

Multiparameter Persistence Landscapes

Oliver Vipond

*Mathematical Institute
University of Oxford
Oxford, OX2 7DT, UK*

VIPOND@MATHS.OX.AC.UK

Editor: Sayan Mukherjee

Abstract

An important problem in the field of Topological Data Analysis is defining topological summaries which can be combined with traditional data analytic tools. In recent work Bubenik introduced the persistence landscape, a stable representation of persistence diagrams amenable to statistical analysis and machine learning tools. In this paper we generalise the persistence landscape to multiparameter persistence modules providing a stable representation of the rank invariant. We show that multiparameter landscapes are stable with respect to the interleaving distance and persistence weighted Wasserstein distance, and that the collection of multiparameter landscapes faithfully represents the rank invariant. Finally we provide example calculations and statistical tests to demonstrate a range of potential applications and how one can interpret the landscapes associated to a multiparameter module.

Keywords: Topological Data Analysis, Multiparameter Persistence, Persistence Landscapes, Machine Learning, Statistical Topology

1. Introduction

Topological and Geometric Data Analysis (TGDA) describes an emerging set of analytic tools which leverage the underlying shape data to produce topological summaries. These techniques have been particularly successful at providing new insight for high dimensional data sets, topological data structures and biological data sets (Gameiro et al., 2015; Kelnin and Guo-Wei, 2015; Nicolau et al., 2011). An ideal topological summary should discriminate well between different spaces, be stable to perturbations of the initial data, and be amenable to statistical analysis.

Persistent homology (PH) has become a ubiquitous tool in the TGDA arsenal. PH studies the homology groups of a family of topological spaces built upon data. Homology is an invariant from algebraic topology which detects different dimensional holes in a topological space. Degree zero homology detects connected components, degree one homology detects loops, degree two homology detects voids and higher degree homology detects higher dimensional cavities.

Persistent homology tracks how these homological features persist throughout the family of topological spaces. The associated topological summary for a 1-parameter family is given in terms of a persistence diagram marking the parameter values for births and deaths of homological features. The persistence diagram may equivalently be thought of as a multiset

of points in \mathbb{R}^2 . For an introduction to persistent homology see Edelsbrunner and Harer (2010) and Oudot (2015).

Qualitatively one considers long-lived homological features detected as inherent to the data set and the short-lived features as noise. Nevertheless in various applications it has been revealed that the short-lived features provide important discriminating information in classification (Bubenik et al., 2019). For example in analysing the topology of brain arteries it was found the 28th longest persisting feature provided the most useful discriminating information (Bendich et al., 2016). We would like therefore to have a statistical framework and topological summary in which one can detect statistically significant topological features of large, medium and small persistence, and in particular does not discard short-lived features as noise.

One can equip the collection of persistence diagrams with a natural pseudo-metric known as the bottleneck distance. The resulting metric space of persistence diagrams does not enjoy desired properties for traditional statistical analysis. For example, a collection of persistence diagrams may not have a well defined mean. As a result there have been various attempts to vectorize the persistence diagram, in order that the summary is more amenable to statistical analysis and machine learning techniques.

The article Bubenik (2015) introduces a stable vectorization of the persistence diagram, the persistence landscape, a function in Lebesgue p -space. The article Bubenik (2018) explores some further properties of the persistence landscape. The persistence landscape naturally enjoys unique means and one can perform traditional hypothesis tests upon the landscape function and numerical statistics derived from this function. Another popular stable vectorization of the persistence diagram is the persistence image (Adams et al., 2017) which has been shown to produce favourable classification results when combined with machine learning techniques.

There are several natural situations where we may wish to build a richer structure of topological spaces on a data set and track the changes to the homology whilst varying multiple parameters. For example in (Keller et al., 2018) the topology of chemical compounds is studied using 2-parameter filtrations. The PH theory becomes wildly more complicated when the family of topological spaces on our data set is indexed by multiple parameters. The associated family of homology groups is known as a multiparameter persistence module. The theory of multiparameter persistence modules is presented in Carlsson and Zomorodian (2009) and unlike the single parameter case where we may associate a persistence diagram to a module, there is not an analogous complete discrete invariant in the multiparameter setting.

There exist various approaches to define invariants for multiparameter persistence modules in the literature. The rank invariant of a module has been studied in the context of H_0 -modules and shape matching and has been shown to be stable when endowed with a matching distance (Frosini and Landi, 2011; Cerri et al., 2013). An alternative approach uses algebraic geometry to construct numeric functions on multiparameter persistence modules (Skryzalin and Carlsson, 2017), generalizing the ring of algebraic functions on barcode space for the single parameter case (Adcock et al., 2016). However, as noted by Kališnik (2019), the disadvantage of both of these approaches is that equipped with their vector space norms, these invariants are unstable with respect to the natural distance to consider on multiparameter modules, the interleaving distance. Another approach is to study

algebraic invariants associated to the multigraded algebra structure of multiparameter persistence modules (Harrington et al., 2019; Miller, 2017).

In this article we introduce a family of new stable invariants for multiparameter persistence modules, naturally extending the results of Bubenik (2015) from the setting of single parameter persistence modules to multiparameter persistence modules. We shall consider persistence modules indexed by continuous real parameters. Our incomplete invariants, the *multiparameter persistence landscapes*, are derived from the rank invariant associated to a multiparameter persistence module, and they are continuous functions in Lesbegue p -space. The multiparameter persistence landscape reduces to a family of single parameter persistence landscapes: the multiparameter landscape of a persistence module M indexed over \mathbb{R}^n is characterized by the fact that it restricts to the single parameter landscape function of the restriction of M to every line parallel to the main diagonal vector $\mathbf{1} = (1, 1, \dots, 1)$.

Thus, whilst the single parameter persistence landscape faithfully represents the persistence diagram of a single parameter module (Betthausen et al., 2019), the multiparameter persistence landscape encodes the persistence diagrams of the family of single parameter persistence submodules which lie on the lines of slope $\mathbf{1}$ through the parameter space.

Living in Lesbegue p -space, the multiparameter landscapes are naturally endowed with a distance function and there is a uniquely defined mean associated to multiple landscapes. Considered as Banach space valued random variables, multiparameter landscapes satisfy the Strong Law of Large Numbers and the Central Limit Theorem and are thus well suited to statistical analysis. The natural inner-product structure on the landscape functions gives rise to a positive-definite kernel, which can be leveraged by machine learning algorithms.

The multiparameter landscape functions are sensitive to homological features of large, medium and small persistence. The landscapes also have the advantage of being interpretable since they are closely related to the rank invariant. Moreover one can derive stable \mathbb{R} -valued numeric invariants from the landscape functions using the linear functionals in the dual space. We can produce confidence intervals and perform hypothesis tests on these numeric invariants which are viewed as \mathbb{R} -valued random variables.

In the 2-parameter case we visualise a multiparameter landscapes as a surface $\lambda : \mathbb{R}^2 \rightarrow \mathbb{R}$. We shall present computational examples in the 2-parameter case using the RIVET software presented in Lesnick and Wright (2015). These examples will serve a range of potential applications, demonstrating that the landscapes are sensitive to both the topology and geometry of data.

2. Multiparameter Persistence

The central objects of interest in multiparameter persistence theory are multiparameter persistence modules. These are algebraic objects which can be used to encode rich topological and geometric properties of a data set. The following example gives a construction of a multiparameter persistence module.

Example 1 (*Sublevel-set Multiparameter Persistence Module*) *Let X be a topological space and $f : X \rightarrow \mathbb{R}^n$, called a filtering function. We can associate a family of topological subspaces indexed by vectors $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{R}^n$ induced by f :*

$$X_{\mathbf{a}} = \{x \in X : f(x)_i < a_i \ \forall i = 1, \dots, n\}$$

this is known as the sublevel-set filtration. For any $\mathbf{b} \in \mathbb{R}^n$ such that $a_i \leq b_i$ for all $i = 1, \dots, n$, we have an inclusion map $X_{\mathbf{a}} \hookrightarrow X_{\mathbf{b}}$. If we let H denote a singular homology functor with coefficients in a field then applying this functor to the collection $\{X_{\mathbf{a}}\}_{\mathbf{a} \in \mathbb{R}^n}$ and the appropriate inclusion maps gives rise to a family of vector spaces $\{H(X_{\mathbf{a}})\}_{\mathbf{a} \in \mathbb{R}^n}$ and linear maps $\{H(X_{\mathbf{a}}) \rightarrow H(X_{\mathbf{b}})\}_{\mathbf{a} \leq \mathbf{b}}$ known as a Sublevel-set Multiparameter Persistence Module.

In the case of point cloud data in Euclidean space we can capture the spatial distribution of the points with our filtering function. Let $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^N = X$ denote a point cloud in Euclidean space. We may define our filtering function from this point cloud $f_P : \mathbb{R}^N \rightarrow \mathbb{R}^n$ to have the first parameter be the distance function from the point cloud $f_P(\mathbf{x})_1 = \min_{p \in P} \|x - p\|$. The remaining coordinates can be used to filter our space by other parameters of interest. For example, one could choose the second parameter to be a density function in order to introduce robustness to outliers of the point cloud. Alternatively a second parameter could be chosen to track a numeric property associated to the points of the point cloud such as charge or mass in the case that the points in the point cloud represent atoms (Keller et al., 2018).

The resulting sublevel-set multiparameter persistence module will then encode the topology of the spatial distribution of the point cloud and its interdependence with other chosen filtration parameters. We provide such constructions and example computations in Section 6.

We can give a compact mathematical description of multiparameter persistence modules as objects of functor categories or graded modules over polynomial rings.

Let P_n denote the monoid ring of the monoid $([0, \infty)^n, +)$ over a field \mathbb{F} . Equivalently one may think of P_n as a pseudo-polynomial ring $\mathbb{F}[x_1, \dots, x_n]$ in which exponents are only required to be non-negative and can be non-integral. Let A_n denote the polynomial ring $\mathbb{F}[x_1, \dots, x_n]$ or analogously the monoid ring of $(\mathbb{N}^n, +)$ over \mathbb{F} .

Let \mathbf{P} denote the category associated to the poset P , so that \mathbf{R}^n and \mathbf{Z}^n denote the categories associated to the posets (\mathbb{R}^n, \leq) and (\mathbb{Z}^n, \leq) under the standard coordinate-wise partial orders. Let \mathbf{Vect} denote the category of vector spaces and linear maps over \mathbb{F} , and \mathbf{vect} denote the subcategory of finite dimensional vector spaces. Moreover, for a category \mathbf{C} , let us denote the \mathbf{C} -valued functors on \mathbf{P} by $\mathbf{C}^{\mathbf{P}}$.

Definition 1 (Multiparameter Persistence Module) *Let M be a module over the ring P_n . We say M is a persistence module if M is an \mathbb{R}^n -graded P_n -module. That is to say M has a decomposition as a \mathbb{F} -vector space $M = \bigoplus_{\mathbf{a} \in \mathbb{R}^n} M_{\mathbf{a}}$ compatible with the action of P_n :*

$$m \in M_{\mathbf{a}} \implies \mathbf{x}^{\mathbf{b}} \cdot m \in M_{\mathbf{a}+\mathbf{b}}$$

We require a morphism of \mathbb{R}^n -graded modules $f : M \rightarrow N$ to be compatible with the module structure $f(\mathbf{x}^{\mathbf{b}} \cdot m) = \mathbf{x}^{\mathbf{b}} \cdot f(m)$ and respect the grading of the modules so that $m \in M_{\mathbf{a}}$ implies $f(m) \in N_{\mathbf{a}}$.

In the setting of a sublevel-set persistence module the vector space at each grade is $H(X_{\mathbf{a}})$, and the action of $\mathbf{x}^{\mathbf{b}}$ on $H(X_{\mathbf{a}})$ is given by the linear map of homology groups induced by the inclusion $X_{\mathbf{a}} \hookrightarrow X_{\mathbf{a}+\mathbf{b}}$.

Definition 2 (*Multiparameter Persistence Module*) Let M be an element of the functor category $\mathbf{Vect}^{\mathbf{R}^n}$. We say M is a multiparameter persistence module. A morphism of persistence modules is a natural transformation $M \Rightarrow N$.

Definition 1 and Definition 2 are equivalent as realised by an equivalence of categories between \mathbb{R}^n -graded $P_n\text{-Mod}$ and $\mathbf{Vect}^{\mathbf{R}^n}$ (Lesnick, 2015).

The theory of persistent homology is well developed for topological spaces filtered by a single parameter. Under appropriate finiteness conditions, one can record the homological features of the filtered topological space captured by the single parameter persistence module in a multiset of intervals known as the barcode. Moreover the Krull-Schmidt Decomposition Theorem establishes the barcode as a complete invariant (Crawley-Boevey, 2015). No such discrete complete invariant exists for multiparameter modules. However, associated to a pointwise finite dimensional multiparameter module is a family of single parameter modules whose collection of barcodes is known as the fibered barcode (Lesnick and Wright, 2015). Since we are working with persistence modules induced by real world finite data, the pointwise finite dimensionality assumption we require for the fibered barcode to exist is not a serious restriction. We shall see later that we will be able to reduce the computation of our multiparameter module invariant, the multiparameter persistence landscapes, to queries of the fibered barcode.

Definition 3 (*Fibered Barcode*) (Lesnick and Wright, 2015) Let \mathbf{L} denote the subposet of \mathbf{R}^n corresponding to a positively sloped line $L \subset \mathbb{R}^n$. Let $\iota_L : (\mathbf{R}, \|\cdot\|_\infty) \rightarrow (\mathbf{R}^n, \|\cdot\|_\infty)$ denote the isometric embedding with $\iota_L(\mathbf{R}) = \mathbf{L}$ and $\iota_L(0) \in \{x_n = 0\}$. For $M \in \mathbf{vect}^{\mathbf{R}^n}$ the composite $M^L = M \circ \iota_L$ is a pointwise finite-dimensional single parameter persistence module, and thus has an associated barcode $\mathcal{B}(M^L)$. Let \mathcal{L} denote the set of positively sloped lines. The collection $\{\mathcal{B}(M^L) : L \in \mathcal{L}\}$ is known as the fibered barcode of M .

In view of applications of multiparameter persistence to data analysis we would like a distance function to compare multiparameter persistence modules. The interleaving distance, d_I , is a pseudo-metric on multiparameter persistence modules and provides a notion of algebraic similarity. Details of the interleaving distance can be found in the Appendix A.2. This pseudo-metric has been shown to be the most discriminating, stable metric on multiparameter persistence modules (Lesnick, 2015). A comprehensive mathematical formulation of metrics for generalized persistence modules is presented in Bubenik et al. (2015).

Despite the strong theoretical properties which make the interleaving distance a good candidate to compare multiparameter persistence modules, it has a couple of undesirable properties for data analysis. The interleaving distance has been shown to be NP-hard to compute (Bjerkevik and Botnan, 2018; Bjerkevik et al., 2019) and so is not computable in applications. Moreover the interleaving distance has the characteristics of an L^∞ -style distance, in that it measures the worst place two modules match up over the parameter space. In contrast, an L^p -style distance, for $p \in [1, \infty)$, is sensitive to the difference of two modules over the whole parameter space.

2.1. Interval Decomposable Modules

Given the complicated nature of unconstrained multiparameter persistence modules, it is common to consider subclasses of multiparameter persistence modules. These modules admit a discrete complete description in analogy with single parameter modules. We can use the decomposition of this subclass of modules into simple summands to define matching distances between these modules. Recall that we denote the category of finite dimensional vector spaces as \mathbf{vect} .

Definition 4 (*Interval Decomposable Modules*) *We define a subposet $I \leq \mathbb{R}^n$ to be an interval if $s, t \in I, s \leq r \leq t$ implies $r \in I$ and for any $s, t \in I \exists r_i \in I$ connecting s and $t, s = r_0 \leq r_1 \geq r_2 \leq r_3 \geq \dots \leq r_n = t$. The interval module $\mathbb{1}_I \in \mathbf{vect}^{\mathbb{R}^n}$ associated to an interval I has a one dimensional vector space at each $a \in I$ and internal isomorphisms given by the identity wherever possible. We say a module $M \in \mathbf{vect}^{\mathbb{R}^n}$ is interval decomposable if $M \cong \bigoplus_{j \in \mathcal{J}} \mathbb{1}_{I_j}$, for some indexed set of intervals $\{I_j : j \in \mathcal{J}\}$.*

The Krull-Schmidt-Remak-Azumaya Theorem guarantees that the decomposition of an interval decomposable module is unique up to reordering. We can thus assign the indexed set of intervals in the decomposition of a module M to be the barcode $\mathcal{B}(M) = \{I_j : j \in \mathcal{J}\}$.

Definition 5 (ε -Matching) *Let $\{I_j : j \in \mathcal{J}\}$ and $\{J_k : k \in \mathcal{K}\}$ be indexed sets of intervals. We say a partial bijection $\sigma : \mathcal{J} \rightarrow \mathcal{K}$ is an ε -matching if $d_I(\mathbb{1}_{I_j}, \mathbb{1}_{J_{\sigma(j)}}) \leq \varepsilon$ for matched intervals and $d_I(\mathbb{1}_{I_j}, 0), d_I(0, \mathbb{1}_{J_k}) \leq \varepsilon$ for unmatched intervals.*

Definition 6 (*Bottleneck Distance*) *Let M and N be interval decomposable modules. The bottleneck distance between the modules is given by*

$$d_B(M, N) = \inf\{\varepsilon \geq 0 : \mathcal{B}(M), \mathcal{B}(N) \text{ admit an } \varepsilon\text{-matching}\}.$$

One would hope to attain a result analogous to the isometry theorem for ordinary one dimensional persistent homology relating the bottleneck distance and the interleaving distance. In the single parameter case an ε -interleaving induces an ε -matching between summands (Bauer and Lesnick, 2014). In contrast, the interleaving distance and bottleneck distance do not coincide for multiparameter interval decomposable modules. Certainly the bottleneck distance provides an upper bound on the interleaving distance. However general interleavings of interval decomposable multiparameter modules do not necessarily induce a matching of interval summands. This is best illustrated by an example provided in Bjerkevik (2016) for which the optimal matching between 1-interleaved modules is a 3-matching.

We can further define a Wasserstein distance for interval decomposable modules.

Definition 7 (p -Wasserstein Distance) *Let M, N be interval decomposable persistence modules with barcodes $\{I_j : j \in \mathcal{J}\}$ and $\{J_\kappa : \kappa \in \mathcal{K}\}$ respectively. Append a collection of empty intervals to each barcode of cardinality the size of the other barcode. For a matching $\sigma : \mathcal{J} \rightarrow \mathcal{K}$, let $\varepsilon_j = d_I(\mathbb{1}_{I_j}, \mathbb{1}_{J_{\sigma(j)}})$. The p -Wasserstein distance is given by*

$$d_{W_p}(M, N) = \inf_{\sigma: \mathcal{J} \rightarrow \mathcal{K}} \left[\sum_{\mathcal{J}} \varepsilon_j^p \right]^{\frac{1}{p}}.$$

The bottleneck distance is simply the ∞ -Wasserstein distance. If we wish to place extra emphasis on intervals with large persistence we may use the persistence weighted p -Wasserstein distance (Definition 9). In order to ensure the persistence weighted p -Wasserstein distance is well defined we need to check that intervals are Lebesgue measurable.

Proposition 8 *If $I \subset \mathbb{R}^n$ is an interval of the partially ordered set \mathbb{R}^n , then I is Lebesgue measurable.*

Proof Recall the following characterisation of Lebesgue measurable sets. A set $I \subset \mathbb{R}^n$ is Lebesgue measurable if and only if for all $\varepsilon > 0$ there exists an open set U and a closed set C such that $U \subseteq I \subseteq C$, and the Lebesgue measure $\mu(C \setminus U) < \varepsilon$.

To establish intervals are measurable it suffices to show that for any interval $I \subset \mathbb{R}^n$ there is some open $U \subset \mathbb{R}^n$ and some closed set $C \subset \mathbb{R}^n$ (open and closed under the standard Euclidean topology) such that $U \subseteq I \subseteq C$ and $C \setminus U$ has zero measure. We shall first show that the interior of the closure of an interval $(\bar{I})^\circ$ coincides with the interior of that interval $(\bar{I})^\circ = I^\circ$. If $\mathbf{x} \in (\bar{I})^\circ$, then there is some ∞ -norm ball contained in \bar{I} containing \mathbf{x} in its interior. Thus there exist $\mathbf{a}, \mathbf{b} \in \bar{I}$ with $a_i < x_i < b_i$ for all $i = 1, \dots, n$. The interval I is dense in \bar{I} so we can perturb \mathbf{a} and \mathbf{b} to $\mathbf{a}', \mathbf{b}' \in I$ with $a'_i < x_i < b'_i$. Since I is an interval containing \mathbf{a}', \mathbf{b}' we have that the open set $\{\mathbf{y} \in \mathbb{R}^n : a'_i < y_i < b'_i\}$ is contained in I , and so $\mathbf{x} \in I^\circ$.

The set \bar{I} is closed and so it is Lebesgue measurable and in particular $\bar{I} \setminus (\bar{I})^\circ$ has zero measure. Thus $\bar{I} \setminus I^\circ = \bar{I} \setminus (\bar{I})^\circ$ also has zero measure and we have found the requisite $U = I^\circ$ and $C = \bar{I}$ witnessing that I is measurable. \blacksquare

Definition 9 (*Persistence Weighted p -Wasserstein Distance*) *Let M, N be interval decomposable persistence modules with barcodes $\{I_j : j \in \mathcal{J}\}$ and $\{J_\kappa : \kappa \in \mathcal{K}\}$. For a union of a pair of intervals $I \cup J \subset \mathbb{R}^n$ let $|I \cup J|$ denote the Lebesgue measure. The persistence weighted p -Wasserstein distance is given by*

$$d_{\overline{W}_p}(M, N) = \inf_{\sigma: \mathcal{J} \rightarrow \mathcal{K}} \left[\sum_{\mathcal{J}} |I_j \cup J_{\sigma(j)}| \varepsilon_j^p \right]^{\frac{1}{p}}.$$

The p -landscape distance we introduce in the following section is similar to the persistence weighted p -Wasserstein Distance and can be defined for persistence modules which do not admit an interval decomposition.

In Section 4 we will show that our invariant is stable with respect to interleaving distance and the persistence weighted Wasserstein distance. In particular the L^p vector space norms on multiparameter landscapes provide stable, computable distance functions for multiparameter persistence modules.

3. Persistence Landscapes

In this section we shall recall the definition of the single parameter persistence landscape and its properties. We shall generalize the definition to multiparameter persistence modules and show which properties of the single parameter persistence landscape are preserved.

From this point onward all single parameter persistence modules we consider shall be pointwise finite dimensional in order that they admit an interval decomposition (Crawley-Boevey, 2015). The multiparameter persistence modules we consider will be pointwise finite dimensional, but will not necessarily admit an interval decomposition.

3.1. Single Parameter Persistence Landscapes

The *persistence landscape* associated to a single parameter persistence module is defined in Bubenik (2015). The persistence landscape is derived from the rank invariant of a module.

Definition 10 (*Rank Invariant*) *Let $M \in \mathbf{vect}^{\mathbb{R}}$ be a persistence module then for $a \leq b$ the function $\beta^{\cdot, \cdot}$ giving the corresponding Betti number is the rank invariant of M :*

$$\beta^{a,b} = \dim(\mathrm{Im}(M_a \rightarrow M_b)).$$

Definition 11 (*Rank Function*) *The rank function $\mathrm{rk} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is given by*

$$\mathrm{rk}(b, d) = \begin{cases} \beta^{b,d} & \text{if } b \leq d \\ 0 & \text{otherwise.} \end{cases}$$

Definition 12 (*Rescaled Rank Function*) *The rescaled rank function $r : \mathbb{R}^2 \rightarrow \mathbb{R}$ is supported on the upper half plane:*

$$r(m, h) = \begin{cases} \beta^{m-h, m+h} & \text{if } h \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Observe that the rank function has support contained in the upper triangular half of the plane with the coordinates corresponding to “births” and “deaths”, whilst the rescaled rank function has support contained in the upper half plane with coordinates corresponding to “midpoints” and “half-lives”.

Definition 13 (*Persistence Landscape*) (Bubenik, 2015) *The persistence landscape is a function $\lambda : \mathbb{N} \times \mathbb{R} \rightarrow \overline{\mathbb{R}}$, where $\overline{\mathbb{R}}$ denotes the extended real numbers, $[-\infty, \infty]$. The function $\lambda(k, t) : \mathbb{R} \rightarrow \overline{\mathbb{R}}$ is defined by*

$$\lambda(k, t) = \sup\{h \geq 0 : \beta^{t-h, t+h} \geq k\}.$$

The value $\lambda(k, t)$ gives the maximal radius of an interval centred at t that is contained in at least k intervals of the barcode. The persistence landscape and persistence diagram of a suitably well-behaved single parameter module carry the same information (Betthausen et al., 2019).

Alternatively the persistence landscape of a single parameter module can be derived from the landscape functions of the modules’ interval summands; see Figure 1.




Figure 1: We show the persistence diagram of a single parameter module M on the left and the associated persistence landscapes on the right. One can see that $\lambda(M)(k, t)$ is the k max of the landscape functions $\lambda((b_j, d_j))(1, t)$ of the interval summands of M .

Definition 14 (*Persistence Landscape*) (Bubenik, 2015) Let M be a single parameter persistence module with associated persistence diagram given by the indexed set $\{(b_j, d_j) : j \in \mathcal{J}\}$. The persistence landscape of M may be equivalently defined as:

$$\lambda(k, t) = \text{kmax}_{j \in \mathcal{J}} \{\lambda((b_j, d_j))(1, t)\}$$

where kmax denotes the k^{th} largest value of the indexed set and $\lambda((b_j, d_j))$ is the landscape associated to the interval module $\mathbb{1}_{(b_j, d_j)}$.

Lemma 15 (Bubenik, 2015) The persistence landscape has the following properties:

1. $\lambda(k, t) \geq 0$.
2. $\lambda(k, t) \geq \lambda(k + 1, t)$.
3. $\lambda(k, t)$ is 1-Lipschitz.

The first two properties are immediate from the definition and the third property is proved by Bubenik (2015).

3.2. Multiparameter Persistence Landscapes

Let us define the multiparameter persistence landscape in analogy with the single parameter case. The rank invariant, rank function and rescaled rank function defined above generalise naturally to multiparameter persistence modules:

Definition 16 (*Rank Invariant*) Let $M \in \mathbf{vect}^{\mathbb{R}^n}$ be a multiparameter persistence module, then for $\mathbf{a} \leq \mathbf{b}$ the function $\beta^{\cdot, \cdot}$ giving the corresponding Betti number is the rank invariant of M :

$$\beta^{\mathbf{a}, \mathbf{b}} = \dim(\text{Im}(M_{\mathbf{a}} \rightarrow M_{\mathbf{b}})).$$

Definition 17 (*Multiparameter Rank Function*) The rank function $\text{rk} : \mathbb{R}^{2n} \rightarrow \mathbb{R}$ is given by

$$\text{rk}(\mathbf{b}, \mathbf{d}) = \begin{cases} \beta^{\mathbf{b}, \mathbf{d}} & \text{if } \mathbf{b} \leq \mathbf{d} \\ 0 & \text{otherwise.} \end{cases}$$

Definition 18 (*Rescaled Multiparameter Rank Function*) The rescaled rank function $r : \mathbb{R}^{2n} \rightarrow \mathbb{R}$

$$r(\mathbf{m}, \mathbf{h}) = \begin{cases} \beta^{\mathbf{m}-\mathbf{h}, \mathbf{m}+\mathbf{h}} & \text{if } \mathbf{h} \geq \mathbf{0} \\ 0 & \text{otherwise.} \end{cases}$$

One could perform statistical analysis directly to the rank function and rescaled rank function. Endowed with the L^p normed vector space structure these functions are not stable with respect to the interleaving distance for all $p \in [1, \infty]$.

Example 2 (*Rank Function Instability*) For $\varepsilon > 0$ and $N \in \mathbb{N}$ consider the multiparameter persistence module M with presentation $M = \langle \{(a_i, \mathbf{0})\}_{i=1, \dots, N} \mid \{\mathbf{x}^\varepsilon \cdot a_i\}_{i=1, \dots, N}\rangle$. The persistence module M is such that $\|\text{rk}_M\|_\infty = N, \|\text{rk}_M\|_p = \infty$ for all $p \in [1, \infty)$, and $d_I(M, 0) = \varepsilon$. Details of presentations and the interleaving distance d_I for multiparameter persistence modules can be found in Appendix A.

We wish to define a stable invariant and we derive a landscape function from the rank invariant.

Definition 19 (*Multiparameter Persistence Landscape*) The multiparameter persistence landscape considers the maximal radius over which k features persist in every (positive) direction through \mathbf{x} in the parameter space $\lambda : \mathbb{N} \times \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$.

$$\lambda(k, \mathbf{x}) = \sup\{\varepsilon \geq 0 : \beta^{\mathbf{x}-\mathbf{h}, \mathbf{x}+\mathbf{h}} \geq k \text{ for all } \mathbf{h} \geq \mathbf{0} \text{ with } \|\mathbf{h}\|_\infty \leq \varepsilon\}.$$

It is worth noting that when restricted to a single parameter persistence module this definition coincides with the single parameter persistence landscape (Bubenik, 2015). If there are multiple persistence modules under consideration we shall denote the landscape associated to module M as $\lambda(M)(k, \mathbf{x})$.

Lemma 20 *The multiparameter persistence landscape has the following properties:*

1. $\lambda(k, \mathbf{x}) \geq 0$.
2. $\lambda(k, \mathbf{x}) \geq \lambda(k+1, \mathbf{x})$.
3. $\lambda(k, \mathbf{x})$ is 1-Lipschitz.

Proof The first two properties follow immediately from the definition. Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and let $k \in \mathbb{N}$. Without loss of generality assume that $\lambda(k, \mathbf{x}) \geq \lambda(k, \mathbf{y})$ and that also $r = \lambda(k, \mathbf{x}) \geq \|\mathbf{x} - \mathbf{y}\|_\infty = \delta$. We seek to show that $\lambda(k, \mathbf{y}) \geq \lambda(k, \mathbf{x}) - \|\mathbf{x} - \mathbf{y}\|_\infty$.

For any $\varepsilon \geq \mathbf{0}$ such that $\|\varepsilon\|_\infty \leq r - \delta$ let us define $\mathbf{h} = (|x_i - y_i| + \varepsilon_i)_i$. We observe that $\|\mathbf{h}\|_\infty \leq \|\mathbf{x} - \mathbf{y}\|_\infty + \|\varepsilon\|_\infty \leq r$ and in particular that:

$$\mathbf{x} - \mathbf{h} \leq \mathbf{y} - \varepsilon \leq \mathbf{y} + \varepsilon \leq \mathbf{x} + \mathbf{h}.$$

Thus $\|\mathbf{h}\|_\infty \leq r$ means the map $M(\mathbf{x} - \mathbf{h} \leq \mathbf{x} + \mathbf{h})$ which factors through $M(\mathbf{y} - \varepsilon \leq \mathbf{y} + \varepsilon)$ has rank at least k . Since ε was arbitrary we see that $\lambda(k, \mathbf{y}) \geq r - \delta$. \blacksquare

In extending the persistence landscape definition to multiparameter persistence modules we encountered a choice of p -norm for the ball in \mathbb{R}^n over which we ask features persist. Whilst Lemma 20 holds for all choices of p -norm it transpires that the most natural choice is the ∞ -norm. This choice considerably simplifies computation of the multiparameter persistence landscape and gives rise to a number of further properties which we explore in this section.

Lemma 21 *Let M be a multiparameter persistence module with rank invariant β_M^\cdot . For all $\mathbf{h} \geq \mathbf{0}$ with $\|\mathbf{h}\|_\infty = h$ we have that $\beta_M^{\mathbf{x}-h\mathbf{1}, \mathbf{x}+h\mathbf{1}} \leq \beta_M^{\mathbf{x}-\mathbf{h}, \mathbf{x}+\mathbf{h}}$.*

Proof For all $\mathbf{h} \geq \mathbf{0}$ with $\|\mathbf{h}\|_\infty = h$ we have that $\mathbf{x} - h\mathbf{1} \leq \mathbf{x} - \mathbf{h} \leq \mathbf{x} + \mathbf{h} \leq \mathbf{x} + h\mathbf{1}$. Hence the linear map $M(\mathbf{x} - h\mathbf{1} \leq \mathbf{x} + h\mathbf{1})$ factors through the map $M(\mathbf{x} - \mathbf{h} \leq \mathbf{x} + \mathbf{h})$ and thus the result follows. \blacksquare

In view of Lemma 21, we see that in order to compute the value of a multiparameter persistence landscape $\lambda(k, \mathbf{x})$ at the point $\mathbf{x} = \mathbf{x}_0$, we only need to compute $\sup\{\varepsilon \geq 0 : \beta^{\mathbf{x}_0 - \varepsilon\mathbf{1}, \mathbf{x}_0 + \varepsilon\mathbf{1}} \geq k\}$. Thus we only need to compute a single barcode in the fibered barcode to compute the landscape value at a point.

Proposition 22 *Let M be a persistence module and let \mathbf{L} be a line of slope $\mathbf{1}$ through the parameter space \mathbf{R}^n . Let $\iota_L : (\mathbf{R}, \|\cdot\|_\infty) \rightarrow (\mathbf{R}^n, \|\cdot\|_\infty)$ denote the isometric embedding with $\iota_L(\mathbf{R}) = \mathbf{L}$ and $\iota_L(0) \in \{x_n = 0\}$. The restriction of the multiparameter landscape of M to L and the single parameter persistence landscape of the persistence module M^L coincide, $\lambda(M^L)(k, t) = \lambda(M)(k, \iota_L(t))$.*

Proof Following the landscape definitions we see that if $\lambda(M)(k, \iota_L(t)) > h$ then we have that $\beta_M^{\iota_L(t)-h\mathbf{1}, \iota_L(t)+h\mathbf{1}} = \beta_{M^L}^{t-h, t+h} \geq k$ and so $\lambda(M^L)(k, t) > h$ thus $\lambda(M^L)(k, t) \geq \lambda(M)(k, \iota_L(t))$. Conversely, if $\lambda(M^L)(k, t) > h$ then $\beta_{M^L}^{t-h, t+h} \geq k$, and so by Lemma 21 we have that $\beta_M^{\iota_L(t)-\mathbf{h}, \iota_L(t)+\mathbf{h}} \geq \beta_M^{\iota_L(t)-h\mathbf{1}, \iota_L(t)+h\mathbf{1}} \geq k$ for all $\mathbf{h} \geq \mathbf{0}$ with $\|\mathbf{h}\|_\infty = h$ and thus $\lambda(M^L)(k, t) \leq \lambda(M)(k, \iota_L(t))$. \blacksquare

An immediate consequence of Proposition 22 is that like the single parameter persistence landscape, the multiparameter persistence landscape also admits a decomposition as the kmax of a series of simple landscape functions when our module is interval decomposable.

Proposition 23 *The multiparameter persistence landscape of an interval decomposable module $M \cong \bigoplus_j \mathbb{1}_{I_j}$ can be equivalently defined as:*

$$\lambda(M)(k, \mathbf{x}) = \text{kmax}_j \{\lambda(\mathbb{1}_{I_j})(1, \mathbf{x})\}.$$

Proposition 24 follows from Proposition 22 and the fact that the map from persistence diagram to persistence landscape is invertible for finitely presented single parameter persistence modules (Betthausen et al., 2019). For completeness, we provide a proof of Proposition 24. See Appendix A.1 and Definition 45 for multiparameter persistence module presentations.




Figure 2: We illustrate a pair of interval decomposable multiparameter persistence modules, $M, N \in \mathbf{vect}^{\mathbf{R}^2}$, which have distinct rank invariants but the same multiparameter persistence landscapes. Each module is the direct sum of two rectangular intervals. The first summand of each module is shaded with green dots and the second summand is shaded in solid red.

Proposition 24 *Let $M \in \mathbf{vect}_{fin}^{\mathbf{R}^n}$ be a finitely presented multiparameter persistence module and let $\lambda : \mathbb{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the associated multiparameter persistence landscape. Using λ we can recover the rank invariant of M on the set of pairs of parameter values which lie on lines of slope $\mathbf{1}$, $\{(\mathbf{a}, \mathbf{a} + r\mathbf{1}) : \mathbf{a} \in \mathbb{R}^n, r \geq 0\}$.*

Proof The rank invariant on the set $\{(\mathbf{a}, \mathbf{a} + r\mathbf{1}) : \mathbf{a} \in \mathbb{R}^n, r \geq 0\}$ is derived from the landscape as follows:

$$\beta^{\mathbf{a}, \mathbf{a} + r\mathbf{1}} = \limsup_{\delta \rightarrow 0^+} \max \left\{ k \in \mathbb{N} : \lambda(k, \mathbf{a} + \frac{(r + \delta)}{2}\mathbf{1}) \geq \frac{r + \delta}{2} \right\}.$$

Consider a pair of parameter values $(\mathbf{a}, \mathbf{a} + r\mathbf{1})$. If $\beta^{\mathbf{a}, \mathbf{a} + r\mathbf{1}} \geq k$, since M is finitely presented, $\beta^{\mathbf{a}, \mathbf{a} + (r + \delta)\mathbf{1}} \geq k$ for all sufficiently small δ . Let $\mathbf{c}_\delta = \mathbf{a} + \frac{(r + \delta)}{2}\mathbf{1}$. By Lemma 21 $\beta_M^{\mathbf{c}_\delta - \mathbf{h}, \mathbf{c}_\delta + \mathbf{h}} \geq k$ for all $\mathbf{h} \geq \mathbf{0}$ with $\|\mathbf{h}\|_\infty \leq \frac{(r + \delta)}{2}$, and thus we have that $\lambda(k, \mathbf{c}_\delta) \geq \frac{(r + \delta)}{2}$ for sufficiently small δ .

Conversely, suppose that $\lambda(k, \mathbf{c}_\delta) \geq \frac{(r + \delta)}{2}$ for some small positive δ . This implies that $\beta^{\mathbf{a}, \mathbf{a} + r\mathbf{1}} \geq \beta^{\mathbf{a}, \mathbf{a} + (r + \frac{\delta}{2})\mathbf{1}} \geq k$. ■

The following example illustrates two modules with distinct rank invariants which are not distinguished by their multiparameter persistence landscapes.

Example 3 *Denote by $[\mathbf{a}, \mathbf{b}]$ the rectangular interval module with opposite vertices \mathbf{a}, \mathbf{b} . Let $M = [(0, 1), (10, 2)] \oplus [(4, 1), (6, 2)]$ and $N = [(0, 1), (6, 2)] \oplus [(4, 1), (10, 2)]$ be interval decomposable 2-parameter persistence modules. These two modules have different rank invariant $(\beta_M^{(0,1), (10,2)} = 1 \neq 0 = \beta_N^{(0,1), (10,2)})$ but the same multiparameter persistence landscapes, see Figure 2.*

One would hope that a multiparameter module invariant could distinguish the modules in the previous example. The multiparameter persistence landscape fails to distinguish

these modules since the rank invariants of these modules coincide on all pairs of parameter values lying on lines of slope $\mathbf{1}$. This in turn occurs since the overlap between the summands in the x_1 -coordinate is greater than their significance in the x_2 -coordinate. Thus altering x_1, x_2 at the same rate we cannot detect the interaction between the summands with the multiparameter persistence landscape. A simple reparametrisation scaling parameters x_i appropriately would allow us to distinguish these modules and motivates the following definition.

Definition 25 (*\mathbf{w} -Weighted Persistence Landscape*) Let $\mathbf{w} \in \{\mathbf{u} \in \mathbb{R}^n : u_i > 0, \|\mathbf{u}\|_\infty = 1\}$ be a weighting vector corresponding to a rescaling of the parameter space \mathbb{R}^n . Define the \mathbf{w} -weighted infinity norm to be $\|\mathbf{h}\|_\infty^{\mathbf{w}} = \|(w_i h_i)_i\|_\infty$. The \mathbf{w} -Weighted Persistence Landscape is a function $\lambda_{\mathbf{w}} : \mathbb{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$.

$$\lambda_{\mathbf{w}}(k, \mathbf{x}) = \sup\{\varepsilon \geq 0 : \beta^{\mathbf{x}-\mathbf{h}, \mathbf{x}+\mathbf{h}} \geq k \text{ for all } \mathbf{h} \geq \mathbf{0} \text{ with } \|\mathbf{h}\|_\infty^{\mathbf{w}} \leq \varepsilon\}.$$

Remark 26 The ordinary multiparameter persistence landscape is the $\mathbf{1}$ -weighted persistence landscape.

The example illustrated in Figure 2 exhibits the dependence on the relative scaling of the significant parameter values in a multiparameter persistence module and highlights the importance of normalisation or choice of weighting vectors \mathbf{w} in practical applications. Consider the weighted landscape with weighting $\mathbf{w} = (\frac{1}{10}, 1)$ for the examples in Figure 2. This weighted landscape distinguishes the two modules, indeed $\lambda_{\mathbf{w}}(M)(1, (5, 1.5)) = 0.5$ and $\lambda_{\mathbf{w}}(N)(1, (5, 1.5)) = 0.1$.

The weighting hyperparameter for multiparameter persistence landscapes is analogous to the Poisson parameter used for the Poisson-weighted persistence landscape kernel introduced in Bubenik (2018), and the Weighting Function for persistence images introduced in Adams et al. (2017). Determining an appropriate choice of these hyperparameters is dependent on a given applied setting.

Definition 27 (*\mathbf{w} -Rescaling*) Let $\varphi_{\mathbf{w}} \in \text{Aut}(\mathbb{R}^n)$ denote the invertible rescaling $\varphi_{\mathbf{w}}(\mathbf{x}) = (w_i x_i)_i$ for $\mathbf{w} \in \{\mathbf{u} \in \mathbb{R}^n : u_i > 0\}$.

The following proposition makes precise the relationship between the weighted landscape and rescaling the parameter space.

Proposition 28 Let \mathbf{w} be a rescaling vector and let $\lambda_{\mathbf{w}}$ denote the function taking a module to its \mathbf{w} -weighted persistence landscape. Let $(\varphi_{\mathbf{w}})^*$ denote the pull back of $\varphi_{\mathbf{w}}$. The \mathbf{w} -weighted persistence landscape is given by $\lambda_{\mathbf{w}} = (\text{id} \times \varphi_{\mathbf{w}^{-1}})^* \circ \lambda_{\mathbf{1}} \circ (\varphi_{\mathbf{w}})^*$ and so the following diagram commutes:

$$\begin{array}{ccc} \text{vect}^{\mathbb{R}^n} & \xleftarrow{(\varphi_{\mathbf{w}})^*} & \text{vect}^{\mathbb{R}^n} \\ \downarrow \lambda_{\mathbf{1}} & & \downarrow \lambda_{\mathbf{w}} \\ L^p(\mathbb{N} \times \mathbb{R}^n) & \xrightarrow{(\text{id} \times \varphi_{\mathbf{w}^{-1}})^*} & L^p(\mathbb{N} \times \mathbb{R}^n). \end{array}$$

Proof Let $M \in \mathbf{vect}^{\mathbf{R}^n}$. It is a straight forward definition chase to see the diagram commutes. For convenience of notation let us use shorthand notation for component wise multiplication $\mathbf{a} \odot \mathbf{b} = (a_i b_i)_i$. Direct computation yields the result:

$$\begin{aligned}
\lambda_{\mathbf{w}}(M)(k, \mathbf{x}) &= \sup\{\varepsilon \geq 0 : \beta_M^{\mathbf{x}-\mathbf{h}, \mathbf{x}+\mathbf{h}} \geq k, \text{ for all } \mathbf{h} \geq \mathbf{0} \text{ with } \|\mathbf{h}\|_{\infty}^{\mathbf{w}} \leq \varepsilon\} \\
&= \sup\{\varepsilon \geq 0 : \beta_{M \circ \varphi_{\mathbf{w}}}^{\mathbf{w}^{-1} \odot (\mathbf{x}-\mathbf{h}), \mathbf{w}^{-1} \odot (\mathbf{x}+\mathbf{h})} \geq k, \text{ for all } \mathbf{h} \geq \mathbf{0} \text{ with } \|\mathbf{w}^{-1} \odot \mathbf{h}\|_{\infty}^{\mathbf{w}} \leq \varepsilon\} \\
&= \sup\{\varepsilon \geq 0 : \beta_{M \circ \varphi_{\mathbf{w}}}^{\mathbf{w}^{-1} \odot \mathbf{x}-\mathbf{t}, \mathbf{w}^{-1} \odot \mathbf{x}+\mathbf{t}} \geq k, \text{ for all } \mathbf{t} \geq \mathbf{0} \text{ with } \|\mathbf{t}\|_{\infty} \leq \varepsilon\} \\
&= \lambda_1(M \circ \varphi_{\mathbf{w}})(k, \mathbf{w}^{-1} \odot \mathbf{x}) = (\text{id} \times \varphi_{\mathbf{w}^{-1}})^* \circ \lambda_1 \circ (\varphi_{\mathbf{w}})^*(M)(k, \mathbf{x}). \quad \blacksquare
\end{aligned}$$

4. Stability and Injectivity

In this section we shall show that the multiparameter landscapes are stable with respect to the interleaving distance and persistence-weighted Wasserstein distance. We will then provide an injectivity result that shows the collection of weighted persistence landscapes derived from a persistence module contains almost all the information in the rank invariant of that module.

For each $p \in [1, \infty]$ let us define a p -distance on the space of multiparameter landscapes completely analogously with the definition as in Bubenik (2015) where we implicitly are viewing our landscapes as elements of Lebesgue space $L^p(\mathbb{N} \times \mathbb{R}^n)$. Our landscapes are all measurable since they are continuous; however they may be unbounded. We can either choose to permit infinite distances or alternatively truncate our landscapes to a bounded region if we wish to ensure that our distances are finite.

Definition 29 (*p-Landscape Distance*) Let M, N be multiparameter persistence modules. The p -landscape distance between M and N , $d_{\lambda}^{(p)}(M, N)$, is defined to be:

$$d_{\lambda}^{(p)}(M, N) = \|\lambda(M) - \lambda(N)\|_p.$$

4.1. Stability

Unlike the infinity norm of rank function and rank invariant, the infinity norm of multiparameter persistence landscapes is stable with respect to the interleaving distance.

Theorem 30 (*Multiparameter Persistence Landscape Stability*) If $M, N \in \mathbf{vect}^{\mathbf{R}^n}$ are multiparameter persistence modules then the ∞ -landscape distance of the multiparameter persistence landscapes is bounded by the interleaving distance d_I .

$$d_{\lambda}^{(\infty)}(M, N) \leq d_I(M, N).$$

Proof Suppose M, N are ε -interleaved. Let $\mathbf{x} \in \mathbb{R}^n$ and assume without loss of generality that $r = \lambda(M)(k, \mathbf{x}) \geq \lambda(N)(k, \mathbf{x})$ and that also $\lambda(M)(k, \mathbf{x}) \geq \varepsilon$.

For any $\mathbf{h} \geq \mathbf{0}$ with $\|\mathbf{h}\|_\infty < r - \varepsilon$ we have that $\|\mathbf{h} + \varepsilon\mathbf{1}\|_\infty < r$. Since $r = \lambda(M)(k, \mathbf{x})$ we know that the map $M(\mathbf{x} - (\mathbf{h} + \varepsilon\mathbf{1}) \leq \mathbf{x} + (\mathbf{h} + \varepsilon\mathbf{1}))$ has rank at least k .

An ε -interleaving gives rise to commutative diagram:

$$\begin{array}{ccc} M(\mathbf{x} - (\mathbf{h} + \varepsilon\mathbf{1})) & \xrightarrow{\hspace{10em}} & M(\mathbf{x} + (\mathbf{h} + \varepsilon\mathbf{1})) \\ & \searrow & \nearrow \\ & N(\mathbf{x} - \mathbf{h}) \longrightarrow N(\mathbf{x} + \mathbf{h}) & \end{array}$$

Thus we see that the map $N(\mathbf{x} - \mathbf{h} \leq \mathbf{x} + \mathbf{h})$ has rank at least k . Hence $\lambda(N)(k, x) \geq r - \varepsilon$ and thus the infinity norm distance between $\lambda(M)$ and $\lambda(N)$ is at most ε . \blacksquare

Corollary 31 (*Multiparameter Sublevel-set Landscape Stability Theorem*) *Let $f, g : X \rightarrow \mathbb{R}^n$ be filtering functions and let $M(f), M(g)$ denote the induced sublevel-set multiparameter persistence modules. The sublevel-set persistence modules satisfy:*

$$d_\lambda^{(\infty)}(M(f), M(g)) \leq \|f - g\|_\infty.$$

Proof $M(f), M(g)$ are $\|f - g\|_\infty$ -interleaved. \blacksquare

In practical applications we may wish to truncate our landscapes to a bounded region of the parameter space $R \subset \mathbb{R}^n$. Since $\|h\mathbf{1}_R\|_p \leq |R|\|h\mathbf{1}_R\|_\infty$, the ∞ -landscape distance stability result yields a coarse bound for the p -landscape distance.

Corollary 32 *Let M, N be multiparameter persistence modules and $R \subset \mathbb{N} \times \mathbb{R}^n$ a Lebesgue measurable subset of the parameter space. The p -landscape distance restricted to the region R is stable with respect to the interleaving distance:*

$$\|(\lambda(M) - \lambda(N))\mathbf{1}_R\|_p \leq |R|d_I(M, N).$$

The weighted landscapes also satisfy stability with respect to the interleaving distance. This can be shown directly or using Proposition 28 and the stability result in the unweighted case.

Corollary 33 (*Multiparameter \mathbf{w} -Weighted Landscape Stability*) *Let M, N be multiparameter persistence modules. For unit weightings $\mathbf{w} \in \{\mathbf{u} \in \mathbb{R}^n : u_i > 0, \|\mathbf{u}\|_\infty = 1\}$ the ∞ -Landscape distance of the weighted Multiparameter Persistence Landscapes is bounded by the interleaving distance.*

$$\|\lambda_{\mathbf{w}}(M) - \lambda_{\mathbf{w}}(N)\|_\infty \leq d_I(M, N).$$

Proof Suppose that M, N are ε -interleaved and that \mathbf{w} is a unit weighting. Since \mathbf{w} is a unit weighting $(\varphi_{\mathbf{w}})^*(M), (\varphi_{\mathbf{w}})^*(N)$ are also ε -interleaved. We attain that:

$$\begin{aligned} \|\lambda_{\mathbf{w}}(M) - \lambda_{\mathbf{w}}(N)\|_{\infty} &= \|(\text{id} \times \varphi_{\mathbf{w}^{-1}})^* \circ \lambda_{\mathbf{1}} \circ (\varphi_{\mathbf{w}})^*(M) - (\text{id} \times \varphi_{\mathbf{w}^{-1}})^* \circ \lambda_{\mathbf{1}} \circ (\varphi_{\mathbf{w}})^*(N)\|_{\infty} \\ &= \|\lambda_{\mathbf{1}} \circ (\varphi_{\mathbf{w}})^*(M) - \lambda_{\mathbf{1}} \circ (\varphi_{\mathbf{w}})^*(N)\|_{\infty} \\ &\leq d_I((\varphi_{\mathbf{w}})^*(M), (\varphi_{\mathbf{w}})^*(N)) \leq d_I(M, N). \end{aligned}$$

The first equality is a result of Proposition 28 and the penultimate inequality is a direct application of Theorem 30. \blacksquare

The p -landscape distance restricted to interval decomposable modules is stable with respect to the persistence weighted p -Wasserstein distance.

Proposition 34 (*p -Landscape Distance Stability of Interval Decomposable Modules*) *Let M, N be interval decomposable multiparameter modules with barcodes consisting of finitely many intervals. The p -landscape distance is stable with respect to the persistence weighted p -Wasserstein distance:*

$$d_{\lambda}^{(p)}(M, N) \leq d_{\overline{W}_p}(M, N).$$

Proof Let us use the shorthand notation $\lambda_M = \lambda(M)$, and suppose M, N have barcodes $\{I_j : j \in \mathcal{J}\}$ and $\{J_{\kappa} : \kappa \in \mathcal{K}\}$ with equal cardinality by appending a set of empty intervals of the cardinality of the other set. Recall that the landscape for M can be expressed as a pointwise maximum, $\lambda_M(k, \mathbf{x}) = \text{kmax}_{\mathcal{J}} \lambda_{\mathbb{1}_{I_j}}(1, \mathbf{x})$. Let $\sigma : \mathcal{J} \rightarrow \mathcal{K}$ be any bijection realising the persistence weighted p -Wasserstein distance, and let $\varepsilon_j = d_I(\mathbb{1}_{I_j}, \mathbb{1}_{J_{\sigma(j)}})$.

$$\begin{aligned} d_{\lambda_{\infty}}^{(p)}(M, N)^p &= \|\lambda_M - \lambda_N\|_p^p = \sum_{k=1}^{\infty} \int_{\mathbb{R}^n} |\lambda_M(k, \mathbf{x}) - \lambda_N(k, \mathbf{x})|^p d\mu \\ &= \int_{\mathbb{R}^n} \sum_{k=1}^{\infty} |\text{kmax}_{\mathcal{J}} \lambda_{\mathbb{1}_{I_j}}(1, \mathbf{x}) - \text{kmax}_{\mathcal{K}} \lambda_{\mathbb{1}_{J_{\kappa}}}(1, \mathbf{x})|^p d\mu \\ &\leq \int_{\mathbb{R}^n} \sum_{j \in \mathcal{J}} |\lambda_{\mathbb{1}_{I_j}}(1, \mathbf{x}) - \lambda_{\mathbb{1}_{J_{\sigma(j)}}}(1, \mathbf{x})|^p d\mu \\ &\leq \int_{\mathbb{R}^n} \sum_{j \in \mathcal{J}} \varepsilon_j^p \mathbb{1}_{\{I_j \cup J_{\sigma(j)}\}} d\mu \\ &= \sum_{j \in \mathcal{J}} |I_j \cup J_{\sigma(j)}| \varepsilon_j^p = d_{\overline{W}_p}^p(M, N)^p. \end{aligned}$$

The inequality between the second and third line follows from the general fact that for any $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ the sum $\sum |u_i - v_i|^q$ is minimized by ordering the components of each tuple. The fourth line bounds the third line by Theorem 30 applied to the matched interval summands. \blacksquare

4.2. Injectivity

We now show that the collection of weighted landscapes associated to a module preserves almost all the information contained in the rank invariant of finitely presented persistence modules. See Appendix A.1 and Definition 45 for multiparameter persistence module presentations.

Definition 35 (*Weighted Landscape Space*) Let $\mathbf{W} := \{\mathbf{u} \in \mathbb{R}^n : u_i > 0, \|\mathbf{u}\|_\infty = 1\}$ denote the set of unit weights. Each multiparameter module M gives rise to a function $\lambda_{\mathbf{W}}(M) : \mathbf{W} \rightarrow L^\infty(\mathbb{N} \times \mathbb{R}^n)$ mapping each unit weight to the associated weighted landscape of that module. We define weighted landscape space to be this function space equipped with the metric:

$$d_{\mathbf{W}}(\lambda_{\mathbf{W}}(M), \lambda_{\mathbf{W}}(N)) = \|\lambda_{\mathbf{W}}(M) - \lambda_{\mathbf{W}}(N)\|_\infty = \sup_{\mathbf{w} \in \mathbf{W}} \{\|\lambda_{\mathbf{w}}(M) - \lambda_{\mathbf{w}}(N)\|_\infty\}.$$

Proposition 36 Let $\mathbf{vect}_{fin}^{\mathbb{R}^n}$ denote the space of finitely presented multiparameter persistence modules. Let us define an equivalence relation on $\mathbf{vect}_{fin}^{\mathbb{R}^n}$ identifying $M \sim N$ if the rank invariant of M and N coincide. Denote the quotient space of finitely presented multiparameter persistence modules under this equivalence relation by $\mathbf{vect}_{fin}^{\mathbb{R}^n} / \sim$. The map $\lambda_{\mathbf{W}} : M \mapsto \{(\mathbf{w}, \lambda_{\mathbf{w}}(M)) : w_i > 0, \|\mathbf{w}\|_\infty = 1\}$ is injective and 1-Lipschitz on $\mathbf{vect}_{fin}^{\mathbb{R}^n} / \sim$, where we equip the quotient space with the distance induced by the interleaving distance, and the weighted landscape space with the metric $d_{\mathbf{W}}$.

Proof We shall first show that the map λ is injective. Let $M \in \mathbf{vect}_{fin}^{\mathbb{R}^n}$ and recall Proposition 24. For all $\mathbf{a} < \mathbf{b}$ there is some unit rescaling vector \mathbf{w} such that the pair of parameter values $(\varphi_{\mathbf{w}}(\mathbf{a}), \varphi_{\mathbf{w}}(\mathbf{b}))$ lie on a line of slope $\mathbf{1}$. Thus we can recover the rank invariant of M from the collection $\{(\mathbf{w}, \lambda_{\mathbf{w}}(M))\}$. The fact that the map $\lambda_{\mathbf{W}}$ is 1-Lipschitz is an immediate consequence of Corollary 33 Multiparameter \mathbf{w} -Weighted Landscape Stability. ■

Since $\lambda_{\mathbf{W}}$ is 1-Lipschitz we can compute a lower bound on the interleaving distance between modules from the collection of weighted landscapes. We would be interested to investigate further the relationship between the landscape distance and the interleaving distance to understand when the landscape distance provides a good lower bound for the interleaving distance.

5. Statistics on Multiparameter Landscapes

A principal advantage of working with landscapes as a summary statistic for our data is that we are always able to take the pointwise mean of a collection of landscapes. Whilst the mean of a collection of landscapes is not necessarily the landscape of some module, one can still interpret the mean landscape. The location of local maxima of the mean landscape correspond to the parameter values above which significant features of the sample landscapes live, which in turn correspond to significant topological features of the data set from which the sample landscapes have been derived.

The space of persistence landscapes endowed with the p -landscape distance naturally spans a subspace of Lebesgue space, a Banach space. We would like to perform statistical

analysis on a set landscapes produced from data sets to distinguish significant topological signals from sampling noise. In Appendix B we review relevant results from the theory of Banach Space valued random variables. In this section we apply these results to multiparameter persistence landscapes. We attain the same collection of results enjoyed by the single parameter persistence landscape established in Bubenik (2015).

5.1. Convergence Results for Multiparameter Landscapes

We shall take the same probabilistic approach as in Bubenik (2015) in viewing multiparameter landscapes derived from a data set as a Banach space valued random variable. The model for applying statistical analysis to persistence landscapes will trace the following general setup.

Suppose X is a Borel measurable random variable on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ thought of as sampling data from some distribution. Further let $\Lambda = \Lambda(X)$ denote the multiparameter persistence landscape associated to some multifiltration of the data X , so that in summary $\Lambda : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow L^p(\mathbb{N} \times \mathbb{R}^n)$ for $1 \leq p < \infty$ is a random variable taking values in a real, separable Banach Space.

Let $\{X_i\}$ be i.i.d copies of X and $\{\Lambda_i\}$ their associated landscapes. Denoting the point-wise mean of the first n landscapes by $\bar{\Lambda}^n$ and applying the general theory of probability in Banach spaces presented in Appendix B we attain several results. Observe that in practice we may be required to truncate our multiparameter landscapes to a bounded region in order to satisfy the finiteness criteria in the convergence results.

Associated to a well-behaved Banach space valued random variable $\Lambda : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow L^p(\mathbb{N} \times \mathbb{R}^n)$ is a set function $I_\Lambda : \mathcal{F} \rightarrow \mathcal{B}$ called the Pettis Integral of Λ . This can be thought of as the expectation of a Banach space valued random variable. For more details see Appendix B and Ledoux and Talagrand (2011).

Theorem 37 (*Strong Law of Large Numbers*) *With our notation as in the above discussion $\bar{\Lambda}^n \rightarrow I_\Lambda(\Omega)$ almost surely if and only if $\mathbb{E}[\|\Lambda\|] < \infty$.*

Theorem 38 (*Central Limit Theorem*) *Let us consider the landscapes endowed with the p -landscape distance for $p \geq 2$. If $\mathbb{E}[\|\Lambda\|] < \infty$ and $\mathbb{E}[\|\Lambda^2\|] < \infty$, then $\sqrt{n}(\bar{\Lambda}^n - I_\Lambda(\Omega))$ converges weakly to a Gaussian random variable $G(\Lambda)$ with the same covariance structure as Λ .*

The central limit theorem for the landscapes induces a central limit theorem for associated real valued random variables and facilitates the computation of approximate confidence intervals.

Corollary 39 *Let us consider the landscapes endowed with the p -landscape distance for $p \geq 2$. Suppose $\mathbb{E}[\|\Lambda\|] < \infty$ and $\mathbb{E}[\|\Lambda^2\|] < \infty$. If $f \in L^p(\mathbb{N} \times \mathbb{R}^n)^*$, so that $Y = f(\Lambda)$ is a real valued random variable, then $\sqrt{n}(\bar{Y}^n - \mathbb{E}[Y]) \rightarrow \mathcal{N}(0, \text{Var}(Y))$ converges in distribution.*

Corollary 40 (*Approximate Confidence Intervals*) *Suppose Y is a real-valued random variable attained from a functional applied to the multiparameter landscape Λ satisfying the conditions of Corollary 39. Let $\{Y_i\}_{i=1}^n$ be i.i.d. instances of this random variable and*

$S_n^2 = \frac{1}{n-1} \sum_{i=1}^n (Y_i - \bar{Y}_n)^2$ the sample variance. An approximate $(1 - \alpha)$ confidence interval for $\mathbb{E}[Y]$ is given by: $[\bar{Y}_n - z_{\frac{\alpha}{2}} \frac{S_n}{\sqrt{n}}, \bar{Y}_n + z_{\frac{\alpha}{2}} \frac{S_n}{\sqrt{n}}]$, where $z_{\frac{\alpha}{2}}$ is the $\frac{\alpha}{2}$ critical value for the normal distribution.

In practice, a functional of choice could be given by integrating the landscapes over a subset R of the parameter domain $f_R(\Lambda) = \int_R \Lambda d\mu$. These functionals can be used to establish the significance of homological features in different regions of the parameter space.

We remark that recent work has attained confidence bands for single parameter persistence landscapes (Chazal et al., 2014, 2015). It is shown in Chazal et al. (2014) that the single parameter persistence landscapes satisfy a uniform central limit theorem, and moreover the rate of convergence is computed. It would be interesting to see similar analysis performed in the multiparameter setting.

6. Example Computations and Machine Learning Applications

In this section we shall present example computations of multiparameter persistence landscapes and demonstrate a simple application of machine learning to the persistence landscapes. We use the RIVET software (The RIVET Developers, 2018) for computations of 2-parameter persistence modules presented in Lesnick and Wright (2015). RIVET supports the fast computation of multigraded Betti numbers (see Definition 46) and an interactive visualisation for 2-parameter persistence modules. The software computes a data structure associated to a module which facilitates real time queries of the fibered barcode. As far as we know, RIVET is the only publicly available TDA software package supporting multiparameter persistent homology calculations.

The software supports a range of input formats including: point cloud, metric space, algebraic chain complex, and explicit bifiltered complex. In particular we shall use the software to calculate and query the fibered barcode associated to a module along a selection of one dimensional slices of the parameter space.

In order to reduce computational cost, RIVET approximates multiparameter modules with a discretization. These approximations can be taken to arbitrary accuracy with respect to the interleaving distance, see Appendix A.3. Discretization is not the only approach one can take in computations involving continuous multiparameter modules. In contrast Miller (2017) develops a primary decomposition of modules which facilitates a finite description of a wide class of persistence modules which would require infinitely many generators if discretized. Our only obstruction to using this approach rather than discretization is the lack of available software to cope with these presentations.

Computation of the module with RIVET is the most computationally expensive procedure in our calculations. Details of the time and space complexity of the algorithm may be found in Lesnick and Wright (2015), loosely if m denotes the size of the filtered complex associated to the input data, in the worst case one requires time $O(m^5)$ and space $O(m^5)$ to compute the data structure which admits fast queries of the fibered barcode. Further details of the software and complexity may be found in Lesnick and Wright (2015).

In theory, since our landscape is derived solely from the rank invariant, we need not calculate the full module and fibered barcode. Recall that the value of the multiparameter persistence landscape at each point can be calculated using the single parameter persistence

landscape associated to the line of slope $\mathbf{1}$ passing through that point. Thus we could reduce the computation of the multiparameter landscape in any dimension to repeated single parameter persistent homology calculations. This reduction would be highly parallelizable and likely to provide significant speedup.

Proposition 41 *Let $M \in \mathbf{vect}^{\mathbf{R}^2}$ be a multiparameter persistence module derived from a simplicial complex with m simplices. Let $\varepsilon > 0$ be some tolerance value and $[0, R] \times [0, R] \subset \mathbb{R}^2$ a subset of our parameter space. We can compute an ε -approximation $\lambda(M)^{(\varepsilon)}$ to the persistence landscape $\lambda(M)$ of M on the region $[0, R] \times [0, R]$ in time $O(m^3 \frac{R}{\varepsilon})$. Our approximation is with respect to the infinity norm, that is $\|\lambda(M)^{(\varepsilon)} - \lambda(M)\|_\infty \leq \varepsilon$.*

Proof Divide the region $[0, R] \times [0, R]$ into a grid of spacing ε . It suffices to calculate the values of the landscape on this grid since the landscape functions are 1-Lipschitz and so we can extend the grid values to an ε -approximate function on $[0, R] \times [0, R]$. Thus we reduce our computation to the computation of $\frac{2R}{\varepsilon}$ single parameter landscapes corresponding to the collection of $\frac{2R}{\varepsilon}$ slope $\mathbf{1}$ lines passing through the points of the grid.

Given birth-death pairs, Bubenik and Dłotko (2017) provide an algorithm to compute the persistence landscapes in time $O(m^2)$. It is well known from Edelsbrunner and Harer (2010) that one can produce birth-death pairs from a filtration of size m in time $O(m^3)$. Hence the result follows. \blacksquare

It is possible that the above time estimate for the landscape computation could be improved by using vineyard style updates between the single parameter landscapes (Cohen-Steiner et al., 2006). Moreover it may be that in practical applications, computing the module with RIVET is significantly faster than the worst time bound $O(m^5)$ and thus using the fibered barcode queries will be faster than the computation of a series of single parameter landscapes. Note also that the $\frac{2R}{\varepsilon}$ single parameter landscape calculations are independent and so can be computed in parallel. We postpone comparisons of different computational algorithms, benchmarking, and efficient implementation to follow up work.

One may want to utilise machine learning algorithms with landscape functions as a collection of features for a data set. Recall that if we consider the persistence landscapes associated with the 2-landscape distance then we are naturally in the setting of a Hilbert Space. The inner product on this space is positive definite on the space of persistence landscapes. As such we may use this kernel to learn non-linear relationships in our data and then apply convex optimisation techniques to an SVM.

Another point to note is that we can discretize the landscape to give an n -dimensional array as a summary of our data to which one could apply a convolutional neural network. This transform from landscape to multidimensional array will satisfy stability with respect to the landscape distance. A similar approach is used by Adams et al. (2017) to produce a *persistence image* from a persistence diagram.

We provide three computational examples together with the application of a basic statistical test and standard SVM classifier. Our examples demonstrate that the multiparameter landscape is sensitive to both topology and geometry. We do not claim that the multiparameter landscape is the optimal analytic tool to perform the various tasks in our examples, rather we demonstrate a range of potential applications.

6.1. Concentric Circles

Our first example will look at points sampled from densities concentrated around a pair of concentric circles with radii 1 and 3 respectively. We colour the points from each circle in two distinct ways. Colouring A assigns the large circle colour parameter 0.5 and the small circle colour parameter 1.5. Colouring B assigns the small circle colour parameter 0.5 and the large circle colour parameter 1.5. We examine how the multiparameter landscapes differ depending on the colouring of the circles. For each colouring we perform 30 samples, each sample consisting of 100 points uniformly sampled from each circle Figure 3a.


We produce a filtration on each pointcloud with the Rips filtration in the first parameter and the colour parameter in the second parameter. Thus at parameter value $(r, c) \in \mathbb{R}^2$, we have the space $X_{(r,c)} = \text{VR}(\mathcal{P}_c, r)$ where \mathcal{P}_c denotes the sampled points with colour parameter no more than c . We apply the degree one homology functor H_1 to produce a multiparameter persistence module which detects the loops in our filtered topological space.

We compute the average landscapes of the H_1 -modules for the two different colourings, Figure 3. When the large circle has the smaller colour parameter value, the first landscape ($k = 1$) can detect the large circle Figure 3b. We see the large circle in the first landscape as the large mountain spanning the parameter subspace $[1, 5.4] \times [0.5, 1.5]$. When the large circle has the higher parameter value, the persistence in the Rips filtration parameter is diminished by the presence of the small circle with smaller colour parameter. In both colourings, the second landscape ($k = 2$) exhibits the range of parameter values for which both circles are detected Figure 3c.


We test the robustness of the landscape by repeating the sampling this time with only 50 points per circle and perturbing both the location and colour of the sampled points with the addition of i.i.d. normals $\mathcal{N}(0, 0.3)$, Figure 4a. We illustrate in Figure 4b and Figure 4c the average landscapes taken over 30 noisy samples. The resulting landscapes are similar to those of the larger samples without noise.

Let us perform a statistical test to determine whether the multiparameter landscapes can detect that the noisy samples are drawn from different distributions. Consider the functional $f_R(\lambda) = \int_R \lambda d\mu$. Using the results of Section 5 we find approximate confidence intervals for $f_R(\lambda)$ with $R = \{1\} \times ([2, 6] \times [0, 1.5]) \subset \mathbb{N} \times \mathbb{R}^2$. We attain approximate 99%-confidence intervals on the noisy samples: for Colouring A $[0.400, 0.474]$, and for Colouring B $[0.00556, 0.00809]$. A two sample t -test on the values of this functional on the two sets of colourings attains a p -value of 0.00629. Thus we reject the null hypothesis that the functional values on the landscapes of the two colourings have the same mean.


We also perform a permutation test for the test statistic $\|\bar{\lambda}^A - \bar{\lambda}^B\|_4$ (with norm taken for the first landscape over the whole parameter space $[0, 6] \times [0, 3]$). We apply 10,000 random permutations to the collection of landscapes randomly assigning 30 landscapes to group A and 30 landscapes to group B and then computing the test statistic. Each of the sampled permutations attained a smaller test statistic than the observed test statistic, (indicating an approximate p -value of 0), and identifying that the landscapes for Colouring A have been drawn from a distinct distribution than the landscapes for Colouring B.



(a) An example point cloud sample from each colouring.




(b) The mean first landscape for each colouring taken over the 30 samples, $\bar{\lambda}(1, \mathbf{x})$.




(c) The mean second landscape for each colouring taken over the 30 samples, $\bar{\lambda}(2, \mathbf{x})$.

Figure 3: The first column shows the plots for Colouring A and the second column Colouring B.


MULTIPARAMETER PERSISTENCE LANDSCAPES



(a) An example point cloud sample from each colouring with noise added.



(b) The mean first landscape $\bar{\lambda}(1, \mathbf{x})$ taken over the 30 noisy samples.



(c) The mean second landscape $\bar{\lambda}(2, \mathbf{x})$ taken over the 30 noisy samples.

Figure 4: The first column shows the plots for Colouring A and the second column Colouring B.

6.2. Modal Estimation

For this example we work on meteorite data which we have lifted from Good and Gaskins (1980). The data set consists of values of the proportion of silica measured in 22 samples. Our task is to infer how many modes there are in the distribution from which this data has been sampled.

A standard approach to this task is kernel density estimation (KDE). With data $\{x_i\} \subset \mathbb{R}^n$ one estimates the probability density function (pdf) of the distribution using a sum of normalized kernels:

$$f(x) = \frac{1}{n} \sum_{i=1}^n K_\sigma(x - x_i)$$

Here K_σ is a density function with mass concentrated about the origin, for example a Gaussian centred at the origin. There are two natural parameters in this KDE setup. The *log bandwidth* parameter σ of the kernel function K_σ , and a *threshold* parameter which dictates how large a peak in the estimated distribution must be to be considered a mode. The choice of these parameters will dramatically alter our inferred number of modes (see Figure 5).

Figure 6 is a surface plot of the KDEs ranging over various bandwidth parameters, demonstrating the change in the number of modes as we change the bandwidth. The surface has been triangulated using a triangulation subordinate to a regular grid on our parameter space. To each 2-simplex τ in the triangulation we attach two parameters; the mean bandwidth $\sigma(\tau)$, and the mean probability density value $p(\tau)$, (averages taken over the vertices of the simplex). We produce a bifiltration by taking the simplicial closure of the 2-simplices with appropriate parameter values, $X_{(\sigma_0, p_0)} = \text{SC}(\{\tau : \sigma(\tau) \leq \sigma_0, p(\tau) \geq 1 - p_0\})$.

The multiparameter landscape detects that three modes appear in the KDEs for a range of parameter values. Looking at the landscapes associated to the H_0 -module we see that the infinity norm of the first three landscapes is constant but decreases significantly between the third and fourth landscapes, Figure 7. This indicates that within this setup, three modes are seen across a significantly wider range of parameter values than four modes, suggesting the data is drawn from a tri-modal distribution which coincides with our expected result. Whilst in this simple example one could suggest there are three modes from inspection,




Figure 5: We plot kernel density estimates on the meteorite data (red) for a range of bandwidth parameters. As we increase the bandwidth parameter we yield fewer modes in our kernel density estimate.




Figure 6: A triangulated surface plot of the KDE for a range of bandwidth parameters. We observe three modes in the KDE estimate for a large range of bandwidth values.

the landscape analysis can equally be applied to higher dimensional data sets for which visualisation is not possible.

This basic example can be generalized to detect other properties of KDEs robust to changes in parameter values. For example one could detect significant i -dimensional holes in the distribution by considering the H_i module in a similar setup. For related work see Persistence Terraces (Moon et al., 2018).




Figure 7: The first to fifth multiparameter persistence landscapes associated to the H_0 module for the KDE surface. The infinity norm of the first three landscapes is constant $\|\lambda(1, \mathbf{x})\|_\infty = \|\lambda(2, \mathbf{x})\|_\infty = \|\lambda(3, \mathbf{x})\|_\infty = 0.283$ but drops significantly between the third and fourth landscape $\|\lambda(4, \mathbf{x})\|_\infty = 0.120$, $\|\lambda(5, \mathbf{x})\|_\infty = 0.106$. This indicates that there are three modes in the KDE for a wide range of bandwidth parameters.

6.3. Curvature

In this subsection we shall work with a synthetic data set sampled from spaces of different curvature. Single parameter persistence landscapes have been used to detect curvature (Bubenik et al., 2019). This example is used to emphasise the ability of the multiparameter landscapes to detect geometric differences between point samples. The samples consist of 100 points chosen uniformly with respect to the volume measure from discs of radius 1 in the hyperbolic plane, the surface of the unit sphere and Euclidean space so that the spaces have constant curvature of $-1, 1, 0$ respectively. Topologically these disks are all trivial, our landscapes are detecting geometric differences induced by the distribution of points.

We would like to show that the multiparameter landscape is able to detect the curvature of the space from which a sample is drawn given only the pairwise distances between points. A multifiltered complex is built on the sampled points by filtering the Rips complex with the third nearest neighbour density function ρ on the points. Explicitly, if \mathcal{P} denotes our sampled points and $(r, \rho_0) \in \mathbb{R}^2$ then $X_{(r, \rho_0)} = \text{VR}(\mathcal{P}_{\rho_0}, r)$ where $\mathcal{P}_{\rho_0} = \{p \in \mathcal{P} : \rho(p) \leq \rho_0\}$ for the third nearest neighbour density function ρ . We take 100 samples of 100 points in each space and investigate the resulting multiparameter landscapes for dimension 1 homology.

We plot the average first multiparameter landscapes in Figure 8a and the differences between the average landscapes in Figure 8b. We observe that the persistence of cycles is affected by the curvature of the space. The more negative the curvature the longer the one dimensional cycles persist.


Let us now apply a simple machine learning algorithm to the multiparameter landscapes to see if we can reliably distinguish the curvature of the space from which our small samples have been drawn.

Using the Python package LinearSVC, we train a Support Vector Machine (SVM) with linear kernel on discretizations of the first 10 landscapes for the samples of the hyperbolic discs and elliptic discs, using l^2 penalty and squared hinge loss function. We randomly partition our samples into 160 training samples and 40 test samples and evaluate the accuracy by the proportion of test samples correctly classified. Repeating this process 100 times we attain an average classification score of 85.78%. Thus we see that the multiparameter landscapes are able to reliably detect curvature given a relatively small local sample. It is possible that alternative choices of filtration parameters may be better suited to detecting curvature.

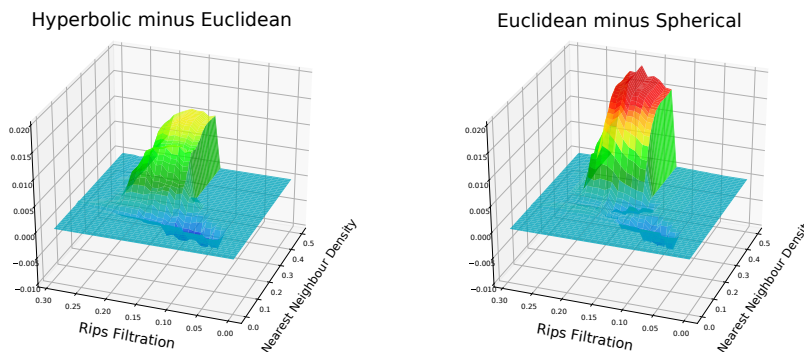
7. Discussion

Multiparameter persistence landscapes provide a stable representation of the rank invariant of a persistence module whilst retaining the discriminating power of the rank invariant. Moreover the landscape distance provides a computable lower bound for the optimal stable distance on persistence modules, the interleaving distance.

The multiparameter landscape also offers a bridge from topological data analysis to machine learning and statistical analysis of multiparameter modules. The multiparameter landscapes, although hard to visualize in dimensions higher than 2, are interpretable in any dimension with large landscape values indicating features robust to changes in the filtration parameters, and non-zero landscapes for large k indicating a large number of homological features.



(a) The mean first landscape $\bar{\lambda}(1, \mathbf{x})$ of the H_1 module for the hyperbolic, Euclidean and elliptic discs taken over 100 samples.



(b) The pointwise difference between the mean landscapes $\bar{\lambda}(1, \mathbf{x})$.

Figure 8: Multiparameter persistence landscapes for uniform samples of discs of unit radius from spaces of different curvature.

An important question to address in practical applications of multiparameter persistence landscapes is how to determine the optimal weighting \mathbf{w} of the parameter space to be used in any given application. The optimal choice of weighting is likely to be highly dependent on the application in mind. The need to choose a hyperparameter is not unique to multiparameter persistence and similar challenges are presented in choosing hyperparameters for various vectorizations of single parameter persistence applications (Bubenik, 2018; Adams et al., 2017; Chazal et al., 2014). We note that multiparameter persistence landscapes depend continuously on the choice of weighting vector \mathbf{w} , and so given sufficient training data one could tune this parameter to maximize discriminatory power.

The multiparameter landscapes highlight several open questions and challenges in the development of the theory and applications of multiparameter persistent homology that we would be interested to see addressed.

1. We would like to understand the relationship between the interleaving distance and landscape distance associated to modules to understand when the landscape distance provides a good lower bound estimate.
2. We have restricted our invariant to the discriminating power of the rank invariant. We would be interested to see if we could combine our landscapes with invariants

that capture the more subtle relationships between features born at incomparable parameter values.

3. With respect to applications there is a question as to whether the increased complexity of multiparameter persistence is necessary, or whether single parameter persistence may be sufficient. One could possibly substitute a multiparameter persistence landscape for a vectorization of a well-chosen single parameter persistence module without significant loss of information.
4. The Bootstrap Method has been used to compute confidence bands for single parameter persistence landscapes (Chazal et al., 2014). We would be interested in applying similar analysis for multiparameter landscapes.

Finally it is worth remarking that the construction of multiparameter persistence landscapes from multiparameter persistence modules can be generalized to produce stable invariants of generalized persistence modules indexed over other posets. Providing the indexing poset P is equipped with a superlinear family of translations Ω , one can derive a landscape function $\lambda : \mathbb{N} \times P \rightarrow \overline{\mathbb{R}}$ from the rank function $\text{rk} : P \times P \rightarrow \mathbb{N}$. This landscape equipped with the supremum norm is stable with respect to the interleaving distance induced by the superlinear family, and provides an interpretable, stable representation of the rank function. This vectorization may prove a useful invariant should the computation of generalized persistence modules be developed in future work.

Acknowledgements

The author would like to give recognition to the *Theory and Foundations of TGDA* workshop hosted at Ohio State University which facilitated useful conversations with experts in TGDA. The author wishes to thank his supervisors Ulrike Tillmann and Vidit Nanda for their guidance and support with this project, and Peter Bubenik for helpful suggestions. The author would like to thank the anonymous reviewers for their detailed feedback on an earlier draft of this paper. The author gratefully acknowledges support from EPSRC studentship EP/N509711/1 and EPSRC grant EP/R018472/1.

Appendix A. Multiparameter Persistence Theory

For completeness we present the multiparameter persistence theory, defining presentations and interleavings of multiparameter persistence modules.

A.1. Presentations

Definition 42 (*Translation Endofunctors*) (Bubenik et al., 2015) Let \mathbf{P} be the category associated to a preordered set (proset) and let $\Gamma : \mathbf{P} \rightarrow \mathbf{P}$ be an endofunctor. We say that Γ is a translation. Since Γ is a functor Γ is monotone $x \leq y$ implies that $\Gamma(x) \leq \Gamma(y)$. We say that Γ is increasing if $x \leq \Gamma(x)$ for all $x \in \mathbf{P}$.

Let \mathbf{Transp} denote the set of increasing translations of \mathbf{P} and observe that \mathbf{Transp} is a monoid with respect to composition.

It is straight forward to see that \mathbf{Transp} also has a natural proset structure with preorder $\Gamma \leq K \Leftrightarrow \Gamma(x) \leq K(x)$ for all x . This preorder is compatible with the monoid structure of \mathbf{Transp} and $\Gamma \leq K$ implies there is a unique natural transformation $\eta_K^\Gamma : \Gamma \Rightarrow K$. If \mathbf{P} is a poset then so is \mathbf{Transp} .

Definition 43 (*Shift Functor*) Let \mathbf{C} be a category, F an element of the functor category $\mathbf{C}^{\mathbf{P}}$, and Γ a translation endofunctor. Let $F(\Gamma)$ denote $F \circ \Gamma \in \mathbf{C}^{\mathbf{P}}$, we call this functor the Γ shift of F .

For a multiparameter persistence module $M \in \mathbf{Vect}^{\mathbf{R}^n}$ we shall write $M(\mathbf{a})$ to denote the shift by the translation in \mathbb{R}^n , $\Gamma_{\mathbf{a}}(\mathbf{x}) = \mathbf{x} + \mathbf{a}$. We define an \mathbb{R}^n -graded set to be some set \mathcal{X} together with a map $\text{gr} : \mathcal{X} \rightarrow \mathbb{R}^n$. For an element $j \in \mathcal{X}$, we shall refer to $\text{gr}(j) = \mathbf{a} \in \mathbb{R}^n$ as the grade of j . Let P_n denote the \mathbb{R}^n -graded monoid ring of the monoid $([0, \infty)^n, +)$ over a field \mathbb{F} .

Definition 44 (*Free Module*) Let \mathcal{X} be an \mathbb{R}^n -graded set. The free module on \mathcal{X} is denoted as $\text{Free}[\mathcal{X}]$, and defined to be:

$$\text{Free}[\mathcal{X}] = \bigoplus_{j \in \mathcal{X}} P_n(-\text{gr}(j)).$$

Note that each element of the graded set gives rise to an independent copy of the \mathbb{R}^n -graded ring P_n , and the negative shift determines that the elements of grade \mathbf{a} in P_n are shifted to be elements of grade $\mathbf{a} + \text{gr}(j)$ in the shifted copy $P_n(-\text{gr}(j))$. The notion of a free module on a graded set can equivalently be defined using a universal property characterisation.

We say a subset $\mathcal{R} \subset M$, of a persistence module is homogeneous if $\mathcal{R} \subset \cup_{\mathbf{a} \in \mathbb{R}^n} M_{\mathbf{a}}$, that is to say each element has a well defined grade.

Definition 45 (*Presentations*) Let \mathcal{X} be a graded set and \mathcal{R} a homogeneous subset of the free module on \mathcal{X} generating the submodule $\langle \mathcal{R} \rangle$. We say that a persistence module M has presentation $\langle \mathcal{X} | \mathcal{R} \rangle$ if:

$$M \cong \frac{\text{Free}[\mathcal{X}]}{\langle \mathcal{R} \rangle}.$$

We say that a presentation is finite if both \mathcal{X} and \mathcal{R} are finite. Let I denote the ideal of P_n generated by the elements $\{\mathbf{x}^{\mathbf{a}} : \mathbf{a} > 0\}$ and let $\Phi_{\langle \mathcal{X} | \mathcal{R} \rangle} : \text{Free}[\mathcal{R}] \rightarrow \text{Free}[\mathcal{X}]$ be the map induced by the inclusion $\mathcal{R} \hookrightarrow \text{Free}[\mathcal{X}]$. We say that a presentation of M is minimal if $\mathcal{R} \subset I \cdot \text{Free}[\mathcal{X}]$ and $\ker \Phi_{\langle \mathcal{X} | \mathcal{R} \rangle} \subset I \cdot \text{Free}[\mathcal{R}]$.

Definition 46 (Multigraded Betti Numbers) Let M be a persistence module. The associated multigraded Betti numbers are maps $\xi_i(M) : \mathbb{R}^n \rightarrow \mathbb{N}$ defined by:

$$\xi_i(M)(\mathbf{a}) = \dim_{\mathbb{F}}(\text{Tor}_i^{P_n}(M, P_n/IP_n)_{\mathbf{a}}).$$

Standard homological algebra arguments establish that the multigraded Betti numbers are well defined (see (Lesnick and Wright, 2015) for details). If $\langle \mathcal{X} | \mathcal{R} \rangle$ is a minimal presentation for M then $\xi_0(M)(\mathbf{a}) = |\text{gr}_{\mathcal{X}}^{-1}(\mathbf{a})|$ and $\xi_1(M)(\mathbf{a}) = |\text{gr}_{\mathcal{R}}^{-1}(\mathbf{a})|$, where $\text{gr}_{\mathcal{X}} : J \rightarrow \mathbb{R}^n$, so that $|\text{gr}_{\mathcal{X}}^{-1}(\mathbf{a})|$ gives the cardinality of the collection of elements of \mathcal{X} with grading \mathbf{a} .

The multigraded Betti numbers are related to the initial topological space with $\xi_0(M)$ marking the filtration values for the birth of homological features, and $\xi_1(M)$ marking the filtration values for relations between features.

A.2. Generalized Interleavings

We adopt the notion of a generalized interleaving from (Bubenik et al., 2015) to define interleaving distances on multiparameter persistence modules.

Definition 47 (Interleaving) (Bubenik et al., 2015) Let \mathbf{P} be a proset and \mathbf{C} be a category. Let $F, G \in \mathbf{C}^{\mathbf{P}}$ be modules and $\Gamma, K \in \mathbf{Transp}$. We say that F, G are (Γ, K) -interleaved if there exist natural transformations $\varphi : F \Rightarrow G\Gamma$, $\psi : G \Rightarrow FK$ satisfying the coherence criteria that $(\psi\Gamma)\varphi = F\eta_{K\Gamma}^{\text{id}}$, $(\varphi K)\psi = G\eta_{\Gamma K}^{\text{id}}$ where $\eta_{\alpha}^{\text{id}}$ denotes the unique natural transformation between the translations $\text{id} \leq \alpha$.

An interleaving may be thought of as an approximate isomorphism. Indeed if we take $\Gamma = K = \text{id}$ then F, G are (Γ, K) -interleaved if and only if F, G are isomorphic. By warping the proset with translations Γ, K we admit flexibility to the rigid notion of isomorphism. In order to introduce an associated distance we must assign a weight to the translations to quantify how close the interleaving is to an isomorphism.

Definition 48 (Superlinear Family) (Bubenik et al., 2015) Let $\Omega : [0, \infty) \rightarrow \mathbf{Transp}$ be a superlinear function: $\Omega_{\varepsilon_1 + \varepsilon_2} \geq \Omega_{\varepsilon_1} \Omega_{\varepsilon_2}$. We say that Ω is a superlinear family.

Definition 49 (ε -Interleaving) (Bubenik et al., 2015) Let $F, G \in \mathbf{C}^{\mathbf{P}}$ be modules and Ω a superlinear family. We say F and G are ε -interleaved with respect to Ω if they are $(\Omega_{\varepsilon}, \Omega_{\varepsilon})$ -interleaved.

Proposition 50 (Induced Interleaving Distance) (Bubenik et al., 2015) Given a superlinear family Ω and modules $F, G \in \mathbf{C}^{\mathbf{P}}$, we have an induced interleaving distance given by:

$$d^{\Omega}(F, G) = \inf\{\varepsilon \geq 0 : F, G \text{ are } \varepsilon\text{-interleaved with respect to } \Omega\}.$$




Figure 9: A pair of interval decomposable modules M, N with interleaving distance $d_I(M, N) = 1$.

The most common superlinear family to consider for persistence modules indexed by \mathbb{R}^n is given by the translation in the diagonal direction, $\Omega_\varepsilon(\mathbf{x}) = \mathbf{x} + \varepsilon \mathbf{1}$. We refer to the interleaving distance induced by this superlinear family as simply *the interleaving distance* and denote this distance by d_I .

Example 4 Let $M, N \in \mathbf{vect}^{\mathbb{R}^2}$ be interval decomposable modules with $M = [(1, 1), (4, 4)] \oplus [(2, 0), (4, 2)]$ and $N = [(2, 2), (3, 3)]$. There is a $1 + \delta$ -interleaving between M and N for all $\delta > 0$ obtained by matching the first summand of M with the summand of N . Explicitly, the natural transformation $\varphi : M \rightarrow N\Omega_{1+\delta}$ is zero everywhere except at grades $\mathbf{a} \in \mathbb{R}^2$ such that \mathbf{a} lies in the support of the first summand of M and $\Omega_{1+\delta}(\mathbf{a})$ lies in the support of N , for which it is an isomorphism. The natural transformation $\psi : N \rightarrow M\Omega_{1+\delta}$ is defined similarly. There is no 1-interleaving between this pair of modules. For example, the morphism $M_{(2,0)} \rightarrow M_{(4,2)}$ is non-zero and so cannot factor through $N_{\Omega_1(2,0)} = N_{(3,1)} = 0$. Hence the interleaving distance between these modules is $d_I(M, N) = 1$.

A.3. Discretization and Continuous Extension

In Section 6 our computations were simplified by restricting a continuous module to a finite grid and then dealing with the continuous extension of this discretization. We will show that restricting to a finite grid gives us a suitable approximation to our module with respect to the interleaving distance between modules.

Definition 51 (*Grid Function*) Let $\mathcal{G} : \mathbb{Z}^n \rightarrow \mathbb{R}^n$ be defined by component-wise strictly increasing functions $\mathcal{G}_i : \mathbb{Z} \rightarrow \mathbb{R}$ with $\sup \mathcal{G}_i = \sup -\mathcal{G}_i = \infty$. We say \mathcal{G} is a grid function. Let us define the size of \mathcal{G} to be

$$|\mathcal{G}| = \max_{i:1 \leq i \leq n} \sup_{z \in \mathbb{Z}} |\mathcal{G}_i(z) - \mathcal{G}_i(z+1)|.$$

Definition 52 (*Discretization*) Let $M \in \mathbf{Vect}^{\mathbb{R}^n}$ be a persistence module and \mathcal{G} a grid function. We say that $M \circ \mathcal{G} \in \mathbf{Vect}^{\mathbb{Z}^n}$ is the \mathcal{G} -discretization of M .

Definition 53 (*Continuous Extension*)(Lesnick and Wright, 2015) Let $Q \in \mathbf{Vect}^{\mathbf{Z}^n}$ be a discrete persistence module and \mathcal{G} a grid function. For $\mathbf{x} \in \mathbb{R}^n$ let us define floor and ceiling functions:

$$\lfloor \mathbf{x} \rfloor_{\mathcal{G}} = \max\{\mathbf{z} \in \text{Im } \mathcal{G} : \mathbf{z} \leq \mathbf{x}\} \text{ and } \lceil \mathbf{x} \rceil_{\mathcal{G}} = \min\{\mathbf{z} \in \text{Im } \mathcal{G} : \mathbf{z} \geq \mathbf{x}\}.$$

We define the continuous extension $E_{\mathcal{G}}(Q) \in \mathbf{Vect}^{\mathbf{R}^n}$ to be the persistence module with:

$$E_{\mathcal{G}}(Q)_{\mathbf{a}} = Q_{\mathcal{G}^{-1}(\lfloor \mathbf{a} \rfloor_{\mathcal{G}})} \text{ and } E_{\mathcal{G}}(Q)(\mathbf{a} \leq \mathbf{b}) = Q(\mathcal{G}^{-1}(\lfloor \mathbf{a} \rfloor_{\mathcal{G}}) \leq \mathcal{G}^{-1}(\lfloor \mathbf{b} \rfloor_{\mathcal{G}}))$$

with the obvious action on the morphisms of $\mathbf{Vect}^{\mathbf{Z}^n}$. We have defined a functor $E_{\mathcal{G}} : \mathbf{Vect}^{\mathbf{Z}^n} \rightarrow \mathbf{Vect}^{\mathbf{R}^n}$.

The following proposition shows that discretization is stable with respect to the interleaving distance, and so we may produce an arbitrarily close approximation to a persistence module by restricting the module to a grid of sufficiently small size.

Proposition 54 Let $M \in \mathbf{Vect}^{\mathbf{R}^n}$ be a persistence module and \mathcal{G} a grid function, we have that:

$$d_I(M, E_{\mathcal{G}}(M \circ \mathcal{G})) \leq |\mathcal{G}|.$$

Proof The modules $M, E_{\mathcal{G}}(M \circ \mathcal{G})$ are $(\lceil \cdot \rceil_{\mathcal{G}}, \text{id})$ -interleaved with natural transformations given by the appropriate internal morphisms of M . ■

Appendix B. Probability in Banach Spaces

We shall present the general theory of Probability for Banach space valued random variables from which we derive the convergence results of Section 5. Let us begin by defining some notation. Let $(\mathcal{B}, \|\cdot\|)$ denote a real, separable Banach Space with topological dual space \mathcal{B}^* . Let $V : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathcal{B}$ denote a Borel measurable random variable. The covariance structure of such a random variable is defined to be the set of expectations

$$\{\mathbb{E}[(f(V) - \mathbb{E}[f(V)])(g(V) - \mathbb{E}[g(V)])] : f, g \in \mathcal{B}^*\}.$$

In order to take expected values of Banach valued random variables we require the notion of the Pettis Integral, which is an extension of the Lebesgue integral to functions on measure spaces taking values in normed spaces. We shall briefly introduce the properties of this integral and existence criteria, the essence of which is built upon reducing the problem to integrability of \mathbb{R} -valued functions.

Definition 55 (*Scalarly Integrable*)(Geitz, 1981) A function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is scalarly integrable if for all $f \in \mathcal{B}^*$ we have that $f(V)$ is measurable and $f(V) \in L^1(\mu)$.

Definition 56 (*Pettis Integrable*)(Geitz, 1981) A scalarly integrable function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is Pettis integrable if for all $E \in \mathcal{F}$ there is an element $I_V(E) \in \mathcal{B}$ such that:

$$\int_E f(V) d\mu = f(I_V(E)) \text{ for all } f \in \mathcal{B}^*.$$

The set function $I_V : \mathcal{F} \rightarrow \mathcal{B}$ is called the Pettis Integral of V with respect to μ . We may also refer to $I_V(\Omega)$ as the Pettis Integral of V and denote this by I_V .

Definition 57 (*Strong and Weak Measurability*)(Musiał, 1991) A function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is simple if there exist $v_1, \dots, v_m \in \mathcal{B}$ and $E_1, \dots, E_m \in \mathcal{F}$ such that $f = \sum_{i=1}^m v_i \mathbb{1}_{E_i}$.

A function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is strongly μ -measurable if there exists a sequence of simple functions $V_n : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ such that $\lim_{n \rightarrow \infty} \|V_n(\omega) - V(\omega)\| = 0$ μ -almost everywhere.

A function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is weakly μ -measurable if for all functionals in the dual space $v^* \in \mathcal{B}^*$ the real-valued function $v^*V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathbb{R}$ is μ -measurable.

We shall suppress the prefix μ if the measure is clear from context.

Theorem 58 (*Pettis Measurability Theorem*)(Musiał, 1991)(Theorem 3.1) A function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ is strongly measurable if and only if:

1. V is weakly measurable.
2. There exists a μ -null set E such that $V(\Omega \setminus E)$ is a separable subset of \mathcal{B} .

Proposition 59 (*Musiał, 1991*)(Proposition 5.1) A strongly measurable \mathcal{B} -valued function $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ with $\mathbb{E}^\mu[\|V\|] < \infty$ is Pettis Integrable.

A consequence of the Pettis Measurability Theorem is that if the codomain is a separable Banach Space then the notions of weak and strong measurability coincide. Thus Proposition 59 gives a sufficient criterion for Pettis Integrability in the setting of multiparameter persistence landscapes endowed with the p -norm for $p \in [1, \infty)$ for which the underlying Banach space is separable.

Corollary 60 Let $V : (\Omega, \mathcal{F}, \mu) \rightarrow \mathcal{B}$ be measurable with \mathcal{B} real and separable. If $\mathbb{E}^\mu[\|V\|] < \infty$, then V is Pettis Integrable and $\|I_V(\Omega)\| \leq \mathbb{E}^\mu[\|V\|]$.

Proof By assumption V is measurable. Moreover, $\mathbb{E}^\mu[\|V\|] < \infty$ implies that for any linear functional $v^* \in \mathcal{B}^*$ we have $\mathbb{E}^\mu[\|v^*V\|] \leq \|v^*\| \mathbb{E}^\mu[\|V\|] < \infty$ and so V is scalarly integrable. Hence by Proposition 59 V is Pettis Integrable. A straight forward application of the Hahn-Banach theorem establishes that $\|I_V(\Omega)\| \leq \mathbb{E}^\mu[\|V\|]$. \blacksquare

Theorem 61 (*Strong Law of Large Numbers*)(Ledoux and Talagrand, 2011)(Collary 7.10) Let V_i be i.i.d copies of $V : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathcal{B}$ and let $S_n = \sum_{i=1}^n V_i$. We have that $\mathbb{E}[\|V\|] < \infty$ if and only if:

$$\frac{S_n}{n} \rightarrow I_V(\Omega) \text{ almost surely as } n \rightarrow \infty.$$

Definition 62 We say a \mathcal{B} -valued random variable X is Gaussian if for all $f \in \mathcal{B}^*$ the real valued random variable $f(X)$ is Gaussian with mean zero. Note that such a Gaussian random variable is determined by its covariance structure (Ledoux and Talagrand, 2011).

The next result only applies for a certain class of Banach spaces. The type and cotype of a Banach space can be thought loosely of as a measure of how close that Banach space is to a Hilbert space, see (Ledoux and Talagrand, 2011) for more details. For $p \in [1, 2]$ the Lebesgue space L^p has type p and cotype 2, and for $p \in [2, \infty)$ the Lebesgue space L^p has type 2 and cotype p .

Theorem 63 (*Central Limit Theorem*)(Hoffmann-Jørgensen and Pisier, 1976) Let \mathcal{B} be a Banach space of type 2 and $V : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathcal{B}$. If $I_V = 0$ and $\mathbb{E}[\|V\|^2] < \infty$ then $\frac{1}{\sqrt{n}}S_n$ converges weakly to a Gaussian random variable with the same covariance structure as V .

References

- Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017. URL <http://jmlr.org/papers/v18/16-337.html>.
- Aaron Adcock, Erik Carlsson, and Gunnar Carlsson. The ring of algebraic functions on persistence bar codes. *Homology, Homotopy and Applications*, 18(1):381–402, 2016. doi: 10.4310/hha.2016.v18.n1.a21. URL <https://doi.org/10.4310/hha.2016.v18.n1.a21>.
- Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry (SoCG 2014)*, pages 355–364. ACM, 2014. ISBN 978-1-4503-2594-3. doi: 10.1145/2582112.2582168. URL <http://doi.acm.org/10.1145/2582112.2582168>.
- Paul Bendich, J. S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *Annals of Applied Statistics*, 10(1):198–218, 2016. ISSN 1932-6157. doi: 10.1214/15-AOAS886. URL <https://uncch.pure.elsevier.com/en/publications/persistent-homology-analysis-of-brain-artery-trees>.
- Leo Betthausen, Peter Bubenik, and Parker B. Edwards. Graded persistence diagrams and persistence landscapes. *arXiv e-prints*, art. arXiv:1904.12807, 2019.
- Håvard Bakke Bjerkevik and Magnus Bakke Botnan. Computational complexity of the interleaving distance. In *34th International Symposium on Computational Geometry (SoCG 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 1–15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPIcs.SoCG.2018.13.
- Håvard Bakke Bjerkevik. Stability of higher-dimensional interval decomposable persistence modules. *arXiv e-prints*, abs/1609.02086, 2016. URL <http://arxiv.org/abs/1609.02086>.
- Håvard Bakke Bjerkevik, Magnus Bakke Botnan, and Michael Kerber. Computing the interleaving distance is NP-hard. *Foundations of Computational Mathematics*, 2019. ISSN 1615-3383. doi: 10.1007/s10208-019-09442-y. URL <https://doi.org/10.1007/s10208-019-09442-y>.
- Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research*, 16(1):77–102, 2015. ISSN 1532-4435. URL <http://jmlr.org/papers/v16/bubenik15a.html>.
- Peter Bubenik. The persistence landscape and some of its properties. *arXiv e-prints*, art. arXiv:1810.04963, 2018.
- Peter Bubenik and Paweł Dłotko. A persistence landscapes toolbox for topological statistics. *Journal of Symbolic Computation*, 78:91–114, 2017. ISSN 0747-7171. doi: <https://doi.org/10.1016/j.jsc.2016.03.009>. URL <http://www.sciencedirect.com/science/article/pii/S0747717116300104>.

- Peter Bubenik, Vin de Silva, and Jonathan Scott. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, 15(6):1501–1531, 2015. ISSN 16153383. doi: 10.1007/s10208-014-9229-5.
- Peter Bubenik, Michael Hull, Dhruv Patel, and Benjamin Whittle. Persistent homology detects curvature. *Inverse Problems*, 2019. URL <http://iopscience.iop.org/10.1088/1361-6420/ab4ac0>.
- Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete and Computational Geometry*, 42(1):71–93, 2009. ISSN 01795376. doi: 10.1007/s00454-009-9176-0.
- Andrea Cerri, Barbara Di Fabio, Massimo Ferri, Patrizio Frosini, and Claudia Landi. Betti numbers in multidimensional persistent homology are stable functions. *Mathematical Methods in the Applied Sciences*, 36(12):1543–1557, 2013. ISSN 1099-1476. doi: 10.1002/mma.2704. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mma.2704>.
- Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, and Larry A. Wasserman. Stochastic convergence of persistence landscapes and silhouettes. *Journal of Computational Geometry*, 6:140–161, 2014.
- Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Aarti Singh, and Larry A. Wasserman. On the bootstrap for persistence diagrams and landscapes. *Modeling and Analysis of Information Systems*, 20(6):111–120, 2015. ISSN 2313-5417, 1818-1015. doi: 10.18255/1818-1015-2013-6-111-120. URL <https://www.mais-journal.ru/jour/article/view/162>.
- David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the Twenty-second Annual Symposium on Computational Geometry (SoCG 2016)*, pages 119–126. ACM, 2006. ISBN 1-59593-340-9. doi: 10.1145/1137856.1137877. URL <http://doi.acm.org/10.1145/1137856.1137877>.
- William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015. doi: 10.1142/S0219498815500668. URL <https://www.worldscientific.com/doi/abs/10.1142/S0219498815500668>.
- Herbert Edelsbrunner and John Harer. *Computational Topology - an Introduction*. American Mathematical Society, 2010. ISBN 978-0-8218-4925-5.
- Patrizio Frosini and Claudia Landi. Persistent Betti numbers for a noise tolerant shape-based approach to image retrieval. In *Computer Analysis of Images and Patterns*, pages 294–301. Springer-Verlag Berlin Heidelberg, August 2011. ISBN 978-3-642-23671-6 978-3-642-23672-3. doi: 10.1007/978-3-642-23672-3_36. URL https://link.springer.com/chapter/10.1007/978-3-642-23672-3_36.
- Marcio Gameiro, Yasuaki Hiraoka, Shunsuke Izumi, Miroslav Kramar, Konstantin Mischakow, and Vidit Nanda. A topological measurement of protein compressibility. *Japan*

- Journal of Industrial and Applied Mathematics*, 32(1):1–17, 2015. ISSN 0916-7005, 1868-937X. doi: 10.1007/s13160-014-0153-5. URL <https://link.springer.com/article/10.1007/s13160-014-0153-5>.
- Robert F. Geitz. Pettis Integration. *Proceedings of the American Mathematical Society*, 82(1):81, 1981. ISSN 00029939. doi: 10.2307/2044321. URL <http://www.jstor.org/stable/2044321?origin=crossref>.
- Irving J. Good and Ray A. Gaskins. Density estimation and bump-hunting by the penalized likelihood method exemplified by scattering and meteorite data. *Journal of the American Statistical Association*, 75(369):42–56, 1980. ISSN 01621459. URL <http://www.jstor.org/stable/2287377>.
- Heather A. Harrington, Nina Otter, Hal Schenck, and Ulrike Tillmann. Stratifying multiparameter persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 3(3):439–471, 2019. doi: 10.1137/18M1224350. URL <https://doi.org/10.1137/18M1224350>.
- Jørgen Hoffmann-Jørgensen and Gilles Pisier. The Law of Large Numbers and the Central Limit Theorem in Banach Spaces. *The Annals of Probability*, 4(4):587–599, 1976.
- Sara Kališnik. Tropical coordinates on the space of persistence barcodes. *Foundations of Computational Mathematics*, 19(1):101–129, 2019. ISSN 1615-3383. doi: 10.1007/s10208-018-9379-y. URL <https://doi.org/10.1007/s10208-018-9379-y>.
- Xia Kelin and Wei Guo-Wei. Multidimensional persistence in biomolecular data. *Journal of Computational Chemistry*, 36(20):1502–1520, July 2015. ISSN 1096-987X. doi: 10.1002/jcc.23953. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.23953>.
- Bryn Keller, Michael Lesnick, and Theodore L. Willke. PHoS: Persistent homology for virtual screening. 2018. doi: 10.26434/chemrxiv.6969260.v1. URL https://chemrxiv.org/articles/PHoS_Persistent_Homology_for_Virtual_Screening/6969260.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces, Isoperimetry and Processes*. Springer-Verlag Berlin Heidelberg, 2011. ISBN 978-3-642-20212-4.
- Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, Jun 2015. ISSN 1615-3383. doi: 10.1007/s10208-015-9255-y. URL <https://doi.org/10.1007/s10208-015-9255-y>.
- Michael Lesnick and Matthew Wright. Interactive visualization of 2-D persistence modules. *arXiv e-prints*, art. arXiv:1512.00180, December 2015.
- Ezra Miller. Data structures for real multiparameter persistence modules. *arXiv e-prints*, art. arXiv:1709.08155, September 2017.
- Chul Moon, Noah Giansiracusa, and Nicole A. Lazar. Persistence terrace for topological inference of point cloud data. *Journal of Computational and Graphical Statistics*, 27(3):576–586, 2018. doi: 10.1080/10618600.2017.1422432. URL <https://doi.org/10.1080/10618600.2017.1422432>.

- Kazimierz Musiał. Topics in the theory of Pettis integration. 1991. ISSN 0049-4704.
- Monica Nicolau, Arnold J. Levine, and Gunnar Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences of the United States of America*, 108(17):7265–7270, April 2011. ISSN 0027-8424. doi: 10.1073/pnas.1102826108. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3084136/>.
- Steve Y. Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*. Number 209 in Mathematical Surveys and Monographs. American Mathematical Society, 2015. URL <https://hal.inria.fr/hal-01247501>.
- Jacek Skryzalin and Gunnar Carlsson. Numeric invariants from multidimensional persistence. *Journal of Applied and Computational Topology*, 1(1):89–119, 2017. ISSN 2367-1726. doi: 10.1007/s41468-017-0003-z. URL <http://link.springer.com/10.1007/s41468-017-0003-z>.
- The RIVET Developers. Rivet, 2018. URL <http://rivet.online>.