wang, *A kernel for multi-parameter persistent homology*, Computers & graphics: X 2 (2019), 100005.
- 27. W. Crawley-Boevey, *Decomposition of pointwise finite-dimensional persistence modules*, Journal of Algebra and its Applications **14** (2015), no. 05, 1550066.
- B. Di Fabio and M. Ferri, Comparing persistence diagrams through complex vectors, Image Analysis and Processing — ICIAP 2015, Springer International Publishing, Berlin, Heidelberg, 2015, pp. 294– 305
- P. Diaconis, S. Holmes, and M. Shahshahani, Sampling from a manifold, Advances in Modern Statistical Theory and Applications: A Festschrift in honor of Morris L. Eaton, Institute of Mathematical Statistics, 2013, pp. 102–125.
- 30. V. Divol and T. Lacombe, *Understanding the topology and the geometry of the space of persistence diagrams via optimal partial transport*, Journal of Applied and Computational Topology **5** (2021), no. 1, 1–53.
- 31. P. Donatini, P. Frosini, and A. Lovato, *Size functions for signature recognition*, Vision Geometry VII (San Diego, CA, United States) (R. A. Melter, A. Y. Wu, and L. J. Latecki, eds.), SPIE, 1998.
- 32. J. P. Eckmann and D. Ruelle, *Ergodic theory of chaos and strange attractors*, Rev. Mod. Phys. **57** (1985), 617–656.
- 33. B. T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, Sivaraman Balakrishnan, and Aarti Singh, *Confidence sets for persistence diagrams*, Annals of Statistics **42** (2014), no. 6, 2301–2339.
- M. Ferri, P. Frosini, A. Lovato, and C. Zambelli, *Point selection: A new comparison scheme for size functions (with an application to monogram recognition)*, Proceedings of the Third Asian Conference on Computer Vision-Volume I Volume I (Berlin, Heidelberg), ACCV '98, Springer-Verlag, 1998, p. 329–337.
- G. A. Gottwald and I. Melbourne, A new test for chaos in deterministic systems, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 460 (2004), no. 2042, 603–611.



- G. A. Gottwald and I. Melbourne, On the validity of the 0-1 test for chaos, Nonlinearity 22 (2009), no. 6, 1367.
- 37. G. A. Gottwald and I. Melbourne, *The 0-1 test for chaos: A review*, Chaos Detection and Predictability (C. Skokos, G. A. Gottwald, and J. Laskar, eds.), Springer, Berlin, Germany, 2016, pp. 221–247.
- 38. M. Henon, On the numerical computation of Poincaré maps, Physica D: Nonlinear Phenomena 5 (1982), no. 2, 412 414.
- 39. S. Kališnik, *Tropical coordinates on the space of persistence barcodes*, Foundations of Computational Mathematics (2018).
- 40. G. Kusano, K. Fukumizu, and Y. Hiraoka, Kernel method for persistence diagrams via kernel embedding and weight factor, Journal of Machine Learning Research 18 (2018), no. 189, 1–41.
- 41. G. Kusano, Y. Hiraoka, and K. Fukumizu, *Persistence weighted gaussian kernel for topological data analysis*, International Conference on Machine Learning, 2016, pp. 2004–2013.
- 42. G. Kusano, Y. Hiraoka, and K. Fukumizu, *Persistence weighted gaussian kernel for topological data analysis*, ICML, 2016.
- R. Kwitt, S. Huber, M. Niethammer, W. Lin, and U. Bauer, Statistical topological data analysis a kernel perspective, Advances in Neural Information Processing Systems 28 (C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, eds.), Curran Associates, Inc., Montreal, Quebec, Canada, 2015, pp. 3052–3060.
- T. Le and M. Yamada, Persistence Fisher kernel: A Riemannian manifold kernel for persistence diagrams, 32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada., 2018.
- 45. M. Lesnick, *The theory of the interleaving distance on multidimensional persistence modules*, Foundations of Computational Mathematics **15** (2015), no. 3, 613–650 (English).
- 46. C. Li, M. Ovsjanikov, and F. Chazal, *Persistence-based structural recognition*, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 1995–2002.
- 47. M. McCullough, M. Small, T. Stemler, and H. Ho-Ching Iu, *Time lagged ordinal partition networks for capturing dynamics of continuous dynamical systems*, Chaos: An Interdisciplinary Journal of Nonlinear Science **25** (2015), no. 5, 053101.
- 48. Y. Mileyko, S. Mukherjee, and J. Harer, *Probability measures on the space of persistence diagrams*, Inverse Problems **27** (2011), no. 12, 124007.
- 49. E. Munch, K. Turner, P. Bendich, S. Mukherjee, J. Mattingly, and J. Harer, *Probabilistic fréchet means for time varying persistence diagrams*, Electron. J. Statist. 9 (2015), 1173–1204.
- D. Pachauri, C. Hinrichs, M. K. Chung, S. C. Johnson, and V. Singh, *Topology-based kernels with application to inference problems in alzheimer's disease*, IEEE Transactions on Medical Imaging 30 (2011), no. 10, 1760–1770.
- 51. T. Padellini and P. Brutti, *Persistence flamelets: Multiscale persistent homology for kernel density exploration*, arXiv preprint arXiv:1709.07097 (2017).
- 52. P. Palaniyandi, On computing Poincaré map by Hénon method, Chaos, Solitons & Fractals 39 (2009), no. 4, 1877 1882.
- 53. D. Pickup, X. Sun, P. L. Rosin, R. R. Martin, Z. Cheng, Z. Lian, M. Aono, A. Ben Hamza, A. Bronstein, M. Bronstein, S. Bu, U. Castellani, S. Cheng, V. Garro, A. Giachetti, A. Godil, J. Han, H. Johan, L. Lai, B. Li, C. Li, H. Li, R. Litman, X. Liu, Z. Liu, Y. Lu, A. Tatsuma, and J. Ye, SHREC'14 track: Shape retrieval of non-rigid 3d human models, Proceedings of the 7th Eurographics workshop on 3D Object Retrieval, EG 3DOR'14, Eurographics Association, 2014.
- L. Polanco and J. A. Perea, Adaptive template systems: Data-driven feature selection for learning with persistence diagrams, 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), IEEE, 2019, pp. 1115–1121.
- 55. J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt, *A stable multi-scale kernel for topological machine learning*, The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2015.
- 56. D. Rouse, A. Watkins, D. Porter, J. Harer, P. Bendich, N. Strawn, E. Munch, J. DeSena, J. Clarke, J. Gilbert, P. Chin, and Andrew Newman, *Feature-aided multiple hypothesis tracking using topological and statistical behavior classifiers*, Signal Processing, Sensor/Information Fusion, and Target Recognition XXIV (Baltimore, Maryland, United States) (I. Kadar, ed.), SPIE, may 2015.
- 57. W. Rudin, Real and complex analysis, Tata McGraw-Hill Education, New York, NY, USA, 2006.
- M. Sandri, Numerical calculation of lyapunov exponents, The Mathematica Journal 6 (1986), no. 3, 78–84.



- N. Singh, H. D. Couture, J. S. Marron, C. Perou, and M. Niethammer, *Topological descriptors of histology images*, Machine Learning in Medical Imaging: 5th International Workshop, MLMI 2014, Held in Conjunction with MICCAI 2014, Boston, MA, USA, September 14, 2014. Proceedings (Boston, MA, USA) (G. Wu, D. Zhang, and L. Zhou, eds.), Springer International Publishing, 2014, pp. 231–239.
- J. Sun, M. Ovsjanikov, and L. Guibas, A concise and provably informative multi-scale signature based on heat diffusion, Proceedings of the Symposium on Geometry Processing (Aire-la-Ville, Switzerland, Switzerland), SGP '09, Eurographics Association, 2009, pp. 1383–1392.
- L. N. Trefethen, Approximation theory and approximation practice (applied mathematics), SIAM, Philadelphia, PA, USA, 2012.
- 62. K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer, *Fréchet means for distributions of persistence diagrams*, Discrete & Computational Geometry **52** (2014), no. 1, 44–70 (English).
- 63. S. Tymochko, E. Munch, and F. A. Khasawneh, *Adaptive partitioning for template functions on per*sistence diagrams, 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), IEEE, 2019, pp. 1227–1234.
- 64. A. Wagner, Nonembeddability of persistence diagrams with p>2 wasserstein metric, arXiv preprint arXiv:1910.13935 (2019).
- 65. M. C. Yesilli, F. A. Khasawneh, and A. Otto, *Topological feature vectors for chatter detection in turning processes*, The International Journal of Advanced Manufacturing Technology (2022), 1–27.
- M. C. Yesilli, S. Tymochko, F. A. Khasawneh, and E. Munch, hatter diagnosis in milling using supervised learning and topological features vector, 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), IEEE, 2019, pp. 1211–1218.
- Q. Zhao and Y. Wang, Learning metrics for persistence-based summaries and applications for graph classification, Proceedings of the 33rd International Conference on Neural Information Processing Systems (Red Hook, NY, USA), Curran Associates Inc., 2019, pp. 9859–9870.
- 68. X. Zhu, A. Vartanian, M. Bansal, D. Nguyen, and L. Brandl, Stochastic multiresolution persistent homology kernel., IJCAI, 2016, pp. 2449–2457.
- B. Zieliński, M. Lipiński, M. Juda, M. Zeppelzauer, and P. Dłotko, *Persistence bag-of-words for topological data analysis*, Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI'19, AAAI Press, 2019, p. 4489–4495.
- B. Zieliński, M. Lipiński, M. Juda, M. Zeppelzauer, and P. Dłotko, Persistence codebooks for topological data analysis, Artificial Intelligence Review 54 (2021), no. 3, 1969–2009.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

