



A Framework for Differential Calculus on Persistence Barcodes

Jacob Leygonie¹ · Steve Oudot² · Ulrike Tillmann¹

Received: 27 April 2020 / Revised: 25 February 2021 / Accepted: 30 April 2021 /
Published online: 13 July 2021
© The Author(s) 2021

Abstract

We define notions of differentiability for maps from and to the space of persistence barcodes. Inspired by the theory of diffeological spaces, the proposed framework uses lifts to the space of ordered barcodes, from which derivatives can be computed. The two derived notions of differentiability (respectively, from and to the space of barcodes) combine together naturally to produce a chain rule that enables the use of gradient descent for objective functions factoring through the space of barcodes. We illustrate the versatility of this framework by showing how it can be used to analyze the smoothness of various parametrized families of filtrations arising in topological data analysis.

Keywords Persistent homology · Persistence barcodes · Optimization

Mathematics Subject Classification 55N31 · 62R40

Contents

1 Introduction	1070
1.1 Motivation	1070
1.2 Related Work	1071

Communicated by Herbert Edelsbrunner.

✉ Jacob Leygonie
jacob.leygonie@maths.ox.ac.uk

Steve Oudot
steve.oudot@inria.fr

Ulrike Tillmann
tillmann@maths.ox.ac.uk

¹ Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

² INRIA Saclay, Île-de-France, 1, rue Honoré d'Estienne d'Orves, 91120 Palaiseau, France

1.3 Contributions and Outline of the Paper	1073
2 Preliminary Notions	1076
2.1 Persistence Modules and Persistent Homology	1076
2.2 Persistence Barcodes/Diagrams	1078
2.3 Morse Functions	1079
2.4 Diffeology Theory	1080
2.5 Stratified Manifolds	1081
3 Differentiability for Maps from or to the Space of Barcodes	1082
3.1 Differentiability of Barcode Valued Maps	1082
3.2 Differentiability of Maps Defined on Barcodes	1085
3.3 Chain Rule	1086
3.4 Higher-Order Derivatives	1087
3.5 The Space of Barcodes as a Diffeological Space	1088
4 The Case of Barcode Valued Maps Derived from Real Functions on a Simplicial Complex	1089
4.1 Generic Smoothness of the Barcode Valued Map	1091
4.2 Differential of the Barcode Valued Map	1094
4.3 Directional Differentiability of the Barcode Valued Map Along Strata	1095
4.4 The Barcode Valued Map as a Permutation Map	1097
5 Application to Common Simplicial Filtrations	1100
5.1 Lower Star Filtrations	1100
5.2 Rips Filtrations of Point Clouds	1102
5.3 Rips Filtrations of Clouds of Ellipsoids	1104
5.4 Arbitrary Filtrations of a Simplicial Complex	1107
6 The Case of Barcode Valued Maps Derived from Real Functions on a Manifold	1108
6.1 Smoothness of the Barcode Valued Map	1109
6.2 Discussion: Generic Differentiability	1112
6.3 A Simple Example	1113
7 The Case of Maps on Barcodes Derived from Vectorizations and Loss Functions	1114
7.1 The Differentiability of Persistence Images	1115
7.2 Linear Representations of Barcodes	1117
7.3 Semi-algebraic and Subanalytic Functions on Barcodes	1118
7.4 The Bottleneck Distance to a Diagram	1120
A Local Isometry of the Barcode on a Simplicial Complex	1121
B The Bottleneck Distance to a Fixed Diagram: The General Case	1123
References	1129

1 Introduction

1.1 Motivation

Barcodes have been introduced in topological data analysis (TDA) as a means to encode the topological structure of spaces and real-valued functions. They have been shown to provide complementary information compared to classical geometric or statistical methods, which explains their interest for applications. However, so far they have been essentially used as an alternative representation of the input, engineered by the user, as opposed to optimized to fit the problem best.

Optimizing barcodes using, e.g., gradient descent requires to differentiate objective functions that factor through the space Bar of barcodes:

$$\mathcal{M} \longrightarrow Bar \longrightarrow \mathbb{R}, \quad (1)$$

where \mathcal{M} is a parameter space equipped with a differential structure, typically a smooth finite-dimensional manifold. A compelling example arises in the context of supervised learning, where the barcodes can be used as features for data, generated by using some *filter function* $f : K \rightarrow \mathbb{R}$ on a fixed graph or simplicial complex K . Instead of considering f as a hyperparameter, it can be beneficial to optimize it among a family $\{f_\theta : K \rightarrow \mathbb{R}\}_{\theta \in \mathcal{M}}$ parametrized by a smooth map which we call the *parametrization*:

$$F : \theta \in \mathcal{M} \mapsto f_\theta \in \mathbb{R}^K.$$

Post-composing F with the *persistent homology* operator Dgm_p in homology degree p yields a map $\text{Dgm}_p \circ F : \mathcal{M} \rightarrow \text{Bar}$. Given a loss function $\mathcal{L} : \text{Bar} \rightarrow \mathbb{R}$, the goal is then to minimize the functional

$$\mathcal{M} \xrightarrow{\text{Dgm}_p \circ F} \text{Bar} \xrightarrow{\mathcal{L}} \mathbb{R} \tag{2}$$

using variational approaches, which are standard in large-scale learning applications. In order to do so, we need to put a sensible smooth structure on Bar and to derive an analogue of the chain rule, so that we can compute the differential of $\mathcal{L} \circ \text{Dgm}_p \circ F$ as the composition of the differentials of \mathcal{L} and $\text{Dgm}_p \circ F$. The difficulty arises as Bar is not a manifold and so far has not been given a structure in which the above makes sense.

Beyond optimization, we want to be able to address other types of applications where differential calculus is involved. For this, a variety of potential scenarios must be considered, e.g., when the filter functions are defined on a fixed smooth manifold, or when the second arrow in (1) takes its values in \mathbb{R}^n or more generally in some smooth finite-dimensional manifold. The goal of our study is to provide a unified framework that accounts for all these scenarios.

1.2 Related Work

Despite the lack of a smooth structure on the space Bar , developing heuristic methods to differentiate the composition in Eq. (2) has been an active direction of research lately, leading to innovative computational applications. In Table 1, we specify, for each of these contributions, the choice of parametrization F and of loss function \mathcal{L} , the optimization problem under consideration, and the sufficient conditions worked out to guarantee the differentiability of the composition in (2).

In the context of point cloud inference considered by [27], the positions of points in a fixed Euclidean space form the parameter space \mathcal{M} , and the resulting Rips filtration (resp. Alpha filtration) of the total complex on the point cloud is the parametrization F . The loss function \mathcal{L} is given by the least-squares approximation of a fixed barcode. By developing a clear functional point of view on the connection between the barcode of the Rips or Alpha filtration and the positions of the points in the cloud, based on lifts to Euclidean space, the authors show that \mathcal{L} is differentiable wherever the pairwise distances between points in the cloud are distinct. The approach is further refined by [19], where it is observed that the parametrization F is a subanalytic map,

which implies that the barcode-valued map admits subanalytic (hence generically differentiable) lifts. In turn, this fact is leveraged to show that any probability measure with a density w.r.t. the Hausdorff measure on \mathcal{M} induces an expected persistence diagram (viewed as a measure in the plane) with a density w.r.t. the Lebesgue measure.

In many applications, F parametrizes *lower-star filtrations*, i.e., filter functions induced by their restrictions to the vertices of K [3, 14, 29, 30, 32, 43]. In [43], the problem of shape matching is cast into an optimization problem involving the barcodes of the shapes. [14] uses the degree-0 persistent homology as a regularizer for classifiers. Similarly, [32] proposes a persistence-based regularization as an additional loss for deep learning models in the context of image segmentation. In [30], a dataset of graphs is seen as part of a bigger common simplicial complex, which allows to learn a filter function which is shared across the whole dataset. These contributions require the differentiability of (2), and they show that it holds whenever the filter function f_θ is injective over the vertex set.

Functions on a grid are used in [3] to tackle the problem of surface reconstruction. These functions are sums of Gaussians whose means and variances are parameters one wants to optimize according to an objective/loss that depends on the degree-1 persistent homology of the functions. [29] considers optimization problems involving persistence with many useful applications as in generative modeling, classification robustness, and adversarial attacks. Both contributions need to take the derivative of (2), and to do so, they require the existence of an inverse map taking interval endpoints in the persistence diagram $\text{Dgm}_p(f_\theta)$ to the corresponding vertices of K . This is a strictly weaker requirement than the injectivity of f_θ , as used in the previous contributions, because an inverse map always exists (provided for instance by the standard reduction algorithm for persistent homology). However, per se, it does not guarantee the differentiability of the composition—see, e.g., [30] for a counter-example.

This variety of applications motivates the search for a unified framework for expressing the differentiability of the arrows in diagrams of the form:

$$\mathcal{M} \longrightarrow \text{Bar} \longrightarrow \mathbb{R}. \quad (3)$$

Since the first appearance of this paper as a preprint, there have been novel applications of persistence differentiability in optimization. For instance, the first author has developed a graph classification framework based on the Laplacian operator [48], applying the differentiability of the persistence map (Theorem 4.9) to the case of extended persistence. In addition, new heuristics to smooth and regularize loss functions as in Eq. (3) improved the optimization procedure for specific data science problems [11, 46]. Another strong guarantee is provided when the loss in Eq. (3) is semi-algebraic (and more generally subanalytic or definable in some o-minimal structure), as then the classic stochastic gradient descent (SGD) algorithm converges to critical points [20]. The bridge between this result in non-smooth analysis and persistence based optimization problems is made in [9], where sufficient conditions for loss functions as in Eq. (3) to be semi-algebraic are given. The main results of [9] also derive from our general framework, see Remark 4.25.

Table 1 Current frameworks for differentiating the composition in (2)

Targeted Application	Loss function \mathcal{L}	Parametrization F	Conditions for Differentiability of (2)
Topological regularizer for classification [14]	Penalty on the degree-0 persistent homology	Parametrizations of lower star filtrations	Injectivity of the filter function over the vertices
Topological regularizer for image segmentation [32]	Penalty on the degrees 0 and 1 persistent homology		
Graph classification [30]	Any loss of supervised learning		
Shape matching [43]	Bottleneck or Wasserstein distance		
Surface reconstruction [3]	Sum of persistences distance		Local correspondence between barcodes interval endpoints and vertices in the simplicial complex
Generative modeling, Adversarial attacks, Regularization [29]	Weighted sum of persistences		
Point cloud continuation [27]	Least-squares approximation of a fixed barcode	Point clouds determining a Rips filtration	Distinct pairwise distances between points

The first column lists the targeted applications. The second and third columns show the choices of loss function \mathcal{L} and parametrization F . The differentiability of $\mathcal{L} \circ \text{Dgm}_p \circ F$ is guaranteed under the conditions listed in the fourth column

1.3 Contributions and Outline of the Paper

Ultimately, our framework should make it possible to determine when and how maps between smooth manifolds \mathcal{M} and \mathcal{N} that factor through the space of barcodes can be differentiated:

$$\mathcal{M} \xrightarrow{B} \text{Bar} \xrightarrow{V} \mathcal{N} .$$

To achieve this goal, in Sect. 3 we define differentiability via lifts in full generality, thereby extending the approach initially proposed by [27] for the specific case of parametrizations by Rips filtrations. Here, we provide some of the details. As a space of multi-sets (assumed by default to have finitely many off-diagonal points), Bar does not naturally come equipped with a differential structure. However, it is covered by maps of the form:

$$\begin{array}{c} \mathbb{R}^{2m} \times \mathbb{R}^n \\ \downarrow Q_{m,n} \\ Bar \end{array}$$

where $\mathbb{R}^{2m} \times \mathbb{R}^n$ can be thought of as the space of ordered barcodes with fixed number m (resp. n) of finite (resp. infinite) intervals, and where $Q_{m,n}$ is the quotient map modulo the order—turning vectors into multi-sets (Definition 3.1). Then, the map $B : \mathcal{M} \rightarrow Bar$ is said to be r -differentiable at parameter $\theta \in \mathcal{M}$ if it admits a local C^r lift \tilde{B} into $\mathbb{R}^{2m} \times \mathbb{R}^n$ for some $m, n \in \mathbb{N}$:

$$\begin{array}{ccc} & \mathbb{R}^{2m} \times \mathbb{R}^n & \\ & \searrow \tilde{B} & \downarrow Q_{m,n} \\ \theta \in U \subset \mathcal{M} & \xrightarrow{B} & Bar \end{array} \tag{4}$$

This means that the map \tilde{B} tracks smoothly and consistently the points in the barcodes $B(\theta')$, for θ' ranging over some open neighborhood U of θ . Dually, the map $V : Bar \rightarrow \mathcal{N}$ is r -differentiable at $D \in Bar$ if for every possible choice of m, n , the composition $V \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathcal{N}$ is C^r on an open neighborhood of every pre-image \tilde{D} of D :

$$\begin{array}{ccc} \mathbb{R}^{2m} \times \mathbb{R}^n & & \\ \downarrow Q_{m,n} & & \\ Bar & \xrightarrow{V} & \mathcal{N} \end{array} \tag{5}$$

The choice of m, n and pre-image \tilde{D} of D should be thought of as the type of perturbation we allow around D . Thus, essentially, V is asked to be smooth with respect to any finite perturbation of D . In Sect. 3.5, we connect these definitions to the theory of diffeological spaces, showing that our two definitions of differentiability for maps B and V are dual to each other and make the barcode space Bar a diffeological space.

We then define the differentials of the maps B and V , given simply by the differentials of the lift $\tilde{B} : \mathcal{M} \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ (for B) and of the composition $V \circ Q_{m,n}$ on the pre-image $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^n$ (for V). Although these differentials taken individually are not defined uniquely, their corresponding diagrams (4) and (5) combine together as follows:

$$\begin{array}{ccccc} & & \mathbb{R}^{2m} \times \mathbb{R}^n & & \\ & & \searrow \tilde{B} & \downarrow Q_{m,n} & \\ \theta \in U \subset \mathcal{M} & \xrightarrow{B} & Bar & \xrightarrow{V} & \mathcal{N} \end{array}$$

implying that the composition $V \circ B = (V \circ Q_{m,n}) \circ \tilde{B}$ is a C^r map between smooth manifolds, whose derivative is obtained by composing the differentials of B and V , and this regardless of the choice of lift and pre-image. This is our analogue of the chain rule in ordinary differential calculus (Proposition 3.14).

In Sects. 4 and 6, we focus on barcode-valued maps $B : \mathcal{M} \rightarrow \text{Bar}$ arising from filter functions on fixed smooth manifolds or simplicial complexes. These maps are usually not differentiable everywhere on their domain. However, motivated by the aforementioned applications, we seek conditions under which B is differentiable almost everywhere on \mathcal{M} . A natural approach for this would be to use Rademacher’s theorem [24, Thm. 3.1.6], as we know that B is Lipschitz continuous by the stability theorem of persistent homology [5,12,16]. However, this approach has several important shortcomings:

- it depends on a choice of measure on \mathcal{M} ;
- it calls for a generalization of Rademacher’s theorem to maps taking values in arbitrary metric spaces, and to the best of our knowledge, existing generalizations only provide directional metric differentials (see, e.g., [41]);
- more fundamentally, it is not constructive and therefore does not provide formulae for the differentials;
- finally, in the context of optimization, it is important to guarantee the existence of differentials/gradients in an open neighborhood of the considered parameter θ , and not just in a full-measure subset.

We therefore propose to follow a different approach, seeking conditions that ensure the differentiability of B on a generic (i.e., open and dense) subset of \mathcal{M} , with explicit differential.

Our first scenario (Sect. 4) considers a parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ of filter functions on a fixed simplicial complex K . Given a homology degree $p \leq d$, where d is the maximal simplex dimension in K , the barcode-valued map B decomposes as $B = \text{Dgm}_p \circ F$, and in Theorem 4.9 we show that B is r -differentiable on a generic subset of \mathcal{M} whenever F is C^r over \mathcal{M} or a generic subset thereof. The proof relies on the fact that the pre-order on the simplices of K induced by the values assigned by the filter function $F(\theta)$ is generically constant around θ in \mathcal{M} . We then relate the differential of B to those of F in Proposition 4.14, yielding a closed formula that can be leveraged in practical implementations. Finally, we study the behavior of B at singular points by means of a stratification of the parameter space \mathcal{M} , whereby the top-dimensional strata are the locations where B is differentiable, and the lower-dimensional strata characterize the defect of differentiability of B . We show in Theorem 4.19 that we can define *directional derivatives* along each incident stratum at any given point $\theta \in \mathcal{M}$. We also show that the barcode valued map can be globally lifted and expressed as a permutation map on each stratum (Corollary 4.24).

In Sect. 5, we illustrate the impact of our framework on a series of examples of parametrizations coming from earlier work, including lower-star filtrations, Rips filtrations and some of their generalizations. For each example, we examine the differentiability of the barcode-valued map and, whenever readily computable, we give the expressions of its differential. This allows us to recover the differentiability results from earlier work in a principled way.

Our second scenario (Sect. 6) considers a parametrization $F : \mathcal{M} \rightarrow C^\infty(\mathcal{X}, \mathbb{R})$ of smooth filter functions on a fixed smooth compact d -dimensional manifold \mathcal{X} . In this scenario, given a parameter $\theta \in \mathcal{M}$, the barcode-valued map B computes all the barcodes of f_θ at once, and collates them in a vector of barcodes:

$$B : \theta \in \mathcal{M} \mapsto (\text{Dgm}_0(f_\theta), \dots, \text{Dgm}_d(f_\theta)) \in \text{Bar}^{d+1}.$$

We show that B is ∞ -differentiable at any parameter θ such that f_θ is Morse with distinct critical values (Theorem 6.1). The key insights are: on the one hand, that at any such parameter θ the implicit function theorem allows us to smoothly track the critical points of $f_{\theta'}$ as θ' ranges over a small enough open neighborhood around θ ; on the other hand, that the stability theorem provides a consistent correspondence between the critical points of $f_{\theta'}$ and the interval endpoints in its barcodes.

In Sect. 7, we look at examples of classes of maps $V : \text{Bar} \rightarrow \mathcal{N}$. We first consider persistence images [1] and more generally linear representations of barcodes, as an illustration of our framework on barcode vectorizations. We show that persistence images and linear representations are ∞ -differentiable under suitable choices of weighting function (Propositions 7.3 and 7.5). We then consider the case where $V : \text{Bar} \rightarrow \mathbb{R}$ is the bottleneck or Wasserstein distance to a fixed barcode and show it is semi-algebraic in a suitable sense (Proposition 7.7), which is useful in a context of optimization. We then focus on the bottleneck distance to a fixed barcode D_0 , which we believe can be of interest in the context of inverse problems. We show that this distance is differentiable on a generic subset of Bar (Propositions 7.9 and B.1).

Finally, throughout the paper we sprinkle our exposition with examples of parametrizations and loss functions that illustrate our results and demonstrate their potential for applications.

2 Preliminary Notions

Throughout the paper, vector spaces and homology groups are taken over a fixed field \mathbb{k} , omitted in our notations whenever clear from the context. As much as possible, we keep separate terminologies for different notions of differentiability, for instance: maps from or to the space of barcodes are called *r-differentiable* when maps between manifolds are simply called C^r . The only exception to this rule is the term *smooth* for maps, which has a versatile meaning that should nonetheless always be clear from the context.

2.1 Persistence Modules and Persistent Homology

Definition 2.1 A *persistence module* \mathbb{V} is a functor from the poset (\mathbb{R}, \leq) to the category $\mathbf{Vect}_{\mathbb{k}}$ of vector spaces over \mathbb{k} .

In other words, a persistence module is a collection $\mathbb{V} = \{V_t, v_{s,t} : V_s \rightarrow V_t\}_{(s,t) \in \mathbb{R}^2, s \leq t}$ of vector spaces V_t and linear maps $v_{s,t}$, such that $v_{t,t} = \text{id}_{V_t}$ for all $t \in \mathbb{R}$ and $v_{s,t} \circ v_{r,s} = v_{r,t}$ for all $r \leq s \leq t \in \mathbb{R}$. We say that \mathbb{V} is *pointwise*

finite-dimensional (or *pdf* for short) if every V_t is finite-dimensional. Unless otherwise stated, persistence modules in the following will be pdf.

Definition 2.2 A *morphism* $\eta : \mathbb{V} \rightarrow \mathbb{W}$ between two persistence modules is a natural transformation between functors.

In other words, writing $\mathbb{V} = \{V_t, v_{s,t}\}_{s \leq t}$ and $\mathbb{W} = \{W_t, w_{s,t}\}_{s \leq t}$, a morphism $\eta : \mathbb{V} \rightarrow \mathbb{W}$ is a collection of linear maps $\{\eta_t : V_t \rightarrow W_t\}_{t \in \mathbb{R}}$ such that the following diagram commutes for all $s \leq t$:

$$\begin{array}{ccc}
 V_s & \xrightarrow{v_{s,t}} & V_t \\
 \eta_s \downarrow & & \downarrow \eta_t \\
 W_s & \xrightarrow{w_{s,t}} & W_t
 \end{array}$$

We say that η is an *isomorphism* of persistence modules if all the η_t are isomorphisms of vector spaces. We denote by **Pers** the category of persistence modules. **Pers** is an abelian category, so it admits kernels, cokernels, images and direct sums, which are defined pointwise. By Crawley-Boevey’s theorem [7], we know that persistence modules essentially uniquely decompose as direct sums of elementary modules called *interval modules*. The interval module \mathbb{I}_J associated with an interval J of \mathbb{R} is defined as the module with copies of the field \mathbb{k} over J and zero spaces elsewhere, the copies of \mathbb{k} being connected by identity maps.

Theorem 2.3 For any persistence module \mathbb{V} , there is a unique multi-set \mathcal{J} of intervals of \mathbb{R} such that

$$\mathbb{V} \cong \bigoplus_{J \in \mathcal{J}} \mathbb{I}_J, \tag{6}$$

Persistence modules of particular interest are the ones induced by the sub-level sets of real-valued functions.

Definition 2.4 Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a real-valued function on a topological space. Write $\mathcal{X}^t := f^{-1}((-\infty, t])$ for the closed sublevel set of f at level $t \in \mathbb{R}$. Given $p \in \mathbb{N}$, the *sublevel set persistent homology* of f in degree p is the (non-necessarily pdf) persistence module $\mathbf{H}_p(f)$ defined by:

- the vector spaces $\{H_p(\mathcal{X}^t)\}_{t \in \mathbb{R}}$, where H_p is the singular homology functor in degree p with coefficients in \mathbb{k} ;
- the linear maps $\{v_{s,t} : H_p(\mathcal{X}^s) \rightarrow H_p(\mathcal{X}^t)\}_{s \leq t}$ induced by inclusions $\mathcal{X}^s \hookrightarrow \mathcal{X}^t$.

In the following, we restrict our focus to *finite-type* persistence modules induced by *tame* functions, defined as follows:

Definition 2.5 A persistence module \mathbb{V} is of *finite type* if it admits a decomposition into finitely many interval modules.

Definition 2.6 A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is *tame* if its persistent homology modules in any degree are of finite type.

In particular, *filter functions* on a finite simplicial complex (see below) and *Morse functions* on a smooth manifold (see Sect. 2.3) are tame.

Definition 2.7 Let K be a finite simplicial complex. A *filter function* $f : K \rightarrow \mathbb{R}$ is a function that is monotonous with respect to inclusions of faces in K , i.e., $f(\sigma) \leq f(\sigma')$ for all $\sigma \subseteq \sigma' \in K$. This implies in particular that every sublevel set $K^t := \{\sigma \in K \mid f(\sigma) \leq t\}$ is a sub-complex of K .

2.2 Persistence Barcodes/Diagrams

Given a decomposition of a finite-type persistence module \mathbb{V} as in (6), the (finite) multi-set \mathcal{J} is called the *barcode* of \mathbb{V} . An alternate representation is as a (finite) multi-set B of points in the plane, where each interval $J \in \mathcal{J}$ is mapped to the point $(\inf J, \sup J)$. To this multi-set of points, we add Δ^∞ , that is the multi-set containing countably many copies of the diagonal $\Delta := \{(b, b) \mid b \in \mathbb{R}\}$, to obtain the so-called *persistence diagram* of \mathbb{V} . When \mathbb{V} is the sublevel set persistent homology of a tame function f in degree p , we denote by $\text{Dgm}_p(f)$ its persistence diagram. Persistence diagrams can also be defined independently of persistence modules as follows:

Definition 2.8 A *persistence diagram* is the union $B \cup \Delta^\infty$ of a finite multi-set B of elements in $\mathbb{R} \times \bar{\mathbb{R}}$, where $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$, with countably many copies of the diagonal Δ . The set of persistence diagrams is denoted by Bar .

From now on, we also use the terminology *barcodes* for persistence diagrams. Following this terminology, we also call *intervals* the points in a persistence diagram. Points lying on the diagonal Δ are qualified as *diagonal*, the others are qualified as *off-diagonal*.

Remark 2.9 In the above definitions, we follow the literature on extended persistence, in which persistence diagrams can have points everywhere in the extended plane $\mathbb{R} \times \bar{\mathbb{R}}$. This is because our framework extends naturally to that setting.

Note also that, in the literature, the diagonal is sometimes not included in the diagrams. Here, we are including it with infinite multiplicity. This is in the spirit of taking the quotient category of observable persistence modules, as defined by [8].

Definition 2.10 Given two barcodes $D, D' \in \text{Bar}$, viewed as multi-sets, a *matching* is a bijection $\gamma : D \rightarrow D'$. The *cost* of γ is the quantity

$$c(\gamma) := \sup_{x \in D} \|x - \gamma(x)\|_\infty \in \bar{\mathbb{R}}.$$

We denote by $\Gamma(D, D')$ the set of all matchings between D and D' .

Definition 2.11 The *bottleneck distance* between two barcodes $D, D' \in \text{Bar}$ is

$$d_\infty(D, D') := \inf_{\gamma \in \Gamma(D, D')} c(\gamma)$$

Given $q \in \mathbb{R}_+^*$, a slight modification of the matching cost yields the q -th *Wasserstein distance* on barcodes as introduced in [17]:

$$d_q(D, D') := \inf_{\gamma \in \Gamma(D, D')} \left(\sum_{x \in D} \|x - \gamma(x)\|_\infty^q \right)^{\frac{1}{q}} \tag{7}$$

Since we include all points in the diagonal with infinite multiplicity in our definition of barcodes, d_∞ is a true metric¹ and not just a pseudo-metric. Indeed, for any $D, D' \in \text{Bar}$, we have $d_\infty(D, D') = 0 \Rightarrow D = D'$. We call *bottleneck topology* the topology induced by d_∞ , which by the previous observation makes *Bar* a Hausdorff space.

A key fact is the Lipschitz continuity of the barcode function, known as the *stability theorem* [5,12,16]:

Theorem 2.12 *Let $f, g : \mathcal{X} \rightarrow \mathbb{R}$ be two real-valued functions with well-defined barcodes. Then,*

$$d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) \leq \|f - g\|_\infty.$$

Note that the assumptions in the theorem are quite general and hold in our cases of interest: tame functions on a compact manifold, and filter functions on a simplicial complex.

2.3 Morse Functions

Morse functions are a special type of tame functions, for which there is a bijective correspondence between critical points in the domain and interval endpoints in the barcode. This correspondence, detailed in Proposition 2.14, will be instrumental in the analysis of Sect. 6. For a proper introduction to Morse theory, we refer the reader to [39].

Definition 2.13 Given a smooth d -dimensional manifold \mathcal{X} , a smooth function $f : \mathcal{X} \rightarrow \mathbb{R}$ is called *Morse* if its Hessian at critical points (i.e., points where the gradient of f vanishes) is non-degenerate.

Note that we do not assume a priori that the values of f at critical points (called *critical values*) are all distinct. For such a value a , we call *multiplicity* of a the number of critical points in the level-set $f^{-1}(a)$. We also introduce the notation $\text{Crit}(f)$ to refer to the set of critical points, which is discrete in \mathcal{X} . In particular, if \mathcal{X} is compact, which will be the case in this paper, $\text{Crit}(f)$ is finite. The number of negative eigenvalues of f at a critical point x is called the *index* of x .

Proposition 2.14 *Assume \mathcal{X} is compact and all the critical values of f have multiplicity 1. Denote by $E(f)$ the multi-set of finite endpoints of off-diagonal intervals (including the left endpoints of infinite intervals) of $\text{Dgm}_0(f) \sqcup \dots \sqcup \text{Dgm}_d(f)$. Then, f induces a bijection $\text{Crit}(f) \rightarrow E(f)$.*

¹ In fact an extended metric as it can take infinite values.

This result is folklore, and we give a proof only for completeness.

Proof Let $a \leq b$ be real numbers. Write \mathcal{X}^a for the sublevel set $f^{-1}((-\infty, a])$. If $[a, b]$ contains a unique critical value c of f , then \mathcal{X}^b has the homotopy type of \mathcal{X}^a glued together with a cell e_p of dimension p , where p is the index of the unique critical point x associated with c [39]. Therefore, $H_*(\mathcal{X}^b, \mathcal{X}^a)$ is trivial except for $* = p$ where it is spanned by the homology class of e_p . This does not depend on the choice of a, b surrounding c and sufficiently close to it. Then, using the long exact sequence in homology, we deduce that either there is one birth in degree p at value c in the persistent homology module, or there is one death in degree $p-1$. Hence, c is either a left endpoint of an interval of $\text{Dgm}_p(f)$, or a right endpoint of an interval of $\text{Dgm}_{p-1}(f)$. In either case, we can define the map $x \mapsto f(x)$ for any $x \in \text{Crit}(f)$, and we have just shown that its codomain is indeed $E(f)$. The map is injective because the critical values of f have multiplicity 1 by assumption. We now show it is onto. Let $a \in \mathbb{R}$ be a non-critical value of f . For any (small enough) $\varepsilon, \eta > 0$, the interval $[a - \eta, a + \varepsilon]$ contains no critical value of f , therefore $\mathcal{X}^{a+\varepsilon}$ deformation retracts onto $\mathcal{X}^{a-\eta}$, thus implying that the inclusions $H_p(\mathcal{X}^{a-\eta}) \rightarrow H_p(\mathcal{X}^{a+\varepsilon})$ are identity maps for any homology degree p . By the decomposition Theorem 2.3, this implies that a cannot be an endpoint of an interval summand, i.e., $a \notin E(f)$. \square

The assumption that each critical value of f has multiplicity 1 is superfluous in Proposition 2.14, if we allow the correspondence map to match trivial intervals. Let $[a, b]$ be an interval containing a unique critical value c . One can still use Morse theory and glue as many critical cells e_p to \mathcal{X}^a as there are critical points in $f^{-1}(c)$ in order to obtain a CW structure on \mathcal{X}^b from the one of \mathcal{X}^a . Considering the different critical cells, we know exactly the ranks of the morphisms $H_p(\mathcal{X}^a) \rightarrow H_p(\mathcal{X}^b)$ induced by inclusions in each homology degree p .

2.4 Diffeology Theory

Diffeology theory provides a principled approach to equip a set with a smooth structure. We use some concepts of the theory in Sect. 3.5, where we equip the set *Bar* of barcodes with a diffeology and identify the resulting smooth maps. We refer the reader to [33] for a detailed introduction to the material presented below. In the following, we call *domain* any open set in any arbitrary Euclidean space.

Definition 2.15 Given a non-empty set S , a *diffeology* is a collection \mathcal{D} of pairs (U, P) , called *plots*, where U is a domain and $P : U \rightarrow S$ is a map from U to S , satisfying the following axioms:

(Covering) For any element $s \in S$ and any integer $n \in \mathbb{N}$, the constant map $x \in \mathbb{R}^n \mapsto s \in S$ is a plot.

(Locality) If for a pair (U, P) we have that, for any $x \in U$ there exists an open neighborhood $U' \subseteq U$ of x such that the restriction $(U', P|_{U'})$ is a plot, then (U, P) itself is a plot.

(Smoothness compatibility) For any plot (U, P) and any smooth map $F : W \rightarrow U$ where W is a domain, the composition $(W, P \circ F)$ is a plot.

If a set S comes equipped with a diffeology \mathcal{D} , then it is called a *diffeological space*. We think of a diffeological space S as a space where we impose which functions, the plots, from a manifold to S , are smooth. Notice that any set can be made a diffeological space by taking all possible maps as plots. This is the *coarsest* diffeology on S , where \mathcal{D} is said to be *finer* than the diffeology \mathcal{D}' if $\mathcal{D} \subset \mathcal{D}'$, and *coarser* if the converse inclusion holds.² The prototypical diffeological space is the Euclidean space \mathbb{R}^n with the usual smooth maps from domains to \mathbb{R}^n as plots.

Definition 2.16 A *morphism* $f : S \rightarrow S'$, or *smooth map*, between two diffeological spaces S and S' , is a map such that for each plot P of S , $f \circ P$ is a plot of S' . f is called a *diffeomorphism* if it is a bijection and $f^{-1} : S' \rightarrow S$ is smooth. A map $f : A \rightarrow S'$, where $A \subseteq S$, is *locally smooth* if for any plot P of S , $f \circ P|_{P^{-1}(A)}$ is a plot of S' . f is a *local diffeomorphism* if it is a bijection onto its image and if f^{-1} is locally smooth as a map $S' \supseteq f(A) \rightarrow S$.

Obviously, identities are smooth, and smooth maps compose together into smooth maps; therefore, we can consider the category **Diffeo** of diffeological spaces. Finite-dimensional smooth manifolds with or without boundaries and corners, Fréchet manifolds and Frélicher spaces, viewed as diffeological spaces with their usual smooth maps, form strict subcategories of **Diffeo**. In fact, finite-dimensional smooth manifolds can be defined in the context of diffeology as follows:

Definition 2.17 A diffeological space \mathcal{M} is a *n-dimensional diffeological manifold* if it is locally diffeomorphic to \mathbb{R}^n at every point in \mathcal{M} .

Theorem 2.18 [33, § 4.3] *Every n-dimensional smooth manifold \mathcal{M} is an n-dimensional diffeological manifold once equipped with the diffeology given by the smooth maps $U \rightarrow \mathcal{M}$ from arbitrary domains U . Conversely, every n-dimensional diffeological manifold is an n-dimensional smooth manifold.*

One appealing feature of **Diffeo**, compared to the category of smooth manifolds for instance, is that it is closed under usual set operations—here we only consider coproducts and quotients:

Definition 2.19 For an arbitrary family of diffeological spaces $\{(S_j, \mathcal{D}_j)\}_{j \in \mathcal{J}}$, the *sum diffeology* on $\bigsqcup_{j \in \mathcal{J}} S_j$ is the finest diffeology making the injections $S_i \rightarrow \bigsqcup_{j \in \mathcal{J}} S_j$ smooth.

Definition 2.20 For a diffeological space (S, \mathcal{D}) and an equivalence relation \sim on S , the *quotient diffeology* on S/\sim is the finest diffeology making the quotient map $S \rightarrow S/\sim$ smooth.

2.5 Stratified Manifolds

Stratified manifolds play a role in Sect. 4.3 of this paper. For background material on the subject, see, e.g., [37].

² This terminology is the opposite to the one used when comparing topologies.

Definition 2.21 Let \mathcal{M} be a smooth d -dimensional manifold. A *Whitney stratification* $\mathcal{S}_{\mathcal{M}}$ of \mathcal{M} is a collection of connected smooth submanifolds (not necessarily closed) of \mathcal{M} , called *strata*, satisfying the following axioms:

(Partition) The strata partition \mathcal{M} .

(Locally finite) Each point of \mathcal{M} has an open neighborhood meeting with finitely many strata.

(Frontier) For each stratum $\mathcal{M}' \in \mathcal{S}_{\mathcal{M}}$, the set $\overline{\mathcal{M}'} \setminus \mathcal{M}'$ is a union of strata, where $\overline{\mathcal{M}'}$ is the closure of \mathcal{M}' in \mathcal{M} .

(Condition b) Consider a pair of strata $(\mathcal{M}', \mathcal{M}'')$ and an element $\theta \in \mathcal{M}'$. If there are sequences of points $(\theta'_k)_{k \in \mathbb{N}}$ and $(\theta''_k)_{k \in \mathbb{N}}$ lying in \mathcal{M}' and \mathcal{M}'' , respectively, both converging to θ , such that the line (θ'_k, θ''_k) (defined in some local coordinate system around θ) converges to some line l and $T_{\theta'_k} \mathcal{M}''$ converges to some flat, then this flat contains l .

Stratified maps are those that behave nicely with respect to stratifications. Here, we only use a subset of the axioms they satisfy; hence, we talk about *weakly stratified maps*.

Definition 2.22 Let \mathcal{M}, \mathcal{N} be stratified manifolds. A map $f : \mathcal{M} \rightarrow \mathcal{N}$ is *weakly stratified* if the pre-images $f^{-1}(\mathcal{N}')$, for any stratum $\mathcal{N}' \in \mathcal{S}_{\mathcal{N}}$, is a union of strata in $\mathcal{S}_{\mathcal{M}}$.

3 Differentiability for Maps from or to the Space of Barcodes

In Sect. 3.1, we provide a general framework for studying the differentiability of maps from a smooth manifold to *Bar*. Then, in Sect. 3.2 we provide the analogue for maps with *Bar* as domain and a smooth manifold as co-domain. Both frameworks are in some sense dual to each other, and inspired by the theory of diffeological spaces—we develop this connection in Sect. 3.5. We then derive a chain rule in Sect. 3.3: if a map between manifolds factors through *Bar*, then it is smooth whenever both terms in the factorization are smooth according to our definitions, and in this case its differential can be computed explicitly.

3.1 Differentiability of Barcode Valued Maps

Throughout this section, \mathcal{M} denotes a smooth finite-dimensional manifold without boundary, which may or may not be compact. Our approach to characterizing the smoothness of a barcode valued map is to factor it through the bundle of *ordered barcodes*:

Definition 3.1 For each choice of nonnegative integers m, n , the space of *ordered barcodes* with m finite bars and n infinite ones is $\mathbb{R}^{2m} \times \mathbb{R}^n$, equipped with the Euclidean norm and the resulting smooth structure. The corresponding *quotient map* $Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \text{Bar}$ quotients the space by the action³ of the

³ \mathfrak{S}_m acts on \mathbb{R}^{2m} by permutation of pairs of adjacent coordinates while \mathfrak{S}_n acts on \mathbb{R}^n by permutation of coordinates.

product of symmetric groups $\mathfrak{S}_m \times \mathfrak{S}_n$, that is: for any ordered barcode $\tilde{D} = (b_1, d_1, \dots, b_m, d_m, v_1, \dots, v_n) \in \mathbb{R}^{2m} \times \mathbb{R}^n$,

$$Q_{m,n}(\tilde{D}) := \{(b_i, d_i)\}_{i=1}^m \cup \{(v_j, +\infty)\}_{j=1}^n \cup \Delta^\infty.$$

One can think of an ordered barcode $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^n$ as a vector describing a persistence diagram with at most m bounded off-diagonal points and exactly n unbounded points. The former have their coordinates encoded in the adjacent pairs of the $2m$ first components in \tilde{D} , while the latter have the abscissa of their left endpoint encoded in the last n components of \tilde{D} . The quotient map $Q_{m,n}$ forgets about the ordering of the bars in the barcodes. So far $Q_{m,n}$ is merely a map between sets, and it is natural to ask whether it is regular in some reasonable sense:

Proposition 3.2 *For any $m, n \in \mathbb{N}^2$, $Q_{m,n}$ is 1-Lipschitz when Bar is equipped with the bottleneck topology.*

Proof For any two elements $\tilde{D}_1, \tilde{D}_2 \in \mathbb{R}^{2m} \times \mathbb{R}^n$, there is an obvious matching γ on their images $Q_{m,n}(\tilde{D}_1), Q_{m,n}(\tilde{D}_2)$ given by matching the components of the vectors \tilde{D}_1 and \tilde{D}_2 entry-wise. The cost of this matching is then bounded above by the supremum norm of $\tilde{D}_1 - \tilde{D}_2$, by the definition of the matching cost $c(\gamma)$. In turn, the supremum norm is bounded above by the ℓ^2 norm. \square

We then say that a barcode valued map is smooth if it admits a smooth lift into the space of ordered barcodes for some choice of m, n :

Definition 3.3 Let $B : \mathcal{M} \rightarrow Bar$ be a barcode valued map. Let $x \in \mathcal{M}$ and $r \in \mathbb{N} \cup \{+\infty\}$. We say that B is r -differentiable at x if there exists an open neighborhood U of x , integers $m, n \in \mathbb{N}$ and a map $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ of class C^r such that $B = Q_{m,n} \circ \tilde{B}$ on U . For an integer $d \in \mathbb{N}$, a function $\mathcal{B} : \mathcal{M} \rightarrow Bar^{d+1}$ is r -differentiable at $x \in \mathcal{M}$ if each of its $d + 1$ components is. We call \tilde{B} a local lift of B .

Remark 3.4 (Locally finite number of off-diagonal points) If a function B as above is r -differentiable at $x \in \mathcal{M}$, then locally for any x' around x we can upper-bound the number of off-diagonal points arising in $B(x')$ by $m + n$. Notice that off-diagonal points can possibly appear in $B(x')$ and become part of the diagonal Δ in $B(x)$, which is to say that Definition 3.3 does not restrict the function B to locally consist in a fixed number of off-diagonal points. Informally, in analogy with the fact that a barcode has finitely many off-diagonal points, our definition of smoothness allows finitely many appearances or disappearances of off-diagonal points in the neighborhood of a barcode.

Remark 3.5 (0-differentiability is stronger than bottleneck continuity) If $B : \mathcal{M} \rightarrow Bar$ is 0-differentiable, then B is continuous when Bar is given the bottleneck topology. This comes from the Lipschitz continuity of $Q_{m,n}$ (Proposition 3.2) and the fact that continuity is stable under composition. The converse is false, because, on the one hand, if B is 0-differentiable then locally the number of off-diagonal points in the image of B is uniformly bounded (see the previous remark), while, on the other hand,

the number of off-diagonal points appearing in barcodes in any given open bottleneck ball is arbitrarily large.

Definition 3.6 Let $B : \mathcal{M} \rightarrow Bar$ be 1-differentiable at some x , and $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ be a C^1 lift of B defined on an open neighborhood U of x . The differential (or derivative) $d_{x, \tilde{B}}B$ of B at x with respect to \tilde{B} is defined to be the differential of \tilde{B} at x :

$$T_x \mathcal{M} \xrightarrow{d_x \tilde{B}} \mathbb{R}^{2m} \times \mathbb{R}^n.$$

Post-composing with the quotient map, we can see $Q_{m,n} \circ d_{x, \tilde{B}}B : T_x \mathcal{M} \rightarrow Bar$ as a multi-set of co-vectors, one above each off-diagonal point of $B(x)$ (plus some distinguished diagonal points), describing linear changes in the coordinates of the points of $B(x)$ under infinitesimal perturbations of x . In this respect, the spaces of ordered barcodes \mathbb{R}^{2m+n} play the role of tangent spaces over Bar . For practical computations, it can be convenient to work with an alternate yet equivalent notion of differentiability, based on point trackings:

Definition 3.7 Let $B : \mathcal{M} \rightarrow Bar$ be a barcode valued map. Let $x \in \mathcal{M}$ and $r \in \mathbb{N} \cup \{+\infty\}$. A C^r local coordinate system for B at x is a collection of maps $\{b_i, d_i : U \rightarrow \mathbb{R}\}_{i \in I}$ and $\{v_j : U \rightarrow \mathbb{R}\}_{j \in J}$ for finite sets I, J defined on an open neighborhood U of x , such that:

- (Smooth) the maps b_i, d_i, v_j are of class C^r ;
- (Tracking) For any $x' \in U$, we have the multi-set equality

$$B(x') = \{(b_i(x'), d_i(x'))\}_{i \in I} \cup \{(v_j(x'), +\infty)\}_{j \in J} \cup \Delta^\infty.$$

Thus, in a local coordinate system, we have maps b_i, d_i (resp. v_j) that track the endpoints of bounded (resp. unbounded) intervals in the image barcode through B . We will often abbreviate the data of a local coordinate system of B at x by $\mathcal{T} = (U, \{b_i, d_i\}_{i \in I}, \{v_j\}_{j \in J})$.

Our two notions of differentiability are indeed equivalent:

Proposition 3.8 Let $B : \mathcal{M} \rightarrow Bar$ be a barcode valued map and $x \in \mathcal{M}$. Then, B is r -differentiable at x if and only if it admits a C^r local coordinate system at x . Specifically, post-composing a C^r local lift $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ around x with the quotient map $Q_{m,n}$ yields a C^r local coordinate system, and conversely, fixing an order on the functions of a C^r local coordinate system yields a C^r local lift.

Proof (\Rightarrow) Let $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ be a C^r local lift of B at x . Extract the components of $(b_1(x'), d_1(x'), \dots, b_m(x'), d_m(x'), v_1(x'), \dots, v_n(x')) := \tilde{B}(x')$ to get a local coordinate system, which is C^r over U as \tilde{B} is. (\Leftarrow) Let $\mathcal{T} = (U, \{b_i, d_i\}_{i \in I}, \{v_j\}_{j \in J})$ be a C^r local coordinate system for B at x . Set $m = |I|$ and $n = |J|$, and fix two arbitrary bijections $s : \{1, \dots, m\} \rightarrow I$ and $t : \{1, \dots, n\} \rightarrow J$. Then, the map $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ defined as:

$$\tilde{B}(x') := [b_{s(1)}(x'), d_{s(1)}(x'), \dots, b_{s(m)}(x'), d_{s(m)}(x'), v_{t(1)}(x'), \dots, v_{t(n)}(x')]$$

is a lift of B . As a map valued in a Euclidean space, \tilde{B} is C^r because all its coordinate functions are. □

Remark 3.9 (*Non-uniqueness of differentials*) It is important to keep in mind that the differential of B at x is not uniquely defined, as it depends on the choice of local lift. Indeed, for two distinct lifts \tilde{B}, \tilde{B}' of B at x , we usually get distinct differentials $dB_{x, \tilde{B}}, dB_{x, \tilde{B}'}$. For instance, if \tilde{B}' is obtained from \tilde{B} by appending an extra pair of coordinates of the form (f, f) , where f is a smooth real function, then $dB_{x, \tilde{B}'}$ takes its values in a different codomain than that of $dB_{x, \tilde{B}}$. Note that this will not be an issue in the rest of the paper, as any choice of differential will yield a valid chain rule (Sect. 3.3).

3.2 Differentiability of Maps Defined on Barcodes

Let \mathcal{N} be a smooth finite-dimensional manifold without boundary. Our notion of differentiability for maps $V : Bar \rightarrow \mathcal{N}$ is in some sense dual to the one for maps $B : \mathcal{M} \rightarrow Bar$, as will be justified formally in the next section.

Definition 3.10 Let $V : Bar \rightarrow \mathcal{N}$ be a map on barcodes. Let $D \in Bar$ and $r \in \mathbb{N} \cup \{+\infty\}$. V is said to be r -differentiable at D , if for all integers m, n and all vectors $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^n$ such that $Q_{m,n}(\tilde{D}) = D$, the map $V \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathcal{N}$ is C^r on an open neighborhood of \tilde{D} .

Notice that for each choice of m, n we have a unique map $V \circ Q_{m,n}$, and we must check its differentiability at all the (possibly many) distinct pre-images \tilde{D} of D and for all m, n . One can think of a choice of m, n and pre-image \tilde{D} of D as a choice of tangent space of Bar at D .

Example 3.11 (*Total persistence function*) Let $V : Bar \rightarrow \mathbb{R}$ be defined as the sum, over bounded intervals (b, d) in a barcode D , of the length $(d - b)$. Given $D \in Bar$ and an ordered barcode $\tilde{D} \in \mathbb{R}^{2m+n}$ such that $Q_{m,n}(\tilde{D}) = D$, the map $V \circ Q_{m,n}$ is a linear form and in particular is of class C^∞ at \tilde{D} . Explicitly, we have

$$V \circ Q_{m,n} : (b_1, d_1, \dots, b_m, d_m, v_1, \dots, v_n) \in \mathbb{R}^{2m+n} \mapsto \sum_{i=1}^m d_i - b_i \in \mathbb{R}$$

Therefore, V is ∞ -differentiable everywhere on Bar .

The relationship between 0-differentiability and the bottleneck continuity for maps V is the opposite to the one that holds for maps B (recall Remark 3.5):

Remark 3.12 (*Bottleneck continuity is stronger than 0-differentiability*) If $V : Bar \rightarrow \mathcal{N}$ is continuous when Bar is equipped with the bottleneck topology, then V is 0-differentiable. This is because the quotient map $Q_{m,n}$ is continuous (Proposition 3.2) and the composition of continuous maps is continuous. The converse is false, as seen, for instance, when taking V to be the total persistence function: although 0-differentiable (because ∞ -differentiable) on Bar , V is not continuous in the bottleneck topology as it is unbounded in any open bottleneck ball.

Definition 3.13 Let $V : Bar \rightarrow \mathcal{N}$ be 1-differentiable at $D \in Bar$, and $\tilde{D} \in \mathbb{R}^{2m+n}$ be a pre-image of D via $Q_{m,n}$. The differential (or derivative) of V at D with respect to \tilde{D} is the map

$$d_{D,\tilde{D}}V : \mathbb{R}^{2m+n} \xrightarrow{d_{\tilde{D}}(V \circ Q_{m,n})} T_{V(D)}\mathcal{N}.$$

3.3 Chain Rule

We now combine the previous definitions to produce a chain rule.

Proposition 3.14 Let $B : \mathcal{M} \rightarrow Bar$ be r -differentiable at $x \in \mathcal{M}$, and $V : Bar \rightarrow \mathcal{N}$ be r -differentiable at $B(x)$. Then:

- (i) $V \circ B : \mathcal{M} \rightarrow \mathcal{N}$ is C^r at x as a map between smooth manifolds;
- (ii) If $r \geq 1$, then for any local C^1 lift $\tilde{B} : U \rightarrow \mathbb{R}^{2m+n}$ of B around x we have:

$$d_x(V \circ B) = d_{B(x),\tilde{B}(x)}V \circ d_{x,\tilde{B}}B.$$

The meaning of this formula is that even though the differentials of B and of V may depend on the choice of lift $\tilde{B} : \mathcal{M} \rightarrow \mathbb{R}^{2m+n}$, their composition does not, and in fact it matches with the usual differential of $V \circ B$ as a map between smooth manifolds.

Proof Since B is r -differentiable at x , there exists an open neighborhood U of x and a local C^r lift $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ for some integers m, n , such that $B|_U = Q_{m,n} \circ \tilde{B}$. Meanwhile, since V is r -differentiable at $B(x)$, the map $V \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathcal{N}$ is C^r at $\tilde{B}(x)$. This implies that the composition $V \circ B|_U = (V \circ Q_{m,n}) \circ \tilde{B}$ is C^r at x and therefore that $V \circ B$ itself is C^r at x since U is open. This proves (i). The formula of (ii) follows then from applying the usual chain rule to $(V \circ Q_{m,n})$ and \tilde{B} , which are C^1 maps between smooth manifolds without boundary. \square

Example 3.15 In [30], given a C^∞ neural network architecture $F_0 : \mathbb{R}^N \rightarrow \mathbb{R}^{K_0}$ valued in the set of functions over the vertices of a fixed graph K , the optimization pipeline requires taking the gradient of the following loss function:

$$\mathcal{L} : \theta \in \mathbb{R}^N \mapsto \sum_{(b,d) \in \text{Dgm}_p(F_0(\theta)) \setminus \Delta \text{ bounded}} s(b, d) \in \mathbb{R},$$

where $s : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a fixed smooth map, and $\text{Dgm}_p(F_0(\theta))$ is the degree- p persistence diagram associated with the lower star filtration induced by $F_0(\theta)$ on K (see Sect. 5.1 dedicated to the full analysis of lower star filtrations). We may see \mathcal{L} as the composition:

$$\mathcal{L} : \theta \in \mathbb{R}^N \xrightarrow{B} \text{Dgm}_p(F_0(\theta)) \in Bar \xrightarrow{V} V(\text{Dgm}_p(F_0(\theta))) \in \mathbb{R},$$

where $V : D \in Bar \mapsto \sum_{(b,d) \in D \setminus \Delta \text{ bounded}} s(b, d) \in \mathbb{R}$. On the one hand, B is ∞ -differentiable at every θ where $F_0(\theta)$ is injective, as will be detailed in Sect. 5.1.

On the other hand, V is ∞ -differentiable everywhere on Bar , a fact obtained exactly as in the case of the total persistence function of Example 3.11. By the chain rule (Proposition 3.14), we deduce that the loss \mathcal{L} is smooth at every θ where $F_0(\theta)$ is injective. Thus we recover the differentiability result of [30]. In fact, the upcoming Theorem 4.9 ensures that B is ∞ -differentiable over an open dense subset of \mathbb{R}^N , and therefore so is \mathcal{L} by the chain rule.

3.4 Higher-Order Derivatives

The notions of derivatives introduced in Definitions 3.6 and 3.13 extend naturally to higher orders. For simplicity, we place ourselves in the Euclidean setting, letting $\mathcal{M} = \mathbb{R}^N$ and $\mathcal{N} = \mathbb{R}^{N'}$ for some $N, N' \in \mathbb{N}$.

Definition 3.16 Let $B : \mathbb{R}^N \rightarrow Bar$ be r -differentiable at some x , and $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ be a C^r lift of B defined on an open neighborhood U of x . The r -th differential (or derivative) of B at x with respect to \tilde{B} is defined to be the r -th Fréchet differential of \tilde{B} at x :

$$d_x^r \tilde{B} : (\mathbb{R}^N)^r \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n.$$

Dually:

Definition 3.17 Let $V : Bar \rightarrow \mathbb{R}^{N'}$ be r -differentiable at $D \in Bar$, and $\tilde{D} \in \mathbb{R}^{2m+n}$ be a pre-image of D via $Q_{m,n}$. The r -th differential (or derivative) of V at D with respect to \tilde{D} is the r -th Fréchet differential of $V \circ Q_{m,n}$ at \tilde{D} :

$$d_{\tilde{D}}^r (V \circ Q_{m,n}) : (\mathbb{R}^{2m+n})^r \rightarrow \mathbb{R}^{N'}.$$

Note that, given maps $B : \mathbb{R}^N \rightarrow Bar$ and $V : Bar \rightarrow \mathbb{R}^{N'}$ that are r -differentiable at x and $B(x)$, respectively, the chain rule of Sect. 3.3 adapts readily to higher-order derivatives of $B \circ V$ at x .

Meanwhile, we get a natural Taylor expansion of B at x with respect to \tilde{B} :

$$T_{x, \tilde{B}}^r B : h \in \mathbb{R}^N \mapsto \tilde{B}(x) + d_x \tilde{B}(h) + \dots + \frac{1}{r!} d_x^r \tilde{B}(h, \dots, h) \in \mathbb{R}^{2m} \times \mathbb{R}^n.$$

Proposition 3.18 Let $B : \mathbb{R}^N \rightarrow Bar$ be r -differentiable at some x , and $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ be a C^r lift of B defined on an open neighborhood U of x . Then,

$$d_\infty(B(x + h), (Q_{m,n} \circ T_{x, \tilde{B}}^r B)(h)) = o(\|h\|^r).$$

Proof This follows from applying the standard Taylor–Young theorem to \tilde{B} , then post-composing by $Q_{m,n}$ —which is 1-Lipschitz by Proposition 3.2. □

To our knowledge, there is in general no equivalent of this result for the map V , due to the lack of a Lipschitz-continuous section of $Q_{m,n}$.

3.5 The Space of Barcodes as a Diffeological Space

In this subsection, we detail how Bar , when viewed as the quotient of a disjoint union of Euclidean spaces, is canonically made into a diffeological space, as defined in Sect. 2.4. We then show that the resulting notions of diffeological smooth maps from and to Bar coincide with the definitions 3.3 and 3.10 of differentiability we chose for maps from and to Bar in the previous sections, thus making these two definitions dual to each other.

As a set, Bar is isomorphic to $(\bigsqcup_{m,n \in \mathbb{N}} \mathbb{R}^{2m+n}) / \sim$, where \sim is the transitive closure of the following relations for m, n ranging over \mathbb{N} :

- For any permutations π, τ of $\{1, \dots, m\}$ and $\{1, \dots, n\}$, respectively,

$$[(b_i, d_i)_{i=1}^m, (v_j)_{j=1}^n] \sim [(b_{\pi(i)}, d_{\pi(i)})_{i=1}^m, (v_{\tau(j)})_{j=1}^n],$$

which indicates that persistence diagrams are multi-sets (i.e., intervals are not ordered);

- Any element $[(b_i, d_i)_{i=1}^m, (v_j)_{j=1}^n] \in \mathbb{R}^{2m+n}$ such that one of the first m adjacent pairs (b_i, d_i) satisfies $b_i = d_i$ is equivalent to the element of $\mathbb{R}^{2(m-1)+n}$ obtained by removing (b_i, d_i) . These identifications correspond to quotienting multi-sets by the diagonal Δ .

Since the Euclidean spaces \mathbb{R}^{2m+n} are equipped with their Euclidean diffeologies, we obtain a canonical diffeology $\mathcal{D}(Bar)$ over Bar from Definitions 2.19 and 2.20. The plots of $\mathcal{D}(Bar)$ can be concretely characterized as follows:

Proposition 3.19 *Let $U \subseteq \mathbb{R}^d$ be open and $B : U \rightarrow Bar$. Then, B is a plot in $\mathcal{D}(Bar)$ if and only if, for every $x \in U$, there exists an open neighborhood $V \subseteq U$ of x and a C^∞ lift $\tilde{B} : V \rightarrow \mathbb{R}^{2m+n}$ such that $B|_V = Q_{m,n} \circ \tilde{B}$.*

In other words, a plot in $\mathcal{D}(Bar)$ is an ∞ -differentiable map from a domain U to Bar .

Proof Note that the characterization of the quotient diffeology, as given in Definition 2.20, is in fact the characterization of the so-called *push-forward diffeology* induced by the quotient map—see [33, § 1.43]. According to that characterization, $B : U \rightarrow Bar$ is a plot if and only if, for every element $z \in U$, there exists an open neighborhood $W \subseteq U$ of z such that the restriction $B|_W$ admits a lift⁴ $\tilde{B} : W \rightarrow \bigsqcup_{m,n \in \mathbb{N}} \mathbb{R}^{2m+n}$, i.e., a plot \tilde{B} of $\bigsqcup_{m,n \in \mathbb{N}} \mathbb{R}^{2m+n}$ that matches with $B|_W$ once post-composed with the quotient map modulo \sim . In turn, by the characterization of the sum diffeology in [33, § 1.39], \tilde{B} is a plot of $\bigsqcup_{m,n \in \mathbb{N}} \mathbb{R}^{2m+n}$ if and only if, for any $x \in W$, there is an open neighborhood $V \subseteq W$ of x and a pair of indices (m, n) such that the restriction $\tilde{B}|_V$ maps into \mathbb{R}^{2m+n} and is in fact a plot of \mathbb{R}^{2m+n} . Equivalently, we have $B|_V = Q_{m,n} \circ \tilde{B}|_V$, where $\tilde{B}|_V$ is of class C^∞ (since the spaces of ordered barcodes are equipped with their canonical Euclidean diffeologies). \square

⁴ Strictly speaking, according to [33, § 1.43] there is also the alternative that the restriction $B|_W$ be constant, but in this case it also admits a lift to $\bigsqcup_{m,n \in \mathbb{N}} \mathbb{R}^{2m+n}$. Indeed, calling D the unique barcode in the image of $B|_W$, we can choose one pre-image \tilde{D} of D in one of the spaces of ordered barcodes \mathbb{R}^{2m+n} , then take \tilde{B} to be the constant map $W \rightarrow \{\tilde{D}\}$.

Corollary 3.20 *The smooth maps in **Diffeo** from a smooth manifold \mathcal{M} without boundary (equipped with the diffeology from Theorem 2.18) to the diffeological space Bar are exactly the ∞ -differentiable maps from \mathcal{M} to Bar .*

Proof Let $B : \mathcal{M} \rightarrow Bar$ be a smooth map in **Diffeo**. For any plot $\phi : U \rightarrow \mathcal{M}$, the composition $B \circ \phi$ is a plot in $\mathcal{D}(Bar)$, therefore it locally rewrites as $Q_{m,n} \circ \tilde{B}$ for some C^∞ lift \tilde{B} , by Proposition 3.19. Choosing ϕ to be a local coordinate chart, we then locally have $B = Q_{m,n} \circ \tilde{B} \circ \phi^{-1}$, which means that B is ∞ -differentiable. Conversely, if B is ∞ -differentiable, it locally rewrites as $B = Q_{m,n} \circ \tilde{B}$, hence for any plot $\phi : U \rightarrow \mathcal{M}$ the composition $B \circ \phi$ locally rewrites as $Q_{m,n} \circ \tilde{B} \circ \phi$ and therefore is a plot in $\mathcal{D}(Bar)$ by Proposition 3.19. \square

Dually:

Corollary 3.21 *The smooth maps in **Diffeo** from the diffeological space Bar to a smooth manifold \mathcal{N} without boundary (equipped with the diffeology from Theorem 2.18) are exactly the ∞ -differentiable maps from Bar to \mathcal{N} .*

Proof Let $V : Bar \rightarrow \mathcal{N}$ be a smooth map in **Diffeo**. By Proposition 3.19, any ∞ -differentiable map $B : U \rightarrow Bar$ defined on a domain U is a plot; therefore, the composition $V \circ B : U \rightarrow \mathcal{N}$ is a plot hence C^∞ . In particular, the map $Q_{m,n} = Q_{m,n} \circ Id_{\mathbb{R}^{2m+n}} : \mathbb{R}^{2m+n} \rightarrow Bar$ is ∞ -differentiable; therefore, $V \circ Q_{m,n}$ is C^∞ . This shows that V is ∞ -differentiable. Conversely, if V is ∞ -differentiable, the maps $V \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathcal{N}$, for varying integers m, n , are C^∞ . By Proposition 3.19, if $B : U \rightarrow Bar$ is a plot, then it locally rewrites as $Q_{m,n} \circ \tilde{B}$ for some C^∞ lift \tilde{B} , therefore $V \circ B$ is locally of the form $(V \circ Q_{m,n}) \circ \tilde{B}$, which is of class C^∞ as a map between manifolds by the chain rule. Thus, $V \circ B$ is a plot, and therefore V is smooth in **Diffeo**. \square

Conceptually, we have made Bar into a diffeological space by viewing it as the quotient of the direct limit of the spaces of ordered barcode. Then, ∞ -differentiable maps are simply morphisms in **Diffeo** from or to smooth manifolds, rather than maps satisfying the a priori unrelated Definitions 3.3 and 3.10. More generally, by seeing Bar as one object in **Diffeo** where morphisms can come in or out, we have notions of smooth maps from or to Bar with respect to any other diffeological space. For instance, a map $f : Bar \rightarrow Bar$ is smooth if and only if all the maps $f \circ Q_{m,n}$, for varying integers m, n , are ∞ -differentiable (the proof is left as an exercise to the reader). Note, however, that diffeology does not characterize the r -differentiable maps for finite r nor the maps that are differentiable only locally, two concepts that are prominent in our analysis.

4 The Case of Barcode Valued Maps Derived from Real Functions on a Simplicial Complex

In this section, we consider barcode valued maps $B_p : \mathcal{M} \rightarrow Bar$ that factor through the space \mathbb{R}^K of real functions on a fixed finite abstract simplicial complex K :

$$B_p : \mathcal{M} \xrightarrow{F} \mathbb{R}^K \xrightarrow{\text{Dgm}_p} \text{Bar}$$

In other words, we consider barcodes derived from real functions on K . Note that Dgm_p , the barcode map in degree p , is only defined on the subspace of filter functions, i.e., functions $K \rightarrow \mathbb{R}$ that are monotonous with respect to inclusions of faces in K . This subspace is a convex polytope bounded by the hyperplanes of equations $f(\sigma) = f(\sigma')$ for $\sigma \subsetneq \sigma' \in K$. From now on, we consistently assume that F takes its values in this polytope.

Example 4.1 (*Height filters*) Given an embedded simplicial complex $K \subseteq \mathbb{R}^d$, let $\mathcal{M} = \mathbb{S}^{d-1}$ and $F : \theta \mapsto (\sigma \in K \mapsto \max_{x \in \sigma} \langle \theta, x \rangle)$. The filter functions considered here are the height functions on K , parametrized on the unit sphere \mathbb{S}^{d-1} by the map F .

By analogy with the previous example, we generally call F the *parametrization* associated with B , although it may not always be a topological embedding of \mathcal{M} into \mathbb{R}^K (it may not even be injective). We also call \mathcal{M} the *parameter space* and use the generic notation θ to refer to an element in \mathcal{M} .

As we shall see in Sect. 4.1, a local coordinate system for the map B_p at $\theta \in \mathcal{M}$ can be derived when the order of the values of the filter function $F(\theta)$ remains constant locally around θ . For this purpose, we introduce the following equivalence relation on filter functions $K \rightarrow \mathbb{R}$:

Definition 4.2 Given a filter function $f : K \rightarrow \mathbb{R}$, the increasing order of its values induce a pre-order on the simplices of K . Two filter functions f, g are said to be *ordering equivalent*, written $f \sim g$, if they induce the same pre-order on K . This relation is an equivalence relation on filter functions, and we denote by $[f]$ the equivalence class of f . The (finite) set of equivalence classes is denoted by $\Omega(\mathbb{R}^K)$.

In order to compare barcodes across an entire equivalence class of functions, we introduce *barcode templates* as follows:

Definition 4.3 Given a filter function $f \in \mathbb{R}^K$ and a homology degree $0 \leq p \leq d$, a *barcode template* (P_p, U_p) is composed of a multi-set P_p of pairs of simplices in K , together with a multi-set U_p of simplices in K , such that:

$$\text{Dgm}_p(f) = \{(f(\sigma), f(\sigma'))\}_{(\sigma, \sigma') \in P_p} \cup \{(f(\sigma), +\infty)\}_{\sigma \in U_p} \cup \Delta^\infty \tag{8}$$

Note that we do not require a priori that $\dim \sigma = p$ and $\dim \sigma' = p + 1$.

Proposition 4.4 For any filter function $f \in \mathbb{R}^K$ and homology degree $0 \leq p \leq d$, there exists a barcode template (P_p, U_p) of f .

Proof Consider the interval decomposition $\mathbf{H}_p(f) \cong \bigoplus_{J \in \mathcal{J}} \mathbb{I}_J$ of the p -th persistent homology module of f . Note that every interval endpoint in the decomposition corresponds to the f -value of some simplex of K (since the persistent homology module has internal isomorphisms in-between these values). For every bounded interval J with endpoints $b, d \in \mathbb{R}$ choose an element (σ_J, σ'_J) in $f^{-1}(b) \times f^{-1}(d) \subseteq K \times K$,

then form the multi-set $P_p := \{(\sigma_J, \sigma'_J) \mid J \in \mathcal{J} \text{ bounded}\}$. Meanwhile, for every unbounded interval J with finite endpoint $v \in \mathbb{R}$ choose an element σ_J in $f^{-1}(v)$, then form the multi-set $U_p := \{\sigma_J \mid J \in \mathcal{J} \text{ unbounded}\}$. \square

Barcode templates get their name from the fact that they are an invariant of the ordering equivalence relation \sim :

Proposition 4.5 *If f, f' are ordering equivalent filter functions, then any barcode template of f is also a barcode template of f' and vice-versa.*

The proof, detailed hereafter, relies on the following elementary lemma.

Lemma 4.6 *Let \mathbb{V} be a persistence module, and $h : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous increasing function. Denote by \mathbb{V}_h the shift of \mathbb{V} by h , i.e., for any $s \leq t$, $\mathbb{V}_{h,t} := \mathbb{V}_{h(t)}$ and $v_{s,t}^{\mathbb{V}_h} := v_{h(s),h(t)}^{\mathbb{V}}$. If \mathbb{V} decomposes as $\mathbb{V} \cong \bigoplus_{J \in \mathcal{J}} \mathbb{I}_J$, then $\mathbb{V}_h \cong \bigoplus_{J \in \mathcal{J}} \mathbb{I}_{h^{-1}(J)}$.*

Proof The operation that takes a persistence module to its shift by h is an endofunctor of **Pers** which commutes with direct sums. In particular, it preserves isomorphisms. \square

Proof of Proposition 4.5 Let f, f' be two ordering equivalent filter functions. Since $f \sim f'$, we have $f(\sigma) = f(\sigma') \Rightarrow f'(\sigma) = f'(\sigma')$ for any pair of simplices $\sigma, \sigma' \in K$. Therefore, the map $h : f(\sigma) \in f(K) \mapsto f'(\sigma) \in f'(K)$ is well defined. Furthermore, h is an increasing function and we extend it monotonously and continuously over all \mathbb{R} . Then, by the reparametrization Lemma 4.6, any barcode template of f is also a barcode template of f' . \square

4.1 Generic Smoothness of the Barcode Valued Map

We now state our first significant results (one local and the other global) about the differentiability of the map B_p in the context of this section. Equipping \mathbb{R}^K with the usual Euclidean norm, we assume that the parametrization F is of class C^r as a map $\mathcal{M} \rightarrow \mathbb{R}^K$. Under this hypothesis, we show that B_p is r -differentiable in the sense of Definition 3.3 on a generic (open and dense) subset of \mathcal{M} . The intuition behind these results is that, whenever the filter functions $F(\theta')$ are all ordering equivalent in a neighborhood of θ , we can pick a barcode template that is consistent across all filter functions $F(\theta')$ in this neighborhood (by Propositions 4.4 and 4.5) and Eq. (8) then behaves like a local coordinate system for B at θ .

Here is our local result:

Theorem 4.7 (Local discrete smoothness) *Let $\theta \in \mathcal{M}$. Suppose the parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ is of class C^r ($r \geq 0$) on some open neighborhood U of θ and that $F(\theta') \sim F(\theta)$ for all $\theta' \in U$. Then, B_p is r -differentiable at θ .*

Proof Note that, as an open set, U is an open submanifold of \mathcal{M} of same dimension. By Proposition 4.4, we can pick a barcode template (P_p, U_p) for $F(\theta)$. By Proposition 4.5, this barcode template is consistent for all $F(\theta')$ where $\theta' \in U$. Therefore, we can locally write:

$$\forall \theta' \in U, B_p(\theta') = \{(F(\theta')(\sigma), F(\theta')(\sigma'))\}_{(\sigma, \sigma') \in P_p} \cup \{(F(\theta')(\sigma), +\infty)\}_{\sigma \in U_p} \cup \Delta^\infty$$

which is a local coordinate system for B_p at θ . This local coordinate system is C^r because F itself is C^r over U . As a result, B_p is r -differentiable at θ , by Proposition 3.8. \square

Corollary 4.8 *Let $\theta \in \mathcal{M}$. Suppose that the parametrization F is of class C^r ($r \geq 0$) on some open neighborhood of θ , and that the filter function $F(\theta)$ is injective. Then, B_p is r -differentiable at θ .*

Proof For such a θ , all the quantities $F(\theta)(\sigma) - F(\theta)(\sigma')$ for $\sigma \neq \sigma' \in K$ are either strictly positive or strictly negative. Therefore, by continuity they keep their sign in an open neighborhood of θ , over which all filter functions are thus ordering equivalent. The result follows then from Theorem 4.7. \square

Here is our global result:

Theorem 4.9 (Global discrete smoothness) *Suppose the parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ is continuous over \mathcal{M} and of class C^r ($r \geq 0$) on some open subset U of \mathcal{M} . Then, B_p is r -differentiable on the set $U \cap \tilde{\mathcal{M}}$, where*

$$\tilde{\mathcal{M}} := \{\theta \in \mathcal{M} \mid \exists \text{ open neighborhood } U_\theta \text{ of } \theta \text{ s.t. } F(\theta') \sim F(\theta) \text{ for all } \theta' \in U_\theta\}, \tag{9}$$

which is generic (i.e., open and dense) in \mathcal{M} . In particular, if F is C^r on some generic subset of \mathcal{M} in the first place, then so is B_p (on some possibly smaller generic subset).

Proof Observe that $\tilde{\mathcal{M}}$ is open in \mathcal{M} . As a consequence, for every $\theta \in U \cap \tilde{\mathcal{M}}$ there is some open neighborhood on which F is C^r and all the filter functions $F(\theta')$ are ordering equivalent, which by Theorem 4.7 implies that B_p is r -differentiable at θ . Thus, all that remains to be shown is that $\tilde{\mathcal{M}}$ is dense in \mathcal{M} , which is the subject of Lemma 4.10 below. \square

Lemma 4.10 *If a parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ is continuous, then the set $\tilde{\mathcal{M}}$ (as defined in Eq. (9)) is dense in \mathcal{M} .*

Proof Let $h : \mathcal{M} \rightarrow \mathbb{R}$ be a continuous function. Consider the boundary of the zero-level set $h^{-1}(0)$:

$$\partial h^{-1}(0) = \overline{h^{-1}(0)} \setminus (h^{-1}(0))^o.$$

Since h is continuous, $h^{-1}(0)$ is closed in \mathcal{M} , therefore $\partial h^{-1}(0)$ is closed with empty interior, i.e., its complement $(\partial h^{-1}(0))^c$ in \mathcal{M} is open and dense.

Consider now the case of function $h_{\sigma, \sigma'} : \theta \in \mathcal{M} \mapsto F(\theta)(\sigma) - F(\theta)(\sigma') \in \mathbb{R}$ for some fixed simplices $\sigma \neq \sigma'$ of K . The map $h_{\sigma, \sigma'}$ is continuous by continuity of the parametrization F ; therefore, the previous paragraph implies that $(\partial h_{\sigma, \sigma'}^{-1}(0))^c$ is generic in \mathcal{M} . Hence, the finite intersection

$$\hat{\mathcal{M}} := \bigcap_{\sigma \neq \sigma' \in K} (\partial h_{\sigma, \sigma'}^{-1}(0))^c$$

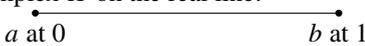
is also generic in \mathcal{M} . We now show that $\hat{\mathcal{M}}$ is a subspace of $\tilde{\mathcal{M}}$.

Let $\theta \in \hat{\mathcal{M}}$ and $\sigma \neq \sigma' \in K$. If $h_{\sigma,\sigma'}(\theta) > 0$, then by continuity we have $h_{\sigma,\sigma'} > 0$ over some open neighborhood $V_{\sigma,\sigma'}$ of θ . Similarly, if $h_{\sigma,\sigma'}(\theta) < 0$. And if $h_{\sigma,\sigma'}(\theta) = 0$, then, since $\theta \in \hat{\mathcal{M}}$, θ lies in the interior of the level set $h_{\sigma,\sigma'}^{-1}(0)$, and therefore there is also an open neighborhood $V_{\sigma,\sigma'}$ of θ over which $h_{\sigma,\sigma'} = 0$. Let V be the finite intersection $\bigcap_{\sigma \neq \sigma' \in K} V_{\sigma,\sigma'}$, which is open and non-empty in \mathcal{M} . For every $\sigma \neq \sigma' \in K$, the sign $F(\theta')(\sigma) - F(\theta')(\sigma')$ is constant over all $\theta' \in V$, where by sign we really distinguish between three possibilities: negative, positive, null. Therefore, the pre-order on the simplices of K induced by $F(\theta')$ is constant over the $\theta' \in V$. In other words, all the $F(\theta')$ are ordering equivalent. Therefore, $\theta \in \tilde{\mathcal{M}}$. Since this is true for any $\theta \in \hat{\mathcal{M}}$, we conclude that $\hat{\mathcal{M}} \subseteq \tilde{\mathcal{M}}$, and so the latter is also dense in \mathcal{M} . □

Example 4.11 (Height functions again) Let us reconsider the scenario of Example 4.1. The parametrization F of height filters is C^0 on the entire sphere \mathbb{S}^{d-1} . Moreover, F is smooth at every direction $\theta \in \mathbb{S}^{d-1}$ that is not orthogonal to some difference $v - v'$ of vertices $v \neq v' \in K_0$ in \mathbb{R}^d . The set U of such directions is generic in \mathbb{S}^{d-1} ; therefore, B_p is ∞ -differentiable over the generic subset $U \cap \tilde{\mathbb{S}}^{d-1}$ by Theorem 4.9, with $\tilde{\mathbb{S}}^{d-1}$ defined as in Eq. (9). In fact, we have $U \cap \tilde{\mathbb{S}}^{d-1} = U$ in this case. Indeed, for any direction $\theta \in U$, the values of the height function h_θ at the vertices of K are pairwise distinct, and by continuity this remains true in a neighborhood of θ . The pre-order on the simplices of K induced by the height function is then constant over this neighborhood.

In Theorems 4.7 and 4.9, one cannot avoid the condition that filter functions are locally ordering equivalent. Indeed, in the next examples, we highlight that there is generally no hope for the barcode valued map B_p to be differentiable everywhere, even if the parametrization F is. This is because, essentially, the time of appearance of a simplex is a maximum of smooth functions, which can be non-smooth at a point where two functions achieve the maximum. The condition that the induced pre-order is locally constant around θ is only a sufficient condition though, because a maximum of two smooth functions can still be smooth at a point where the maximum is attained by the two functions. We provide a second example to illustrate this fact.

Example 4.12 (Singular parameter) Let us consider the following geometric simplicial complex K on the real line:



That is, K has vertices $K_0 = \{a, b\}$ with respective coordinates $\{0, 1\}$, and edges $K_1 = \{ab\}$. Consider the parametrization that filters the complex according to the squared Euclidean distance to a point, i.e., $F : \theta \in \mathbb{R} \mapsto (\sigma \in K \mapsto \max_{x \in \sigma} (x - \theta)^2)$. The map B_0 is then essentially a real function that tracks the squared Euclidean distance of the vertex closest to θ , specifically:

$$B_0(\theta) = \{(\min(\theta^2, (1 - \theta)^2), +\infty)\} \cup \Delta^\infty.$$

Hence, B_0 is not differentiable at $\theta = \frac{1}{2}$ since $\frac{1}{2}$ is a singular point of the map $\theta \mapsto \min(\theta^2, (1 - \theta)^2)$. Meanwhile, for $\theta < \frac{1}{2}$, we have $F(\theta)(a) < F(\theta)(b)$,

whereas whenever $\theta > \frac{1}{2}$, we have $F(\theta)(a) > F(\theta)(b)$. In particular, the pre-order induced by the filter functions $F(\theta)$ is not constant around $\theta = \frac{1}{2}$, and so $\frac{1}{2} \notin \tilde{\mathbb{R}}$.

Example 4.13 (*Only sufficient condition*) We remove the edge ab from the geometric complex K in the previous example, and we see the points a and b as lying on the x -axis of \mathbb{R}^2 . Consider the parametrization of height filters $F : \theta \in \mathbb{S}^1 \mapsto (\sigma \in K \mapsto \max_{x \in \sigma} \langle \theta, x \rangle)$. The map B_p is then trivial for each degree p except 0, where it writes as follows:

$$B_0(\theta) = \{(\langle \theta, a \rangle, +\infty), (\langle \theta, b \rangle, +\infty)\} \cup \Delta = \{(0, +\infty), (\langle \theta, (1, 0) \rangle, +\infty)\} \cup \Delta^\infty.$$

We see that we have a valid local coordinate system given by the two smooth maps $\theta \mapsto 0$ and $\theta \mapsto \langle \theta, (0, 1) \rangle$, so the map B_0 is ∞ -differentiable everywhere on \mathbb{S}^1 by Proposition 3.8. Meanwhile, we have $F(\theta)(a) < F(\theta)(b)$ whenever $\langle \theta, (1, 0) \rangle > 0$, and $F(\theta)(a) > F(\theta)(b)$ whenever $\langle \theta, (1, 0) \rangle < 0$, therefore the pre-order induced by the filter functions $F(\theta)$ is not constant around $\theta = (0, 1)$ and $v = (0, -1)$, hence $(0, 1), (0, -1) \notin \tilde{\mathbb{R}}$.

4.2 Differential of the Barcode Valued Map

Given a continuous parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ of class C^1 on some open set $U \subseteq \mathcal{M}$, Theorem 4.9 guarantees that a barcode template, through Equation (8), provides a C^1 local coordinate system for B_p around each point $\theta \in U \cap \tilde{\mathcal{M}}$. In turn, by Proposition 3.8, any arbitrary ordering on the functions of this local coordinate system induces a C^1 local lift of B_p . Hence, we have the following formula for the corresponding differential:

Proposition 4.14 *Given $\theta \in U \cap \tilde{\mathcal{M}}$ and a barcode template (P_p, U_p) of $F(\theta)$, for any choice of ordering $(\sigma_1, \sigma'_1), \dots, (\sigma_m, \sigma'_m), \tau_1, \dots, \tau_n$ of (P_p, U_p) , the map*

$$\tilde{B}_p : \theta' \mapsto \left[(F(\theta')(\sigma_i), F(\theta')(\sigma'_i))_{i=1}^m, (F(\theta')(\tau_j))_{j=1}^n \right]$$

is a local C^1 lift of B_p around θ , and the corresponding differential for B_p at θ is:

$$d_{\theta, \tilde{B}_p} B_p(\cdot) = \left[(d_\theta F(\cdot)(\sigma_i), d_\theta F(\cdot)(\sigma'_i))_{i=1}^m, (d_\theta F(\cdot)(\tau_j))_{j=1}^n \right].$$

Remark 4.15 (*Algorithm for computing derivatives*) Suppose we are given a parametrization F whose differential we can compute. Let $\theta \in \mathcal{M}$. If the barcode of $F(\theta)$ is given to us, then the proof of Proposition 4.4 provides an algorithm to build a barcode template (P_p, U_p) for $F(\theta)$. If the barcode of $F(\theta)$ is not given in the first place, then the matrix reduction algorithm for computing persistence [23,49] outputs both the barcode and a barcode template. In both scenarios, Proposition 4.14 gives a formula to compute a differential of B_p at θ from the barcode template (P_p, U_p) . The optimization pipelines mentioned in Introduction [3,14,27,30,43] apply this strategy to compute differentials.

4.3 Directional Differentiability of the Barcode Valued Map Along Strata

In this section, we define directional derivatives for the barcode valued map $B_p : \mathcal{M} \rightarrow \text{Bar}$ at points where it may not be differentiable in the sense of Definition 3.3. For this, we stratify the parameter space \mathcal{M} in such a way that B_p is differentiable on the top-dimensional strata, then we define its derivatives on lower-dimensional strata via directional lifts. Intuitively, the strata in \mathcal{M} are prescribed by the ordering equivalence classes in \mathbb{R}^K , as we know from Theorem 4.7 that the pre-order on simplices plays a key role in the differentiability of B_p .

Formally, consider the stratification of \mathbb{R}^K formed by the collection $\Omega(\mathbb{R}^K)$ of ordering equivalence classes. This is a Whitney stratification, obtained by cutting \mathbb{R}^K with the hyperplanes $\{f(\sigma) = f(\sigma')\}$ for varying simplices $\sigma \neq \sigma' \in K$. We look for stratifications of \mathcal{M} that make the parametrization F weakly stratified (in the sense of Definition 2.22) and smooth on each stratum. Here are typical scenarios where such stratifications exist:

Proposition 4.16 *Let $F : \mathcal{M} \rightarrow \mathbb{R}^K$ be a continuous parametrization. Suppose that, either*

- (i) \mathcal{M} is a semi-algebraic set in \mathbb{R}^N and F is a semi-algebraic map, or
- (ii) \mathcal{M} is a compact subanalytic set in a real analytic manifold and F is a subanalytic map.

Then, there is a Whitney stratification of \mathcal{M} , made of semi-algebraic (resp. subanalytic) strata, such that F is weakly stratified with C^∞ restrictions to each stratum.

Proof This is Section I.1.7 of [28], after observing that the stratification $\Omega(\mathbb{R}^K)$ is made of semi-algebraic strata. □

Example 4.17 We consider the parametrization F of height filters on the sphere \mathbb{S}^{d-1} from Example 4.11. By Proposition 4.16, there is a stratification of \mathbb{S}^{d-1} that makes F weakly stratified and C^∞ on each stratum. To be more specific, such a stratification is obtained by taking the pre-images⁵ of the strata of $\Omega(\mathbb{R}^K)$ via F . Figure 1 illustrates the result in the case $d = 3$, where the obtained stratification of \mathbb{S}^2 is made of an arrangement of great circles, each circle being the pre-image of a set $\{F(\theta)(v) = F(\theta)(v')\}$ for vertices $v \neq v'$.

Once a stratification $\mathcal{S}_\mathcal{M}$ of \mathcal{M} is given, we can introduce a notion of derivative for B_p at $\theta \in \mathcal{M}$ in the direction of an *incident* stratum \mathcal{M}' , i.e., a stratum whose closure in \mathcal{M} contains θ .

Definition 4.18 Let $B : \mathcal{M} \rightarrow \text{Bar}$ be a map defined on a stratified space $(\mathcal{M}, \mathcal{S}_\mathcal{M})$. Let $\theta \in \mathcal{M}$, and let $\mathcal{M}' \in \mathcal{S}_\mathcal{M}$ be a stratum incident to θ . The map B is *r-differentiable* at θ along \mathcal{M}' if there is an open neighborhood U of θ in \mathcal{M} and a C^r map $\tilde{B} : U \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ for some integers m, n such that $B = Q_{m,n} \circ \tilde{B}$ on $U \cap \mathcal{M}'$. The differential $d_\theta \tilde{B}$ is called a *directional derivative* of B at θ along \mathcal{M}' .

⁵ This is called the *pull-back stratification*. In fact, for any smooth map $F : \mathcal{M} \rightarrow \mathbb{R}^K$ that is transverse with respect to $\Omega(\mathbb{R}^K)$ and to any stratification of \mathcal{M} (e.g., the trivial one), the pull-back of $\Omega(\mathbb{R}^K)$ via F makes the latter weakly stratified and C^∞ on each stratum—see, e.g., [28, § I.1.3].

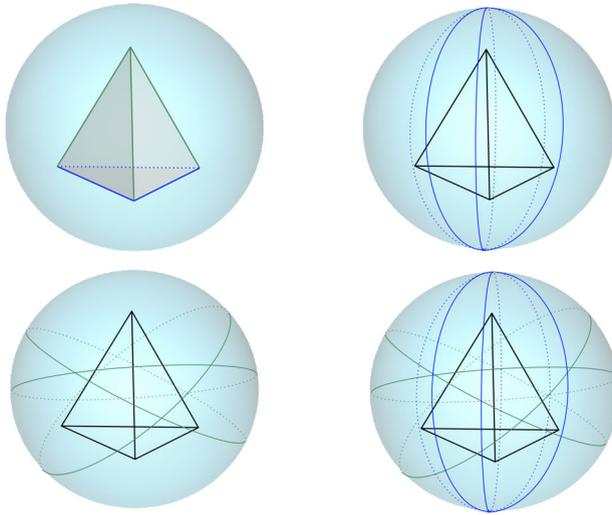


Fig. 1 The regular tetrahedron K embedded in \mathbb{R}^3 (top left) induces the stratification of \mathbb{S}^2 for the parametrization of height filters (bottom right). The intersections of great circles are 0-dimensional strata, the parts of the great circles that do not intersect with each other are 1-dimensional strata, and the rest of \mathbb{S}^2 forms the two-dimensional strata. Each great circle corresponds to unit vectors that are orthogonal to a given edge of the simplicial complex. The edges joining the vertices in the base of the tetrahedron produce the blue circles (top right), while the edges joining the apex of the tetrahedron to the base face produce the green circles (bottom left) (Color figure online)

This definition agrees with the notions of r -differentiability and derivatives introduced in Sect. 3 when \mathcal{M}' contains an open neighborhood around θ , i.e., for θ located in a top-dimensional stratum \mathcal{M}' . When θ is located in some lower-dimensional stratum, it admits finitely many incident strata \mathcal{M}' (possibly not top-dimensional), each one of which yields a specific directional derivative at θ . The definition of each derivative involves a local C^r lift \tilde{B} of B near θ in \mathcal{M}' . This lift is required to extend smoothly over an open neighborhood U in \mathcal{M} , to ensure that \tilde{B} and its derivatives have well-defined limits at θ .

Theorem 4.19 (Discrete smoothness along strata) *Let $r \in \mathbb{N}$ and $F : \mathcal{M} \rightarrow \mathbb{R}^K$. Suppose $\mathcal{S}_{\mathcal{M}}$ is a Whitney stratification of \mathcal{M} such that:*

- (i) F is a weakly stratified map with respect to $\mathcal{S}_{\mathcal{M}}$ and $\Omega(\mathbb{R}^K)$, and
- (ii) the restriction of F to each stratum of $\mathcal{S}_{\mathcal{M}}$ is C^r , and
- (iii) for every $\theta \in \mathcal{M}$ and every incident stratum $\mathcal{M}' \in \mathcal{S}_{\mathcal{M}}$, there is an open neighborhood U of θ in \mathcal{M} such that $F|_{\mathcal{M}' \cap U}$ extends to a C^r map $U \rightarrow \mathbb{R}^K$.

Then, at every $\theta \in \mathcal{M}$, the barcode valued map $B_p : \mathcal{M} \rightarrow \text{Bar}$ is r -differentiable along each stratum incident to θ . In particular, B_p is r -differentiable in the sense of Definition 3.3 inside each top-dimensional stratum.

Proof Let $\theta \in \mathcal{M}$ and \mathcal{M}' a stratum incident to θ . By (i), combined with Propositions 4.4 and 4.5, there exists a barcode template (P_p, U_p) that is consistent across all

$F(\theta')$ for $\theta' \in \mathcal{M}'$. Therefore, for all $\theta' \in \mathcal{M}'$:

$$B_p(\theta') = \left\{ (F(\theta')(\sigma), F(\theta')(\sigma')) \right\}_{(\sigma, \sigma') \in P_p} \cup \left\{ (F(\theta')(\sigma), +\infty) \right\}_{\sigma \in U_p} \cup \Delta^\infty, \tag{10}$$

which by (ii) provides a C^r local coordinate system for $B_{p|_{\mathcal{M}'}}$. Then, by Proposition 3.8, there is a C^r lift of $B_{p|_{\mathcal{M}'}}$, whose coordinate functions are of the form $\theta' \mapsto F(\theta')(\sigma)$. Using (iii), we extend each coordinate function of this lift (hence the lift itself) to an open neighborhood U of θ in \mathcal{M} . □

Combining Proposition 4.16 with Theorem 4.19 yields the following:

Corollary 4.20 *Under the hypotheses of Proposition 4.16, there is a Whitney stratification of \mathcal{M} , made of semi-algebraic (resp. subanalytic) strata, such that B_p is ∞ -differentiable on the top-dimensional strata (whose union is generic in \mathcal{M}). If furthermore F is globally C^r , then B_p is everywhere r -differentiable along incident strata.*

Example 4.21 Consider again the setup of Example 4.12. We stratify \mathbb{R} by the point $\{\frac{1}{2}\}$ and the half-lines $(-\infty; \frac{1}{2})$ and $(\frac{1}{2}; +\infty)$. The parametrization F is C^∞ and sends strata into strata; therefore, by Theorem 4.19 the barcode valued map B_0 admits directional derivatives everywhere on \mathbb{R} . More precisely, recall that we have a lift $\tilde{B}_0 : \theta \mapsto \min(\theta^2, (1 - \theta)^2)$, which is smooth in the top-dimensional strata, while at $\theta = \frac{1}{2}$ it admits directional derivatives along the two half-lines, whose values are 1 and -1 , respectively, and thus do not agree.

Example 4.22 Consider again the stratification $\mathcal{S}_{\mathbb{S}^{d-1}}$ by the great circles of the parameter space \mathbb{S}^{d-1} associated with the parametrization of height filters (Example 4.17). By Corollary 4.20, we know that there exists a refinement $\mathcal{S}'_{\mathbb{S}^{d-1}}$ of $\mathcal{S}_{\mathbb{S}^{d-1}}$ such that B_p admits directional derivatives along incident strata of $\mathcal{S}'_{\mathbb{S}^{d-1}}$ at every point $\theta \in \mathbb{S}^{d-1}$. In fact, we can even take $\mathcal{S}'_{\mathbb{S}^{d-1}}$ to be $\mathcal{S}_{\mathbb{S}^{d-1}}$ itself. Indeed, all the directions in a given stratum $\mathcal{M}' \in \mathcal{S}_{\mathbb{S}^{d-1}}$ induce the same pre-order on the simplices of K , therefore

- the restriction $F|_{\mathcal{M}'}$ is valued in a stratum of $\Omega(\mathbb{R}^K)$, and
- for every simplex $\sigma \in K$, there is a vertex $\bar{v}(\sigma)$ such that $F|_{\mathcal{M}' }(\cdot)(\sigma) = \langle \cdot, \bar{v}(\sigma) \rangle$.

Consequently, the assumptions of Theorem 4.19 hold, and the barcode valued map B_p admits directional derivatives along incident strata of $\mathcal{S}_{\mathbb{S}^{d-1}}$ at every point $\theta \in \mathbb{S}^{d-1}$.

4.4 The Barcode Valued Map as a Permutation Map

In this section, we work out a global lift of the barcode valued map, which restricts nicely to each stratum of a stratification of \mathcal{M} . To do so, we first focus on the map Dgm which, given a filter function $f \in \mathbb{R}^K$ on a fixed simplicial complex K of dimension d , returns the vector of all its barcodes $(\text{Dgm}_p(f))_{p=0}^d$. We observe that Dgm admits a global Euclidean lift, and furthermore, that this lift is essentially a permutation map on each stratum of $\Omega(\mathbb{R}^K)$. Throughout, we fix an ordering of the simplices of K , so

that the canonical basis of \mathbb{R}^K turns into a basis of $\mathbb{R}^{\#K}$, and we let $\phi : \mathbb{R}^K \rightarrow \mathbb{R}^{\#K}$ be the corresponding isomorphism.

Proposition 4.23 *There exist integers m_p, n_p for $0 \leq p \leq d$ such that $\sum_{p=0}^d (2m_p + n_p) = \#K$, and a map $\text{Perm} : \mathbb{R}^K \rightarrow \prod_{p=0}^d \mathbb{R}^{2m_p} \times \mathbb{R}^{n_p} \cong \mathbb{R}^{\#K}$ whose restriction $\text{Perm}|_S$ to each ordering equivalence class $S \in \Omega(\mathbb{R}^K)$ is a permutation matrix, and such that the following diagram commutes:⁶*

$$\begin{array}{ccc}
 \mathbb{R}^{\#K} & \xrightarrow{\text{Perm}} & \prod_{p=0}^d \mathbb{R}^{2m_p} \times \mathbb{R}^{n_p} \\
 \uparrow \phi & & \downarrow Q := \prod_{p=0}^d Q_{m_p, n_p} \\
 \mathbb{R}^K & \xrightarrow{\text{Dgm}} & \text{Bar}^{d+1}
 \end{array} \tag{11}$$

For simplicity, from now on we identify $f \in \mathbb{R}^K$ with its image in $\mathbb{R}^{\#K}$ without explicitly mentioning the map ϕ .

Proof Given a filter function $f \in \mathbb{R}^K$, we define a total barcode template (P, U) for f to be the data of $d + 1$ barcode templates (P_p, U_p) for f in each homology degree, such that each simplex of K appears exactly once, in a unique P_p or U_p . We further require that the pairs (σ, σ') appearing in P_p consist of a p -dimensional simplex σ and a $(p + 1)$ -dimensional simplex σ' , while the unpaired simplices appearing in U_p must be p -dimensional. A simplex σ is then labeled *positive* if it appears as the first component of a pair in some P_p or U_p , and *negative* otherwise.

Note that total barcode templates always exist, by an argument similar to (yet somewhat more involved than) the one used in the proof of Proposition 4.4. Alternatively, note that applying the matrix reduction algorithm for computing persistence [23,49] to the sublevel-sets filtration of f produces a total barcode template. By Proposition 4.5, total barcode templates are invariant under ordering equivalences. We therefore fix a unique total barcode template $(P(S), U(S))$ per ordering equivalence class $S \in \Omega(\mathbb{R}^K)$ (there are only finitely many such classes), and we denote by $m_p(S) := \#P_p(S), n_p(S) := \#U_p(S)$ their sizes in each homology degree p .

Since the barcode templates $(P(S), U(S))$ are total, we have $\sum_{p=0}^d (2m_p(S) + n_p(S)) = \#K$. Besides, since the number of infinite intervals in the barcode of a filter function is given by the Betti numbers of the simplicial complex K , an easy induction on the homology degree shows that the number of positive (resp. negative) simplices in each homology degree is independent of the choice of filter function and of total barcode template. Therefore, the integers $m_p(S), n_p(S)$ do not depend on the stratum S .

For each stratum $S \in \Omega(\mathbb{R}^K)$ and homology degree p , we pick arbitrary orderings $(\sigma_{k,S}, \sigma'_{k,S})_{k=0}^{m_p}$ of $P_p(S)$ and $(\tau_{k,S})_{k=0}^{n_p}$ of $U_p(S)$. Any filter function $f \in S$ admits $(P(S), U(S))$ as total barcode template, therefore we get that $\text{Dgm}_p(f) =$

⁶ Strictly speaking, like Dgm, this diagram applies only to the set of filter functions in \mathbb{R}^K .

$Q_{m_p, n_p}((f(\sigma_{k,S}), f(\sigma'_{k,S}))_{k=0}^{m_p}, (f(\tau_{k,S}))_{k=0}^{n_p})$ in every homology degree p . We simply set $\text{Perm}(f) := [(f(\sigma_{k,S}), f(\sigma'_{k,S}))_{k=0}^{m_p}, (f(\tau_{k,S}))_{k=0}^{n_p}]_{p=0}^d \in \prod_{p=0}^d \mathbb{R}^{2m_p} \times \mathbb{R}^{n_p}$, which ensures the commutativity of (11). Since each simplex of K appears exactly once in $(P(S), U(S))$, the vector $\text{Perm}(f)$ is a re-ordering of the coordinates of f (i.e., of its values on the simplices) and therefore $\text{Perm}|_S$ is a permutation matrix. \square

We now turn to the parametrized barcode valued map

$$B : \theta \in \mathcal{M} \xrightarrow{F} F(\theta) \in \mathbb{R}^K \xrightarrow{\text{Dgm}} [\text{Dgm}_p(F(\theta))]_{p=0}^d \in \text{Bar}^{d+1}$$

determined by a parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ of filter functions. We show that if \mathcal{M} admits a Whitney stratification $\mathcal{S}_{\mathcal{M}}$ satisfying the assumptions of Theorem 4.19, then B admits a global lift \tilde{B} that acts as a permutation of F -values on each stratum.

Corollary 4.24 *Using the same notations as in Proposition 4.23, the map*

$$\tilde{B} : \theta \in \mathcal{M} \mapsto \text{Perm}(F(\theta)) \in \prod_{p=0}^d \mathbb{R}^{2m_p} \times \mathbb{R}^{n_p}$$

is a global lift of B , i.e., $Q \circ \tilde{B} = B$ everywhere on \mathcal{M} . If moreover \mathcal{M} admits a Whitney stratification $\mathcal{S}_{\mathcal{M}}$ satisfying the assumptions of Theorem 4.19, then $\tilde{B} = \text{Perm}_{\mathcal{M}'} \circ F$ for some permutation matrix $\text{Perm}_{\mathcal{M}'}$ over each stratum $\mathcal{M}' \in \mathcal{S}_{\mathcal{M}}$. Consequently, B is r -differentiable along incident strata everywhere on \mathcal{M} , with directional derivatives given by the ones of \tilde{B} .

The last part of the statement expresses the fact that directional derivatives of B are simply given by permuting the directional derivatives of the coordinate functions of F .

Proof The first part of the statement is a straight consequence of Proposition 4.23. Let $\mathcal{S}_{\mathcal{M}}$ be a stratification satisfying the assumptions of Theorem 4.19. As F is weakly stratified with respect to $\mathcal{S}_{\mathcal{M}}$ and $\Omega(\mathbb{R}^K)$, it sends strata into strata and therefore by Proposition 4.23 we have $\tilde{B} = \text{Perm}_{\mathcal{M}'} \circ F$ for some permutation matrix $\text{Perm}_{\mathcal{M}'}$ over each stratum $\mathcal{M}' \in \mathcal{S}_{\mathcal{M}}$. Then, since F admits local smooth extensions over each stratum \mathcal{M}' of $\mathcal{S}_{\mathcal{M}}$, so do its coordinate functions and in turn so does $\tilde{B} = \text{Perm}_{\mathcal{M}'} \circ F$. These local extensions of \tilde{B} yield directional derivatives for B along incident strata. \square

Remark 4.25 Recall that the map Perm is a linear map when restricted to the strata of $\Omega(\mathbb{R}^K)$, which are simply polyhedra in \mathbb{R}^K . Therefore, if \mathcal{M} is a semi-algebraic set (resp. subanalytic set or definable set in an o-minimal structure) and F is a semi-algebraic (resp. subanalytic or definable) map, then the global lift $\tilde{B} = \text{Perm} \circ F$ of Corollary 4.24 is itself a semi-algebraic (resp. subanalytic or definable) map. Thus, we recover Proposition 3.2 and Corollary 3.3 of [9]. Meanwhile, the differentiability of \tilde{B} on top-dimensional strata (as per Corollary 4.20) recovers their Proposition 3.4.

We conclude this section with a side result whose proof (deferred to “Appendix A”) relies on Proposition 4.23. This result states that Dgm is locally an isometry on top-dimensional strata of $\Omega(\mathbb{R}^K)$. It involves the distance $d_0(f)$ of any filter function $f \in$

\mathbb{R}^K to the union of strata of $\Omega(\mathbb{R}^K)$ of codimension at least 1:

$$d_0(f) = \frac{1}{2} \min_{\sigma \neq \sigma'} |f(\sigma) - f(\sigma')|.$$

Proposition 4.26 *Let $f, g \in \mathbb{R}^K$ be two filter functions that are located in the closure of a common top-dimensional stratum $S \in \Omega(\mathbb{R}^K)$. Then:*

$$\max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) \geq \min(\|f - g\|_\infty, \max(d_0(f), d_0(g))). \tag{12}$$

In particular, for any filter function $f \in \mathbb{R}^K$ located in a top-dimensional stratum, the map Dgm is a local isometry in a closed ball of radius $d_0(f)$ around f , specifically:

$$\begin{aligned} \forall g \in \mathbb{R}^K, \|f - g\|_\infty \leq d_0(f) &\implies \\ \max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) &= \|f - g\|_\infty \end{aligned} \tag{13}$$

$$\begin{aligned} \forall g, h \in \mathbb{R}^K, \max(\|f - g\|_\infty, \|f - h\|_\infty) \leq \frac{d_0(f)}{3} &\implies \\ \max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(g), \text{Dgm}_p(h)) &= \|g - h\|_\infty. \end{aligned} \tag{14}$$

5 Application to Common Simplicial Filtrations

In this section, we leverage Theorems 4.7 and 4.9 in the case of a few important classes of parametrizations of filter functions on a simplicial complex K of dimension d . In each case, we derive a characterization of the parameter values where B_p is differentiable, and whenever possible we provide an explicit differential of B_p using Proposition 4.14. In the following, we fix a homology degree $0 \leq p \leq d$.

5.1 Lower Star Filtrations

Parametrizations of lower star filtrations are involved in most practical scenarios [3, 14,29,30,43]; here, we provide a common analysis of their differentiability.

Definition 5.1 Given a function $f : K_0 \rightarrow \mathbb{R}$ defined on the vertices of K , we extend it to each simplex σ of K by its highest value on the vertices of σ . The sub-level sets of this function together form the *lower-star filtration* of K induced by f .

One interest of lower-star filtrations is that any parametrization $\mathcal{M} \rightarrow \mathbb{R}^{K_0}$ on the vertex set of K induces a valid parametrization $\mathcal{M} \rightarrow \mathbb{R}^K$ on K itself. Sufficient conditions for the differentiability of such parametrizations are easy to work out thanks to the following observation:

Proposition 5.2 *Let $F_0 : \mathcal{M} \rightarrow \mathbb{R}^{K_0}$ be a C^r parametrization of filter functions on the vertices of K . Then, the induced parametrization $F : \mathcal{M} \rightarrow \mathbb{R}^K$ is C^r at each $\theta \notin \text{Sing}(F_0)$, where $\text{Sing}(F_0)$ is the boundary of the set:*

$$\{\theta \in \mathcal{M}, \exists(v, v') \in K_0, F_0(\theta)(v) = F_0(\theta)(v')\}.$$

Specifically, for every $\theta \notin \text{Sing}(F_0)$, letting

$$\bar{v} : \sigma \in K \mapsto \operatorname{argmax}_{v \text{ vertex in } \sigma} F_0(\theta)(v) \in K_0$$

by breaking ties wherever necessary, there is an open neighborhood U of θ such that $F(\theta')(\sigma) = F_0(\theta')(\bar{v}(\sigma))$ for every $\theta' \in U$ and $\sigma \in K$, from which follows that F is C^r at θ .

Proof The continuity of F comes from the continuity of F_0 and of the max function. If $\theta \in \mathcal{M} \setminus \text{Sing}(F_0)$, then the pre-order on K_0 induced by $F_0(\cdot)$ is constant in an open neighborhood U of θ . We want to check that F is C^r at θ , i.e., that all maps $\theta' \mapsto F(\theta')(\sigma)$ are C^r at θ , for a fixed simplex $\sigma \in K$. For σ a vertex of K , this is true by assumption because $F(\cdot)(\sigma) = F_0(\cdot)(\sigma)$. For an arbitrary simplex σ , $F(\cdot)(\sigma) = \max_{v \text{ vertex in } \sigma} F_0(\cdot)(v)$. Since the pre-order induced on K_0 by F_0 is constant over U , the maximum above is attained at vertex $\bar{v}(\sigma)$, and this fact holds for all θ' in U . Thus, $F(\cdot)(\sigma)|_U = F_0(\cdot)(\bar{v}(\sigma))|_U$, which allows us to conclude. \square

Remark 5.3 Recall that $\text{Sing}(F_0)$ is by definition the boundary of $\{\theta \in \mathcal{M}, \exists(v, v') \in K_0, F_0(\theta)(v) = F_0(\theta)(v')\}$, whose complement may not be generic (in fact it may even be empty, e.g., when $F_0 = 0$). This shows the interest of working with locally constant pre-orders on vertices, and not just with locally injective parametrizations as in the works of [3,14,29,30,43].

Defining $\text{Sing}(F_0)$ and \bar{v} as in Proposition 5.2, and combining this result with Proposition 4.14, we deduce the following result on the differentiability of B_p , which only relies on the differentiability of F_0 :

Corollary 5.4 *For any C^r parametrization $F_0 : \mathcal{M} \rightarrow \mathbb{R}^{K_0}$ on the vertices of K , the induced barcode valued map $B_p : \theta \in \mathcal{M} \mapsto \text{Dgm}_p(F(\theta)) \in \text{Bar}$ is r -differentiable outside $\text{Sing}(F_0)$. Moreover, at $\theta \in \mathcal{M} \setminus \text{Sing}(F_0)$, for any barcode template (P_p, U_p) of $F(\theta)$ and any choice of ordering $(\sigma_1, \sigma'_1), \dots, (\sigma_m, \sigma'_m), \tau_1, \dots, \tau_n$ of (P_p, U_p) , the map $\tilde{B}_p : \mathcal{M} \rightarrow \mathbb{R}^m \times \mathbb{R}^n$ defined by:*

$$\tilde{B}_p : \theta' \mapsto \left[(F_0(\theta')(\bar{v}(\sigma_i)), F_0(\theta')(\bar{v}(\sigma'_i)))_{i=1}^m, (F_0(\theta')(\bar{v}(\sigma'_j)))_{j=1}^n \right]$$

is a local C^r lift of B_p around θ . The corresponding differential for B_p at θ is:

$$d_{\theta, \tilde{B}_p} B_p(\cdot) = \left[(d_{\theta} F_0(\cdot)(\bar{v}(\sigma_i)), d_{\theta} F_0(\cdot)(\bar{v}(\sigma'_i)))_{i=1}^m, (d_{\theta} F_0(\cdot)(\bar{v}(\sigma'_j)))_{j=1}^n \right].$$

Proof For $\theta \in \mathcal{M} \setminus \text{Sing}(F_0)$, the pre-order on the vertices K_0 induced by F_0 is constant in an open neighborhood U of θ . By Proposition 5.2, each $F(\theta')(\sigma)$ rewrites as $F_0(\theta')(\bar{v}(\sigma))$ for $\theta' \in U$, which implies that the pre-order on the simplices of K induced by F is also constant over U . The fact that B_p is r -differentiable at θ follows then from Theorem 4.7, since F itself is C^r on an open neighborhood of θ (again by Proposition 5.2, and by the fact that $\text{Sing}(F_0)$ is closed). The rest of the corollary is an immediate consequence of Proposition 4.14. \square

Example 5.5 Consider our running example of parametrization of height filtrations $F_0(\theta) = h_\theta : v \in K_0 \mapsto \langle v, \theta \rangle \in \mathbb{R}$, where K is a fixed geometric simplicial complex in \mathbb{R}^d and $\theta \in \mathbb{S}^{d-1}$. In this case, we know from Example 4.11 that B_p is generically ∞ -differentiable. Corollary 5.4 provides another proof of this fact: since F_0 is C^∞ , B_p is ∞ -differentiable outside $\text{Sing}(F_0)$, which has generic complement in \mathbb{S}^{d-1} . Moreover, the components of the differential of B_p at $\theta \in \mathbb{S}^{d-1} \setminus \text{Sing}(F_0)$ are the $d_\theta F_0(\cdot)(v)$, whose corresponding gradients (in the tangent space $T_\theta \mathbb{S}^{d-1}$ equipped with the Riemannian structure inherited from \mathbb{R}^d) are $v - \langle v, \theta \rangle \theta$.

5.2 Rips Filtrations of Point Clouds

Given a finite point cloud $P = (p_1, \dots, p_n) \in \mathbb{R}^{nd}$, the *Rips* filtration of P is a filtration of the total complex $K := 2^{\{1, \dots, n\}} \setminus \{\emptyset\}$ with $n := \#P$ vertices, where the time of appearance of a simplex $\sigma \subseteq \{1, \dots, n\}$ is $\max_{i, j \in \sigma} \|p_i - p_j\|_2$. [27] optimize the positions of the points of P in \mathbb{R}^d so that the barcode of the Rips filtration reaches some target barcode. Here, we see \mathbb{R}^{nd} as our parameter space \mathcal{M} , and we consider the parametrization

$$F(P)(\sigma) := \max_{i, j \in \sigma} \|p_i - p_j\|_2.$$

The differentiability result of [27] can be expressed as a result on the differentiability of the barcode-valued map $B_p = \text{Dgm}_p \circ F$ using our framework. We require that the points of P lie in general position as defined hereafter:

Definition 5.6 [27] P is in *general position* if the following two conditions hold:

- (i) $\forall i \neq j \in \{1, \dots, n\}, p_i \neq p_j$;
- (ii) $\forall \{i, j\} \neq \{k, l\}, \text{ where } i, j, k, l \in \{1, \dots, n\}, \|p_i - p_j\|_2 \neq \|p_k - p_l\|_2$.

We denote by $\tilde{\mathcal{P}} \subseteq \mathbb{R}^{nd}$ the subspace of point clouds in general position.

Proposition 5.7 $\tilde{\mathcal{P}}$ is generic in \mathbb{R}^{nd} .

Proof The set of point clouds P such that $p_i \neq p_j$ for all $1 \leq i \neq j \leq n$ is clearly generic in \mathbb{R}^{nd} . Moreover, the maps $P = (p_1, \dots, p_n) \mapsto \|p_i - p_j\|_2^2 - \|p_k - p_l\|_2^2$ are smooth everywhere and are submersions on a generic subset of \mathbb{R}^{nd} ; therefore, their 0-sets have generic complements, whose (finite) intersection is also generic. \square

We next observe that the parametrization F is C^∞ at point clouds P in general position.

Proposition 5.8 *The parametrization $F : \mathbb{R}^{nd} \rightarrow \mathbb{R}^K$ is C^∞ over $\tilde{\mathcal{P}}$. Specifically, given $P \in \tilde{\mathcal{P}}$, letting $\{\tilde{v}(\sigma), \tilde{w}(\sigma)\} = \operatorname{argmax}_{i,j \in \sigma} \|p_i - p_j\|_2$ for every $\sigma \in K$, there is an open neighborhood U of P such that $F(P')(\sigma) = \|p'_{\tilde{v}(\sigma)} - p'_{\tilde{w}(\sigma)}\|_2$ for every $P' = (p'_1, \dots, p'_n) \in U$ and $\sigma \in K$, from which follows that F is C^∞ at P .*

Proof The continuity of F follows from the continuity of the Euclidean norm and max function. Assuming P is in general position, the distances $\|p_i - p_j\|_2$, for $i \neq j$ ranging in $\{1, \dots, n\}$, are strictly ordered. By continuity of F , this order remains the same over an open neighborhood U of P in \mathbb{R}^{nd} . Therefore, every $P' = (p'_1, \dots, p'_n) \in U$ is also in general position, and $F(P')(\sigma) = \|p'_{\tilde{v}(\sigma)} - p'_{\tilde{w}(\sigma)}\|_2$ for all $\sigma \in K$. Now, the map $P' \mapsto \|p'_{\tilde{v}(\sigma)} - p'_{\tilde{w}(\sigma)}\|_2$ is C^∞ at P for each σ because $p'_{\tilde{v}(\sigma)} \neq p'_{\tilde{w}(\sigma)}$. This implies that F is C^∞ at P . \square

Defining \tilde{v}, \tilde{w} as in Proposition 5.8, and combining this result with Proposition 4.14, we deduce the following differential of B_p , which only relies on derivatives of the Euclidean distance between points:

Corollary 5.9 *The barcode valued map $B_p : P \in \mathbb{R}^{nd} \mapsto \operatorname{Dgm}_p(F(P)) \in \operatorname{Bar}$ is ∞ -differentiable in $\tilde{\mathcal{P}}$. Moreover, at $P \in \tilde{\mathcal{P}}$, for any barcode template (P_p, U_p) of $F(P)$ and any choice of ordering $(\sigma_1, \sigma'_1), \dots, (\sigma_m, \sigma'_m), \tau_1, \dots, \tau_n$ of (P_p, U_p) , the map \tilde{B}_p defined on a point cloud $P' = (p'_1, \dots, p'_n)$ by:*

$$\tilde{B}_p(P') = \left[\left(\|p'_{\tilde{v}(\sigma_i)} - p'_{\tilde{w}(\sigma_i)}\|_2, \|p'_{\tilde{v}(\sigma'_i)} - p'_{\tilde{w}(\sigma'_i)}\|_2 \right)_{i=1}^m, \left(\|p'_{\tilde{v}(\tau_j)} - p'_{\tilde{w}(\tau_j)}\|_2 \right)_{j=1}^n \right]$$

is a local C^∞ lift of B_p around P . The corresponding differential $d_{P, \tilde{B}_p} B_p : \mathbb{R}^{nd} \rightarrow \mathbb{R}^{2m} \times \mathbb{R}^n$ is defined on a tangent vector $u \in \mathbb{R}^{nd}$ by:

$$d_{P, \tilde{B}_p} B_p(u) = \left[\left(\langle P_{\tilde{v}(\sigma_i), \tilde{w}(\sigma_i)}, u \rangle, \langle P_{\tilde{v}(\sigma'_i), \tilde{w}(\sigma'_i)}, u \rangle \right)_{i=1}^m, \left(\langle P_{\tilde{v}(\tau_j), \tilde{w}(\tau_j)}, u \rangle \right)_{j=1}^n \right],$$

where $P_{i,j}$ denotes the vector with $\frac{p_i - p_j}{\|p_i - p_j\|_2}$ as i -th component (resp. $\frac{p_j - p_i}{\|p_i - p_j\|_2}$ as j -th component) and 0 as other components.

This result implies in particular that B_p is generically ∞ -differentiable, since by Proposition 5.7 the set of point clouds in general position is generic in \mathbb{R}^{nd} .

Proof By Proposition 5.8, F is C^∞ in $\tilde{\mathcal{P}}$, which is open by Proposition 5.7. Given P in general position, the distances $\|p_i - p_j\|_2$, for $i \neq j$ ranging in $\{1, \dots, n\}$, are strictly ordered, and this order remains the same over an open neighborhood U of P in \mathbb{R}^{nd} by continuity. By Proposition 5.8 again, we have $F(P')(\sigma) = \|p'_{\tilde{v}(\sigma)} - p'_{\tilde{w}(\sigma)}\|_2$ for every $P' = (p'_1, \dots, p'_n) \in U$ and $\sigma \in K$. Therefore, the pre-order induced by F on the simplices of K is constant over U . Consequently, B_p is ∞ -differentiable at P by Theorem 4.7. The rest of the statement is an immediate consequence of Proposition 4.14. \square

We conclude this section by considering a parametrization that constrains the points p_1, \dots, p_n to evolve along smooth submanifolds $\mathcal{M}_1, \dots, \mathcal{M}_n$ of \mathbb{R}^d :

Proposition 5.10 *Let $\mathcal{M}_1, \dots, \mathcal{M}_n$ be smooth submanifolds of \mathbb{R}^d . Denoting by $\iota : \mathcal{M}_1 \times \dots \times \mathcal{M}_n \hookrightarrow \mathbb{R}^{nd}$ the inclusion map, the barcode valued map $B_p = \text{Dgm}_p \circ F \circ \iota$ is generically ∞ -differentiable.*

Proof Let $\mathcal{M} := \mathcal{M}_1 \times \dots \times \mathcal{M}_n$. The parametrization $F \circ \iota$ is C^∞ at parameters $\theta \in \mathcal{M}$ such that, locally, for all θ' in a sufficiently small open neighborhood around θ , the following quantities are constant:

- (i) the indices of the points at distance 0 of each other in $\iota(\theta')$, and
- (ii) the pre-order on the pairwise distances in $\iota(\theta')$.

Note that, in this case, the point clouds $\iota(\theta')$ are not necessarily in general position, but the way they violate conditions (i) and (ii) of Definition 5.6 is constant. Let U' (resp. U) denote the set of points in \mathcal{M} where (i) (resp. (ii)) is satisfied. From the above, $F \circ \iota$ is C^∞ over $U \cap U'$. We now show that $U \cap U'$ is generic in \mathcal{M} .

Calling U_{ijkl} the quadric $\{P \in \mathbb{R}^{nd} \mid \|p_i - p_j\|_2 = \|p_k - p_l\|_2\}$, and U'_{ij} the hyperplane $\{P \in \mathbb{R}^{nd} \mid p_i = p_j\}$, for i, j, k, l ranging in $\{1, \dots, n\}$, we have:

$$U = \bigcap_{\{i,j\} \neq \{k,l\}} \mathcal{M} \setminus \partial \iota^{-1}(U_{ijkl})$$

$$U' = \bigcap_{i \neq j} \mathcal{M} \setminus \partial \iota^{-1}(U'_{ij}).$$

Indeed, for any $\{i, j\} \neq \{k, l\}$, the order between $\|p_i - p_j\|_2$ and $\|p_k - p_l\|_2$ in $\iota(\theta)$ is strict when θ is in the (open) complement of $\iota^{-1}(U_{ijkl})$, constantly an equality when θ is inside the (open) interior $\iota^{-1}(U_{ijkl})^\circ$, and not locally constant when θ lies on the boundary $\partial \iota^{-1}(U_{ijkl})$, hence the formula for U . The formula for U' follows from the same argument.

The sets $\partial \iota^{-1}(U_{ijkl})$ and $\partial \iota^{-1}(U'_{ij})$ are boundaries of closed sets, and thus, their complements in \mathcal{M} are generic. As finite intersections of generic sets, U and U' are themselves generic. Theorem 4.9 allows us to conclude. □

5.3 Rips Filtrations of Clouds of Ellipsoids

As pointed out by [4], in some cases, growing isotropic balls around the points of $P = (p_1, \dots, p_n) \in \mathbb{R}^{nd}$ may result in a loss of geometric information. It is then advised to grow rather ellipsoids with distinct covariance matrices around each point, to account for the local anisotropy of the problem. Formally, the *Ellipsoid-Rips* filtration of P with respect to the vector of covariance matrices $A = (A_1, \dots, A_n) \in (S_{d,+}(\mathbb{R}))^n$ is a filtration of the total complex $K := 2^{\{1, \dots, n\}} \setminus \{\emptyset\}$ with $n := \#P$ vertices, in which the

time of appearance of a simplex $\sigma \subseteq \{1, \dots, n\}$ is given by:

$$\max_{i,j \in \sigma} r_{i,j}(A) \text{ where } r_{i,j}(A) := \left\| \frac{p_i - p_j}{\frac{1}{2} \left(\sqrt{q_i \left(\frac{p_i - p_j}{\|p_i - p_j\|_2} \right)} + \sqrt{q_j \left(\frac{p_j - p_i}{\|p_i - p_j\|_2} \right)} \right)} \right\|_2,$$

where the $q_i : x \in \mathbb{R}^d \mapsto \langle A_i x, x \rangle$ are the quadrics determined by the positive definite matrices A_i ⁷. Here, we see the space $(S_{d,+}(\mathbb{R}))^n$ as our parameter space \mathcal{M} , whose smooth structure is inherited from that of $(\mathbb{R}^{\frac{d(d+1)}{2}})^n$, and we consider the parametrization:

$$F(A)(\sigma) := \max_{i,j \in \sigma} r_{i,j}(A).$$

We are then interested in the differentiability of the barcode valued map $B_p = \text{Dgm}_p \circ F$. Inspired by the case of isotropic Rips filtrations, we require that the covariance matrices in A lie in general position as defined hereafter:

Definition 5.11 The pair (A, P) is in *general position* if the two following conditions hold:

- all points in P are distinct, i.e., $p_i \neq p_j$ whenever $1 \leq i \neq j \leq n$;
- all pairwise “ellipsoidal” distances are distinct, i.e., $r_{i,j}(A) \neq r_{k,l}(A)$ whenever $\{i, j\} \neq \{k, l\} \subseteq \{1, \dots, n\}$.

Proposition 5.12 Assume the points of P to be pairwise distinct. Then, the set of vectors of covariance matrices A such that (A, P) is in general position is generic in $S_{d,+}(\mathbb{R}^d)^n$.

Proof First, we claim that the sets $O_{ijkl} := \{A \in S_{d,+}(\mathbb{R}^d)^n \mid r_{i,j}(A) = r_{k,l}(A)\}$, for $\{i, j\} \neq \{k, l\}$, are level-sets of some smooth real valued functions on $S_{d,+}(\mathbb{R}^d)^n$ whose gradients are nowhere zero. To prove this fact, we introduce the quantities $C := \frac{\|p_i - p_j\|_2}{\|p_k - p_l\|_2}$ and $(x, y) := (\frac{p_i - p_j}{\|p_i - p_j\|_2}, \frac{p_k - p_l}{\|p_k - p_l\|_2})$. Then:

$$\begin{aligned} A = (A_1, \dots, A_n) \in O_{ijkl} &\Leftrightarrow r_{i,j}(A) = r_{k,l}(A) \\ &\Leftrightarrow \frac{\sqrt{q_i(x)} + \sqrt{q_j(x)}}{\sqrt{q_k(y)} + \sqrt{q_l(y)}} = C \\ &\Leftrightarrow \frac{\sqrt{\langle A_i x, x \rangle} + \sqrt{\langle A_j x, x \rangle}}{\sqrt{\langle A_k y, y \rangle} + \sqrt{\langle A_l y, y \rangle}} = C. \end{aligned}$$

Note that x, y are nonzero because points in P are distinct. Therefore, the map $f_{ijkl} := A \in S_{d,+}(\mathbb{R})^d \mapsto \frac{\sqrt{\langle A_i x, x \rangle} + \sqrt{\langle A_j x, x \rangle}}{\sqrt{\langle A_k y, y \rangle} + \sqrt{\langle A_l y, y \rangle}} \in \mathbb{R}$ is well defined and smooth on $S_{d,+}(\mathbb{R})^n$

⁷ The quantity $r_{i,j}(A)$ serves as a proxy for the intersection of the two ellipsoids with covariance matrices A_i and A_j centered at p_i and p_j suggested by [4], as the problem of computing intersections of quadrics is in general NP-hard.

as the two inner products in the denominator are always strictly positive. We want to compute $\nabla f_{ijkl} = (\nabla_{A_1} f_{ijkl}, \dots, \nabla_{A_n} f_{ijkl})$ where $\nabla_{A_t} f_{ijkl}$ is the gradient of f_{ijkl} with respect to the t -th component of A . For $t = i$:

$$\nabla_{A_i} f_{ijkl} = \frac{1}{\sqrt{\langle A_k y, y \rangle + \langle A_l y, y \rangle}} \times \frac{1}{2\sqrt{\langle A_i x, x \rangle}} \times \nabla_{A_i} \langle A_i x, x \rangle .$$

The first two factors are strictly positive scalars for any $A \in S_{d,+}(\mathbb{R}^d)$. The last factor is the gradient of a nonzero linear map, so it is nonzero. As a consequence, the gradient $\nabla_A f_{ijkl}$ is nowhere zero, which proves our claim.

Then, by the constant rank theorem, each O_{ijkl} is a smooth sub-manifold of $S_{d,+}(\mathbb{R}^d)^n$ of dimension strictly lower than that of $S_{d,+}(\mathbb{R}^d)^n$. Taking their (finite) union allows us to conclude. □

From this point, the same chain of arguments as in the isotropic case allows us to show that the parametrization F is C^∞ at vectors of covariance matrices A in general position, and to express the differential of B_p at A . Assume the points of P to be pairwise distinct, and denote by $\tilde{\mathcal{A}} \subseteq S_{d,+}(\mathbb{R}^d)^n$ the subspace of covariance matrices A such that (A, P) is in general position.

Proposition 5.13 *The parametrization $F : S_{d,+}(\mathbb{R}^d)^n \rightarrow \mathbb{R}^K$ is C^∞ over $\tilde{\mathcal{A}}$. Specifically, given $A \in \tilde{\mathcal{A}}$, letting $\{\bar{v}(\sigma), \bar{w}(\sigma)\} = \operatorname{argmax}_{i,j \in \sigma} r_{i,j}(A)$ for every $\sigma \in K$, there is an open neighborhood U of A such that $F(A')(\sigma) = r_{\bar{v}(\sigma), \bar{w}(\sigma)}(A')$ for every $A' = (A'_1, \dots, A'_n) \in U$ and $\sigma \in K$, from which follows that F is C^∞ at A .*

Proof Let $A \in \tilde{\mathcal{A}}$. Then, the maps $r_{i,j}$ are C^∞ because the points of P are pairwise distinct, and furthermore the quantities $r_{i,j}(A)$, for $i \neq j$ ranging in $\{1, \dots, n\}$, are strictly ordered. By continuity, this order remains the same over an open neighborhood U of A in $S_{d,+}(\mathbb{R}^d)^n$. Therefore, for every $A' \in U$, for all $\sigma \in K$, we have $F(A')(\sigma) = r_{\bar{v}(\sigma), \bar{w}(\sigma)}(A')$. This implies that F is C^∞ at A . □

Defining \bar{v}, \bar{w} as in Proposition 5.13, and combining this result with Proposition 4.14, we deduce the following formula for the differential of B_p , which only rely on derivatives of the maps $r_{i,j}$:

Corollary 5.14 *The barcode valued map $B_p : A \in S_{d,+}(\mathbb{R}^d)^n \mapsto \operatorname{Dgm}_p(F(A)) \in \operatorname{Bar}$ is ∞ -differentiable over $\tilde{\mathcal{A}}$. Moreover, at $A \in \tilde{\mathcal{A}}$, for any barcode template (P_p, U_p) of $F(A)$ and any choice of ordering $(\sigma_1, \sigma'_1), \dots, (\sigma_m, \sigma'_m), \tau_1, \dots, \tau_n$ of (P_p, U_p) , the map \tilde{B}_p defined by:*

$$A' = (A'_1, \dots, A'_n) \mapsto \left[\left(r_{\bar{v}(\sigma_i), \bar{w}(\sigma_i)}(A'), r_{\bar{v}(\sigma'_i), \bar{w}(\sigma'_i)}(A') \right)_{i=1}^m, \left(r_{\bar{v}(\tau_j), \bar{w}(\tau_j)}(A') \right)_{j=1}^n \right]$$

is a local C^∞ lift of B_p around P , whose differential provides a closed formula for $d_{A, \tilde{B}_p} B_p$.

This result implies in particular that B_p is generically ∞ -differentiable, since by Proposition 5.12 the set of vectors of covariance matrices in general position is generic in $S_{d,+}(\mathbb{R}^d)^n$ (provided the points of P are pairwise distinct).

Proof By Proposition 5.13, F is C^∞ in $\tilde{\mathcal{A}}$, which is open by Proposition 5.12. Given $A \in \tilde{\mathcal{A}}$, the quantities $r_{i,j}(A)$, for $i \neq j$ ranging in $\{1, \dots, n\}$, are strictly ordered, and this order remains the same over an open neighborhood U of A in $S_{d,+}(\mathbb{R}^d)^n$ by continuity. By Proposition 5.13 again, we have $F(A')(\sigma) = r_{\tilde{v}(\sigma), \tilde{w}(\sigma)}(A')$ for every $A' = (A'_1, \dots, A'_n) \in U$ and $\sigma \in K$. Therefore, the pre-order induced by F on the simplices of K is constant over U . Consequently, B_p is ∞ -differentiable at A by Theorem 4.7. The rest of the statement is an immediate consequence of Proposition 4.14. \square

Remark 5.15 Corollaries 5.9 and 5.14 can be combined together to generically differentiate the barcode valued map B_p with respect to both the point positions and the covariance matrices. The corresponding parameter space is $\mathbb{R}^{nd} \times S_{d,+}(\mathbb{R}^d)^n$.

5.4 Arbitrary Filtrations of a Simplicial Complex

In certain scenarios, the optimization takes place in the entire space of filter functions $\text{Filt}(K)$ on a fixed simplicial complex K . For instance, in the context of topological simplification of a filter function f_0 , as described by [2,23], one looks for a filter function $f \in \mathbb{R}^K$ which is ε -close to f_0 in supremum norm and whose diagram $\text{Dgm}_p(f)$ equals $\text{Dgm}_p(f_0) \setminus \Delta_\varepsilon$, where Δ_ε is the set of intervals of $\text{Dgm}_p(f_0)$ that are ε -close to the diagonal. One way to formalize this question is as a soft-constrained optimization problem, whereby the bottleneck distance to the simplified barcode is to be minimized in tandem with the supremum-norm distance to the original function:

$$\min_{f \in \text{Filt}(K)} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(f_0) \setminus \Delta_\varepsilon) + \lambda \|f - f_0\|_\infty,$$

for some fixed mixing parameter λ . This optimization problem can be tackled using a variational approach, for which it is more convenient to work in the manifold \mathbb{R}^K containing $\text{Filt}(K)$. However, in order to avoid leaving $\text{Filt}(K)$, we consider the parametrization of \mathbb{R}^K given by the indicator function of $\text{Filt}(K)$:

$$F := \mathbb{1}_{\text{Filt}(K)} : \mathbb{R}^K \rightarrow \mathbb{R}^K$$

$$f \mapsto \begin{cases} f & \text{if } f \in \text{Filt}(K) \\ 0 & \text{otherwise,} \end{cases}$$

which is smooth generically. The optimization becomes then:

$$\min_{f \in \mathbb{R}^K} d_\infty(\text{Dgm}_p(F(f)), \text{Dgm}_p(f_0) \setminus \Delta_\varepsilon) + \lambda \|F(f) - f_0\|_\infty. \tag{15}$$

Implementing a variational approach such as gradient descent requires both terms in (15) to be differentiable. The second term is generically differentiable, as the parametrization F and the norm $\| \cdot \|_\infty$ are. The first term is the composition

$$f \in \mathbb{R}^K \longrightarrow \text{Dgm}_p(F(f)) \in \text{Bar} \longrightarrow d_\infty(\text{Dgm}_p(F(f)), \text{Dgm}_p(f_0) \setminus \Delta_\epsilon) \in \bar{\mathbb{R}}, \tag{16}$$

which by the chain rule (Proposition 3.14) is differentiable as long as both arrows are. Since F is generically differentiable, so is the first arrow by Theorem 4.9. The second arrow is the bottleneck distance to a fixed diagram and therefore also generically differentiable, as will be argued in Sect. 7. There, we also view Eq. (15) as an instance of semi-algebraic loss function, which can be minimized via stochastic gradient descent (SGD).

6 The Case of Barcode Valued Maps Derived from Real Functions on a Manifold

In this section, we consider barcode valued maps that factor through the space $\mathbb{R}^{\mathcal{X}}$ of real functions on a fixed smooth compact d -manifold \mathcal{X} without boundary. Since we seek statements about the differentiability of B , we restrict the focus to maps that factor through $C^\infty(\mathcal{X}, \mathbb{R})$ equipped with the standard Whitney C^∞ topology⁸:

$$B : \mathcal{M} \xrightarrow{F} C^\infty(\mathcal{X}, \mathbb{R}) \xrightarrow{\text{Dgm}} \text{Bar}^{d+1} .$$

Here, Dgm is the map that takes a function $f \in C^\infty(\mathcal{X}, \mathbb{R})$ to the vector of its barcodes $(\text{Dgm}_p(f))_{p=0}^d$. It is well defined on $C^\infty(\mathcal{X}, \mathbb{R})$, as continuous functions on triangulable spaces have well-defined persistence diagrams [12]. However, as in the previous sections, we want to work only with barcodes that have finitely many off-diagonal points, therefore we further assume that F takes its values in the subset $\text{Tame}(\mathcal{X})$ of tame C^∞ functions—note that $\text{Tame}(\mathcal{X})$ contains the generic subset of Morse functions on \mathcal{X} [39]. Hence, the factorization:

$$B : \mathcal{M} \xrightarrow{F} \text{Tame}(\mathcal{X}) \xrightarrow{\text{Dgm}} \text{Bar}^{d+1} .$$

As before, we call F the *parametrization* associated with B , and \mathcal{M} the *parameter space*, whose elements are generally referred to as θ . We also denote $F(\theta)$ by f_θ to emphasize the fact that F is valued in a function space. The map Dgm takes f_θ to the vector of its barcodes $(\text{Dgm}_p(f_\theta))_{p=0}^d$, so we can take advantage of the bijective correspondence between the critical points of f_θ (provided f_θ is Morse) and the interval endpoints in this vector (Proposition 2.14).

As in the case of a parametrization valued in the space of filter functions on a simplicial complex, we need F to be smooth in some reasonable sense to ensure that

⁸ This topology coincides with all usual topologies on $C^\infty(\mathcal{X}, \mathbb{R})$ because \mathcal{X} is compact.

the composite B is ∞ -differentiable. For this, we define a curve $c : \mathbb{R} \rightarrow C^\infty(\mathcal{X}, \mathbb{R})$ to be *differentiable* if the limit $\lim_{h \rightarrow 0} \frac{c(t+h) - c(t)}{h}$ exists for all $t \in \mathbb{R}$. The limit can be viewed as a curve, and when iterated limits exist, we say that c is a *smooth* curve. We then say that the parametrization F is *smooth*⁹ if it sends every smooth curve $\theta(t)$ in \mathcal{M} to a smooth curve $F(\theta(t))$ in $C^\infty(\mathcal{X}, \mathbb{R})$. By Corollary 11.9 in [38], if F is smooth, then its uncurried version

$$\tilde{F} : (\theta, x) \in \mathcal{M} \times \mathcal{X} \mapsto F(\theta)(x) \in \mathbb{R} \tag{17}$$

is a smooth map in the usual sense, to which we can therefore apply standard results from differential calculus, typically the implicit function theorem. This will be instrumental in the proof of our main result (Theorem 6.1).

6.1 Smoothness of the Barcode Valued Map

Theorem 6.1 (Continuous smoothness) *Let $F : \mathcal{M} \rightarrow C^\infty(\mathcal{X}, \mathbb{R})$ be a parametrization of class C^∞_c valued in $\text{Tame}(\mathcal{X})$. Let $\theta \in \mathcal{M}$ be a parameter such that f_θ is Morse with critical values of multiplicity 1. Then, B is ∞ -differentiable at θ .*

Proof Since f_θ is a Morse function on a compact manifold, $\text{Crit}(f_\theta)$ is a finite set whose cardinality we denote by N_θ . We will proceed by proving the following statements in sequence:

- (i) There exist an open neighborhood U of θ and smooth maps $\pi_l : U \rightarrow \mathcal{X}$ for $1 \leq l \leq N_\theta$ that track the critical points, that is:

$$\forall \theta' \in U, \text{Crit}(f_{\theta'}) = \{\pi_l(\theta')\}_{1 \leq l \leq N_\theta} \tag{18}$$

- (ii) Shrinking U if necessary, we further have that for any $\theta' \in U$, $f_{\theta'}$ is Morse with critical values of multiplicity 1.
- (iii) Let $\theta' \in U$ and $(b, d) \in \text{Dgm}_p(\mathcal{X}, f_{\theta'}) \setminus \Delta$ for some homology degree p . Then, either $d = +\infty$, in which case there exists a unique $1 \leq l \leq N_\theta$ such that $b = f_{\theta'}(\pi_l(\theta'))$, or $d < +\infty$, in which case there exist unique $1 \leq l \neq l' \leq N_\theta$ such that $(b, d) = (f_{\theta'}(\pi_l(\theta')), f_{\theta'}(\pi_{l'}(\theta')))$.
- (iv) For all $\theta_1, \theta_2 \in U$, $1 \leq l \neq l' \leq N_\theta$, and $0 \leq p \leq d$, we have: $(f_{\theta_1}(\pi_l(\theta_1)), f_{\theta_1}(\pi_{l'}(\theta_1))) \in \text{Dgm}_p(f_{\theta_1})$ (resp. $(f_{\theta_1}(\pi_l(\theta_1)), +\infty) \in \text{Dgm}_p(f_{\theta_1})$) if and only if $(f_{\theta_2}(\pi_l(\theta_2)), f_{\theta_2}(\pi_{l'}(\theta_2))) \in \text{Dgm}_p(f_{\theta_2})$ (resp. $(f_{\theta_2}(\pi_l(\theta_2)), +\infty) \in \text{Dgm}_p(f_{\theta_2})$).
- (v) There exist smooth local coordinate systems for B_p at θ for every $0 \leq p \leq d$. Therefore, by Proposition 3.8, the barcode valued map B is ∞ -differentiable at θ .

The proofs of assertions (i) and (ii) use differential geometry: we show that we can smoothly track the critical points of $f_{\theta'}$ as θ' varies in a neighborhood of θ . The proof of assertion (iii) simply exploits the fact that the endpoints in the barcodes of a

⁹ This notion of smooth map is key in the theory developed in [25,36] for adapting the concepts of differential geometry to a wide category of infinite-dimensional vector spaces, including Banach and Fréchet spaces.

Morse function are its critical values (Proposition 2.14). Assertion (iv) means that the critical points do not exchange their contributions to the persistence diagrams when the parameter is varying. This will be shown using standard tools in persistence theory. Assertion (v) is obtained by re-indexing the set $\{1, \dots, N_\theta\}$ such that, through this re-indexation, the maps $\theta' \mapsto f_{\theta'}(\pi_l(\theta'))$ provide local coordinate systems as defined in Definition 3.6. □

Proof of assertion (i): The tangent bundle $T\mathcal{X} = \bigsqcup_{x \in \mathcal{X}} \{x\} \times T_x\mathcal{X}$ is a smooth manifold of dimension $2d$. Let x_1, \dots, x_{N_θ} be the critical points of f_θ . Locally, in an open neighborhood \mathbb{V} of these critical points, the tangent bundle is parallelizable, i.e., we have a diffeomorphism $T\mathbb{V} \cong \mathbb{V} \times \mathbb{R}^d$ and the projection onto the second component provides a smooth map to \mathbb{R}^d . Consider the map:

$$\partial F : (\theta', x) \in \mathcal{M} \times \mathbb{V} \mapsto \nabla f_{\theta'}(x) \in T_x\mathbb{V} \cong \mathbb{R}^d,$$

which is smooth due to the smoothness of \tilde{F} , see Eq. (17). Then, at the critical points we have $\partial F(\theta, x_l) = \nabla f_\theta(x_l) = 0$. Moreover, because f_θ is Morse, $\nabla_x \partial F(\theta, x_l) = \nabla^2 f_\theta(x_l)$ is invertible, where $\nabla_x \partial F$ denotes the first derivative of ∂F with respect to its second argument. We can then apply the implicit function theorem to ∂F : there exist an open neighborhood U_l of θ , an open neighborhood V_l of x_l (contained in \mathbb{V}) and a smooth diffeomorphism $\pi_l : U_l \rightarrow V_l$ such that

$$\forall (\theta', x) \in U_l \times V_l, \partial F(\theta', x) = 0 \iff x = \pi_l(\theta'). \tag{19}$$

Let $U = \bigcap_{l=1}^{N_\theta} U_l$. After shrinking each V_l so that it equals $\pi_l(U)$, we obtain that (19) holds over $U \times V_l$ for every $1 \leq l \leq N_\theta$. Now, by definition of ∂F and the (\Leftarrow) of (19), we have

$$\forall \theta' \in U, \{\pi_l(\theta')\}_{1 \leq l \leq N_\theta} \subseteq \text{Crit}(f_{\theta'}).$$

We now show the converse inclusion. From the (\Rightarrow) in Eq. (19), it is sufficient to prove that no critical points of $f_{\theta'}$ can be found in the compact set $W := \mathcal{X} \setminus (\bigcup_{l=1}^{N_\theta} V_l)$ when θ' ranges over U . We equip \mathcal{X} with an arbitrary Riemannian metric g , and we consider the smooth map:

$$\partial G : (\theta', x) \in U \times \mathcal{X} \mapsto g(\nabla f_{\theta'}(x), \nabla f_{\theta'}(x)) \in \mathbb{R},$$

where $\nabla f_{\theta'}(x) \in T_x\mathcal{X}$. In particular, $\partial G(\theta', x)$ is zero if and only if x is a critical point of $f_{\theta'}$. As a result, ∂G does not vanish on $\{\theta\} \times W$ since W includes no critical point of f_θ . By the compactness of W and the continuity of ∂G , there exists an open neighborhood U' of θ such that $\partial G|_{U' \times W}$ does not vanish either. Assertion (i) follows after shrinking U to $U \cap U'$. □

Proof of assertion (ii): Let U be as in assertion (i). Since f_θ is Morse, $\nabla_x \partial F(\theta, x_l) = \nabla^2 f_\theta(x_l)$ is invertible for each $l \in \{1, \dots, N_\theta\}$. ∂F is of class C^1 as it is of class C^∞ , so we get open neighborhoods U'_l of θ and V'_l of x_l such that $\nabla_x \partial F$ is invertible over

$U'_l \times V'_l$. We shrink U to $U \cap (\bigcap_{l=1}^{N_\theta} U'_l)$ and each V_l to $V_l \cap V'_l$, so that the critical points of $f_{\theta'}$ are non-degenerate for $\theta' \in U$. Shrinking U further if necessary, a similar argument ensures that the critical values of $f_{\theta'}$ have multiplicity 1 for all $\theta' \in U$. This concludes the proof of assertion (ii).

Proof of assertion (iii): Let $\theta' \in U$. Let $(b, d) \in \text{Dgm}_p(f_{\theta'}) \setminus \Delta$ for some homology degree $0 \leq p \leq d$. We assume that $d < +\infty$. From assertion (ii), $f_{\theta'}$ is Morse with critical values of multiplicity 1. Therefore, by Proposition 2.14, $f_{\theta'}$ induces a bijection between the multi-sets $\text{Crit}(f_{\theta'})$ and $E(f_{\theta'})$. Meanwhile, assertion (i) provides the equality $\text{Crit}(f_{\theta'}) = \{\pi_l(\theta')\}_{1 \leq l \leq N_\theta}$, so $f_{\theta'}$ induces a bijection $\{\pi_l(\theta')\}_{1 \leq l \leq N_\theta} \rightarrow E(f_{\theta'})$. By taking pre-images of b and d which are in $E(f_{\theta'})$, there exist some unique indices $1 \leq l \neq l' \leq N_\theta$ such that $(b, d) = (f_{\theta'}(\pi_l(\theta')), f_{\theta'}(\pi_{l'}(\theta')))$. The case $d = +\infty$ is proven the same way. \square

Proof of assertion (iv): The maps $(\theta_1, \theta_2) \in U^2 \mapsto |f_{\theta_1}(\pi_l(\theta_1)) - f_{\theta_2}(\pi_{l'}(\theta_2))| \in \mathbb{R}_+$, for varying $1 \leq l \neq l' \leq N_\theta$, are continuous. They are strictly positive at (θ, θ) because f_θ has critical values of multiplicity 1, so

$$\inf_{1 \leq l \neq l' \leq N_\theta} |f_\theta(\pi_l(\theta)) - f_\theta(\pi_{l'}(\theta))| > 0.$$

By continuity, shrinking U further if necessary, we have

$$\inf_{1 \leq l \neq l' \leq N_\theta, (\theta_1, \theta_2) \in U^2} |f_{\theta_1}(\pi_l(\theta_1)) - f_{\theta_2}(\pi_{l'}(\theta_2))| > 0.$$

Let ε be a real number such that:

$$0 < \varepsilon < \inf_{1 \leq l \neq l' \leq N_\theta, (\theta_1, \theta_2) \in U^2} |f_{\theta_1}(\pi_l(\theta_1)) - f_{\theta_2}(\pi_{l'}(\theta_2))|. \tag{20}$$

By continuity of \tilde{F} and compactness of \mathcal{X} , we can shrink U further¹⁰ so that $\|f_{\theta_1} - f_{\theta_2}\|_\infty \leq \frac{\varepsilon}{2}$ for any $\theta_1, \theta_2 \in U$. From the stability theorem 2.12 we then have:

$$\forall \theta_1, \theta_2 \in U, \forall 0 \leq p \leq d, \quad d_\infty(\text{Dgm}_p(f_{\theta_1}), \text{Dgm}_p(f_{\theta_2})) \leq \frac{\varepsilon}{2}. \tag{21}$$

Let us fix two parameters $\theta_1, \theta_2 \in U$ and a homology degree p . Let $1 \leq l_1 \neq l'_1 \leq N_\theta$ be such that $(f_{\theta_1}(\pi_{l_1}(\theta_1)), f_{\theta_1}(\pi_{l'_1}(\theta_1))) \in \text{Dgm}_p(f_{\theta_1})$. From Equation (21), there exists a matching $\gamma : \text{Dgm}_p(f_{\theta_1}) \rightarrow \text{Dgm}_p(f_{\theta_2})$ with $\text{cost } c(\gamma) \leq \frac{\varepsilon}{2}$. In particular, if we denote $(b, d) := \gamma(f_{\theta_1}(\pi_{l_1}(\theta_1)), f_{\theta_1}(\pi_{l'_1}(\theta_1))) \in \mathbb{R}^2$, then

$$|f_{\theta_1}(\pi_{l_1}(\theta_1)) - b| \leq \frac{\varepsilon}{2} \quad \text{and} \quad |f_{\theta_1}(\pi_{l'_1}(\theta_1)) - d| \leq \frac{\varepsilon}{2}. \tag{22}$$

¹⁰ This does not prevent our choice of ε from satisfying (20), because shrinking U increases the right-hand side of this equation.

Of course we cannot have $d = +\infty$. Also, we cannot have $(b, d) \in \Delta$, i.e., $b = d$, because then the triangle inequality would imply that $|f_{\theta_1}(\pi_{l_1}(\theta_1)) - f_{\theta_1}(\pi_{l'_1}(\theta_1))| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$, which contradicts (20). Thus, (b, d) is a bounded off-diagonal point of $\text{Dgm}_p(f_{\theta_2})$. By assertion (iii), there exist indices $1 \leq l_2 \neq l'_2 \leq N_\theta$ such that $b = f_{\theta_2}(\pi_{l_2}(\theta_2))$ and $d = f_{\theta_2}(\pi_{l'_2}(\theta_2))$. Equations (22) and (20) together force $l_2 = l_1$ and $l'_2 = l'_1$. Hence, $(f_{\theta_2}(\pi_{l_1}(\theta_2)), f_{\theta_2}(\pi_{l'_1}(\theta_2))) = (b, d) \in \text{Dgm}_p(f_{\theta_2})$, which proves the result. The case of an index $1 \leq l \leq N_\theta$ such that $(f_{\theta_1}(\pi_l(\theta_1)), +\infty) \in \text{Dgm}_p(f_{\theta_1})$ is treated in the same way. \square

Proof of assertion (v): For any homology degree $0 \leq p \leq d$, by assertion (iii), each bounded off-diagonal interval (b, d) in $\text{Dgm}_p(f_\theta) \setminus \Delta$ can be rewritten as $(f_\theta(\pi_{l_{b,p}}(\theta)), f_\theta(\pi_{l_{d,p}}(\theta)))$ for some indices $l_{b,p} \neq l_{d,p}$. Similarly, each interval $(v, +\infty)$ can be rewritten as $(f_\theta(\pi_{l_{v,p}}(\theta)), +\infty)$ for some index $l_{v,p}$. By assertion (iv), for any parameter $\theta' \in U$, $B_p(\theta')$ equals

$$\begin{aligned} & \left\{ (f_{\theta'}(\pi_{l_{b,p}}(\theta')), f_{\theta'}(\pi_{l_{d,p}}(\theta'))) \right\}_{(b,d) \in \text{Dgm}_p(f_\theta) \setminus \Delta} \\ & \cup \left\{ (f_{\theta'}(\pi_{l_{v,p}}(\theta')), +\infty) \right\}_{(v, +\infty) \in \text{Dgm}_p(f_\theta)} \cup \Delta^\infty. \end{aligned}$$

This provides a smooth local coordinate system (see Definition 3.7) for B_p at θ , therefore B_p is ∞ -differentiable at θ by Proposition 3.8. Since this is true for every $0 \leq p \leq d$, B itself is ∞ -differentiable at θ . \square

Remark 6.2 (Multiplicity one) The upcoming Fig. 5 shows how important the assumption that f_θ has critical values of multiplicity 1 is for the conclusion of Theorem 6.1 to hold. Roughly speaking, the assumption implies that the critical points do not exchange their contributions to the persistence diagrams of f_θ under perturbations of θ . We proved this fact using the stability theorem for persistence diagrams (see the proof of assertion (iv) above); however, it is also a consequence of the so-called *structural stability theorem* for dynamical systems [42]. This result implies that the gradient vector field induced by a Morse function f_θ with distinct critical values is *structurally stable*, and as an immediate consequence, that the Morse–Smale complex of f_θ does not change as we smoothly perturb f_θ . The Morse–Smale complex allows us to recover the persistence module completely and, in turn, the barcode of f_θ .

6.2 Discussion: Generic Differentiability

Theorem 6.1 guarantees that B is ∞ -differentiable at parameters θ that produce Morse functions with critical values of multiplicity 1. The set of such functions is a generic subspace of $C^\infty(\mathcal{X}, \mathbb{R})$ [26]. We can also argue that, under some extra conditions on the parametrization F , the set $D(\mathcal{M}, \mathcal{X})$ of parameters $\theta \in \mathcal{M}$ that produce Morse functions f_θ with critical values of multiplicity 1 is generic in \mathcal{M} :

Proposition 6.3 [40] *If F is smooth and generically large, i.e., for generic $x \in \mathcal{X}$ the map $\theta \in \mathcal{M} \mapsto df_\theta(x) \in T_x \mathcal{X}^*$ is a submersion, then $D(\mathcal{M}, \mathcal{X})$ is generic in \mathcal{M} .*

There are important examples where this result applies, such as for instance:

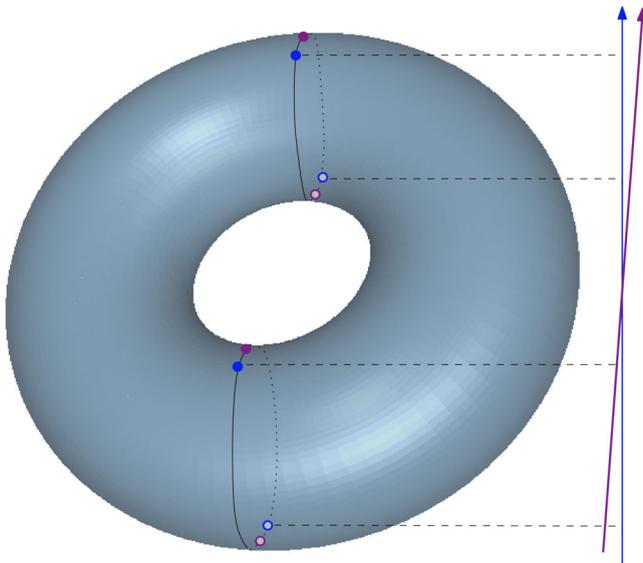


Fig. 2 A torus filtered by a generic height function. The blue arrow indicates the direction θ . By the correspondence of Proposition 2.14, the 4 critical points (blue dots) correspond from bottom to top to an infinite bar in degree 0, an infinite bar in degree 1, another infinite bar in degree 1, and an infinite bar in degree 2 of the resulting barcode $\text{Dgm}(h_\theta)$. The implicit function theorem applied to these critical points allows us to track them smoothly when perturbing the height function (purple arrow). The correspondence to points in the barcode remains unchanged (Color figure online)

Example 6.4 [40] Assume \mathcal{X} is embedded in \mathbb{R}^d and translated so as not to contain the origin. Then, each of the following parametrizations F is smooth and generically large:

$$\begin{aligned}
 v \in \mathbb{R}^d &\mapsto (x \in \mathcal{X} \mapsto \langle v, x \rangle \in \mathbb{R}) \\
 p \in \mathbb{R}^d &\mapsto (x \in \mathcal{X} \mapsto |x - p|^2 \in \mathbb{R}) \\
 A \in S_+(\mathbb{R}^d) &\mapsto (x \in \mathcal{X} \mapsto \frac{1}{2} \langle Ax, x \rangle \in \mathbb{R})
 \end{aligned}$$

6.3 A Simple Example

Take the ground space \mathcal{X} to be the torus $\mathbb{S}^1 \times \mathbb{S}^1$ embedded in \mathbb{R}^3 , the parameter space \mathcal{M} to be the 2-sphere \mathbb{S}^2 , and the parametrization F to be the family of height filtrations, i.e., $F : \theta \in \mathbb{S}^2 \mapsto (x \in \mathcal{X} \mapsto \langle \theta, x \rangle \in \mathbb{R})$. For a generic direction $\theta \in \mathbb{S}^2$, the induced height function, which we denote by h_θ , will be Morse and no two critical points are in the same level set. In this case, we can track the critical points smoothly as we vary θ , and the barcodes $\text{Dgm}_p(h_\theta)$ also evolve smoothly. An example of this situation is given in Fig. 2.

Even in this elementary situation, the singular parameters $\theta \in \mathbb{S}^2$ can exhibit pathological behaviors. There are two specific heights, on opposite sides of the sphere \mathbb{S}^2 ,

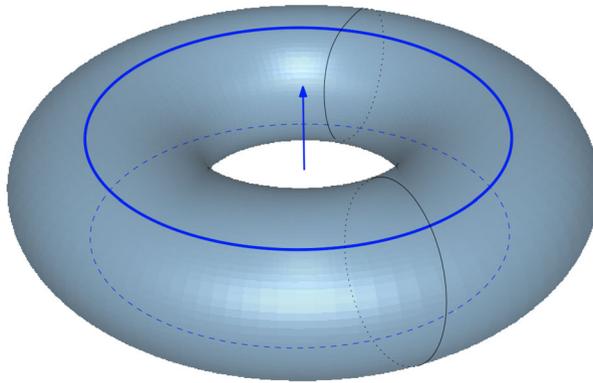


Fig. 3 Horizontal torus filtered by the vertical height function h_θ . The critical sets are the two blue circles, one of which corresponds to both a birth of a connected component and a loop, while the other corresponds to the births of a loop and a 2-cycle. Observe that any slight perturbation of θ results in a valid Morse function with 4 critical points; however, the locations of these points do not vary smoothly, and not even continuously, at θ (Color figure online)

that produce *Morse–Bott* functions. We show one of them in Fig. 3. At such a parameter θ , the critical sets are codimension-1 submanifolds of \mathcal{X} , and smooth perturbations of θ may result in discontinuous changes in the critical set.

There are other directions θ at which the assumptions of Theorem 6.1 are not met, yet the interval endpoints in the barcode can still be tracked smoothly. Such a case is shown in Fig. 4, where the height function h_θ is Morse but with a critical value of multiplicity 2. In this specific case, the implicit function theorem still applies to both critical points and provides a smooth local coordinate system for the barcode of h_θ .

However, in the general case, such a Morse function with two critical points sitting in the same level-set can induce a change in the correspondence with interval endpoints in the barcode, potentially resulting in non-smooth behavior of the barcode valued map B . An example is given in Fig. 5.

7 The Case of Maps on Barcodes Derived from Vectorizations and Loss Functions

We continue on with examples of differentiable maps, this time focusing on maps $V : \text{Bar} \rightarrow \mathcal{N}$ defined on barcodes and valued in a smooth finite-dimensional manifold. There is a plethora of examples of such maps V in the literature on topological data analysis [1,6,10,15,21,34]. Most of them take \mathcal{N} to be a Euclidean or Hilbert space, and they were designed to provide meaningful (e.g., stable, discriminative) representations of barcodes that can be fed to machine learning algorithms. A prototypical example of such a map is the persistence image of [1], which we study in Sect. 7.1. Other maps even have $\mathcal{N} = \mathbb{R}$ as codomain, and they are meant to be used as loss terms in optimization tasks [3,14,32]. Many examples of such vectorizations and loss functions are part of the wide class of *linear representations*, which we study in Sect. 7.2. In

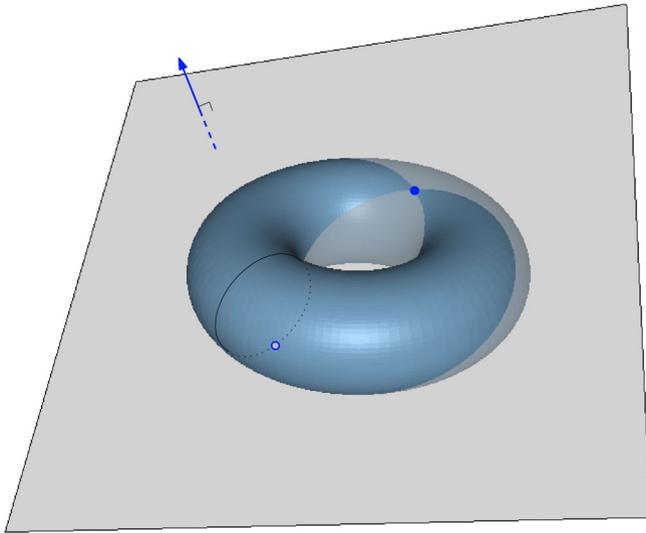


Fig. 4 A height function h_θ (blue arrow) that is Morse with two critical points (blue dots) in the same level set (the hyperplane), producing two distinct loops in the persistence diagram. The critical points can still be tracked smoothly around θ , and no change in the pairing occurs in the barcode $\text{Dgm}(h_\theta)$ (Color figure online)

Sect. 7.4, we study an important example of nonlinear loss, namely the bottleneck distance to a fixed barcode, which we believe can be of interest in the context of inverse problems. The machinery developed in this section is likely to be adaptable to other examples of maps on barcodes, however the purpose of the section is to provide a proof of concept rather than an exhaustive treatment.

7.1 The Differentiability of Persistence Images

Recall that Bar is equipped with the bottleneck topology. Let Bar_n be the subset of Bar containing the barcodes with n infinite intervals. In particular, Bar_0 is the set of barcodes whose intervals are bounded.

Proposition 7.1 *The set of path connected components of Bar is enumerable. More precisely, $\pi_0(\text{Bar}) = \{\text{Bar}_n\}_{n=0}^{+\infty}$.*

Proof Since $\text{Bar} = \bigsqcup_{n=0}^{+\infty} \text{Bar}_n$, we only need to prove that each Bar_n is a maximal connected subset of Bar . First note that Bar_n is path connected, as we can always move n infinite intervals to n other ones continuously, and similarly move the bounded off-diagonal intervals to the diagonal. We now prove the maximality of Bar_n . Let $A \subseteq \text{Bar} \setminus \text{Bar}_n$ be non-empty. Any element in A has infinite bottleneck distance to any element in Bar_n , since their numbers of infinite intervals are different. Therefore, $A \cup \text{Bar}_n$ cannot be path-connected, and so Bar_n is maximal. \square

We view the persistence image as a map $V : \text{Bar}_0 \rightarrow \mathbb{R}^{n^2}$ for some discretization step $n \in \mathbb{N}$:

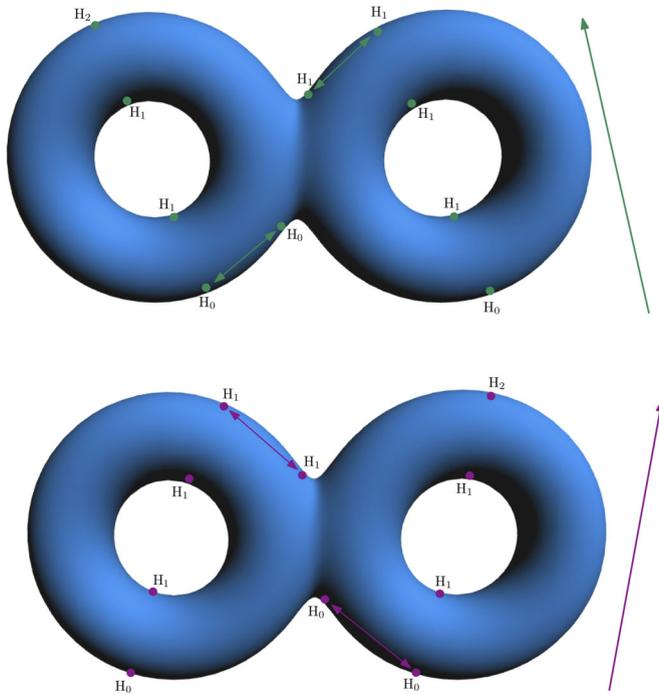


Fig. 5 A 2-torus filtered by two infinitesimal perturbations of the vertical height function, together with their critical points and labels to indicate the dimension of the homology group that they affect. Paired critical points correspond to bounded intervals in the associated barcodes. Here, the vertical height function is Morse and the critical points evolve smoothly. However, the pairing between critical points is not constant, nor their homological dimensions. Therefore, the barcode valued map is not smooth at the vertical direction (Color figure online)

Definition 7.2 Let $D \in Bar_0$. We fix a *weighting function* $\omega : \mathbb{R} \rightarrow \mathbb{R}$ that is zero at the origin. For $(b, d) \in \mathbb{R}^2$, consider the Gaussian

$$g_{b,d} : (x, y) \in \mathbb{R}^2 \mapsto \frac{1}{2\pi\sigma^2} e^{-[(x-b)^2+(y-(d-b))^2]/2\sigma^2}$$

for some fixed variance $\sigma > 0$. The *persistence surface* associated with D is the map

$$\rho_D : (x, y) \in \mathbb{R}^2 \mapsto \sum_{(b,d) \in D} \omega(d-b)g_{b,d}(x, y).$$

Given a square $B \subset \mathbb{R}^2$, we subdivide it into n^2 regular squares $B_{k,l}$ for $1 \leq k, l \leq n$. Then, we define the *persistence image* of D to be the histogram

$$V_{B,n} : D \in Bar_0 \mapsto \left(\int_{(x,y) \in B_{k,l}} \rho_D(x, y) dx dy \right)_{1 \leq k, l \leq n} \in \mathbb{R}^{n^2}$$

Proposition 7.3 *If ω is C^r over \mathbb{R}^2 for some integer $r \in \mathbb{N}$, then $V_{B,n}$ is r -differentiable everywhere in Bar_0 .*

Proof The maps $(b, d) \in \mathbb{R}^2 \mapsto \int_{(x,y) \in B_{k,l}} g_{b,d}(x, y) dx dy \in \mathbb{R}$ are C^∞ for any fixed box $B_{k,l}$. For any space of ordered barcodes $\mathbb{R}^{2m} \times \mathbb{R}^0$ and any $\tilde{D} = (b_1, d_1, \dots, b_m, d_m) \in \mathbb{R}^{2m} \times \mathbb{R}^0$,

$$V_{B,n}(Q_{m,0}(\tilde{D})) = \left(\sum_{1 \leq i \leq m} \omega(d_i - b_i) \int_{(x,y) \in B_{k,l}} g_{b_i,d_i}(x, y) dx dy \right)_{1 \leq k,l \leq n} \in \mathbb{R}^{n^2},$$

which is C^r at every $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^0$. □

In [1], the weighting function ω is chosen to be the ramp function $\omega_t : \mathbb{R} \mapsto \mathbb{R}$ defined as

$$\omega_t(u) = \begin{cases} 0 & \text{if } u \leq 0 \\ \frac{u}{t} & \text{if } 0 \leq u \leq t \\ 1 & \text{if } t \leq u \end{cases} \tag{23}$$

for some parameter $t > 0$. Thus, the ramp function is differentiable everywhere except at 0 and t . This implies that the persistence image $V_{B,n}$ is nowhere differentiable, as every neighborhood of a barcode always contains some neighborhood of the diagonal Δ . Thanks to Proposition 7.3, this issue can be resolved by taking any C^r approximation of the ramp function, which makes the persistence image r -differentiable over Bar_0 .

7.2 Linear Representations of Barcodes

The analysis of persistence images in the previous section can be generalized to the following wide class of vectorizations:

Definition 7.4 Let $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^k$, $\psi : \mathbb{R} \rightarrow \mathbb{R}^k$ and $\omega : \mathbb{R} \rightarrow \mathbb{R}$ be continuous maps such that $\omega(0) = 0$. The associated *linear representation* is the map

$$V : D \in Bar \mapsto \sum_{(b,d) \in D \text{ bounded}} \omega(d - b)\phi(b, d) + \sum_{(v, +\infty) \in D} \psi(v) \in \mathbb{R}^k.$$

Properties of linear representations valued in Banach spaces such as continuity, Lipschitzness and stochastic convergence are analyzed in [19,22]. Many vectorizations in the literature are linear representations, e.g., persistence images [1] and its variations [18,35,44], persistence silhouettes [13] and weighted Betti curves [47].

When $k = 1$, a linear representation may be viewed as a loss function on persistence diagrams. The total persistence in Example 3.11 and more generally the q -Wasserstein

distance to the empty diagram are such loss functions. In addition, the structure elements of [31, Definition 9] form a wide class of parametrized linear losses and linear representations that can be optimized.

In all these examples, the maps ϕ , ψ and ω are not necessarily smooth by design, see, e.g., the ramp function in Eq. (23) for persistence images, but one can always replace them with smooth approximations. We then get r -differentiable maps on barcodes, as expressed in the following result.

Proposition 7.5 *If the maps ϕ, ψ are C^r on generic subsets of \mathbb{R}^2 containing the diagonal Δ , and if ω is C^r on a generic subset of \mathbb{R} containing the origin, then the associated linear representation V is generically r -differentiable. Whenever ϕ, ψ and ω are in fact C^r everywhere, then V is r -differentiable everywhere.*

Proof The subspace of barcodes whose intervals avoid the set of non-differentiability of ϕ, ψ and ω is clearly generic in Bar . Let D be a barcode therein. For any space of ordered barcodes $\mathbb{R}^{2m} \times \mathbb{R}^n$ and pre-image $\tilde{D} = [(b_i, d_i)_{i=1}^m, (v_j)_{j=1}^n] \in \mathbb{R}^{2m} \times \mathbb{R}^n$ of D , we have

$$V(Q_{m,n}(\tilde{D})) = \sum_{1 \leq i \leq m} \omega(d_i - b_i)\phi(b_i, d_i) + \sum_{1 \leq j \leq n} \psi(v_j),$$

which is C^r in a neighborhood of \tilde{D} . □

Let us consider an everywhere r -differentiable linear representation V , and a barcode valued map B on a simplicial complex, which is (generically) differentiable (Theorem 4.9). Using the chain rule 3.14, the composition $V \circ B$ is then itself (generically) differentiable, hence amenable to gradient descent based optimization.

7.3 Semi-algebraic and Subanalytic Functions on Barcodes

We consider another important class of examples arising from loss functions on barcodes that restrict to semi-algebraic maps on the spaces of ordered barcodes. The subanalytic and definable counterparts are analogously defined, and the results of this section are valid in these situations as well. See also [9] for a full treatment of semi-algebraic loss functions in persistence.

Definition 7.6 We say that a map $V : Bar \rightarrow \mathbb{R}$ is *semi-algebraic* if all the precompositions $V \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathbb{R}$ are semi-algebraic.

A prototypical example of semi-algebraic loss on barcodes is the distance to a target barcode D_0 :

$$d_q(D_0, \cdot) : D \in Bar \mapsto d_q(D_0, D) \in \mathbb{R} \cup \{+\infty\}.$$

Here, d_q is the q -th Wasserstein distance on barcodes for any $q \in \mathbb{R}_+^*$ as defined in Eq. (7), and d_∞ is the bottleneck distance.

Proposition 7.7 *For any target barcode D_0 and nonnegative number $q \in \mathbb{R}_+^*$, the map $d_q(D_0, \cdot) : \text{Bar} \rightarrow \mathbb{R}$ is semi-algebraic.*

Proof We consider the case where $q = \infty$, as the same line of arguments works for arbitrary Wasserstein metrics, and rewrite $d_q(D_0, \cdot)$ as d_{D_0} for simplicity. Let $m, n \in \mathbb{N}$. We assume that n is the number of infinite intervals in D_0 , as otherwise the map $d_{D_0} \circ Q_{m,n} : \mathbb{R}^{2m} \times \mathbb{R}^n \rightarrow \mathbb{R}$ takes infinite value everywhere. Then, $d_{D_0} \circ Q_{m,n}$ can be expressed as a minimum of finitely many cost functions, $\min c(\gamma_{m,n})(\cdot)$, each of which is defined in terms of a fixed partial matching $\gamma_{m,n}$ of coordinates in $\mathbb{R}^{2m} \times \mathbb{R}^n$ with interval endpoints of D_0 . As a point-wise maximum of finitely many absolute values, each cost function $c(\gamma_{m,n})(\cdot)$ is semi-algebraic, and so $d_{D_0} \circ Q_{m,n}$ is semi-algebraic. \square

Semi-algebraic functions V on barcodes are particularly useful in the context of optimization when pre-composed with a semi-algebraic parametrization of filter functions $F : \mathcal{M} \rightarrow \mathbb{R}^K$ on a fixed simplicial complex K . Indeed, composition preserves semi-algebraicity, and so from Remark 4.25 the loss function given by the composition

$$\mathcal{L} : \mathcal{M} \xrightarrow{F} \mathbb{R}^K \xrightarrow{\text{Dgm}_p} \text{Bar} \xrightarrow{V} \mathbb{R} \tag{24}$$

is a semi-algebraic map. Then, [20, Corollary 5.9] guarantees that the well-known stochastic gradient descent (SGD) algorithm converges almost surely to critical points of \mathcal{L} .¹¹

This guarantee can be applied to various optimization problems. When choosing the Rips parametrization F of point clouds as in Sect. 5.2, minimizing the loss $\mathcal{L} = d_q(D_0, \cdot) \circ \text{Dgm}_p \circ F$ amounts to solving the problem of point cloud inference originally proposed in [27], see [29] for implementations. Besides, from Sect. 5.4, for F the parametrization of all filter functions on a fixed simplicial complex and an adequate target barcode D_0 , the minimization of \mathcal{L} yields an approach to function simplification. However, when F is not semi-algebraic, typically in the continuous setting developed in Sect. 6, and more generally for an arbitrary barcode valued map $B : \mathcal{M} \rightarrow \text{Bar}$, it is unclear how to perform full-fledged continuous gradient descent to minimize

$$\mathcal{L} : \mathcal{M} \xrightarrow{B} \text{Bar} \xrightarrow{d_q(D_0, \cdot)} \mathbb{R} . \tag{25}$$

While implementing a solution to this problem is beyond the scope of this paper, it serves as a motivation for the next section where we show that the bottleneck distance to D_0 is generically ∞ -differentiable, as then the chain rule of Proposition 3.14 enables the use of gradient descent.

¹¹ The loss \mathcal{L} must also be locally Lipschitz for this result to hold. By the stability theorem [16], Dgm_p is Lipschitz continuous, hence this additional mild requirement is met whenever F and V are locally Lipschitz (for instance, when $V = d_q(D_0, \cdot)$ is the distance to a fixed barcode).

7.4 The Bottleneck Distance to a Diagram

For simplicity, we denote the bottleneck distance to a fixed barcode D_0 by:

$$d_{D_0} : D \in Bar \mapsto d_\infty(D, D_0) \in \mathbb{R}.$$

For ease of exposition, we consider the special case where $D_0 = \Delta^\infty$ is the empty diagram (the diagonal Δ with infinite multiplicity). The analysis of the general case of an arbitrary fixed barcode D_0 is technically more involved and is deferred to “Appendix B.”

Recall that $d_{\Delta^\infty}(D) = +\infty$ for any diagram $D \in Bar$ with infinite bars. Consequently, we consider the restriction of d_{Δ^∞} to the subset Bar_0 introduced in Sect. 7.1. This restriction is valued in the real line: $d_{\Delta^\infty} : Bar_0 \rightarrow \mathbb{R}$. Consider the set Bar_Δ of barcodes which admit a unique point at maximal distance to the diagonal Δ :

$$Bar_\Delta := \left\{ D \in Bar_0 \mid \# \operatorname{argmax}_{(b,d) \in D} \frac{|d-b|}{2} = 1 \right\}. \tag{26}$$

For $D \in Bar_\Delta$, we let $(\bar{b}_D, \bar{d}_D) \in D$ be the unique interval in the set $\operatorname{argmax}_{(b,d) \in D} \frac{|d-b|}{2}$.

Proposition 7.8 *Bar $_\Delta$ is generic in Bar $_0$. Moreover, given $D \in Bar_\Delta$, for $\varepsilon > 0$ small enough, any D' at bottleneck distance less than ε from D satisfies $d_{\Delta^\infty}(D') = \frac{|\bar{d}_{D'} - \bar{b}_{D'}|}{2}$ and $\|(\bar{b}_{D'}, \bar{d}_{D'}) - (\bar{b}_D, \bar{d}_D)\|_\infty < \varepsilon$.*

Proof Given $D \in Bar_0$, consider the set $\operatorname{argmax}_{(b,d) \in D} \frac{|d-b|}{2}$. If this set is not a singleton, then we can move infinitesimally one of its elements away from the diagonal, so as to get a diagram in Bar_Δ . Thus, Bar_Δ is dense in Bar_0 . Let now $D \in Bar_\Delta$, and let δ be the second maximal distance to the diagonal:

$$\delta := \max_{(b,d) \in D \setminus \{(\bar{b}_D, \bar{d}_D)\}} \frac{|d-b|}{2}$$

and $\alpha := \frac{|\bar{d}_D - \bar{b}_D|}{2} - \delta > 0$. Take $\varepsilon \in (0, \frac{\alpha}{4})$. If D' is at bottleneck distance less than ε from D , all the points of D' are within distance less than ε either from the diagonal or from an off-diagonal point of D . As we have picked $\varepsilon < \frac{\alpha}{4}$, there is a unique off-diagonal point (\bar{b}', \bar{d}') of D' that is within distance less than ε from (\bar{b}_D, \bar{d}_D) , and it must be the unique furthest point from Δ in D' . So indeed $D' \in Bar_\Delta$ and $(\bar{b}_{D'}, \bar{d}_{D'}) = (\bar{b}', \bar{d}')$. Therefore, Bar_Δ is open, which concludes the proof. \square

Not surprisingly, d_{Δ^∞} is smooth at every $D \in Bar_\Delta$, with partial derivatives related to the ones of the map $(\bar{b}_D, \bar{d}_D) \mapsto \frac{|\bar{d}_D - \bar{b}_D|}{2}$.

Proposition 7.9 *For any $D \in Bar_\Delta$,*

(i) d_{Δ^∞} is ∞ -differentiable at D , and

(ii) for any $m \in \mathbb{N}$ and $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^0$ such that $Q_{m,0}(\tilde{D}) = D$, there are exactly two nonzero components in the gradient $\nabla_{\tilde{D}}(d_{\Delta^\infty} \circ Q_{m,0})$, one with value $\frac{1}{2}$ and the other with value $-\frac{1}{2}$.

Proof Let $m \in \mathbb{N}$ and $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^0$ be such that $Q_{m,0}(\tilde{D}) = D$. Without loss of generality, we can write $\tilde{D} = (\bar{b}_D, \bar{d}_D, b_2, d_2, \dots, b_m, d_m)$ where (b_i, d_i) is distinct from (\bar{b}_D, \bar{d}_D) for all $2 \leq i \leq m$. By Proposition 3.2, $Q_{m,0}$ is continuous. Therefore, by Proposition 7.8, there is an open neighborhood U of \tilde{D} , such that for any $\tilde{D}' = (\bar{b}_{D'}, \bar{d}_{D'}, b'_2, d'_2, \dots, b'_m, d'_m) \in U$, $Q_{m,0}(\tilde{D}')$ is in Bar_Δ and $d_{\Delta^\infty}(Q_{m,0}(\tilde{D}')) = \frac{|\bar{d}_{D'} - \bar{b}_{D'}|}{2} > 0$. Assertions (i) and (ii) follow. \square

Acknowledgements The authors wish to thank Vidit Nanda and Oliver Vipond for the frequent conversations that influenced this project. The authors are also indebted to the anonymous reviewers for their valuable insights into the final revisions of the manuscript. JL wishes to thank Heather Harrington for general guidance, Yixuan Wang for sharing knowledge on Morse theory and differential geometry, and finally Ambrose Yim and Naya Yerolemu for their feedback. This research has been supported in part by ESPRC Grant EP/R018472/1.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Local Isometry of the Barcode on a Simplicial Complex

Proof of Proposition 4.26 We have several persistence diagrams to compare, so we first simplify the problem as follows. Given two vectors $D = (D_0, \dots, D_d) \in Bar^{d+1}$ and $D' = (D'_0, \dots, D'_d)^{d+1}$ of $d + 1$ barcodes, let $\Gamma(D, D')$ be the set of multi-matchings between D and D' , where a multi-matching is a bijection $\gamma : \bigsqcup_{p=0}^d D_i \rightarrow \bigsqcup_{p=0}^d D'_i$ such that $\gamma(D_p) = D'_p$ for all $0 \leq p \leq d$. The notions of cost $c(\gamma)$ and optimality are the same as for ordinary matchings. Specifically, for an optimal γ in $\Gamma(\text{Dgm}(f), \text{Dgm}(g))$, we have $c(\gamma) = \max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g))$.

We fix an ordering $\sigma_1 < \dots < \sigma_{\#K}$ of the simplices of K , which yields an isomorphism $\phi : \mathbb{R}^K \rightarrow \mathbb{R}^{\#K}$. We denote by f_i the i -th component of f through this isomorphism, i.e., $f_i = f(\sigma_i)$. Let us assume that $f, g \in \mathbb{R}^K$ are two filter functions in a common top dimensional stratum S . If we can prove (12) in this case then it will hold for f, g in the closure of S by a continuity argument. Since f and g are both in S , they induce the same strict order on the simplices, and without loss of generality we can assume that $f_1 < \dots < f_{\#K}$ and $g_1 < \dots < g_{\#K}$. By Proposition 4.23, we can write $\text{Dgm}(f) = Q(P(S)f)$ and $\text{Dgm}(g) = Q(P(S)g)$ for a fixed permutation matrix $P(S)$, which implies that:

$$\forall 1 \leq i \neq j \leq \#K, \forall 0 \leq p \leq d, (f_i, f_j) \in \text{Dgm}_p(f) \Leftrightarrow (g_i, g_j) \in \text{Dgm}_p(g)$$

$$\forall 1 \leq k \leq \#K, \forall 0 \leq p \leq d, (f_k, +\infty) \in \text{Dgm}_p(f) \Leftrightarrow (g_k, +\infty) \in \text{Dgm}_p(g) \tag{27}$$

Let $\gamma \in \Gamma(\text{Dgm}(f), \text{Dgm}(g))$ be optimal. Consider the case where γ sends an off-diagonal point (b, d) of $\text{Dgm}(f)$ onto the diagonal Δ . As (b, d) is of the form (f_i, f_j) (or $(f_i, +\infty)$), this implies that $c(\gamma) \geq \frac{|f_i - f_j|}{2} \geq d_0(f)$. In addition, $\text{Dgm}(f)$ and $\text{Dgm}(g)$ have exactly the same number of bounded and unbounded intervals in each degree, which implies that there exists an off-diagonal interval (b', d') of $\text{Dgm}(g)$ which has pre-image in the diagonal Δ . Again, (b', d') must be of the form (g_k, g_l) (or $(g_k, +\infty)$), so that $c(\gamma) \geq \frac{|g_k - g_l|}{2} \geq d_0(g)$. Therefore, $c(\gamma) \geq \max(d_0(f), d_0(g))$ and we are done.

We now treat the case where all off-diagonal intervals are sent to off-diagonal intervals by γ . We denote by $O(f, g) \subset \Gamma(\text{Dgm}(f), \text{Dgm}(g))$ the set of multi-matchings that send off-diagonal intervals to off-diagonal intervals. By the decomposition $\text{Dgm}(f) = Q(P(S)f)$ (resp. $\text{Dgm}(g) = Q(P(S)g)$) and from the fact that no two values of f (resp. of g) are equal, the bounded end-points of off-diagonal intervals of $\text{Dgm}(f)$ (resp. $\text{Dgm}(g)$) are in bijection with the set $\{f_1, \dots, f_{\#K}\}$ (resp. $\{g_1, \dots, g_{\#K}\}$). Therefore, any multi-matching $v \in O(f, g)$ induces a permutation $\pi(v)$ of $\{1, \dots, \#K\}$. Let us denote by $c(\pi) := \max_i (|f_i - g_{\pi(i)}|)$ the *cost* of a permutation π of $\{1, \dots, \#K\}$. In this formulation, we have:

$$\max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) = c(\gamma) = \min_{v \in O(f, g)} c(v) = \min_{v \in O(f, g)} c(\pi(v)) \tag{28}$$

Consider the following relaxed optimization problem, in which the pairing of coordinates in (27) is ignored:

$$\min_{\pi \text{ permutation of } \{1, \dots, \#K\}} c(\pi) \tag{29}$$

From the fact that $f_1 < \dots < f_{\#K}$ and $g_1 < \dots < g_{\#K}$, the monotonicity of the optimal transport map for the ∞ -Wasserstein distance in \mathbb{R} [45] guarantees that $\pi = \text{Id}$ is a minimizer¹² of (29). Therefore,

$$\begin{aligned} \max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) &= \min_{v \in O(f, g)} c(v) \\ &\geq \min_{\pi \text{ permutation}} c(\pi) = c(\text{Id}) = \|f - g\|_\infty \end{aligned}$$

and we are done.

We now address the second part of the proposition. Let $f \in \mathbb{R}^K$ be a filter function. By the stability Theorem 2.12, showing that Eq. (13) holds amounts to showing that $\max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) \geq \|f - g\|_\infty$. We denote by S the top

¹² Indeed, for a permutation π , let $\text{inv}(\pi)$ denote its number of inversions. Let π be a minimizer (29) with minimal $\text{inv}(\pi)$. Assuming by contradiction that $\text{inv}(\pi) > 0$, there exist $i < j$ such that $\pi(i) > \pi(j)$. Let τ be the transposition that swaps $\pi(i)$ and $\pi(j)$. Since $f_1 < \dots < f_{\#K}$ and $g_1 < \dots < g_{\#K}$, a simple case analysis shows that $c(\tau \circ \pi) \leq c(\pi)$, thus raising a contradiction with the minimality of $\text{inv}(\pi)$. Therefore $\pi = \text{Id}$ is a minimizer of (29).

dimensional stratum S that contains f , and let $g \in \mathbb{R}^K$ be another filter such that $\|f - g\|_\infty \leq d_0(f)$. This implies that g is also in the (closure of the) stratum S . We can then apply (12), and since by assumption $\|f - g\|_\infty \leq d_0(f) \leq \max(d_0(f), d_0(g))$, we have the desired result.

Using similar arguments, we finally prove Eq. (14). Let $f \in \mathbb{R}^K$ be a filter function in some top dimensional stratum S , and $g, h \in \mathbb{R}^K$ be such that $\|f - g\| \leq \frac{d_0(f)}{3}$ and $\|f - h\| \leq \frac{d_0(f)}{3}$. By the stability Theorem 2.12, showing that Eq. (14) holds amounts to showing that $\max_{0 \leq p \leq d} d_\infty(\text{Dgm}_p(g), \text{Dgm}_p(h)) \geq \|g - h\|_\infty$. For every $i \neq j \in \{1, \dots, \#K\}$,

$$\begin{aligned} |g_i - g_j| &= |(f_i - f_j) - [(f_i - g_i) + (g_j - f_j)]| \\ &\geq |f_i - f_j| - 2\|f - g\|_\infty \geq 2d_0(f) - 2\|f - g\|_\infty \geq \frac{4d_0(f)}{3}, \end{aligned}$$

so that $d_0(g) \geq \frac{2d_0(f)}{3}$. Similarly, $d_0(h) \geq \frac{2d_0(f)}{3}$. Meanwhile,

$$\|g - h\|_\infty \leq \|f - g\|_\infty + \|f - h\|_\infty \leq \frac{2d_0(f)}{3}.$$

Therefore, $\|g - h\|_\infty \leq \max(d_0(g), d_0(h))$, and since both g, h are in (the closure of) S , we conclude by using Eq. (12). □

B The Bottleneck Distance to a Fixed Diagram: The General Case

Throughout, we denote by Δ_ϵ the set of elements in \mathbb{R}^2 that are at distance less than $\epsilon > 0$ to the diagonal Δ . We equip, for the purpose of this section only, the spaces of ordered barcodes with the supremum norm $\|\cdot\|_\infty$ rather than the Euclidean norm. Note that (the proof of) Proposition 3.2 ensures that the quotient maps $Q_{m,n}$ are 1-Lipchitz with respect to the metrics in place. We denote by $\mathcal{B}(\cdot, *)$ the ball centered at \cdot with radius $*$ with respect to the supremum norm or bottleneck metric depending on the context.

In this section, we generalize Proposition 7.9; namely, we show the generic differentiability of the bottleneck distance $d_{D_0} : \text{Bar} \rightarrow \mathbb{R} \cup \{+\infty\}$ to an arbitrary fixed diagram $D_0 \in \text{Bar}$.

Proposition B.1 *Let $D_0 \in \text{Bar}$ and n be the number of infinite bars in D_0 . For generic $D \in \text{Bar}_n$, d_{D_0} is ∞ -differentiable at D . Moreover, for any $m \in \mathbb{N}$ and $\tilde{D} \in \mathbb{R}^{2m} \times \mathbb{R}^n$ such that $Q_{m,n}(\tilde{D}) = D$, exactly one of the following possibilities holds:*

- (i) *either the gradient $\nabla_{\tilde{D}}(d_{D_0} \circ Q_{m,n})$ has exactly two nonzero components, one with value $\frac{1}{2}$ and the other with value $-\frac{1}{2}$; or*
- (ii) *the gradient $\nabla_{\tilde{D}}(d_{D_0} \circ Q_{m,n})$ has a unique nonzero component with value 1 or -1 .*

Proposition B.1 states the generic smoothness of d_{D_0} . We first observe that all the compositions $d_{D_0} \circ Q_{m,n}$ are smooth on a generic subset of \mathbb{R}^{2m+n} .

Lemma B.2 For every $m \in \mathbb{N}$, the map

$$d_{D_0} \circ Q_{m,n} : \mathbb{R}^{2m+n} \rightarrow \mathbb{R}$$

is generically smooth, with gradients that are either 0 or as in (i) or (ii) of Proposition B.1.

Proof Let $m \in \mathbb{N}$. Define an ordered matching $\tilde{\gamma} : \mathbb{R}^{2m+n} \rightarrow \mathbb{R}^{2m+n}$ to be an affine map whose first m pairs of coordinate functions (resp. last n coordinate functions) are of the form $\tilde{D} := [(b_i, d_i)_{i=1}^m, (v_j)_{j=1}^n] \mapsto (b_i, d_i) - (b_{0,i}, d_{0,i})$ where $(b_{0,i}, d_{0,i})$ is either an off-diagonal point in D_0 or $(b_{0,i}, d_{0,i}) = (\frac{b_i+d_i}{2}, \frac{b_i+d_i}{2})$ is the orthogonal projection of (b_i, d_i) onto Δ (resp. are of the form $\tilde{D} \mapsto v_j - v_{0,j}$ for some infinite interval $(v_{0,j}, +\infty)$ in D_0). We further require that the collection of intervals $(b_{0,i}, d_{0,i})$ (resp. $(v_{0,j}, +\infty)$) involved in this way are distinct elements in D_0 . We denote by $D_0(\tilde{\gamma})$ the set of bounded off-diagonal intervals $(b_0, d_0) \in D_0$ that are not in the collection $\{b_{0,i}, d_{0,i}\}_{i=1}^m$.

Since the maximum of smooth functions over \mathbb{R}^{2m+n} is smooth¹³ on a generic subset of \mathbb{R}^{2m+n} , the map

$$\tilde{c}(\tilde{\gamma}) : \tilde{D} \in \mathbb{R}^{2m+n} \mapsto \max(\|\tilde{\gamma}(\tilde{D})\|_\infty, \{\frac{|d_0 - b_0|}{2}\}_{(b_0, d_0) \in D_0(\tilde{\gamma})}) \in \mathbb{R} \tag{30}$$

is itself C^∞ on a generic subset of \mathbb{R}^{2m+n} , with gradients either equal to 0 or as in (i) or (ii) of Proposition B.1. Let $\tilde{\Gamma}_m$ be the set of ordered matchings $\tilde{\gamma} : \mathbb{R}^{2m+n} \rightarrow \mathbb{R}^{2m+n}$, which is non-empty and finite. Then, the map

$$\tilde{d}_{D_0,m} : \tilde{D} \in \mathbb{R}^{2m+n} \mapsto \min_{\tilde{\gamma} \in \tilde{\Gamma}_m} \tilde{c}(\tilde{\gamma})(\tilde{D}) \in \mathbb{R}$$

is C^∞ on a generic subset of \mathbb{R}^{2m+n} , with gradients either equal to 0 or as in (i) or (ii) of Proposition B.1.

We will be done if we can show that the two maps $d_{D_0} \circ Q_{m,n}$ and $\tilde{d}_{D_0,m}$ are equal over \mathbb{R}^{2m+n} . Fix an ordered barcode $\tilde{D} \in \mathbb{R}^{2m+n}$ and let $D := Q_{m,n}(\tilde{D})$. Let $\tilde{\gamma} : \mathbb{R}^{2m+n} \rightarrow \mathbb{R}^{2m+n}$ be an ordered matching. The components of $\tilde{\gamma}$ determine a matching γ between D and D_0 , sending (b_i, d_i) onto $(b_{0,i}, d_{0,i})$ and $(v_j, +\infty)$ onto $(v_{0,j}, +\infty)$. By definition of the cost of a matching 2.10 and Equation (30), we have $c(\gamma) = \tilde{c}(\tilde{\gamma})(\tilde{D})$. This yields $\tilde{d}_{D_0,m}(\tilde{D}) \geq d_{D_0}(D) = d_{D_0} \circ Q_{m,n}(\tilde{D})$. Conversely, among the optimal matchings from D to D_0 , it is always possible to find one that sends off-diagonal points of D (and D_0) on the diagonal only by orthogonal projection. This allows us to lift γ at the level of \tilde{D} and to define an ordered matching $\tilde{\gamma}$ such that $\tilde{c}(\tilde{\gamma})(\tilde{D}) = c(\gamma)$. This yields $\tilde{d}_{D_0,m}(\tilde{D}) \leq d_{D_0}(D) = d_{D_0} \circ Q_{m,n}(\tilde{D})$ and therefore $d_{D_0} \circ Q_{m,n} = \tilde{d}_{D_0,m}$ on \mathbb{R}^{2m+n} . \square

¹³ This is the same argument as in the proof of Proposition 4.9. Namely, the set where at least two of the smooth functions involved in the maximum are equal is closed, and therefore the boundary of this set has generic complement. On this complement the maximum of the smooth functions locally equals a unique smooth function.

We cannot directly use Lemma B.2 to prove Proposition B.1. As a matter of fact, by the definition of ∞ -differentiability (Definition 3.10), Proposition B.1 is asking that for generic $D \in \text{Bar}_n$ all the maps $d_{D_0} \circ Q_{m,n}$, for varying $m \in \mathbb{N}$, should be smooth at pre-images of D . However, Lemma B.2 only guarantees that the maps $d_{D_0} \circ Q_{m,n}$, taken individually, are smooth over generic subsets of \mathbb{R}^{2m+n} , and it is not clear a priori how to glue at the level of barcodes these generic subsets lying in different spaces of ordered barcodes \mathbb{R}^{2m+n} . In order to leverage Lemma B.2, we devise intermediate results that infer the smoothness of the maps $d_{D_0} \circ Q_{m',n}$ from the knowledge of the smoothness of a well-chosen map $d_{D_0} \circ Q_{m,n}$. The high-level intuition of each of these intermediate steps is as follows:

1. Infinitesimal perturbations of a given diagram D can be understood as infinitesimal moves of the off-diagonal points of D , together with appearances of small intervals from the diagonal. In Lemma B.3, we devise a generic condition on D ensuring that these new small off-diagonal intervals appearing when perturbing D do not play any role in the bottleneck distance to D_0 .
2. Given a barcode D , we take a pre-image $\tilde{D}_m \in Q_{m,n}^{-1}(D)$ of D which is minimal in the sense that its pairs of adjacent components are not trivial, i.e., not of the form (b, b) . In other words, \tilde{D}_m is an ordering of the endpoints of off-diagonal intervals appearing in D without extra pairs (b, b) lying on the diagonal. Up to an infinitesimal perturbation of \tilde{D}_m , Lemma B.2 ensures that $d_{D_0} \circ Q_{m,n}$ is smooth in an open neighborhood of \tilde{D}_m . It is easy to observe that for any other pre-image $\tilde{D}_{m'}$ of D , the components of the ordered barcode $\tilde{D}_{m'}$ only differ with those of \tilde{D}_m by the addition of trivial pairs of the form (b, b) . According to the previous item, those trivial pairs do not play any role when computing the bottleneck distance to D_0 . Therefore, since $d_{D_0} \circ Q_{m,n}$ is smooth in a neighborhood of \tilde{D}_m , the map $d_{D_0} \circ Q_{m',n}$ is itself smooth in an open neighborhood of $\tilde{D}_{m'}$. We make these intuitions rigorous in Lemma B.4.
3. The previous arguments allow us to construct open balls $\mathcal{B}(\tilde{D}_{m'}, \epsilon)$ of the same radius $\epsilon > 0$ around all pre-images $\tilde{D}_{m'} \in \mathbb{R}^{2m'+n}$ of a generic diagram $D \in \text{Bar}$ over which all maps $d_{D_0} \circ Q_{m',n}$ are smooth. To conclude that d_{D_0} itself is ∞ -differentiable in a neighborhood of D , we show in Lemma B.5 that the ϵ -bottleneck ball around D is covered by the union of the images of the balls $\mathcal{B}(\tilde{D}_{m'}, \epsilon)$.

Let $\hat{\text{Bar}}$ be the set of barcodes $D \in \text{Bar}_n$ such that no intervals of D_0 is at distance $d_\infty(D, D_0)$ to its diagonal projection. It is easy to check that $\hat{\text{Bar}}$ is generic in Bar_n . When perturbing a given barcode D infinitesimally, an arbitrary number of new off-diagonal points may appear from the diagonal. We show that, for $D \in \hat{\text{Bar}}$, these new off-diagonal intervals can be disregarded when computing the bottleneck distance to D_0 .

Lemma B.3 *Let $D \in \hat{\text{Bar}}$. There exists $\epsilon > 0$ such that for any barcode D' which is ϵ -close to D we have that $d_\infty(D', D_0) > \epsilon$, and there exists an optimal matching from D' to D_0 sending $D' \cap \Delta_\epsilon$ (i.e., those points of D' that are ϵ -close to the diagonal), onto the diagonal Δ .*

Proof Let $D \in \hat{B}ar$. Denote by α the minimal gap $|\frac{|d_0-b_0|}{2} - d_\infty(D, D_0)|$ between the distance of off-diagonal intervals (b_0, d_0) of D_0 from their diagonal projections and $d_\infty(D, D_0)$. Since $D \in \hat{B}ar$, α is strictly positive.

We have $d_\infty(D, D_0) > 0$ as otherwise $d_\infty(D, D_0) = 0$ would imply that $D \notin \hat{B}ar$ (as the distance from a diagonal element of D_0 to its diagonal projection would also be 0), so we can pick $\epsilon > 0$ such that

$$\epsilon < \min\left(\frac{d_\infty(D, D_0)}{2}, \frac{\alpha}{2}\right).$$

We now prove that the conclusion of the lemma holds in the bottleneck ball $\mathcal{B}(D, \epsilon)$. Let $D' \in \mathcal{B}(D, \epsilon)$. Since $\epsilon < \frac{d_\infty(D, D_0)}{2}$, we have $d_\infty(D', D_0) > \epsilon$. We assume, seeking contradiction, that there is no optimal matching from D' to D_0 that sends all points of $D' \cap \Delta_\epsilon$ onto Δ .

We restrict our attention to the set $\Gamma^*(D', D_0)$ of optimal matchings from D' to D_0 that are allowed to send off-diagonal points of D' and D_0 to the diagonal only by orthogonal projections. This set is finite and non-empty. We define the Δ -degree of a matching $\gamma \in \Gamma^*(D', D_0)$ to be the number of off-diagonal points of D' and D_0 that are sent to their diagonal projections, and take γ with maximal Δ -degree. By assumption, there exists an off-diagonal point $(b', d') \in D' \cap \Delta_\epsilon$ sent to some off-diagonal point $(b_0, d_0) \in D_0$. Recall that $|\frac{|d_0-b_0|}{2} - d_\infty(D, D_0)| \geq \alpha$. We divide the analysis into two cases: either $\frac{|d_0-b_0|}{2} \geq d_\infty(D, D_0) + \alpha$, or $\frac{|d_0-b_0|}{2} \leq d_\infty(D, D_0) - \alpha$.

In the case where $\frac{|d_0-b_0|}{2} \geq d_\infty(D, D_0) + \alpha$, we have:

$$\begin{aligned} \|(b_0, d_0) - (b', d')\|_\infty &= \|[(b_0, d_0) - (\frac{b'+d'}{2}, \frac{b'+d'}{2})] - [(b', d') - (\frac{b'+d'}{2}, \frac{b'+d'}{2})]\|_\infty \\ &\geq \|(b_0, d_0) - (\frac{b'+d'}{2}, \frac{b'+d'}{2})\|_\infty - \|(b', d') - (\frac{b'+d'}{2}, \frac{b'+d'}{2})\|_\infty \\ &\geq \frac{|d_0 - b_0|}{2} - \frac{|d' - b'|}{2} \\ &\geq d_\infty(D, D_0) + \alpha - \epsilon \\ &\geq d_\infty(D', D_0) + \alpha - 2\epsilon \\ &> d_\infty(D', D_0) \end{aligned}$$

where the first inequality holds by the triangle inequality, the second from the fact that a minimizer of the distance from (b_0, d_0) to the diagonal is the orthogonal projection of (b_0, d_0) onto Δ , the third by assumption on (b_0, d_0) and (b', d') , the fourth by the triangle inequality and the last one by $\epsilon < \frac{\alpha}{2}$. This yields a contradiction as γ is optimal and its cost may not exceed $d_\infty(D', D_0)$.

Consider now the case where $\frac{|d_0 - b_0|}{2} \leq d_\infty(D, D_0) - \alpha$. On the one hand, by the triangle inequality and by the fact that $\epsilon < \alpha$:

$$\frac{|d_0 - b_0|}{2} \leq d_\infty(D, D_0) - \alpha \leq d_\infty(D', D_0) + \epsilon - \alpha < d_\infty(D', D_0)$$

On the other hand, since $\epsilon < \frac{d_\infty(D, D_0)}{2}$,

$$\frac{|d' - b'|}{2} \leq \epsilon < \frac{d_\infty(D, D_0)}{2} \leq \frac{d_\infty(D', D_0) + \epsilon}{2} < d_\infty(D', D_0),$$

where the last inequality comes from the fact that $\epsilon < d_\infty(D', D_0)$ because $\epsilon < \frac{d_\infty(D', D_0) + \epsilon}{2}$. To sum up, both quantities $\frac{|d_0 - b_0|}{2}$ and $\frac{|d' - b'|}{2}$ are upper-bounded by $d_\infty(D', D_0)$. Modifying γ by sending (b_0, d_0) and (b', d') to their diagonal projections, we obtain a matching in $\Gamma^*(D', D_0)$ with Δ -degree strictly higher than that of γ , which contradicts the maximality of the Δ -degree of γ . \square

We say that an ordered barcode $\tilde{D}_m = [(b_i, d_i)_{i=1}^m, (v_j)_{j=1}^n] \in \mathbb{R}^{2m+n}$ is *minimal* if $b_i \neq d_i$ for $1 \leq i \leq m$. This terminology is justified by the fact that the image $D := Q_{m,n}(\tilde{D}_m) \in \text{Bar}_n$ contains exactly m bounded-off diagonal intervals and n unbounded ones, and therefore any other pre-image $\tilde{D}_{m'} \in \mathbb{R}^{2m'+n}$ of D must lie in a space of ordered barcodes of dimension at least $2m + n$ (i.e., $m' \geq m$). We show that under suitable assumptions, the differentiability of all the maps $d_{D_0} \circ Q_{m',n}$ at pre-images $\tilde{D}_{m'}$ of D can be inferred from the differentiability of $d_{D_0} \circ Q_{m,n}$ at the minimal pre-image \tilde{D}_m .

Lemma B.4 *For every $m \in \mathbb{N}$, the set of minimal ordered barcodes in \mathbb{R}^{2m+n} is open. Moreover, given a minimal $\tilde{D}_m \in \mathbb{R}^{2m+n}$ with $D := Q_{m,n}(\tilde{D}_m) \in \hat{B}\text{ar}$, if $d_{D_0} \circ Q_{m,n}$ is C^∞ in an open neighborhood of \tilde{D}_m , then there is an $\epsilon > 0$ such that for all other pre-images $\tilde{D}_{m'}$ of D , the map $d_{D_0} \circ Q_{m',n}$ is C^∞ in $\mathcal{B}(\tilde{D}_{m'}, \epsilon)$, with gradients as in (i) or (ii) of Proposition B.1.*

Proof It is clear that the set of minimal ordered barcodes in \mathbb{R}^{2m+n} is open. We address the second part of the lemma. Let $\tilde{D}_m \in \mathbb{R}^{2m+n}$ be a minimal ordered barcode such that $D := Q_{m,n}(\tilde{D}_m) \in \hat{B}\text{ar}$, and assume there is an open neighborhood U of \tilde{D}_m within which $d_{D_0} \circ Q_{m,n}$ is C^∞ . By continuity of the quotient map and from the fact that $\hat{B}\text{ar}$ is open, we can assume without loss of generality that $Q_{m,n}(U)$ is contained in $\hat{B}\text{ar}$.

For any other pre-image $\tilde{D}_{m'} \in \mathbb{R}^{2m'+n}$ of D , i.e., an ordered barcode such that $Q_{m',n}(\tilde{D}_{m'}) = D = Q_{m,n}(\tilde{D}_m)$, the first m' adjacent pairs of components of $\tilde{D}_{m'}$ must describe in an arbitrary order the m bounded off-diagonal points of D together with $m' - m$ trivial pairs of the form (b, b) . The last n components of $\tilde{D}_{m'}$ must be in correspondence with the left endpoints of infinite intervals in D . In other words, the first $2m'$ components of $\tilde{D}_{m'}$ consist of a re-ordering of the first $2m$ components of \tilde{D}_m , together with $m' - m$ trivial pairs of the form (b, b) . The last n components of $\tilde{D}_{m'}$ consist of a re-ordering of those of \tilde{D}_m .

To every pre-image $\tilde{D}_{m'}$ of D as above, we associate the linear projection $L_{m',m} : \mathbb{R}^{2m'+n} \rightarrow \mathbb{R}^{2m+n}$ that sends $\tilde{D}_{m'}$ to \tilde{D}_m by re-arranging the m non-trivial pairs of components and the n last components, and forgetting the $m' - m$ trivial pairs. Since $D \in \hat{B}ar$, Lemma B.3 provides an $\epsilon > 0$ such that for any $D' \in \mathcal{B}(D, \epsilon)$, the points of D' that are in Δ_ϵ may be sent onto the diagonal when computing the bottleneck distance from D' to D_0 , and furthermore they can be disregarded when computing $d_\infty(D', D_0)$. Therefore, using that the quotient map $Q_{m',n}$ is 1-Lipschitz, we know that for any pre-image $\tilde{D}_{m'}$ of D and $\tilde{D}'_{m'} \in \mathcal{B}(\tilde{D}_{m'}, \epsilon)$, the $m' - m$ pairs of components $\tilde{D}'_{m'}$ with persistence less than ϵ can be disregarded when computing $d_{D_0} \circ Q_{m',n}(\tilde{D}'_{m'})$. Formally, for every $m' \in \mathbb{N}$,

$$\begin{aligned} \forall \tilde{D}'_{m'} \in Q_{m',n}^{-1}(D), \forall \tilde{D}'_{m'} \in \mathcal{B}(\tilde{D}_{m'}, \epsilon), \quad d_{D_0} \circ Q_{m',n}(\tilde{D}'_{m'}) \\ = d_{D_0} \circ Q_{m,n} \circ L_{m',m}(\tilde{D}'_{m'}). \end{aligned} \tag{31}$$

Note that the maps $L_{m',m}$ are 1-Lipschitz. Therefore, we can reduce ϵ in order to ensure that $L_{m',m}(\mathcal{B}(\tilde{D}_{m'}, \epsilon)) \subset U$ for every pre-image $\tilde{D}_{m'}$ of D . Applying the chain rule on $d_{D_0} \circ Q_{m,n}$ and $L_{m',m}$ —which is an affine map hence C^∞ —in Equation (31), we obtain that all the maps $d_{D_0} \circ Q_{m',n}$ are C^∞ in $\mathcal{B}(\tilde{D}_{m'}, \epsilon)$. Also by the chain rule, by definition of $L_{m',m}$, the components of the gradients of the maps $d_{D_0} \circ Q_{m',n}$ are a re-ordering of the components of the gradient of $d_{D_0} \circ Q_{m,n}$. By Lemma B.2, the gradient of the latter is either 0 or as in (i) or (ii) of Proposition B.1. However, the gradient of $d_{D_0} \circ Q_{m,n}$ being 0 at some elements $\tilde{D}'_m \in U$ would mean that the bottleneck distance $d_\infty(Q_{m',n}(\tilde{D}'_m), D_0)$ equals the distance of some off-diagonal interval (b_0, d_0) to its diagonal projection, which is impossible since $Q_{m',n}(\tilde{D}'_m) \in \hat{B}ar$. \square

By means of Lemma B.4, we can deduce at once the differentiability of all the maps $d_{D_0} \circ Q_{m',n}$ over balls of the same radius. We need a last result that connects these balls to an actual open neighborhood of D in Bar_n .

Lemma B.5 *For any $D \in Bar_n$, there exists an $\epsilon > 0$ such that for every $m' \in \mathbb{N}$,*

$$Q_{m',n}^{-1}(\mathcal{B}(D, \epsilon)) \subseteq \bigcup_{\tilde{D}_{m'} \in R^{2m'+n}, Q_{m',n}(\tilde{D}_{m'})=D} \mathcal{B}(\tilde{D}_{m'}, \epsilon).$$

Proof Let $D \in Bar_n$, and $\eta > 0$ be less than all the pairwise distances between geometrically distinct off-diagonal points in D , and less than all the distances from off-diagonal points in D to the diagonal. We take $\epsilon > 0$ such that $\epsilon < \frac{\eta}{2}$. Let $D' \in \mathcal{B}(D, \epsilon)$. Then, for every off-diagonal point (b, d) of D , the number of (off-diagonal) points of D' lying in $\mathcal{B}((b, d), \epsilon)$ equals the multiplicity of (b, d) in D . Let us say that these points in D' are of type **(a)**. The points of D' that are not in the balls $\mathcal{B}((b, d), \epsilon)$, for (b, d) ranging over off-diagonal intervals of D , must be ϵ -close to the diagonal, and we say that these points are of type **(b)**. Note that we can accordingly characterize the components of a pre-image $\tilde{D}'_{m'} \in Q_{m',n}^{-1}(D')$: the pairs of components in $\tilde{D}'_{m'}$ must either be trivial (i.e., of the form (b, b)), or equal to some off-diagonal point of type

(a) or (b). All off-diagonal points of D' , of type (a) or (b), counted with multiplicity, must appear as a pair in $\tilde{D}'_{m'}$.

Given such a pre-image $\tilde{D}'_{m'} \in Q_{m',n}^{-1}(\mathcal{B}(D, \epsilon))$ of D' , we construct another ordered barcode $\tilde{D}_{m'} \in \mathbb{R}^{2m'+n}$ by modifying the components of $\tilde{D}'_{m'}$ at cost less than ϵ (i.e., such that $\|\tilde{D}'_{m'} - \tilde{D}_{m'}\|_\infty < \epsilon$) as follows:

- The last n components of $\tilde{D}'_{m'}$ parametrize the left endpoints of infinite intervals in D' . We change them at cost less than ϵ into the left endpoints of infinite intervals in D .
- If a pair (b', d') among the first m' pairs of components of $\tilde{D}'_{m'}$ is of type (a), it is ϵ -close to a unique off-diagonal point (b, d) of D . We change it into (b, d) .
- If a pair (b', d') among the first m' pairs of components of $\tilde{D}'_{m'}$ is of type (b), it is ϵ -close to the diagonal. We transform it into $(\frac{b'+d'}{2}, \frac{b'+d'}{2})$.
- The remaining pairs in the first m' pairs of components of $\tilde{D}'_{m'}$ must be trivial, and we leave them unchanged.

In this way, we have constructed an ordered barcode $\tilde{D}_{m'}$ such that $\tilde{D}'_{m'} \in \mathcal{B}(\tilde{D}_{m'}, \epsilon)$ and also, by construction, $\tilde{D}_{m'}$ is a pre-image of D , i.e., $Q_{m',n}(\tilde{D}_{m'}) = D$. □

We are now ready to prove Proposition B.1.

Proof of Proposition B.1 Consider the set of barcodes $D \in Bar_n$ that admit an open neighborhood within which d_{D_0} is ∞ -differentiable. By definition, this set is open in Bar_n , and we are left to show that it is also dense. Given an arbitrary $D \in Bar^n$, we will perform a series of infinitesimal perturbations of D , so that there exists a (small) open neighborhood U of D over which d_{D_0} is ∞ -differentiable.

Since \hat{Bar} is generic in Bar_n , up to an infinitesimal perturbation, we can assume that D lies in \hat{Bar} . Let $\tilde{D}_m \in \mathbb{R}^{2m+n}$ be a minimal pre-image of D . By Lemma B.4, the set of minimal ordered barcodes in \mathbb{R}^{2m+n} is open. Moreover, $d_{D_0} \circ Q_{m,n}$ is smooth on a generic subset of \mathbb{R}^{2m+n} by Lemma B.2. Therefore, up to an infinitesimal perturbation of \tilde{D}_m (which results in an infinitesimal perturbation of D by continuity of $Q_{m,n}$), we can further assume that $d_{D_0} \circ Q_{m,n}$ is smooth on a ball $\mathcal{B}(\tilde{D}_m, \epsilon)$ for some $\epsilon > 0$, while \tilde{D}_m remains minimal and D stays in \hat{Bar} .

Reducing ϵ if necessary, by Lemma B.4 all the maps $d_{D_0} \circ Q_{m',n}$ are smooth over $\mathcal{B}(\tilde{D}_{m'}, \epsilon)$, with gradients as in (i) or (ii) of Proposition B.1, where $\tilde{D}_{m'}$ ranges over the pre-images of D . Reducing ϵ further if necessary, we conclude that d_{D_0} is ∞ -differentiable over $\mathcal{B}(D, \epsilon)$ by Lemma B.5. □

References

1. 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. *The Journal of Machine Learning Research*, 18(1):218–252, 2017.
2. Dominique Attali, Marc Glisse, Samuel Hornus, Francis Lazarus, and Dmitriy Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. In *Proc. TopoInVis*, 2009.

3. Rickard Brüel-Gabrielsson, Vignesh Ganapathi-Subramanian, Primož Skraba, and Leonidas J Guibas. Topology-aware surface reconstruction for point clouds. In *Computer Graphics Forum*, volume 39, pages 197–207. Wiley Online Library, 2020.
4. Paul Breiding, Sara Kališnik, Bernd Sturmfels, and Madeleine Weinstein. Learning algebraic varieties from samples. *Revista Matemática Complutense*, 31(3):545–593, 2018.
5. Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *Journal of Computational Geometry*, 6(1):162–191, 2015.
6. Peter Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015.
7. William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *J. Algebra Appl.*, 14(5):art. id. 1550066, 8 pp, 2015.
8. Frédéric Chazal, William Crawley-Boevey, and Vin de Silva. The observable structure of persistence modules. *Homology Homotopy Appl.*, 18(2):247–265, 2016.
9. Mathieu Carriere, Frédéric Chazal, Marc Glisse, Yuichi Ike, and Hariprasad Kannan. Optimizing persistent homology based functions. arXiv preprint [arXiv:2010.08356](https://arxiv.org/abs/2010.08356), 2020. To appear in the proceedings of ICML 2021.
10. Mathieu Carriere, Marco Cuturi, and Steve Oudot. Sliced wasserstein kernel for persistence diagrams. In *International Conference on Machine Learning*, volume 70, pages 664–673. PMLR, 2017.
11. Pádraig Corcoran and Bailin Deng. Regularization of persistent homology gradient computation. In *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*, 2020.
12. Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. SpringerBriefs in Mathematics. Springer, 2016.
13. Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, and Larry Wasserman. Stochastic convergence of persistence landscapes and silhouettes. In *Proceedings of the thirtieth annual symposium on Computational geometry*, pages 474–483, 2014.
14. Chao Chen, Xiuyan Ni, Qinxun Bai, and Yusu Wang. A topological regularizer for classifiers via persistent homology. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2573–2582, 2019.
15. Mathieu Carrière, Steve Oudot, and Maks Ovsjanikov. Stable topological signatures for points on 3d shapes. In *Computer Graphics Forum*, volume 34, pages 1–12. Wiley Online Library, 2015.
16. David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
17. David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have l_p -stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010.
18. Yen-Chi Chen, Daren Wang, Alessandro Rinaldo, and Larry Wasserman. Statistical analysis of persistence intensity functions. arXiv preprint [arXiv:1510.02502](https://arxiv.org/abs/1510.02502), 2015.
19. Vincent Divol and Frédéric Chazal. The density of expected persistence diagrams and its kernel based estimation. *Journal of Computational Geometry*, 10(2):127–153, 2019.
20. Damek Davis, Dmitry Drusvyatskiy, Sham Kakade, and Jason D Lee. Stochastic subgradient method converges on tame functions. *Foundations of Computational Mathematics*, 20(1):119–154, 2020.
21. Barbara Di Fabio and Massimo Ferri. Comparing persistence diagrams through complex vectors. In *International Conference on Image Analysis and Processing*, pages 294–305. Springer, 2015.
22. Vincent Divol and Théo Lacombe. Understanding the topology and the geometry of the space of persistence diagrams via optimal partial transport. *Journal of Applied and Computational Topology*, pages 1–53, 2020.
23. Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
24. Herbert Federer. *Geometric measure theory*. Die Grundlehren der mathematischen Wissenschaften, Band 153. Springer-Verlag New York Inc., New York, 1969.
25. Alfred Frölicher and Andreas Kriegl. *Linear spaces and differentiation theory*. Pure and Applied Mathematics (New York). John Wiley & Sons, Ltd., Chichester, 1988. A Wiley-Interscience Publication.
26. M. Golubitsky and V. Guillemin. *Stable mappings and their singularities*. Springer-Verlag, New York-Heidelberg, 1973. Graduate Texts in Mathematics, Vol. 14.
27. Marcio Gameiro, Yasuaki Hiraoka, and Ipppei Obayashi. Continuation of point clouds via persistence diagrams. *Physica D: Nonlinear Phenomena*, 334:118–132, 2016.
28. Mark Goresky and Robert MacPherson. *Stratified Morse theory*. Ergebnisse der Mathematik, volume 14. Springer-Verlag, Berlin, 1988.



29. Rickard Brüel Gabriëlsson, Bradley J Nelson, Anjan Dwaraknath, and Primoz Skraba. A topology layer for machine learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1553–1563. PMLR, 2020.
30. Christoph Hofer, Florian Graf, Bastian Rieck, Marc Niethammer, and Roland Kwitt. Graph filtration learning. In *International Conference on Machine Learning*, pages 4314–4323. PMLR, 2020.
31. Christoph D Hofer, Roland Kwitt, and Marc Niethammer. Learning representations of persistence barcodes. *Journal of Machine Learning Research*, 20(126):1–45, 2019.
32. Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In *Advances in Neural Information Processing Systems*, pages 5658–5669, 2019.
33. Patrick Iglesias-Zemmour. *Diffeology*. Mathematical Surveys and Monographs, volume 185. American Mathematical Society, Providence, RI, 2013.
34. Sara Kališnik. Tropical coordinates on the space of persistence barcodes. *Found. Comput. Math.*, 19(1):101–129, 2019.
35. Genki Kusano, Yasuaki Hiraoka, and Kenji Fukumizu. Persistence weighted gaussian kernel for topological data analysis. In *International Conference on Machine Learning*, pages 2004–2013, 2016.
36. Andreas Kriegl and Peter W. Michor. *The convenient setting of global analysis*. Mathematical Surveys and Monographs, volume 53. American Mathematical Society, Providence, RI, 1997.
37. John Mather. Notes on topological stability. *Bull. Amer. Math. Soc. (N.S.)*, 49(4):475–506, 2012.
38. Peter W. Michor. *Manifolds of differentiable mappings*. Shiva Mathematics Series, volume 3. Shiva Publishing Ltd., Nantwich, 1980.
39. J. Milnor. *Morse theory*. Annals of Mathematics Studies, No. 51. Princeton University Press, Princeton, N.J., 1963. Based on lecture notes by M. Spivak and R. Wells.
40. Liviu Nicolaescu. *An invitation to Morse theory*. Universitext. Springer, New York, second edition, 2011.
41. Pierre Pansu. Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un. *Annals of Mathematics*, pages 1–60, 1989.
42. Jacob Palis and Stephen Smale. Structural stability theorems. In *Global Analysis (Proc. Sympos. Pure Math., Vol. XIV, Berkeley, Calif., 1968)*, pages 223–231. Amer. Math. Soc., Providence, R.I., 1970.
43. Adrien Poulenard, Primoz Skraba, and Maks Ovsjanikov. Topological function optimization for continuous shape matching. In *Computer Graphics Forum*, volume 37, pages 13–25. Wiley Online Library, 2018.
44. Jan Reininghaus, Stefan Huber, Ulrich Bauer, and Roland Kwitt. A stable multi-scale kernel for topological machine learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4741–4748, 2015.
45. Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.
46. Elchanan Solomon, Alexander Wagner, and Paul Bendich. A fast and robust method for global topological functional optimization. arXiv preprint [arXiv:2009.08496](https://arxiv.org/abs/2009.08496), 2020.
47. Yuhei Umeda. Time series classification via topological data analysis. *Information and Media Technologies*, 12:228–239, 2017.
48. Ka Man Yim and Jacob Leygonie. Optimization of Spectral Wavelets for Persistence-Based Graph Classification. *Frontiers in Applied Mathematics and Statistics*, 7:16, 2021.
49. Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.