

BASIC METRIC GEOMETRY OF THE BOTTLENECK DISTANCE

MAURICIO CHE ^A, FERNANDO GALAZ-GARCÍA ^B, LUIS GUIJARRO ^B, INGRID MEMBRILLO SOLIS ^C,
AND MOTIEJUS VALIUNAS

ABSTRACT. Given a metric pair (X, A) , i.e. a metric space X and a distinguished closed set $A \subset X$, one may construct in a functorial way a pointed pseudometric space $\mathcal{D}_\infty(X, A)$ of persistence diagrams equipped with the bottleneck distance. We investigate the basic metric properties of the spaces $\mathcal{D}_\infty(X, A)$ and obtain characterizations of their metrizability, completeness, separability, and geodesicity.

1. INTRODUCTION

In recent years, persistent homology has received considerable attention as it has shown not only to be a powerful tool in data analysis but to provide new perspectives in the study of abstract metric spaces. In persistent homology, the small, medium and large-scale topological features of geometric objects are encoded in a *persistence diagram*, a multiset of points in the (extended) plane. The construction of such diagrams starts with a filtration of a topological space, and records the exact birth and death times of generators of the persistent homology associated to that particular filtration [20]. Most common filtrations come from the Vietoris–Rips complexes over a given metric space or from sublevel sets of natural functions defined on some topological space. In either case, the obtained persistence diagram contains topological information regarding different levels or scales of the underlying problem of interest [9]. In applications to data analysis, it is crucial to compare the topological features of the analyzed data quantitatively, which in turn might be filtrated or sampled using different methods. Therefore, most applications of persistent homology require not only being able to compute persistence diagrams, but a measure of similarity between such objects.

The bottleneck distance was the first measure of similarity that appeared in the literature to compare persistence diagrams [11]. It is well-known that the bottleneck distance is an (extended) pseudometric [17]. The successful introduction of the bottleneck distance gave rise to alternative similarity measures for persistence diagrams. For instance, the p -Wasserstein metrics, for $1 \leq p < \infty$, form a one-parameter family of similarity measures between persistence diagrams, which, in contrast to the bottleneck distance, are metrics [12]; in the limit, one may think of the bottleneck distance as another member of this family. Out of all the similarity measures that have been defined in persistent homology, only the bottleneck distance has been shown to be equivalent to the interleaving distance, a pseudometric defined for persistent modules, the algebraic analogues of persistence diagrams [17].

Several authors have studied the topological and geometric properties of the spaces defined by the bottleneck distance, $\mathcal{D}_\infty(\mathbb{R}_{\geq 0}^2, \Delta)$, and the p -Wasserstein metrics, $\mathcal{D}_p(\mathbb{R}_{\geq 0}^2, \Delta)$, where $\mathbb{R}_{\geq 0}^2 = \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq y\}$, $\Delta = \{(x, y) \in \mathbb{R}_{\geq 0}^2 : x = y\}$, and $1 \leq p < \infty$, as well as those

Date: January 19, 2024.

2010 Mathematics Subject Classification. 53C23, 55N31, 54F45.

Key words and phrases. bottleneck distance, metric pair, persistence diagram.

^ASupported by CONACYT Doctoral Scholarship No. 769708.

^BSupported by research grants MTM2017–85934–C3–2–P and PID2021-124195NB-C32 from the Ministerio de Economía y Competitividad de España (MINECO), and by ICMAT Severo Ochoa project CEX2019-000904-S(MINECO)

^CSupported by the Leverhulme Trust (grant RPG-2019-055).

of generalizations of these spaces [4, 5, 10, 13]. In [6, 10], the authors study the geometry and topology of spaces of generalized persistence diagrams $\mathcal{D}_p(X, A)$, $1 \leq p < \infty$. These spaces are a generalization of the spaces defined in [16, Definition 4]. When $(X, A) = (\mathbb{R}_{\geq 0}^2, \Delta)$, one recovers the usual space of persistence diagrams, where the coordinates of a point $(x, y) \in \mathbb{R}_{\geq 0}^2$ are interpreted as the birth at scale x and death at scale y of a topological feature of the space under consideration. The spaces of persistence diagrams are constructed in a functorial way: for each $p \in [1, \infty]$, there exists a functor \mathcal{D}_p which assigns to a metric pair (X, A) , i.e. a metric space X and a distinguished non-empty closed subset $A \subset X$, a space of (generalized) persistence diagrams $\mathcal{D}_p(X, A)$ equipped with either the bottleneck distance, when $p = \infty$, or the p -Wasserstein metric, when $1 \leq p < \infty$. The topology and geometry of the spaces $\mathcal{D}_p(\mathbb{R}_{\geq 0}^2, \Delta)$, $1 \leq p < \infty$, were analyzed in [16, 19]. We note, however, that in previous works, the geometric and topological properties of spaces of persistence diagrams equipped with the bottleneck distance had not been completely investigated, since their study requires the use of distinct methods from those used in the study of the p -Wasserstein spaces of persistence diagrams. A further motivation for the study of the spaces $\mathcal{D}_p(X, A)$ is that, when $X = \mathbb{R}_{\geq 0}^{2n} = \{(x_1, y_1, \dots, x_n, y_n) : 0 \leq x_i \leq y_i \text{ for } i = 1, \dots, n\}$ and $A = \Delta_n = \{(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}_{\geq 0}^{2n} : x_i = y_i \text{ for } i = 1, \dots, n\}$ for $n \geq 2$, one may consider the resulting spaces of generalized persistence diagrams as the parameter spaces for rectangle-decomposable persistent modules. These objects arise in multiparameter persistent homology and have received considerable attention in recent years (see, for example, [1, 2, 3, 18]).

The present work aims to fill some gaps in our basic understanding of the topological and geometric properties of the spaces of persistence diagrams equipped with the bottleneck distance. First of all, we note that, given a metric pair (X, A) , the space $\mathcal{D}_\infty(X, A)$ is not necessarily a metric space: the (generalized) bottleneck distance function d_∞ might be a pseudometric that is not a metric. This appears to be the case for most common examples of spaces of (generalized) persistence diagrams equipped with the bottleneck distance, including the usual space of euclidean persistence diagrams.

Theorem A (Metrizability). *The pseudometric space $\mathcal{D}_\infty(X, A)$ is a metric space if and only if $X \setminus A$ with the restricted metric of X is a discrete space.*

If X is complete, then $\mathcal{D}_p(X, A)$ is a complete metric space provided $1 \leq p < \infty$ (see [10, Theorem B (1)]). Here, we obtain the following characterization of completeness for $\mathcal{D}_p(X, A)$ for all $p \in [1, \infty]$.

Theorem B (Completeness). *For any $p \in [1, \infty]$, the space $\mathcal{D}_p(X, A)$ is complete if and only if X/A is complete.*

Corollary C. *For any $p \in [1, \infty]$, if X is complete, then so is $\mathcal{D}_p(X, A)$.*

We also consider the separability of $\mathcal{D}_\infty(X, A)$. For $1 \leq p < \infty$, if (X, A) is separable, then $\mathcal{D}_p(X, A)$ is also separable (see [10, Theorem B (2)]). This is no longer the case for $\mathcal{D}_\infty(X, A)$. Indeed, for instance, the usual space of euclidean persistence diagrams with the bottleneck distance, $\mathcal{D}_\infty(\mathbb{R}_{\geq 0}^2, \Delta)$, is not separable (see Example 5.3), although this particular fact seems to be known already (see [7, Theorem 5]). Here, we obtain the following characterization of separability for $\mathcal{D}_\infty(X, A)$. Given $r > 0$ and a subset $Y \subseteq X$, we let $B_r(Y)$ denote the open r -neighborhood of Y in X , i.e. $B_r(Y) = \{x \in X : d(x, y) < r \text{ for some } y \in Y\}$. Given $y \in X$, we write $B_r(y)$ for $B_r(\{y\})$.

Theorem D (Separability). *The pseudometric space $\mathcal{D}_\infty(X, A)$ is separable if and only if the subset $B_D(A) \setminus B_\delta(A) \subseteq X$ is totally bounded for all $D > \delta > 0$.*

Finally, we give a criterion for $\mathcal{D}_\infty(X, A)$ to be a geodesic pseudometric space when X is proper. Given $x, y \in X$, we call a rectifiable path $\gamma : [0, 1] \rightarrow X$ a *geodesic* from x to y if $\gamma(0) = x$, $\gamma(1) = y$, and γ has length $d(x, y)$. A pseudometric space X is *geodesic* if any two points of X can be joined by a geodesic.

Theorem E (Geodesicity). *Let X be a proper metric space. Then the pseudometric space $\mathcal{D}_\infty(X, A)$ is geodesic if and only if any two points $x, y \in X/A$ such that either $d(x, y) < d(x, A)$ or $y = A$ can be joined by a geodesic in X/A .*

Corollary F. *If X is proper and geodesic, then $\mathcal{D}_\infty(X, A)$ is geodesic.*

Our results are related to some of the theorems in [6], where the authors study the space $\overline{\mathcal{D}}_\infty(X, A)$ of persistence diagrams σ such that the subdiagram $\{x \in \sigma : d(x, A) \geq \delta\}$ is finite for every $\delta > 0$. We point out that $\overline{\mathcal{D}}_\infty(X, A)$ is a closed subspace of $\mathcal{D}_\infty(X, A)$. However, the spaces $\mathcal{D}_\infty(X, A)$ and $\overline{\mathcal{D}}_\infty(X, A)$ do not necessarily satisfy the same metric properties. For instance, if X is a metric space then so is $\overline{\mathcal{D}}_\infty(X, A)$ (see the proof of [10, Lemma 2.4(2)]); the analogous statement for $\mathcal{D}_\infty(X, A)$, however, is far from being true (see Theorem A). In [6, Theorem 6.2], the authors showed that if X is separable then so is $\overline{\mathcal{D}}_\infty(X, A)$, but this is not necessarily true for $\mathcal{D}_\infty(X, A)$ (see Example 5.3). Moreover, if A is distance minimizing (that is, for every $x \in X$ there exists $a \in A$ such that $d(x, a) = d(x, A)$) then there always exists an optimal bijection between two diagrams in $\overline{\mathcal{D}}_\infty(X, A)$ [6, Theorem 4.9], and if in addition X is geodesic then so is $\overline{\mathcal{D}}_\infty(X, A)$ [6, Theorem 5.10]. These two statements are no longer true for $\mathcal{D}_\infty(X, A)$: see Example 6.6. The characterization of completeness of $\overline{\mathcal{D}}_\infty(X, A)$ (see [6, Theorem 6.3 & Proposition 6.10]) does agree with the characterization for $\mathcal{D}_\infty(X, A)$ (see Theorem B); however, the proof given in [6] does not seem to be easy to adapt for the space $\mathcal{D}_\infty(X, A)$.

It is also worth mentioning that in [6] the authors consider extended pseudometric spaces (see Definition 2.1), which are not necessarily metric spaces. Our results also hold if we replace metric pairs with pseudometric ones, and our proofs can be readily adapted with minimal modifications. The sole exception is the proof of Theorem E (see Section 6), which relies on the uniqueness of ultralimits in metric spaces. Ultralimits in a pseudometric space need not be unique. Nevertheless, our arguments remain valid if we select ultralimit points arbitrarily from their possible values. For simplicity, we have stated all our results for metric pairs.

Our article is organized as follows. In Section 2, we recall the definition of the functors \mathcal{D}_p , along with necessary background on persistence diagram spaces. Section 3 contains the proof of Theorem A. We then prove Theorem B and Corollary C in Section 4. Finally, we prove Theorem D in Section 5, and Theorem E together with Corollary F in Section 6.

Acknowledgements. The authors wish to thank the referees for their helpful comments.

2. PRELIMINARIES

In this section, we collect preliminary material that we will use in the rest of the article. The contents of this section are based on [10], where the reader may find further details. We refer the reader to [8] for basic results on the geometry of metric spaces.

Definition 2.1. Let X be a set.

- (1) A map $d : X \times X \rightarrow [0, \infty]$ is an *extended pseudometric* on X if it is symmetric and satisfies the triangle inequality, and a *pseudometric* on X if, in addition, $d(x, y) < \infty$ for $x, y \in X$.
- (2) An (extended) pseudometric d on X such that $d(x, y) = 0$ if and only if $x = y$ is an (*extended*) *metric*. When d is an (extended) pseudometric on X , points at distance zero from each other give a partition of X , and d induces an (extended) metric in the corresponding quotient set.
- (3) An (*extended*) *pseudometric space* (X, d_X) is a set X together with an (extended) pseudometric d_X . An (*extended*) *metric space* (X, d_X) is a set X together with an (extended) metric d_X .

Definition 2.2. A *metric pair* (X, A) consists of a metric space (X, d_X) and a non-empty closed subset $A \subseteq X$. When A is a point, we will say that $(X, \{a_0\})$ is a *pointed metric space with basepoint* a_0 .

We now define the space $\mathcal{D}_\infty(X, A)$ of ∞ -persistence diagrams on a metric pair (X, A) . We will denote multisets by using two curly brackets $\{\{\cdot\}\}$. Persistence diagrams will be denoted by Greek letters.

Let (X, d) be a metric space. We denote by $\tilde{\mathcal{D}}(X)$ the set of countable multisets $\{\{x_1, x_2, \dots\}\}$ of elements of X , and equip it with the *bottleneck distance* \tilde{d}_∞ given by

$$(2.1) \quad \tilde{d}_\infty(\tilde{\sigma}, \tilde{\tau}) = \inf_{\phi: \tilde{\sigma} \rightarrow \tilde{\tau}} \sup_{x \in \tilde{\sigma}} d(x, \phi(x)),$$

where ϕ ranges over all bijections between $\tilde{\sigma}$ and $\tilde{\tau}$ in $\tilde{\mathcal{D}}(X)$. Here, by convention, we set $\inf \emptyset = \infty$, that is, $\tilde{d}_\infty(\tilde{\sigma}, \tilde{\tau}) = \infty$ whenever $\tilde{\sigma}$ and $\tilde{\tau}$ do not have the same cardinality. Observe that \tilde{d}_∞ is an extended pseudometric on $\tilde{\mathcal{D}}(X)$, making $(\tilde{\mathcal{D}}_\infty(X), \tilde{d}_\infty)$ into an extended pseudometric space. We denote the empty multiset by $\tilde{\sigma}_\emptyset$.

We may similarly define the *p-Wasserstein distance* \tilde{d}_p by

$$\tilde{d}_p(\tilde{\sigma}, \tilde{\tau}) = \inf_{\phi: \tilde{\sigma} \rightarrow \tilde{\tau}} \left(\sum_{x \in \tilde{\sigma}} d(x, \phi(x))^p \right)^{1/p}$$

for any $p \in [1, \infty)$, which are extended pseudometrics on $\tilde{\mathcal{D}}(X)$ as well.

Given $\tilde{\sigma}, \tilde{\tau} \in \tilde{\mathcal{D}}(X)$, we define an equivalence relation on $\tilde{\mathcal{D}}(X)$ by setting $\tilde{\sigma} \sim_A \tilde{\tau}$ whenever there exist $\tilde{\alpha}, \tilde{\beta} \in \tilde{\mathcal{D}}(A)$ such that $\tilde{\sigma} \sqcup \tilde{\alpha} = \tilde{\tau} \sqcup \tilde{\beta}$. Let $\mathcal{D}(X, A)$ be the quotient set $\tilde{\mathcal{D}}(X)/\sim_A$. Given $\tilde{\sigma} \in \tilde{\mathcal{D}}(X)$, we let σ be the equivalence class of $\tilde{\sigma}$ in $\mathcal{D}(X, A)$. Observe that σ_\emptyset , the equivalence class of $\tilde{\sigma}_\emptyset$, is also the equivalence class of any $\tilde{\alpha} \in \tilde{\mathcal{D}}(A)$.

The maps \tilde{d}_p induce functions $d_p: \mathcal{D}(X, A) \times \mathcal{D}(X, A) \rightarrow [0, \infty]$ given by

$$(2.2) \quad d_p(\sigma, \tau) = \inf_{\tilde{\alpha}, \tilde{\beta} \in \tilde{\mathcal{D}}(A)} \tilde{d}_p(\tilde{\sigma} \sqcup \tilde{\alpha}, \tilde{\tau} \sqcup \tilde{\beta}).$$

Definition 2.3. For any $p \in [1, \infty]$, the *space of p-persistence diagrams* $\mathcal{D}_p(X, A)$ on the pair (X, A) is the set of all $\sigma \in \mathcal{D}(X, A)$ such that $d_p(\sigma, \sigma_\emptyset) < \infty$.

In particular, we have $d_p(\sigma, \sigma_\emptyset) = (\sum_{x \in \tilde{\sigma}} d(x, A)^p)^{1/p}$ for any $p \in [1, \infty)$ and $d_\infty(\sigma, \sigma_\emptyset) = \sup_{x \in \tilde{\sigma}} d(x, A)$. The inequality $d_p(\sigma, \sigma_\emptyset) < \infty$ may be interpreted as the p -norm of the distances of points in σ to the set A (the total p -persistence, in the original setting of persistence homology) being finite.

Proposition 2.4. For any $p \in [1, \infty)$, the function d_p is an extended metric on $\mathcal{D}(X, A)$ and a metric on $\mathcal{D}_p(X, A)$. On the other hand, the function d_∞ is an extended pseudometric on $\mathcal{D}(X, A)$ and a pseudometric on $\mathcal{D}_\infty(X, A)$.

For simplicity, we will treat elements in $\mathcal{D}(X, A)$ as multisets by themselves, with the understanding that, whenever we do so, we are actually dealing with representatives of such elements in $\tilde{\mathcal{D}}(X, A)$. Thus, for instance, when we write $x \in \sigma$, for $\sigma \in \mathcal{D}(X, A)$, or consider bijections $\phi: \sigma \rightarrow \tau$, for $\sigma, \tau \in \mathcal{D}(X, A)$, we mean $x \in \tilde{\sigma}$ and $\phi: \tilde{\sigma} \rightarrow \tilde{\tau}$ is a bijection for suitable representatives $\tilde{\sigma}, \tilde{\tau} \in \tilde{\mathcal{D}}(X)$. Observe that the spaces $\mathcal{D}_p(X, A)$ and $\mathcal{D}_p(X/A, \{A\})$ are naturally isometric for any $p \in [1, \infty]$.

3. METRIZABILITY: PROOF OF THEOREM A

Theorem A follows from the following two lemmas. Throughout, we let (X, A) be a metric pair.

Lemma 3.1. If $X \setminus A$ is discrete, then $d_\infty(\sigma, \tau) \neq 0$ for any $\sigma, \tau \in \mathcal{D}_\infty(X, A)$ with $\sigma \neq \tau$.

Proof. Since $\sigma \neq \tau$, there exists a point $x \in X \setminus A$ such that if $m_\sigma, m_\tau \in \mathbb{N} \cup \{\aleph_0\}$ are the multiplicities of x in σ, τ (respectively), then $m_\sigma \neq m_\tau$; without loss of generality, $m_\sigma > m_\tau$. As $X \setminus A$ is discrete, $\{x\}$ is open in $X \setminus A$, implying that $B_\varepsilon(x) \subseteq \{x\} \cup A$ for some $\varepsilon > 0$. As $x \notin A$ and A is closed, we have $d(x, A) > 0$ and therefore we may assume (without loss of generality) that $\varepsilon \leq d(x, A)$; it then follows that $B_\varepsilon(x) = \{x\}$.

We claim that $d_\infty(\sigma, \tau) \geq \varepsilon$. Indeed, as $m_\sigma > m_\tau$, any bijection $\gamma: \sigma \rightarrow \tau$ must send a copy of x in σ to a point y in τ such that $x \neq y$ in X , and therefore $d(x, y) \geq \varepsilon$. Thus $\sup_{z \in \sigma} d(z, \gamma(z)) \geq \varepsilon$; taking the infimum over all bijections γ then gives $d_\infty(\sigma, \tau) \geq \varepsilon > 0$, as claimed. \square

Lemma 3.2. *If $X \setminus A$ is not discrete, then there exist $\sigma, \tau \in \mathcal{D}_\infty(X, A)$ such that $\sigma \neq \tau$ and $d_\infty(\sigma, \tau) = 0$.*

Proof. Since $X \setminus A$ is not discrete, there exists a point $x \in X \setminus A$ such that $\{x\}$ is not open in $X \setminus A$. It follows that there exists a sequence $(x_n)_{n=0}^\infty$ in X such that $x_n \neq x$ for any n but $x_n \rightarrow x$ as $n \rightarrow \infty$. Now let $\sigma = \{\{x_n : n \in \mathbb{N}\}\}$ and $\tau = \sigma \sqcup \{\{x\}\}$. Then clearly $\sigma, \tau \in \mathcal{D}_\infty(X, A)$ (as the sequence $(d(x_n, A))_{n=0}^\infty$ in \mathbb{R} converges to $d(x, A)$ and so is bounded), and $\sigma \neq \tau$ since $x \notin \sigma$ but $x \in \tau$.

Now let $\varepsilon > 0$, and let $N = N_\varepsilon \in \mathbb{N}$ be such that $d(x_n, x) \leq \varepsilon/2$ for any $n \geq N$. Define a bijection $\gamma_\varepsilon: \sigma \rightarrow \tau$ by $\gamma_\varepsilon(x_n) = x_n$ for $n < N$, $\gamma_\varepsilon(x_N) = x$, and $\gamma_\varepsilon(x_n) = x_{n-1}$ for $n > N$. We then have $d(x_n, \gamma_\varepsilon(x_n)) = 0$ for $n < N$, $d(x_N, \gamma_\varepsilon(x_N)) = d(x_N, x) \leq \varepsilon/2$, and $d(x_n, \gamma_\varepsilon(x_n)) = d(x_n, x_{n-1}) \leq d(x_n, x) + d(x_{n-1}, x) \leq \varepsilon$ for $n > N$. This implies that $\sup_{y \in \sigma} d(y, \gamma_\varepsilon(y)) \leq \varepsilon$, and therefore $d_\infty(\sigma, \tau) \leq \varepsilon$. As ε was arbitrary, it follows that $d_\infty(\sigma, \tau) = 0$, as required. \square

4. COMPLETENESS: PROOFS OF THEOREM B AND COROLLARY C

Theorem B follows from the following two lemmas.

Lemma 4.1. *For any $p \in [1, \infty]$, if $\mathcal{D}_p(X, \{a_0\})$ is complete, then so is X .*

Proof. Let $(x_n)_{n=0}^\infty$ be a Cauchy sequence in X . Then $(d(x_n, a_0))_{n=0}^\infty$ is a Cauchy sequence in \mathbb{R} , and so converges. If $d(x_n, a_0) \rightarrow 0$ as $n \rightarrow \infty$ then $x_n \rightarrow a_0$ and so the sequence (x_n) converges. Thus, we may assume that $d(x_n, a_0) \rightarrow \delta$ as $n \rightarrow \infty$ for some $\delta > 0$.

Now for each $n \in \mathbb{N}$ let $\sigma_n = \{\{x_n\}\} \in \mathcal{D}_p(X, \{a_0\})$. For each $n, m \in \mathbb{N}$ there is a bijection $\sigma_n \rightarrow \sigma_m$ sending $x_n \mapsto x_m$, implying that $d_p(\sigma_n, \sigma_m) \leq d(x_n, x_m)$. Thus, $(\sigma_n)_{n=0}^\infty$ is a Cauchy sequence in $\mathcal{D}_p(X, \{a_0\})$ and so converges to some $\sigma \in \mathcal{D}_p(X, \{a_0\})$. Now let $\varepsilon \in (0, \delta/2]$. Then there exists $N_\varepsilon \in \mathbb{N}$ such that $d_p(\sigma_n, \sigma) < \varepsilon$, and so there exists a bijection $\gamma_n: \sigma_n \rightarrow \sigma$ such that $d(x', \gamma_n(x')) < \varepsilon \leq \delta/2$ for every $x' \in \sigma_n$. This implies that σ contains a unique point $x \in X \setminus B_{\delta/2}(a_0)$ and that $\gamma_n(x_n) = x$, and hence $d(x_n, x) < \varepsilon$, for every $n \geq N_\varepsilon$. Therefore, the sequence (x_n) converges to $x \in X$, as required. \square

Lemma 4.2. *Let $\ell_\infty(X, a_0)$ be the space of all sequences $(x_n)_{n=0}^\infty$ in X such that $\sup_{n \in \mathbb{N}} d(x_n, a_0) < \infty$, equipped with the (pseudo-)metric $\widehat{d}_\infty((x_n), (y_n)) = \sup_{n \in \mathbb{N}} d(x_n, y_n)$. Let $\Phi: \ell_\infty(X, a_0) \rightarrow \mathcal{D}_\infty(X, \{a_0\})$ be the map sending each sequence to the (equivalence class of) the corresponding multiset. Then the following statements hold:*

- (1) Φ is a continuous map.
- (2) If X is complete then $\ell_\infty(X, a_0)$ is complete.
- (3) Let $(\sigma_n)_{n=0}^\infty$ be a Cauchy sequence in $\mathcal{D}_\infty(X, \{a_0\})$. Then there exists a Cauchy sequence $(\widehat{\sigma}_n)_{n=0}^\infty$ in $\ell_\infty(X, a_0)$ such that $\Phi(\widehat{\sigma}_n) = \sigma_n$ for all n .

Proof. Clearly $d_\infty(\Phi(\widehat{\sigma}), \Phi(\widehat{\tau})) \leq \widehat{d}_\infty(\widehat{\sigma}, \widehat{\tau})$ for any $\widehat{\sigma}, \widehat{\tau} \in \ell_\infty(X, a_0)$; in particular, Φ is continuous. Moreover, it is easy to see that the arguments in the “usual” proof of the fact that $\ell_\infty = \ell_\infty(\mathbb{R}, 0)$ is complete may be adapted to proof that if X is complete, then so is $\ell_\infty(X, a_0)$.

Now let $(\sigma_n)_{n=0}^\infty$ be a Cauchy sequence in $\mathcal{D}_\infty(X, \{a_0\})$. Then for each $k \in \mathbb{N}$ there exists $N_k \in \mathbb{N}$ such that $d_\infty(\sigma_n, \sigma_m) \leq 2^{-k}$ for all $n, m \geq N_k$. In particular, for each $k \in \mathbb{N}$ and $n \geq N_k$ there

exists a bijection $\gamma_{k,n}: \sigma_{N_k} \rightarrow \sigma_n$ such that $d(x, \gamma_{k,n}(x)) \leq 2^{1-k}$ for all $x \in \sigma_{N_k}$. Without loss of generality, assume that $N_k \leq N_\ell$ whenever $k \leq \ell$ and that $N_k \rightarrow \infty$ as $k \rightarrow \infty$.

For each n , let $\tilde{\sigma}_n$ be a countable multiset of points in X representing σ_n such that $\tilde{\sigma}_n$ contains a_0 with infinite multiplicity. Then, for each $k \in \mathbb{N}$ and $n \geq N_k$, there exists a bijection $\tilde{\gamma}_{k,n}: \tilde{\sigma}_{N_k} \rightarrow \tilde{\sigma}_n$ inducing $\gamma_{k,n}$; in particular, we have $d(x, \tilde{\gamma}_{k,n}(x)) \leq 2^{1-k}$ for all $x \in \tilde{\sigma}_{N_k}$. We now construct a sequence $(\hat{\sigma}_n)_{n=0}^\infty$ in $\ell_\infty(X, a_0)$ as follows:

- for $n \leq N_0$, pick an arbitrary bijection $\psi_n: \mathbb{N} \rightarrow \tilde{\sigma}_n$ and let $\hat{\sigma}_n = (\psi_n(m))_{m=0}^\infty$;
- for $n > N_0$, let $k \in \mathbb{N}$ be such that $N_k < n \leq N_{k+1}$ and let

$$\hat{\sigma}_n = ((\tilde{\gamma}_{k,n} \circ \tilde{\gamma}_{k-1,N_k} \circ \cdots \circ \tilde{\gamma}_{1,N_2} \circ \tilde{\gamma}_{0,N_1} \circ \psi_{N_0})(m))_{m=0}^\infty;$$

see Figure 1. It follows from the construction that $\Phi(\hat{\sigma}_n) = \sigma_n$ for all n . We claim that $(\hat{\sigma}_n)_{n=0}^\infty$ is a Cauchy sequence.

Indeed, let $\varepsilon > 0$, and let $k \in \mathbb{N}$ be such that $2^{3-k} \leq \varepsilon$. Note that for any $k', n \in \mathbb{N}$ such that $N_{k'} < n \leq N_{k'+1}$ we have $\hat{d}_\infty(\hat{\sigma}_{N_{k'}}, \hat{\sigma}_n) \leq \sup_{x \in \tilde{\sigma}_{N_{k'}}} d(x, \tilde{\gamma}_{k',n}(x)) \leq 2^{1-k'}$. In particular, if in addition $k' \geq k$ then we have

$$\begin{aligned} \hat{d}_\infty(\hat{\sigma}_{N_k}, \hat{\sigma}_n) &\leq \left[\sum_{\ell=k}^{k'-1} \hat{d}_\infty(\hat{\sigma}_{N_\ell}, \hat{\sigma}_{N_{\ell+1}}) \right] + \hat{d}_\infty(\hat{\sigma}_{N_{k'}}, \hat{\sigma}_n) \leq \sum_{\ell=k-1}^{k'-1} 2^{-\ell} \\ &< \sum_{\ell=k-1}^\infty 2^{-\ell} = 2^{2-k} \leq \frac{\varepsilon}{2} \end{aligned}$$

for any $n > N_k$. Therefore, for all $n, m > N_k$ we have $\hat{d}_\infty(\hat{\sigma}_n, \hat{\sigma}_m) \leq \hat{d}_\infty(\hat{\sigma}_{N_k}, \hat{\sigma}_n) + \hat{d}_\infty(\hat{\sigma}_{N_k}, \hat{\sigma}_m) \leq \varepsilon$. Thus the sequence $(\hat{\sigma}_n)_{n=0}^\infty$ is Cauchy, as claimed. \square




FIGURE 1. The proof of Lemma 4.2: construction of the sequence $(x_n)_{n=0}^\infty = ((\hat{\sigma}_n)_m)_{n=0}^\infty$ of points in X for a fixed $m \in \mathbb{N}$.

Proof of Theorem B. Let $p \in [1, \infty]$ and let (X, A) be a metric pair. In view of the isometry $\mathcal{D}_p(X, A) \cong \mathcal{D}_p(X/A, \{A\})$, $p \in [1, \infty]$, Theorem B is equivalent to the statement that $\mathcal{D}_p(X, \{a_0\})$ is complete if and only if X is complete, where $a_0 \in X$ is any point. Thus, the “only if” implication of Theorem B follows from Lemma 4.1. Suppose now that X is complete. If $1 \leq p < \infty$, then $\mathcal{D}_p(X, A)$ is complete by [10, Theorem B (1)]. If $p = \infty$, then it follows from Lemma 4.2 that $\mathcal{D}_\infty(X, A)$ is complete. Thus, $\mathcal{D}_p(X, A)$ is complete for all $p \in [1, \infty]$. \square

We now prove Corollary C.

Proof of Corollary C. By Theorem B, it is enough to show that if X is complete then so is X/A . Thus, let $(x_n)_{n=0}^\infty$ be a Cauchy sequence in X/A . Then $(d_{X/A}(x_n, A))_{n=0}^\infty$ is a Cauchy sequence in \mathbb{R} , and so converges. If $d_{X/A}(x_n, A) \rightarrow 0$ as $n \rightarrow \infty$ then $x_n \rightarrow A$. Otherwise, $d_{X/A}(x_n, A) \rightarrow \delta$

as $n \rightarrow \infty$ for some $\delta > 0$. As (x_n) is Cauchy, it follows that for each $\varepsilon \in (0, \delta/2)$ there exists N_ε such that for all $n, m \geq N_\varepsilon$ we have $d_{X/A}(x_n, x_m) \leq \varepsilon < \delta/2$, implying that $x_n, x_m \notin B_{\delta/2}(A)$ (in particular, $x_n, x_m \neq A$) and $d_X(x_n, x_m) = d_{X/A}(x_n, x_m)$. In particular, (x_n) is also a Cauchy sequence in X , and so converges to some $x \in X \setminus A$; but then we also have $x_n \rightarrow x$ in X/A as $n \rightarrow \infty$. \square

5. SEPARABILITY: PROOF OF THEOREM D

Theorem D follows from the following two results.

Lemma 5.1. *If $B_D(A) \setminus B_\delta(A)$ is totally bounded for all $D > \delta > 0$, then $\mathcal{D}_\infty(X, A)$ is separable.*

Proof. Let $n \in \mathbb{N}_{\geq 1}$. Then $B_n(A) \setminus B_{1/n}(A)$ is totally bounded, and so there exist $x_1^{(n)}, \dots, x_{k_n}^{(n)} \in X$ such that $B_n(A) \setminus B_{1/n}(A) \subseteq \bigcup_{m=1}^{k_n} B_{1/n}(x_m^{(n)})$. For $m \in \{1, \dots, k_n\}$, we define $U_m^{(n)}$ inductively as

$$U_m^{(n)} = \left[B_n(A) \cap B_{1/n}(x_m^{(n)}) \right] \setminus \left[B_{1/n}(A) \cap B_{1/n}(x_m^{(n)}) \cap \bigcup_{\ell=1}^{m-1} U_\ell^{(n)} \right].$$

We thus get a collection of pairwise disjoint subsets $U_1^{(n)}, \dots, U_{k_n}^{(n)} \subseteq X$ such that $B_n(A) \setminus B_{1/n}(A) = \bigcup_{m=1}^{k_n} U_m^{(n)}$ and $U_m^{(n)} \subseteq B_{1/n}(x_m^{(n)})$.

Now let $\mathcal{F}_n \subseteq \mathcal{D}_\infty(X, A)$ be the set of diagrams in which each point (outside A) is equal to $x_m^{(n)}$ for some m . For each m , the multiplicity of $x_m^{(n)}$ is in $\mathbb{N} \cup \{\aleph_0\}$, leading to $|\mathbb{N} \cup \{\aleph_0\}| = \aleph_0$ choices. Thus the cardinality of \mathcal{F}_n is $\aleph_0^{k_n} = \aleph_0$. Now let $\mathcal{F} = \bigcup_{n=1}^\infty \mathcal{F}_n$, so that \mathcal{F} is a countable union of countable sets and thus countable. We claim that \mathcal{F} is dense in $\mathcal{D}_\infty(X, A)$.

Indeed, let $\sigma \in \mathcal{D}_\infty(X, A)$, and let $\varepsilon > 0$. Pick $n \in \mathbb{N}$ such that $n > d_\infty(\sigma, \sigma_\emptyset)$ and $\frac{1}{n} \leq \varepsilon$. Let $\tau \in \mathcal{F}_n \subseteq \mathcal{F}$ be such that the multiplicity of $x_m^{(n)}$ in τ is equal to the number of points of σ lying inside $U_m^{(n)}$. Then $d_\infty(\sigma, \tau) \leq \frac{1}{n} \leq \varepsilon$: indeed, there is a bijection $\sigma \rightarrow \tau$ sending any point in $U_m^{(n)}$ to $x_m^{(n)}$ and any point in $B_{1/n}(A)$ to a point in A (this leaves no points of σ unmatched as $d_\infty(\sigma, \sigma_\emptyset) < n$ implies that all points of σ lie in $B_n(A)$). Therefore, \mathcal{F} is dense in $\mathcal{D}_\infty(X, A)$, as claimed. \square

Lemma 5.2. *If there exists $D > \delta > 0$ such that $B_D(A) \setminus B_\delta(A)$ is not totally bounded, then $\mathcal{D}_\infty(X, A)$ is not separable.*

Proof. Since $B_D(A) \setminus B_\delta(A)$ is not totally bounded, there exists $\varepsilon > 0$ such that $B_D(A) \setminus B_\delta(A)$ cannot be covered by finitely many ε -balls; without loss of generality, assume that $\varepsilon \leq \delta$. Now let $\mathcal{X} \subseteq B_D(A) \setminus B_\delta(A)$ be a maximal subset such that $d(x, y) \geq \varepsilon$ for all $x, y \in \mathcal{X}$ with $x \neq y$ (such an \mathcal{X} exists by Zorn's Lemma). Then ε -balls around the points of \mathcal{X} cover $B_D(A) \setminus B_\delta(A)$, implying that \mathcal{X} is infinite. In particular, there exists a collection $\{x_i : i \in \mathbb{N}\}$ of points in $B_D(A) \setminus B_\delta(A)$ such that $d(x_i, x_j) \geq \varepsilon$ whenever $i \neq j$.

Let $\{\sigma_i : i \in \mathbb{N}\}$ be a countable subset of $\mathcal{D}_\infty(X, A)$. Define a diagram τ by

$$\tau = \{\{x_i : d(x, x_i) \geq \varepsilon/2 \text{ for all } x \in \sigma_i\}\};$$

since $d(x_i, A) < D$ for all i we have $\tau \in \mathcal{D}_\infty(X, A)$. Note that since $d(x_i, A) \geq \delta > \varepsilon/2$ for each i , the definition of τ is independent of the choice of representative of σ_i (i.e. adding/removing points in A to/from σ_i does not change τ). We now claim that $d_\infty(\tau, \sigma_i) \geq \varepsilon/2$ for each i , so that τ is not in the closure of $\{\sigma_i : i \in \mathbb{N}\}$ in $\mathcal{D}_\infty(X, A)$; this will imply that $\{\sigma_i : i \in \mathbb{N}\}$ is not dense in $\mathcal{D}_\infty(X, A)$, as required.

To prove our claim, let $i \in \mathbb{N}$ and let $\gamma : \tau \rightarrow \sigma_i$ be a bijection.

Case 1: Suppose that $d(x, x_i) \geq \varepsilon/2$ for all $x \in \sigma_i$, and so $x_i \in \tau$. We then have $d(x_i, x) \geq \varepsilon/2$ for all $x \in \sigma_i$ with $x \notin A$, and $d(x_i, A) \geq \delta \geq \varepsilon > \varepsilon/2$ since $x_i \notin B_\delta(A)$. Thus we must have $d(x_i, \gamma(x_i)) \geq \varepsilon/2$.

Case 2: Suppose that $d(x, x_i) < \varepsilon/2$ for some $x \in \sigma_i$, and so $x_i \notin \tau$. Then every point of τ is either in A , or equal to x_j for some $j \neq i$. But we have $d(x, A) \geq d(x_i, A) - d(x, x_i) > \delta - \varepsilon/2 \geq \varepsilon/2$ since $\delta \geq \varepsilon$, and $d(x, x_j) \geq d(x_i, x_j) - d(x, x_i) > \varepsilon - \varepsilon/2 = \varepsilon/2$. Therefore, we must have $d(\gamma^{-1}(x), x) > \varepsilon/2$.

It follows that in either case there exists $y \in \tau$ such that $d(y, \gamma(y)) \geq \varepsilon/2$, and in particular $\sup_{y \in \tau} d(y, \gamma(y)) \geq \varepsilon/2$. Taking the infimum over all bijections $\gamma: \tau \rightarrow \sigma_i$ gives $d_\infty(\tau, \sigma_i) \geq \varepsilon/2$, as claimed. \square

Example 5.3. Consider $\mathcal{D}_\infty(\mathbb{R}_{\geq 0}^2, \Delta)$, the usual space of persistence diagrams. Then for any $D > \delta > 0$, the “strip” $B_D(\Delta) \setminus B_\delta(\Delta)$ is unbounded and, hence, not totally bounded. It follows that $\mathcal{D}_\infty(\mathbb{R}_{\geq 0}^2, \Delta)$ is not separable.

Example 5.4. Consider $\mathcal{D}_\infty([0, \infty), \{0\})$. Then, for any $D > \delta > 0$, the set $B_D(\{0\}) \setminus B_\delta(\{0\})$ is just the interval $[\delta, D)$, which is clearly totally bounded. Thus, $\mathcal{D}_\infty([0, \infty), \{0\})$ is separable.

6. GEODESICITY: PROOF OF THEOREM E

In the proof of Theorem E we use the concepts of ultrafilters and ultralimits, defined as follows.

Definition 6.1. An *ultrafilter* (on \mathbb{N}) is a non-zero function $\omega: \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$, where $\mathcal{P}(\mathbb{N})$ denotes the power set of \mathbb{N} , that is *finitely additive*—that is, $\omega(A) + \omega(B) = \omega(A \cup B)$ for any disjoint subsets $A, B \subseteq \mathbb{N}$. A *principal ultrafilter* is an indicator function $\mathbb{1}_n$ (for some $n \in \mathbb{N}$), defined by setting $\mathbb{1}_n(A) = 1$ if and only if $n \in A$. Given a sequence $(x_n)_{n=1}^\infty$ in a metric space X and an ultrafilter ω , an *ultralimit* $\lim_\omega x_n$ of (x_n) with respect to ω is a point $x \in X$ such that $\omega(\{n \in \mathbb{N} : d(x, x_n) < \varepsilon\}) = 1$ for every $\varepsilon > 0$.

It is easy to check, using finite additivity, that an ultrafilter ω is non-principal if and only if $\omega(F) = 0$ for every finite subset $F \subset \mathbb{N}$. The existence of a non-principal ultrafilter is guaranteed by the axiom of choice [15, §3.1]. We refer the interested reader to [14] for an introduction to the theory of ultrafilters and their use throughout mathematics.

In any metric space X , an ultralimit of a sequence (x_n) is always unique (when it exists): indeed, if $x \neq y$ were two different ultralimits of (x_n) then ω would send both $\{n \in \mathbb{N} : d(x, x_n) < \varepsilon\}$ and $\{n \in \mathbb{N} : d(y, x_n) < \varepsilon\}$ to 1, where $\varepsilon = \frac{1}{2}d(x, y)$, contradicting finite additivity of ω . Moreover, if X is compact then such an ultralimit will always exist [15, §3.1]. If ω is a non-principal ultrafilter, then clearly $\lim_\omega x_n$ will be an accumulation point of the sequence (x_n) .

Our main use of ultrafilters and ultralimits is based on their ability to pick out accumulation points of sequences *consistently*. For instance, roughly speaking, in the proof of Proposition 6.2 below we aim to construct a bijection $\tilde{\gamma}: \tilde{\sigma} \rightarrow \tilde{\tau}$ as a ‘limit’ of a sequence of bijections $\tilde{\gamma}_n: \tilde{\sigma} \rightarrow \tilde{\tau}$. Given a point $x \in \tilde{\sigma}$ it may turn out, for example, that there exists a subset $Y_x \subseteq \tilde{\tau}$ such that $\{n \in \mathbb{N} : \tilde{\gamma}_n(x) = y\}$ is infinite for all $y \in Y_x$; analogously, given $y \in \tilde{\tau}$, there may exist $X_y \subseteq \tilde{\sigma}$ such that $\{n \in \mathbb{N} : \tilde{\gamma}_n^{-1}(y) = x\}$ is infinite for all $x \in X_y$. It is then unclear which (if any) point of Y_x (respectively X_y) should be declared to be $\tilde{\gamma}(x)$ (respectively $\tilde{\gamma}^{-1}(y)$), doing this procedure for all $x \in \tilde{\sigma}$ and $y \in \tilde{\tau}$ in a consistent way: that is, if $y \in Y_x$ is declared to be $\tilde{\gamma}(x)$ then $x \in X_y$ needs to be declared to be $\tilde{\gamma}^{-1}(y)$, and vice versa. Using ultralimits with respect to a non-principal ultrafilter ω allows us to circumvent such an issue—in particular, given $x \in X$ at most one of the subsets $\{n \in \mathbb{N} : \tilde{\gamma}_n(x) = y\}$ for $y \in Y_x$ is sent to 1 by ω , allowing us to pick out the ‘correct’ $y \in Y_x$ unambiguously.

The main ingredient in the proof of Theorem E is the following result.

Proposition 6.2. *Let X be a proper metric space and let $\sigma, \tau \in \mathcal{D}_\infty(X, A)$. Then there exist diagrams $\sigma', \tau' \in \mathcal{D}_\infty(X, A)$ such that $d_\infty(\sigma, \sigma') = d_\infty(\tau, \tau') = 0$ and a bijection $\gamma: \sigma' \rightarrow \tau'$ such that $d_\infty(\sigma, \tau) = \sup\{d(x, \gamma(x)) : x \in \sigma'\}$.*

Proof. By the definition of the pseudometric d_∞ , there is a sequence of bijections $(\gamma_n: \sigma \rightarrow \tau)_{n=1}^\infty$ such that $d(x, \gamma_n(x)) \leq d_\infty(\sigma, \tau) + \frac{1}{n}$ for all n and all $x \in \sigma$. We can lift each γ_n to a bijection $\tilde{\gamma}_n: \tilde{\sigma}_n \rightarrow \tilde{\tau}_n$ for some countable multisets $\tilde{\sigma}_n$ and $\tilde{\tau}_n$ in X representing diagrams σ and τ , respectively. Now let $\tilde{\sigma}$ (respectively $\tilde{\tau}$) be the multiset in X consisting of the multiset $\{\{x \in \sigma : x \notin A\}\}$ (respectively $\{\{x \in \tau : x \notin A\}\}$) together with a countably infinite number of copies of each point $a \in A$ appearing in $\tilde{\sigma}_n \cup \tilde{\tau}_n$ for some $n \geq 1$. It is easy to see that $\tilde{\sigma}$ and $\tilde{\tau}$ are countable. Moreover, since any point $a \in A$ appears in $\tilde{\sigma}$ and in $\tilde{\tau}$ with the same multiplicity (zero or infinite), with a having infinite multiplicity whenever it appears in $\tilde{\sigma}_n$ or in $\tilde{\tau}_n$, it follows that the bijection $\tilde{\gamma}_n: \tilde{\sigma}_n \rightarrow \tilde{\tau}_n$ can be extended to a bijection $\tilde{\gamma}_n: \tilde{\sigma} \rightarrow \tilde{\tau}$ such that $\tilde{\gamma}_n(x) = x$ in X whenever $x \notin \tilde{\sigma}_n$. In particular, such an extended bijection $\tilde{\gamma}_n: \tilde{\sigma} \rightarrow \tilde{\tau}$ satisfies $d(x, \tilde{\gamma}_n(x)) \leq d_\infty(\sigma, \tau) + \frac{1}{n}$ for all $x \in \tilde{\sigma}$.

Now let $\omega: \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$ be a non-principal ultrafilter. We call a point $x \in \tilde{\sigma}$ *convergent* if there exists $y_x \in \tilde{\tau}$ such that $\tilde{\gamma}_n(x) = y_x$ ω -almost surely (that is, ω sends $\{n \in \mathbb{N} : \tilde{\gamma}_n(x) = y_x\}$ to 1), and we let $\tilde{\sigma}_c \subseteq \tilde{\sigma}$ be the multiset of convergent points in $\tilde{\sigma}$. Similarly, we call a point $y \in \tilde{\tau}$ *convergent* if $\tilde{\gamma}_n^{-1}(y) = x_y$ ω -almost surely for some $x_y \in \tilde{\sigma}$, and we let $\tilde{\tau}_c \subseteq \tilde{\tau}$ be the multiset of convergent points in $\tilde{\tau}$. Note that given any $x \in \tilde{\sigma}$ and $y \in \tilde{\tau}$, we have $x = x_y$ if and only if $y = y_x$.

Given $x \in \tilde{\sigma} \setminus \tilde{\sigma}_c$, note that $B_{d_\infty(\sigma, \tau)+1}(x)$ is compact (as X is proper), and so the sequence $(\tilde{\gamma}_n(x))_{n=1}^\infty \subseteq B_{d_\infty(\sigma, \tau)+1}(x)$ has a (unique) ultralimit $y_x^\omega := \lim_\omega \tilde{\gamma}_n(x)$ in X . Similarly, given $y \in \tilde{\tau} \setminus \tilde{\tau}_c$ we can define $x_y^\omega := \lim_\omega \tilde{\gamma}_n^{-1}(y)$; see Figure 2. Now let $\tilde{\sigma}' := \tilde{\sigma} \sqcup \{\{x_y^\omega : y \in \tilde{\tau} \setminus \tilde{\tau}_c\}\}$ and $\tilde{\tau}' := \tilde{\tau} \sqcup \{\{x_x^\omega : x \in \tilde{\sigma} \setminus \tilde{\sigma}_c\}\}$. We define a mapping $\tilde{\gamma}: \tilde{\sigma}' \rightarrow \tilde{\tau}'$ by setting $\tilde{\gamma}(x) = y_x$ for $x \in \tilde{\sigma}_c$, $\tilde{\gamma}(x) = y_x^\omega$ for $x \in \tilde{\sigma} \setminus \tilde{\sigma}_c$, and $\tilde{\gamma}(x_y^\omega) = y$ for $y \in \tilde{\tau} \setminus \tilde{\tau}_c$. Finally, let $\sigma', \tau' \in \mathcal{D}_\infty(X, A)$ be the diagrams represented by $\tilde{\sigma}', \tilde{\tau}'$ (respectively), and let $\gamma: \sigma' \rightarrow \tau'$ be the mapping induced by $\tilde{\gamma}$.

Note that $\tilde{\gamma}$ (and so γ) is a well-defined bijection by construction. Moreover, for any $x \in \tilde{\sigma}'$ we have $d(x, \tilde{\gamma}(x)) \leq d_\infty(\sigma, \tau)$: indeed, we have

$$d(x, \tilde{\gamma}(x)) = \lim_\omega d(x, \tilde{\gamma}_n(x)) \leq \limsup_{n \rightarrow \infty} d(x, \tilde{\gamma}_n(x)) \leq d_\infty(\sigma, \tau)$$

if $x \in \tilde{\sigma}$, and

$$d(x, \tilde{\gamma}(x)) = d(y, x_y^\omega) = \lim_\omega d(y, \tilde{\gamma}_n^{-1}(y)) \leq \limsup_{n \rightarrow \infty} d(y, \tilde{\gamma}_n^{-1}(y)) \leq d_\infty(\sigma, \tau)$$

if $x = x_y^\omega$ for some $y \in \tilde{\tau} \setminus \tilde{\tau}_c$; in both cases, the first inequality follows from the fact that given a sequence $(r_n)_{n=1}^\infty$ in the interval $[0, d_\infty(\sigma, \tau) + 1]$, the ultralimit $\lim_\omega r_n$ is an accumulation point of (r_n) . This implies that $\sup\{d(x, \gamma(x)) : x \in \sigma'\} = \sup\{d(x, \tilde{\gamma}(x)) : x \in \tilde{\sigma}'\} \leq d_\infty(\sigma, \tau)$. Moreover, by the definition of d_∞ we have $\sup\{d(x, \gamma(x)) : x \in \sigma'\} \geq d_\infty(\sigma', \tau')$. Since $|d_\infty(\sigma, \tau) - d_\infty(\sigma', \tau')| \leq d_\infty(\sigma, \sigma') + d_\infty(\tau, \tau')$, it remains to show that $d_\infty(\sigma, \sigma') = d_\infty(\tau, \tau') = 0$.

We now claim that $d_\infty(\sigma, \sigma') = 0$ (the proof that $d_\infty(\tau, \tau') = 0$ is analogous). Indeed, let $\varepsilon > 0$, and let $N \in \mathbb{N}$ be such that $\frac{2}{N} \leq \varepsilon$. We will find a bijection $\tilde{\delta}_N: \tilde{\sigma} \rightarrow \tilde{\sigma}'$ such that $d(x, \tilde{\delta}_N(x)) \leq \frac{2}{N} \leq \varepsilon$ for any $x \in \tilde{\sigma}$.

Let $y \in \tilde{\tau} \setminus \tilde{\tau}_c$. We claim that there exists a sequence $(x_y^{(n)})_{n=1}^\infty$ of points in $\tilde{\sigma}$ such that $x_y^{(n)} \rightarrow x_y^\omega$ as $n \rightarrow \infty$ and $x_y^{(n)} \neq x_y^{(m)}$ (as elements of $\tilde{\sigma}$) when $n \neq m$. Indeed, in order to show this it is enough to show that for each $\varepsilon > 0$, the submultiset $X_\varepsilon = \{\{x \in \tilde{\sigma} : d(x, x_y^\omega) < \varepsilon\}\}$ of $\tilde{\sigma}$ is infinite. Suppose for contradiction that this is not the case, and so $X_\varepsilon = \{\{x_1, \dots, x_k\}\}$ for some $\varepsilon > 0$, and for $1 \leq i \leq k$ let $A_i = \{n \in \mathbb{N} : \tilde{\gamma}_n^{-1}(y) = x_i\}$. Then $\omega(A_i) = 0$ for each i since $y \notin \tilde{\tau}_c$; however, we have $\omega(A_1 \cup \dots \cup A_k) = \omega(\{n \in \mathbb{N} : d(\tilde{\gamma}_n^{-1}(y), x_y^\omega) < \varepsilon\}) = 1$ by the definition of x_y^ω . This contradicts the fact that ω is finitely additive; thus X_ε must be infinite, and so a sequence $(x_y^{(n)})$ with the claimed properties does exist.

Moreover, we can pass each of the sequences $(x_y^{(n)})_{n=1}^\infty$ for $y \in \tilde{\tau} \setminus \tilde{\tau}_c$ to a subsequence so that $x_y^{(n)} \neq x_{y'}^{(m)}$ (as elements of $\tilde{\sigma}$) whenever $(y, n) \neq (y', m)$: for instance, we may pick an enumeration $\tilde{\tau} \setminus \tilde{\tau}_c = \{y_1, y_2, \dots\}$ and inductively build the collection of tuples $\{(x_{y_i}^{(1)}, \dots, x_{y_i}^{(n)}) : 1 \leq i \leq n\}$ with no repeating elements. Finally, after passing each of the sequences $(x_y^{(n)})_{n=1}^\infty$ to a subsequence once again, if necessary, we may assume that $d(x_y^{(n)}, x_y^\omega) \leq \frac{1}{n}$ for all y and n .

We now define the bijection $\tilde{\delta}_N : \tilde{\sigma} \rightarrow \tilde{\sigma}'$ as follows: given $x \in \tilde{\sigma}$,

- if $x = x_y^{(N)}$ for some $y \in \tilde{\tau} \setminus \tilde{\tau}_c$, let $\tilde{\delta}_N(x) = x_y^\omega$;
- if $x = x_y^{(n)}$ for some $y \in \tilde{\tau} \setminus \tilde{\tau}_c$ and some $n \geq N + 1$, let $\tilde{\delta}_N(x) = x_y^{(n-1)}$;
- otherwise, let $\tilde{\delta}_N(x) = x$;

see Figure 2, solid arrows. It is easy to see that $\tilde{\delta}_N$ is a bijection. Moreover, we have $d(x, \tilde{\delta}_N(x)) \leq \frac{2}{N}$ for all $x \in \tilde{\sigma}$: indeed, we have $d(x_y^{(N)}, x_y^\omega) \leq \frac{1}{N}$ and

$$d(x_y^{(n)}, x_y^{(n-1)}) \leq d(x_y^{(n)}, x_y^\omega) + d(x_y^{(n-1)}, x_y^\omega) \leq \frac{1}{n} + \frac{1}{n-1} \leq \frac{2}{N}$$

for $n \geq N + 1$. It follows that $d_\infty(\sigma, \sigma') \leq \sup\{d(x, \tilde{\delta}_N(x)) : x \in \tilde{\sigma}\} \leq \varepsilon$, as required. \square




FIGURE 2. The proof of Proposition 6.2: construction for a point $y \in \tilde{\tau} \setminus \tilde{\tau}_c$. Here $d = d_\infty(\sigma, \tau)$, and the solid arrows represent the bijection $\tilde{\delta}_N : \tilde{\sigma} \rightarrow \tilde{\sigma}'$.

We use the following two lemmas to prove the ‘if’ implication in Theorem E.

Lemma 6.3. *If X is a proper metric space and $A \subseteq X$ is a non-empty closed subset, then any geodesic in X/A with A as an endpoint can be lifted to a geodesic in X .*

Proof. Let $\eta : [0, 1] \rightarrow X/A$ be a constant speed geodesic such that $\eta(0) = A$. Let $x = \eta(1)$ and $r = d(x, A)$; without loss of generality, suppose that $x \notin A$, and therefore $r > 0$ and $\eta(t) \notin A$ for any $t > 0$. Then for any $n \in \mathbb{Z}_{\geq 1}$ there exists $a_n \in A$ such that $d_X(\eta(\frac{1}{n}), a_n) \leq \frac{2r}{n}$. In particular, we have $a_n \in A \cap B$ for each n , where B is the closed ball in X with centre x and radius $2r$. As X is proper and A is closed, $A \cap B$ is compact, implying that a subsequence of the sequence (a_n)

converges to some $a \in A$. We then define a map $\eta': [0, 1] \rightarrow X$ by $\eta'(0) = a$ and $\eta'(t) = \eta(t)$ for $t > 0$. For any $t > 0$ we then have

$$d(a, \eta(t)) \leq \limsup_{n \rightarrow \infty} d(a_n, \eta(t)) \leq \limsup_{n \rightarrow \infty} d(a_n, \eta(1/n)) + \limsup_{n \rightarrow \infty} d(\eta(1/n), \eta(t)) = td(x, A),$$

implying that η' is a geodesic lifting η . \square

Lemma 6.4. *If X is a proper metric space and $A \subseteq X$ is a non-empty closed subset such that any two points $x, y \in X/A$ with either $d(x, y) < d(x, A)$ or $y = A$ can be joined by a geodesic, then the pseudo-metric space $\mathcal{D}_\infty(X, A)$ is geodesic.*

Proof. Let $\sigma, \tau \in \mathcal{D}_\infty(X, A)$, and let $\sigma', \tau' \in \mathcal{D}_\infty(X, A)$ and $\gamma: \sigma' \rightarrow \tau'$ be as given by Proposition 6.2. Without loss of generality, assume that $d(x, \gamma(x)) < \max\{d(x, A), d(\gamma(x), A)\}$ for each $x \in \sigma'$ such that $x \neq A$ and $\gamma(x) \neq A$. Then each $x \in \sigma'$ can be joined by a geodesic in X to $\gamma(x) \in \tau'$ (when $x \in A$ or $\gamma(x) \in A$, this follows from Lemma 6.3); in particular we can pick a constant speed geodesic $\eta_x: [0, 1] \rightarrow X$ from x to $\gamma(x)$. For each $t \in [0, 1]$, define a diagram ρ_t as the equivalence class of the multiset $\{\{\eta_x(t) : x \in \sigma'\}\}$; we have $\rho_t \in \mathcal{D}_\infty(X, \{a_0\})$ since $d(x, \gamma(x)) \leq d(x, a_0) + d(\gamma(x), a_0) \leq d_\infty(\sigma', \sigma_\emptyset) + d_\infty(\tau', \sigma_\emptyset)$ for each $x \in \sigma'$ and therefore

$$d_\infty(\rho_t, \sigma_\emptyset) \leq d_\infty(\sigma', \sigma_\emptyset) + t(d_\infty(\sigma', \sigma_\emptyset) + d_\infty(\tau', \sigma_\emptyset)) \leq 2d_\infty(\sigma', \sigma_\emptyset) + d_\infty(\tau', \sigma_\emptyset) < \infty.$$

Now since $d(\sigma, \sigma') = 0$ we have $d(\sigma, \rho) = d(\sigma', \rho)$ (similarly, $d(\tau, \rho) = d(\tau', \rho)$) for any $\rho \in \mathcal{D}_\infty(X, \{a_0\})$. Moreover, whenever $0 \leq t < u \leq 1$ we have

$$d_\infty(\rho_t, \rho_u) \leq \sup\{d(\eta_x(t), \eta_x(u)) : x \in \sigma'\} = (u - t)d_\infty(\sigma, \tau),$$

and therefore the opposite inequality must also hold since

$$d_\infty(\rho_t, \rho_u) \geq d_\infty(\sigma, \tau) - d_\infty(\sigma, \rho_t) - d_\infty(\rho_u, \tau) \geq [1 - (t - 0) - (1 - u)]d_\infty(\sigma, \tau) = (u - t)d_\infty(\sigma, \tau).$$

It follows that if we define $\xi: [0, 1] \rightarrow \mathcal{D}_\infty(X, \{a_0\})$ by $\xi(0) = \sigma$, $\xi(1) = \tau$, and $\xi(t) = \rho_t$ for $0 < t < 1$, then ξ is a constant speed geodesic in $\mathcal{D}_\infty(X, \{a_0\})$ from σ to τ , as required. \square

In the other direction, we have the following.

Lemma 6.5. *If X and $a_0 \in X$ are such that the pseudo-metric space $\mathcal{D}_\infty(X, \{a_0\})$ is geodesic, then any two points $x, y \in X$ with either $d(x, y) < d(x, a_0)$ or $y = a_0$ can be joined by a geodesic.*

Proof. We say a pair $(x, y) \in X^2$ is *good* if $x \neq y$ and either $d(x, y) < d(x, a_0)$ or $y = a_0$. We aim to show that any good pair of points in X can be joined by a geodesic. For any good pair (x, y) , we pick a constant speed geodesic $\xi_{x,y}: [0, 1] \rightarrow \mathcal{D}_\infty(X, \{a_0\})$ from $\{\{x\}\}$ to $\{\{y\}\}$. Note that we have $d_\infty(\{\{x\}\}, \{\{y\}\}) = d(x, y)$; in particular, if $d(x, y) < d(x, a_0)$ then $d_\infty(\{\{x\}\}, \{\{y\}\}) < d(x, a_0)$, and if $y = a_0$ then $d_\infty(\{\{x\}\}, \{\{y\}\}) = d_\infty(\{\{x\}\}, \sigma_\emptyset) = d(x, a_0)$, implying that we have $d_\infty(\{\{x\}\}, \{\{y\}\}) \leq d(x, a_0)$ in either case.

The idea of the proof is to build a geodesic from x to y for a good pair (x, y) one-quarter of the remaining distance at a time. Any point $v \in \xi_{x,y}(t)$ for $t \leq \frac{1}{4}$ satisfies either $d(v, a_0) < \frac{1}{3}d(x, a_0)$ or $d(v, a_0) > \frac{2}{3}d(x, a_0)$, and such a ‘gap’ ensures that the partial geodesic can be built unambiguously by picking the points satisfying the latter inequality. Moreover, if such a partial geodesic from x to y ends at a point u then one can show that the pair (u, y) is again good, and so the construction can be iterated.

Let $(x, y) \in X^2$ be a good pair. Then for any $t \in [0, \frac{1}{4}]$ there exists a bijection $\gamma: \{\{x\}\} \rightarrow \xi_{x,y}(t)$ such that $d(z, \gamma(z)) < \frac{1}{3}d(x, a_0)$ for every z in (a chosen representative of) $\{\{x\}\}$, implying that there exists a point $u = u(x, y, t) \in \xi_{x,y}(t)$, namely $u(x, y, t) := \gamma(x)$, such that $d(u, a_0) > \frac{2}{3}d(x, a_0)$ and $d(v, a_0) < \frac{1}{3}d(x, a_0)$ for every $v \in \xi_{x,y}(t) \setminus \{\{u\}\}$; see Figure 3. Thus, as for any $0 \leq s < t \leq \frac{1}{4}$ we have $d_\infty(\xi_{x,y}(s), \xi_{x,y}(t)) = (t - s)d(x, y) < \frac{1}{3}d(x, a_0)$, any ‘nearly optimal’ bijection $\gamma: \xi_{x,y}(s) \rightarrow \xi_{x,y}(t)$ must send $u(x, y, s)$ to $u(x, y, t)$ and therefore $d(u(x, y, s), u(x, y, t)) \leq (t - s)d(x, y)$.

Furthermore, we claim that $d(u, y) \leq \frac{3}{4}d(x, y)$ and that (u, y) is a good pair, where $u = u(x, y, \frac{1}{4})$. If $y = a_0$, then $\{\{y\}\} = \sigma_\emptyset$ as diagrams and the claim follows directly from the fact that $d_\infty(\xi_{x,y}(\frac{1}{4}), \sigma_\emptyset) = \frac{3}{4}d(x, y)$ (note that we cannot have $u = y$ as $d(x, u) \leq \frac{1}{4}d(x, y)$). Otherwise, we have $d(x, y) < d(x, a_0)$, and so for ‘nearly optimal’ bijections $\gamma_x: \xi_{x,y}(\frac{1}{4}) \rightarrow \{\{x\}\}$ and $\gamma_y: \xi_{x,y}(\frac{1}{4}) \rightarrow \{\{y\}\}$ we have $d(\gamma_x(z), \gamma_y(z)) \leq d(z, \gamma_x(z)) + d(z, \gamma_y(z)) < d(x, a_0)$ for any $z \in \xi_{x,y}(\frac{1}{4})$. But then we must have $(\gamma_y \circ \gamma_x^{-1})(x) = y$; on the other hand, as u is the only point of $\xi_{x,y}(\frac{1}{4})$ such that $d(u, a_0) > \frac{2}{3}d(x, a_0)$, for ‘nearly optimal’ bijections we have $\gamma_x(u) = x$ and therefore $\gamma_y(u) = y$. Taking infimum over the ‘nearly optimal’ bijections γ_y , this implies that $d(u, y) \leq \frac{3}{4}d(x, y)$. Furthermore, taking the infimum over γ_x we have $d(x, u) \leq \frac{1}{4}d(x, y) < \frac{1}{4}d(x, a_0)$, and therefore $d(u, y) < \frac{3}{4}d(x, a_0) < d(x, a_0) - d(x, u) \leq d(u, a_0)$, implying that (u, y) is a good pair, as claimed.

We now define the sequence $(u_n)_{n=0}^\infty$ inductively, by letting $u_0 = x$ and $u_n = u(u_{n-1}, y, \frac{1}{4})$ for $n \geq 1$. We define a map $\eta: [0, 1] \rightarrow X$ as follows:

$$\eta(t) := \begin{cases} u(u_n, y, 1 - (\frac{4}{3})^n(1-t)) & \text{if } 1 - (\frac{3}{4})^n \leq t < 1 - (\frac{3}{4})^{n+1}, \\ y & \text{if } t = 1. \end{cases}$$

It is straightforward to verify (using the triangle inequality) that $d(\eta(s), \eta(t)) \leq (t-s)d(x, y)$ for all $0 \leq s < t \leq 1$, and therefore that η is a geodesic, as required. \square




FIGURE 3. An example of the situation in the proof of Lemma 6.5; here $d = d(x, a_0)$, the black points represent a diagram $\xi_{x,y}(t)$ for some $t \in [0, \frac{1}{4}]$, and the arrows a ‘nearly optimal’ bijection $\gamma: \{\{x\}\} \rightarrow \xi_{x,y}(t)$.

Proof of Theorem E. The ‘if’ direction follows from Lemma 6.4. The ‘only if’ direction follows from Lemma 6.5 and the isometry $\mathcal{D}_\infty(X, A) \cong \mathcal{D}_\infty(X/A, \{A\})$. \square

Proof of Corollary F. In view of Theorem E, it is enough to show that if X is proper and geodesic, then X/A is geodesic. Thus, let $x, y \in X$, and let $[x], [y] \in X/A$ be their images in X/A . If $d(x, y) < d(x, A) + d(y, A)$, then it is easy to verify that for any geodesic $\eta: [0, 1] \rightarrow X$ from x to y , the composite $q \circ \eta$ is a geodesic in X/A from $[x]$ to $[y]$, where $q: X \rightarrow X/A$ is the quotient map. Otherwise, we have $d_{X/A}([x], [y]) = d(x, A) + d(y, A)$. But then as X is proper and A is closed, there exist $a, b \in A$ such that $d(x, A) = d(x, a)$ and $d(y, A) = d(y, b)$. Therefore, for any geodesics η_x and η_y in X from x to a and from b to y , respectively, the concatenation of $q \circ \eta_x$ and $q \circ \eta_y$ is a geodesic in X/A from $[x]$ to $[y]$, as required. \square

It follows from Lemma 6.5 that the ‘only if’ direction of Theorem E is true even without the assumption that X is proper. However, the same cannot be said about the ‘if’ direction, as the following example shows.

Example 6.6. Let $X = c_0$ (that is, the space of all real sequences $(a_n)_{n=1}^\infty$ converging to zero, equipped with the supremum metric). Let $A = \{(0, 0, \dots)\} \subset X$. Then the space $X = X/A$ is geodesic: indeed, given any $(a_n), (b_n) \in X$ the path $\eta: [0, 1] \rightarrow X$ defined by $\eta(t) = ((1-t)a_n + tb_n)_{n=1}^\infty$ is a geodesic. We claim that $\mathcal{D}_\infty(X, A)$ is not geodesic.

For any finite subset $F \subset \mathbb{Z}_{\geq 1}$, let $\alpha_F = (a_n^F)_{n=1}^\infty \in X$ be such that $a_n^F = 1 + \frac{1}{n}$ if $n \in F$ and $a_n^F = 0$ otherwise. Let \mathcal{F}_+ (respectively \mathcal{F}_-) be the collection of finite subsets of $\mathbb{Z}_{\geq 1}$ of even (respectively odd) cardinality, and consider the diagrams $\sigma_\pm = \{\{\alpha_F : F \in \mathcal{F}_\pm\}\} \in \mathcal{D}_\infty(X, A)$. We claim that there are no geodesics in $\mathcal{D}_\infty(X, A)$ from σ_+ to σ_- .

Note first that $d_\infty(\sigma_+, \sigma_-) = 1$. Indeed, we have $d(\alpha_F, \alpha_{F'}) > 1$ and $d(\alpha_{F'}, A) > 1$ for any $F \in \mathcal{F}_+$ and $F' \in \mathcal{F}_-$, implying that $d_\infty(\sigma_+, \sigma_-) \geq 1$. On the other hand, for any $N \geq 1$ and any $B \subseteq [N]$ the sets $\{F \in \mathcal{F}_+ : F \cap [N] = B\}$ and $\{F \in \mathcal{F}_- : F \cap [N] = B\}$ are both countably infinite (where we write $[N]$ for $\{1, \dots, N\}$), implying that for any N there exists a bijection $\delta_N: \mathcal{F}_+ \rightarrow \mathcal{F}_-$ such that $F \cap [N] = \delta_N(F) \cap [N]$ for any $F \in \mathcal{F}_+$. We can then define a bijection $\gamma_N: \sigma_+ \rightarrow \sigma_-$ by setting $\gamma_N(\alpha_F) = \alpha_{\delta_N(F)}$, and it's easy to check that $d(\alpha, \gamma_N(\alpha)) < 1 + \frac{1}{N}$ for any $\alpha \in \sigma_+$. This implies that $d_\infty(\sigma_+, \sigma_-) = 1$, as claimed.

Now suppose for contradiction that $\xi: [0, 1] \rightarrow \mathcal{D}_\infty(X, A)$ is a constant speed geodesic from σ_+ to σ_- . Then $d_\infty(\sigma_-, \xi(\frac{1}{2})) = \frac{1}{2} < \frac{3}{2}$, implying that there exists a point $\beta = (b_n)_{n=1}^\infty \in \xi(\frac{1}{2})$ (which we fix from now on) such that $|2 - b_1| \leq d(\alpha_{\{1\}}, \beta) < \frac{3}{2}$ and therefore $b_1 > \frac{1}{2}$. Let $F_0 = \{n \in \mathbb{Z}_{\geq 1} : b_n > \frac{1}{2}\}$; as $\beta \in X = c_0$, we know that $|F_0| < \infty$ and that there exists $\varepsilon \in (0, \frac{1}{2})$ such that $b_n \geq \frac{1}{2} + \varepsilon$ for all $n \in F_0$ and $b_n \leq \frac{1}{2} - \varepsilon$ for all $n > \frac{1}{\varepsilon}$. Since $d(\beta, A) \geq \frac{1}{2} + \varepsilon > \frac{1}{2}$ and since $d_\infty(\sigma_\pm, \xi(\frac{1}{2})) = \frac{1}{2}$, there exist subsets $F_\pm \in \mathcal{F}_\pm$ such that $d(\alpha_{F_\pm}, \beta) < \frac{1}{2} + \varepsilon$. Now if $n \in F_0$ then $|0 - b_n| \geq \frac{1}{2} + \varepsilon$ and therefore we must have $a_n^{F_\pm} = 1 + \frac{1}{n}$; on the other hand, if $n \notin F_0$ then $b_n \leq \frac{1}{2}$ and either $\frac{1}{n} \geq \varepsilon$ or $b_n \leq \frac{1}{2} - \varepsilon$, so in either case we have $|1 + \frac{1}{n} - b_n| \geq \frac{1}{2} + \varepsilon$ and therefore we must have $a_n^{F_\pm} = 0$. This implies that $F_- = F_0 = F_+$, which is impossible since $\mathcal{F}_- \cap \mathcal{F}_+ = \emptyset$; thus there are no geodesics in $\mathcal{D}_\infty(X, A)$ from σ_+ to σ_- , as claimed.

Note that, for $\sigma_\pm \in \mathcal{D}_\infty(X, A)$ as in Example 6.6, there are no bijections $\gamma: \sigma_+ \rightarrow \sigma_-$ such that $\sup_{x \in \sigma_+} d(x, \gamma(x)) = d_\infty(\sigma_+, \sigma_-)$. Indeed, if there was such a bijection then the map $[0, 1] \rightarrow \mathcal{D}_\infty(X, A)$ sending t to the diagram $\{(1-t)x + t\gamma(x) : x \in \sigma_+\}$ would be a geodesic from σ_+ to σ_- , but we have shown that no such geodesics exist.

REFERENCES

1. Hideto Asashiba, Mickaël Buchet, Emerson G Escobar, Ken Nakashima, and Michio Yoshiwaki, *On interval decomposability of 2d persistence modules*, Computational Geometry **105** (2022), 101879.
2. Håvard Bakke Bjerkevik, *On the stability of interval decomposable persistence modules*, Discrete Comput. Geom. **66** (2021), no. 1, 92–121. MR 4270636
3. Magnus Bakke Botnan, Vadim Lebovici, and Steve Oudot, *On rectangle-decomposable 2-parameter persistence modules*, Discrete & Computational Geometry **68** (2022), no. 4, 1078–1101.
4. Peter Bubenik and Alex Elchesen, *Universality of persistence diagrams and the bottleneck and Wasserstein distances*, Comput. Geom. **105/106** (2022), Paper No. 101882, 18. MR 4414770
5. ———, *Virtual persistence diagrams, signed measures, Wasserstein distances, and Banach spaces*, J. Appl. Comput. Topol. **6** (2022), no. 4, 429–474. MR 4496687
6. Peter Bubenik and Iryna Hartsock, *Topological and metric properties of spaces of generalized persistence diagrams*, [arXiv:2205.08506](https://arxiv.org/abs/2205.08506) [math.AT], 2022.
7. Peter Bubenik and Tane Vergili, *Topological spaces of persistence modules and their properties*, J. Appl. Comput. Topol. **2** (2018), no. 3-4, 233–269. MR 3927353
8. Dmitri Burago, Yuri Burago, and Sergei Ivanov, *A course in metric geometry*, Graduate Studies in Mathematics, vol. 33, American Mathematical Society, Providence, RI, 2001. MR 1835418
9. Gunnar Carlsson and Mikael Vejdemo-Johansson, *Topological data analysis with applications*, Cambridge University Press, Cambridge, 2022. MR 4346385
10. Mauricio Che, Fernando Galaz-García, Luis Guijarro, and Ingrid Membrillo Solis, *Metric geometry of spaces of persistence diagrams*, [arXiv:2109.14697](https://arxiv.org/abs/2109.14697) [math.MG], 2021.

11. David Cohen-Steiner, Herbert Edelsbrunner, and John Harer, *Stability of persistence diagrams*, Discrete Comput. Geom. **37** (2007), no. 1, 103–120. MR 2279866
12. David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko, *Lipschitz functions have L_p -stable persistence*, Found. Comput. Math. **10** (2010), no. 2, 127–139. MR 2594441
13. Vincent Divol and Théo Lacombe, *Understanding the topology and the geometry of the space of persistence diagrams via optimal partial transport*, J. Appl. Comput. Topol. **5** (2021), no. 1, 1–53. MR 4224153
14. Isaac Goldbring, *Ultrafilters throughout mathematics*, Graduate Studies in Mathematics, no. 220, American Mathematical Society, 2022.
15. Michael Kapovich and Bernhard Leeb, *On asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds*, Geom. Funct. Anal. **5** (1995), 582–603.
16. Yuriy Mileyko, Sayan Mukherjee, and John Harer, *Probability measures on the space of persistence diagrams*, Inverse Problems **27** (2011), no. 12, 124007, 22. MR 2854323
17. Steve Y. Oudot, *Persistence theory: from quiver representations to data analysis*, Mathematical Surveys and Monographs, vol. 209, American Mathematical Society, Providence, RI, 2015. MR 3408277
18. Jacek Skryzalin and Gunnar Carlsson, *Numeric invariants from multidimensional persistence*, J. Appl. Comput. Topol. **1** (2017), no. 1, 89–119. MR 3975550
19. Katharine Turner, Yuriy Mileyko, Sayan Mukherjee, and John Harer, *Fréchet means for distributions of persistence diagrams*, Discrete Comput. Geom. **52** (2014), no. 1, 44–70. MR 3231030
20. Afra Zomorodian and Gunnar Carlsson, *Computing persistent homology*, Discrete Comput. Geom. **33** (2005), no. 2, 249–274. MR 2121296

(Che) DEPARTMENT OF MATHEMATICAL SCIENCES, DURHAM UNIVERSITY, UNITED KINGDOM.

Email address: mauricio.a.che-moguel@durham.ac.uk

(Galaz-García) DEPARTMENT OF MATHEMATICAL SCIENCES, DURHAM UNIVERSITY, UNITED KINGDOM.

Email address: fernando.galaz-garcia@durham.ac.uk

(Guijarro) DEPARTMENT OF MATHEMATICS, UNIVERSIDAD AUTÓNOMA DE MADRID AND ICMAT CSIC-UAM-UC3M, SPAIN

Email address: luis.guijarro@uam.es

(Membrillo Solis) MATHEMATICAL SCIENCES, UNIVERSITY OF SOUTHAMPTON, UNITED KINGDOM

Email address: i.membrillo-solis@soton.ac.uk

Current address: School of Computer Science and Engineering, University of Westminster, United Kingdom

Email address: i.membrillosolis@westminster.ac.uk

(Valiunas) MATHEMATICAL INSTITUTE, UNIVERSITY OF WROCLAW, POLAND

Email address: motiejus.valiunas@math.uni.wroc.pl