

GEOMETRIC BOUNDS FOR PERSISTENCE

ALEXEY BALITSKIY, BARIS COSKUNUZER, AND FACUNDO MÉMOLI

ABSTRACT. In this paper, we offer a new perspective on persistent homology by integrating key concepts from metric geometry. For a given compact subset \mathcal{X} of a Banach space \mathbf{Y} , we analyze the topological features arising in the family $\mathcal{N}_\bullet(\mathcal{X} \subset \mathbf{Y})$ of nested neighborhoods of \mathcal{X} in \mathbf{Y} and provide several geometric bounds on their persistence (lifespans).

We begin by examining the lifespans of these homology classes in terms of their filling radii in \mathbf{Y} , establishing connections between these lifespans and fundamental invariants in metric geometry, such as the Urysohn width. We then derive bounds on these lifespans by considering the ℓ^∞ -principal components of \mathcal{X} , also known as Kolmogorov widths.

Additionally, we introduce and investigate the concept of extinction time of a metric space \mathcal{X} : the critical threshold beyond which no homological features persist in any degree. We propose methods for estimating the Čech and Vietoris–Rips extinction times of \mathcal{X} by relating \mathcal{X} to its convex hull and to its tight span, respectively.

CONTENTS

1. Introduction	8437
2. Background	8443
3. Bounding lifespans via widths	8454
4. Bounding lifespans by spread	8466
5. Bounding extinction times	8469
6. Final remarks	8474
Appendix A. Auxiliary properties of widths	8476
Appendix B. Übercontractible sets	8477
Appendix C. Notation table	8480
Acknowledgment	8480
References	8480

1. INTRODUCTION

Over the past decade, numerous approaches within Topological Data Analysis (TDA) have been developed to uncover patterns across a wide variety of data types.

Received by the editors March 24, 2024, and, in revised form, December 30, 2024.

2020 *Mathematics Subject Classification*. Primary 53C23.

This work was supported by the following grants from the National Science Foundation (Grants # DMS-1926686, DMS-2202584, DMS-2220613, DMS-2229417, DMS-1723003, IIS-1901360, CCF-1740761, and DMS-1547357), from the Simons Foundation (Grant # 579977) and from the BSF (under grant # 2020124). The project has received funding from the HORIZON EUROPE Research and Innovation programme under Marie Skłodowska-Curie grant agreement number 101107896.

Among these, Persistent Homology (PH) has emerged as a cornerstone of TDA, providing a robust multiscale feature extraction framework. This progress has been driven by the development of efficient algorithmic procedures and effective software implementations for its computation (see Section 2.1). Indeed, PH has gained considerable traction in diverse machine learning applications spanning fields such as:

- Bioinformatics and Biomedicine [NLC11, CCR13, GHI⁺15, CW17, AQO⁺20, BHPG⁺21, SL22],
- Finance [GK18, RQD23],
- Materials Science [HNH⁺16, RSDFS16, SDB⁺22, ONH22, LBD⁺18],
- Neuroscience [CI08, SMI⁺08, DMFC12, RNS⁺17], and
- Network Analysis [HKNU17, HKN19, HGR⁺20, CCI⁺20, ZYCW20, AAF19].

In this paper, we explore aspects of the Persistent Homology (PH) methodology through the lens of metric geometry. Our results offer new quantitative interpretations of the PH output, providing deeper insights into its structure and significance.

1.1. Persistent homology (in a nutshell). We recall the basic idea behind persistent homology. Let $\Delta_\bullet = \{\Delta_r \xrightarrow{\iota_{r,s}} \Delta_s\}_{0 < r \leq s}$ be a *filtration*: a *nested* family of topological spaces or simplicial complexes, e.g. obtained via the Vietoris–Rips filtration $\mathcal{V}_\bullet(\mathcal{X})$ or the Čech filtration $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y})$ induced from a compact metric space \mathcal{X} (in the case of the Čech filtration, one typically assumes that \mathcal{X} is a subset of a Banach space \mathbf{Y}). For a nonnegative integer k , let ω be a nontrivial degree- k reduced homology class appearing in the nested family Δ_\bullet , that is, assume that $\omega \in \tilde{H}_k(\Delta_r; \mathbb{F})$ for some $r > 0$.¹ We in fact consider the degree- k *homological spectrum* of the filtration Δ_\bullet , $\text{Spec}_k(\Delta_\bullet)$ to be the collection of all such nonzero homology classes (see equation (1)). The *birth time* b_ω of $\omega \in \tilde{H}_k(\Delta_r; \mathbb{F})$ is the infimal $u > 0$ such that there exists $\omega_u \in \tilde{H}_k(\Delta_u)$ with the property that $(\iota_{u,r})_*(\omega_u) = \omega$. Similarly, we define the *death time* d_ω of ω to be the supremal $v \geq b_\omega$ such that ω does not become homologically trivial in Δ_v , that is $(\iota_{r,v})_*(\omega) \neq 0$. This is informally interpreted as indicating that the nontrivial class ω is “alive” inside the interval $I_\omega = (b_\omega, d_\omega]$.² We call the quantity $d_\omega - b_\omega$ the *lifespan* (or persistence) of the class ω .

The notion of *Persistent Homology* is closely related to but subtly differs from this process of recording birth and death times for individual homology classes described above. The degree- k persistent homology of Δ_\bullet is the directed system of vector spaces $\tilde{H}_k(\Delta_\bullet; \mathbb{F})$. Under suitable tameness assumptions on the family Δ_\bullet , an up-to-isomorphism representation of this directed system can be obtained via its *persistence diagram*, a multiset of intervals I on $\mathbb{R}_{>0}$ supporting certain linearly independent collection of nontrivial homology classes that are alive at all points in I . See Section 2.1 for the precise definition of persistence diagrams (PD).

In many applications, the lifespan of a topological feature ω is critically significant, as it is often interpreted as a measure of the “size” or “importance” of ω . In practice, topological features with long lifespans—those that *persist*—are typically

¹Here \mathbb{F} is a fixed field.

²Whether the interval is left/right open/closed depends on semi-continuity conditions of Δ_\bullet . See Definition 2.18 for the case of neighborhood filtrations, the type of filtrations that we concentrate on in this paper.

considered to represent the primary shape characteristics of a dataset, while features with short lifespans are generally regarded as (topological) noise.³ Therefore, *determining* (or estimating) and *interpreting* the lifespans of these topological features appearing in the persistence diagram has critical significance for applications such as the ones mentioned above.

1.2. Connections with metric geometry and main results. In this paper, we aim to give a *geometric interpretation* of these lifespans by relating them to several notions from Metric Geometry. We study the lifespans of individual homology classes ω appearing throughout the Vietoris–Rips and Čech filtrations by resorting to the notion of *filling radius* and to several notions of *width*. See Section 1.3 for a discussion of the interplay between widths and filling radii in metric geometry. Through these notions of width, our results show that the lifespans of homology classes are controlled by the (geometric) size of their representatives in the filtration thus providing precise (geometric) interpretations of the significance of the features tracked by persistent homology.

We first discuss implications of the absolute (Gromov’s) and relative filling radius of a homology class (Definition 2.31) in our setting. Then, we observe that the Čech lifespan of a homology class ω is equal to its relative filling radius in ambient space (Section 2.3). Next, we give several bounds for the lifespans of individual homology classes by resorting to the notions of Urysohn width, Alexandrov width, and Kolmogorov width. They measure in various ways how well a space can be approximated by a k -dimensional complex, and they are denoted by $\text{UW}_k(\cdot)$, $\text{AW}_k(\cdot)$, $\text{KW}_k(\cdot)$, respectively. Since all these notions of width are monotonically nonincreasing with respect to the dimension parameter (e.g., $\text{UW}_k(\mathcal{X}) \geq \text{UW}_{k+1}(\mathcal{X})$), any degree- k estimate automatically applies to homology classes in higher degrees.

Corollary 3.18 (VR lifespans via Urysohn width). *Let \mathcal{X} be a compact metric space, and let $\omega \in \text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$, $k \geq 1$. Then,*

$$d_\omega - b_\omega \leq \text{UW}_{k-1}(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \text{E}(\mathcal{X}))).$$

In particular,

$$d_\omega - b_\omega \leq \text{UW}_{k-1}(\text{E}(\mathcal{X})).$$

Here and throughout the paper, $\overline{\mathcal{N}}_r(\mathcal{X} \subset \mathcal{Z})$ denotes the closed r -neighborhood of \mathcal{X} inside the metric space \mathcal{Z} and $\text{E}(\mathcal{X})$ denotes the *tight span* of \mathcal{X} (Definition 2.28), a canonically constructed metric space admitting an isometric embedding of \mathcal{X} and enjoying properties reminiscent of (but stronger than) the ones satisfied by the convex hull.

Notice that in a special case, if \mathcal{X} is a closed k -manifold, and $\omega = [\mathcal{X}]$ is its fundamental class, since $b_\omega = 0$ in that case, the result above implies that the Vietoris–Rips lifespan of ω is bounded above by the Urysohn width of \mathcal{X} , i.e., $d_\omega \leq \text{UW}_{k-1}(\mathcal{X})$ (this particular bound goes back to Gromov; it follows from [Gro83, Appendix 1, Example after Lemma (B)] combined with [Gro83, Appendix 1, Proposition (D)]).

³However, short lived, or even ephemeral, topological features also can carry useful information; see Usher and Zhang [UZ16], Bubenik, Hull, Patel and Whittle [BHPW20], and Mémoli and Zhou [MZ24].

Next, we give several bounds for Čech lifespans. The first one is via Alexandrov widths.

Corollary 3.14 (Čech lifespans via Alexandrov width). *Let \mathcal{X} be compact subset of a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 1$. Then,*

$$d_\omega - b_\omega \leq \text{AW}_{k-1}(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

In particular,

$$d_\omega - b_\omega \leq \text{AW}_{k-1}(\overline{\text{conv}}(\mathcal{X}) \subset \mathbf{Y}).$$

Here $\overline{\text{conv}}(\mathcal{X}) \subset \mathbf{Y}$ denotes the closure of the convex hull of $\mathcal{X} \subset \mathbf{Y}$.

Further, we introduce a new notion of width, called *treewidth* (Definition 3.9), in order to obtain a finer estimate on Čech lifespans (Corollary 3.12). Notice that in the estimates above, for each class ω , in order to obtain an upper bound for its lifespan, the corresponding width needs to be calculated for the neighborhood $\mathcal{N}_{b_\omega}(\mathcal{X} \subset \cdot)$ of \mathcal{X} , not for \mathcal{X} itself. In the following, we get rid of this dependency on the neighborhood.

The notion of treewidth permits establishing a certain *multiplicative* bound on lifespans.

Corollary 3.24 (Čech lifespans via treewidth). *Let \mathcal{X} be compact subset in a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, with birth time $b_\omega \geq 1$. Then,*

$$\frac{d_\omega}{b_\omega} \leq C + 1 + \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}).$$

Next, we consider the ℓ^∞ -version of principal component analysis (PCA $_\infty$) for a compact subset \mathcal{X} of a Banach space \mathbf{Y} (e.g., a point cloud in \mathbb{R}^N). By using the estimates in previous sections, we relate the $(k+1)$ -variance $\nu_{k+1}(\mathcal{X} \subset \mathbf{Y})$ with the lifespans of classes ω appearing throughout the Čech filtration (Section 3.5). We note that the variance $\nu_{k+1}(\mathcal{X} \subset \mathbf{Y})$ is also known as the k^{th} *Kolmogorov width* of \mathcal{X} in the approximation theory literature (Remark 3.27).

Corollary 3.29 (Čech lifespans via ℓ^∞ -variance). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 0$. Then,*

$$d_\omega - b_\omega \leq \nu_{k+1}(\mathcal{X} \subset \mathbf{Y}) = \text{KW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Note that by monotonicity of widths (e.g., $\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \geq \text{AW}_{k+1}(\mathcal{X} \subset \mathbf{Y})$), all these bounds apply to homology classes of degree higher than k as well.

While these lifespan bounds depend on the homology degree, we next give a general bound for Čech lifespans which is independent of degree. We achieve this by generalizing Katz's notion of spread [Kat83] to the extrinsic setting and introducing the notion that we call *überspread* (Section 4). Überspread basically measures the Hausdorff distance from the space to the closest *übercontractible* space (a contractible space where all neighborhoods are also contractible, see Definition 4.2) in an ambient space. With this notion, we generalize the existing VR-lifespan estimates via spread obtained by Lim, Mémoli and Okutan [LMO24] to Čech-lifespans in any degree.

Theorem 4.5 (Čech lifespans via überspread). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$ for any $k \geq 0$. Then,*

$$d_\omega - b_\omega \leq 2 \text{ ü-spread}(\mathcal{X} \subset \mathbf{Y}).$$

While the results above are effective for bounding individual lifespans ($d_\omega - b_\omega$), we also attack a more general question: how to obtain a global bound for death times of homology classes across all degrees? To do this we introduce a notion, called *extinction time*, representing the maximal threshold after which there is no nontrivial homology class in any degree $k \geq 0$. We bound both Čech extinction times $\check{\xi}(\mathcal{X} \subset \mathbf{Y})$ and Vietoris–Rips extinction times $\xi(\mathcal{X})$ by relating \mathcal{X} to their convex hulls and tight spans, respectively. In the case of Čech filtrations, we introduce a notion called *convexity deficiency*, $\text{cdef}(\mathcal{X} \subset \mathbf{Y})$, which is the Hausdorff distance of a space \mathcal{X} to its convex hull in \mathbf{Y} (Section 5.1).

Theorem 5.5 (Bounding Čech extinction). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Then,*

$$\check{\xi}(\mathcal{X} \subset \mathbf{Y}) \leq \text{cdef}(\mathcal{X} \subset \mathbf{Y}).$$

In the VR case, we define an analogous notion called *hyperconvexity deficiency*, $\text{hcdef}(\mathcal{X})$, which is the Hausdorff distance between the tight span $\text{E}(\mathcal{X})$ of \mathcal{X} and the isometric copy of \mathcal{X} inside of it (Section 5.2). We then show that a result analogous to Theorem 5.5 is also true in the VR-case (Corollary 5.10).

We highlight the bidirectional relationship between applied topology and metric geometry: on one hand, with the goal of improving their interpretability, we establish upper bounds for crucial quantities which originated in applied topology (e.g., lifespans of homology classes) using concepts from metric geometry; on the other hand, these results yield computational lower bounds or estimates for metric geometry notions inspired by persistent homology; see Remark 2.36. This interplay underscores the synergy between these fields, enabling insights that advance both domains.

Cores. Conceptually, our results establish a relationship between a given space \mathcal{X} and another space $\Lambda_{\mathcal{X}}$, which functions as a *core* for \mathcal{X} . This represents the central thread weaving together the various parts of the paper. Specifically, estimates on the lifespans of homological features that arise as the radius of neighborhoods of \mathcal{X} increases are derived from the distance between \mathcal{X} and $\Lambda_{\mathcal{X}}$:

- In Section 3, where we explore various notions of widths, the role of the core $\Lambda_{\mathcal{X}}$ is, roughly speaking, assumed by a k -dimensional space closest to \mathcal{X} . In this context, the k -width can be interpreted as the distance between \mathcal{X} and this approximate core, $\Lambda_{\mathcal{X}}$. In this section, we first recall the classical notions of Urysohn and Alexandrov width and then introduce a new variant which we call treewidth.
- In Section 4, after recalling the notion of spread, we introduce the notion of überspread, where the core $\Lambda_{\mathcal{X}}$ is treated as an übercontractible space, and the distance between \mathcal{X} and $\Lambda_{\mathcal{X}}$ provides an upper bound on the lifespans in any degree. Notably, the condition imposed on the core in this context is the triviality of its homology groups, rather than any restriction on its dimension.
- In Section 5, where we analyze extinction times, we impose strong geometric conditions—such as convexity or hyperconvexity—on $\Lambda_{\mathcal{X}}$ to derive extinction bounds based on the distance between $\Lambda_{\mathcal{X}}$ and \mathcal{X} . This is done through the concepts of convexity deficiency and hyperconvexity deficiency, which we introduce therein.

1.3. Related work. In this work, we aim to build a bridge between two seemingly disparate fields: applied algebraic topology and metric geometry. Both disciplines address a similar fundamental question regarding the quantification of “shape” using distinct tools:

How to measure the size of a set/space/manifold?

In applied algebraic topology, persistent homology is an effective tool for accomplishing this aim, and the *lifespans* (or *persistence*) of topological features induced by Vietoris–Rips (or Čech) filtrations are used as a measure of the size or importance of the corresponding topological features. In particular, Vietoris–Rips (or Čech) complexes were invented in order to transform a given metric space into a simplicial complex while maintaining its topological information, thus enabling an effective cohomology theory for metric spaces; see the papers by Vietoris [Vie27], Borsuk [Bor48] and Hausmann [Hau95]. Numerous studies in the literature explore Vietoris–Rips complexes and Vietoris–Rips filtrations across various settings; see Latschev [Lat01], Chazal, Cohen-Steiner, de Silva, Guibas, Mémoli, and Oudot [CCSG⁺09, CDSO14], Adamaszek, Adams, Frick, Gillespie, Lim, Mémoli, Moy, Okutan, Reddy, and Wang [AA17, AAR19, AMMW24, AFV23, LMO24, Gil24], Attali, Lieutier and Salinas [ALS13], Rieser, Bubenik and Milicevic [Rie20, BM24], Turner [Tur19], Virk [Vir20, Vir22, Vir21], and Zaremsky [Zar22].

On the other hand, from the metric geometry side, estimating the size of a manifold has been a key problem for several decades. Gromov introduced and studied the notion of *filling radius* in his seminal paper [Gro83], and several other notions of “largeness” in [Gro86]. Before Gromov, certain relative, or extrinsic, versions of the filling radius were studied by Federer and Fleming [FF60], Michael and Simon [MS73] and Bombieri and Simon [BS83] in geometric analysis in connection with the isoperimetric inequality. Gromov brought the filling radius to the realm of systolic geometry, and the study of scalar curvature [Gro83]. Several other measures of size of a given manifold or metric space, nowadays known as *widths*, were also studied and popularized by Gromov [Gro83, Gro88].

The interplay between filling radii, widths, and other metric invariants (including volume) has been an active research area since then. Katz determined the filling radius of spheres and other essential spaces [Kat83, Kat89, Kat90, Kat91]. Several authors have studied the filling radii in comparison with other measures of largeness [Cai94, BH10]. Guth proved some related conjectures of Gromov [Gut11, Gut17]. Sabourau, Nabutovsky, and Rotman related the filling radius with sweep-outs of manifolds [Sab20, NRS21]. Bounds on filling radius in terms of Hausdorff content follow from very general isoperimetric estimates due to Liokumovich, Lishak, Nabutovsky, and Rotman [LLNR22].

In recent years, several articles have explored connections between these two domains, addressing analogous problems with different methodologies. With this aim, Lim, Mémoli and Okutan related the filling radius of a closed manifold to the interval corresponding to the fundamental class in the top VR persistence diagram [LMO24]. As shown in [LMO24, Section 9.3.2], the stability of persistence diagrams of Vietoris–Rips filtrations can be used to obtain lower bounds for the Gromov–Hausdorff distance between spheres through considerations related to their filling radii. These were shown not to be tight by Lim, Mémoli and Smith [LMS23] and the polymath project [ABC⁺22] furthered this line of work; see also Jeffs and Harrison [HJ23] and Rodriguez-Martín [RM24]. Adams and Coskunuzer used a

well-known quantity in metric geometry, Urysohn width, to estimate the lifespans in the persistence diagram of a given space [AC22]. In [Vir22], in the manifold setting, Virk studied the relation between persistence diagrams for large degrees and lower dimensional features.

In this paper, we aim to establish a direct connection between concepts from applied algebraic topology and metric geometry by linking various quantities used in both fields to measure the size of a metric space. In writing this paper, we have prioritized accessibility, aiming to bridge the gap between the two fields and foster greater collaboration and understanding.

2. BACKGROUND

In this section, we provide an overview of the concepts from applied algebraic topology and metric geometry that form the foundation of the paper. We give a summary of our notations in Table 1 in the Appendix.

2.1. Persistent homology. Persistent Homology (PH) is a methodology rooted in Applied Algebraic Topology that captures various structural characteristics of a given topological or metric space. Its development can be traced to the pioneering work of Frosini [Fro90] and Robins [Rob99], with its algorithmic framework later established by Edelsbrunner, Letscher, and Zomorodian [ELZ02]. Earlier manifestations of persistent homology were retrospectively identified in the works of Morse [Mor30], Deheuvels [Deh55], and Barannikov [Bar94].

In the past two decades, PH has been employed as a powerful mathematical machinery for discovering patterns in data in applications within Machine Learning and Data Science. This advancement has been made possible by the development of efficient algorithms capable of computing PH in *polynomial time*. Specifically, the total computational effort is a polynomial function of parameters related to the size of the input simplicial filtration and the maximum homology degree to be computed; see Edelsbrunner, Letscher and Zomorodian [ELZ02], Harker, Mischaikow, Mrozek and Nanda [HMMN14, MN13], and Bauer [Bau21].

For more details on PH and its use in various settings, see Carlsson [Car09], Edelsbrunner and Harer [EH10], Chazal, de Silva, Glisse and Oudot [CDSGO16], Ghrist [Ghr18], Rabadán and Blumberg [RB19], Carlsson and Vejdemo-Johansson [CVJ21], Joharinad and Jost [JJ23], and Polterovich, Rosen, Samvelyan and Zhang [PRSZ20].

Neighborhood notation. Throughout the paper, we use both open and closed neighborhoods and adopt the following notation. In the sequel, given a metric space \mathcal{Z} , a point $z \in \mathcal{Z}$ and $r > 0$, by $B_r(z)$ we will denote the open ball of radius r around z . When \mathcal{X} is a subset of a metric space \mathcal{Z} , through

$$\mathcal{N}_r(\mathcal{X} \subset \mathcal{Z}) := \bigcup_{x \in \mathcal{X}} B_r(x),$$

we will denote the open r -neighborhood of \mathcal{X} in \mathcal{Z} while $\overline{\mathcal{N}}_r(\mathcal{X} \subset \mathcal{Z})$ will denote the similarly defined closed r -neighborhood of \mathcal{X} in \mathcal{Z} .

2.1.1. *Filtrations.* As noted in the previous section, utilizing the PH machinery requires a filtration—a nested family of topological spaces or abstract simplicial complexes—denoted by Δ_\bullet . One of the most natural examples arises by considering nested neighborhoods of a subspace of a metric space, i.e., for \mathcal{X} a subset of a metric space \mathcal{Z} , the family $\{\mathcal{N}_r(\mathcal{X} \subset \mathcal{Z})\}_{r \geq 0}$ defines a filtration. Simplicial constructions are preferred in practical applications and the most common ones are Vietoris–Rips and Čech complexes and the respective filtrations they induce. While our study focuses primarily on these two types of filtrations, most of our results concerning Čech complexes naturally extend to alpha complexes; see Remark 2.9.

Definition 2.1 (Vietoris–Rips complexes). Let $(\mathcal{X}, \mathbf{d}_\mathcal{X})$ be a compact metric space. For $r > 0$, its Vietoris–Rips complex at scale r is the abstract simplicial complex $\mathcal{V}_r(\mathcal{X})$ where a k -simplex $\sigma = [x_{i_0}, x_{i_1}, \dots, x_{i_k}] \in \mathcal{V}_r(\mathcal{X})$ if and only if $\mathbf{d}_\mathcal{X}(x_{i_m}, x_{i_n}) < r$ for any $0 \leq m, n \leq k$.

Definition 2.2 (Čech complexes). Let \mathcal{X} be a compact subset of a metric space \mathcal{Z} . For $r > 0$, the Čech complex at scale r is the abstract simplicial complex $\mathcal{C}_r(\mathcal{X} \subset \mathcal{Z})$ where a k -simplex $\sigma = [x_{i_0}, x_{i_1}, \dots, x_{i_k}] \in \mathcal{C}_r(\mathcal{X} \subset \mathcal{Z})$ if and only if $\bigcap_{m=0}^k B_r(x_{i_m}) \neq \emptyset$ in \mathcal{Z} .

In most scenarios $\mathcal{Z} = \mathbf{Y}$, a Banach space.

Through the geometric realization functor, the nested families of simplicial complexes provided by the Vietoris–Rips and Čech complexes induce filtrations.

Definition 2.3 (Filtration). A *filtration* of a topological space is a collection $\Delta_\bullet = (\Delta_r, \iota_{r,s})_{0 < r \leq s}$ such that for each $0 < r \leq s$, Δ_r is a subset of the given topological space and $\iota_{r,s} : \Delta_r \hookrightarrow \Delta_s$ is the inclusion map.

When there is no risk of confusion we will simply say that Δ_\bullet is a filtration without mentioning the ambient topological space, with the understanding that it can be recovered as the colimit of Δ_\bullet . In the remainder of the paper, we use the notation $\mathcal{V}_\bullet(\mathcal{X})$ and $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y})$ to denote the filtrations induced by the (geometric realizations of the) Vietoris–Rips and Čech complexes of \mathcal{X} , respectively.

Example 2.4 (Neighborhood filtrations). Another example of filtrations arising in geometric scenarios is the following. Let $\mathcal{X} \subset \mathcal{Z}$ be a nonempty compact subset of a metric space $(\mathcal{Z}, \mathbf{d}_\mathcal{Z})$. Then, one considers the filtration $\mathcal{N}_\bullet(\mathcal{X} \subset \mathcal{Z})$ given, for each $r > 0$, by the *open* r -neighborhood $\mathcal{N}_r(\mathcal{X} \subset \mathcal{Z})$ of \mathcal{X} in \mathcal{Z} . We will refer to any filtration arising in that manner as a *neighborhood filtration*. Of particular relevance to this paper will be the case when $\mathcal{Z} = \mathbf{Y}$, a Banach space.

Given the similarities in the definitions of Vietoris–Rips and Čech simplicial complexes, it is natural to expect certain relationships between these two types of complexes. By direct computation, it is straightforward to see that for any compact $\mathcal{X} \subset \mathbf{Y}$,

$$\mathcal{C}_r(\mathcal{X} \subset \mathbf{Y}) \subseteq \mathcal{V}_{2r}(\mathcal{X}) \subseteq \mathcal{C}_{2r}(\mathcal{X} \subset \mathbf{Y}).$$

The Nerve Theorem directly relates the r -neighborhoods of \mathcal{X} in \mathbf{Y} , $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$, with the induced Čech simplicial complexes.

Lemma 2.5 (Nerve theorem; Alexandrov [Ale28] and Borsuk [Bor48]). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . For any $r > 0$, $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$ and $\mathcal{C}_r(\mathcal{X} \subset \mathbf{Y})$ are homotopy equivalent to each other, i.e.,*

$$\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \simeq \mathcal{C}_r(\mathcal{X} \subset \mathbf{Y}).$$

There are “persistent”, or “functorial”, versions of this result; see Bauer, Kerber, Roll and Rolle [BKRR23] for an overview of different variants of the functorial nerve lemma. We will use the version below (see the discussion in [LMO24, Remark 4.4]) to relate the Čech filtration $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y})$ and the Neighborhood filtration $\mathcal{N}_\bullet(\mathcal{X} \subset \mathbf{Y})$.

Theorem 2.6 (Persistent nerve theorem [LMO24, Proposition 4.5]). *There exist homotopy equivalences $\varphi_s : \mathcal{C}_s(\mathcal{X} \subset \mathbf{Y}) \rightarrow \mathcal{N}_s(\mathcal{X} \subset \mathbf{Y})$ for each $s > 0$ such that for each $t > s > 0$ the following diagram commutes up to homotopy:*

$$\begin{array}{ccc} \mathcal{C}_s(\mathcal{X} \subset \mathbf{Y}) & \longrightarrow & \mathcal{C}_t(\mathcal{X} \subset \mathbf{Y}) \\ \varphi_s \downarrow & & \downarrow \varphi_t \\ \mathcal{N}_s(\mathcal{X} \subset \mathbf{Y}) & \longrightarrow & \mathcal{N}_t(\mathcal{X} \subset \mathbf{Y}) \end{array}$$

Note that Theorem 2.6 implies that the persistent homology of $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y})$ is isomorphic to that of $\mathcal{N}_\bullet(\mathcal{X} \subset \mathbf{Y})$, a fact that we will repeatedly use in the sequel.

Since any compact metric space \mathcal{X} can be regarded as a subset of $L^\infty(\mathcal{X})$ (via its Kuratowski embedding, see Definition 2.24), one obtains an analogous result providing a connection between the Vietoris–Rips filtration and the filtration $\mathcal{N}_\bullet(\mathcal{X} \subset L^\infty(\mathcal{X}))$ consisting of nested neighborhoods

$$\{\mathcal{N}_s(\mathcal{X} \subset L^\infty(\mathcal{X})) \subseteq \mathcal{N}_t(\mathcal{X} \subset L^\infty(\mathcal{X}))\}_{0 < s \leq t}$$

of $\mathcal{X} \subset L^\infty(\mathcal{X})$.

Corollary 2.7 ([LMO24, Theorem 4.1]). *There exist homotopy equivalences $\varphi_s : \mathcal{V}_{2s}(\mathcal{X}) \rightarrow \mathcal{N}_s(\mathcal{X} \subset L^\infty(\mathcal{X}))$ for each $s > 0$ such that for each $t > s > 0$ the following diagram commutes up to homotopy:*

$$\begin{array}{ccc} \mathcal{V}_{2s}(\mathcal{X}) & \longrightarrow & \mathcal{V}_{2t}(\mathcal{X}) \\ \varphi_s \downarrow & & \downarrow \varphi_t \\ \mathcal{N}_s(\mathcal{X} \subset L^\infty(\mathcal{X})) & \longrightarrow & \mathcal{N}_t(\mathcal{X} \subset L^\infty(\mathcal{X})) \end{array}$$

Remark 2.8. The result above remains valid if $L^\infty(\mathcal{X})$ is substituted by any other *injective* metric space admitting an isometric embedding of \mathcal{X} ; see Section 2.2 for the definition and Section 5.2, where we in particular utilize the tight span $E(\mathcal{X})$ as one such choice. Also, the proof of Corollary 2.7 yields that $\mathcal{C}_\bullet(\mathcal{X} \subset L^\infty(\mathcal{X}))$ and $\mathcal{V}_{2\bullet}(\mathcal{X})$ are naturally homotopy equivalent; see [LMO24, Section 4].

Remark 2.9 (Alpha complexes). Note that while we only discuss VR and Čech filtered complexes in our paper, our results on Čech lifespans naturally apply to lifespans of homology classes induced by alpha complexes as $\mathcal{A}_r(\mathcal{X} \subset \mathbf{Y}) \simeq \mathcal{C}_r(\mathcal{X} \subset \mathbf{Y})$ where $\mathcal{A}_r(\mathcal{X} \subset \mathbf{Y})$ represents the alpha complex induced by \mathcal{X} with distance threshold $r \geq 0$; see Edelsbrunner and Harer [EH10, III.4].

2.1.2. Persistent homology. Here, we recall basic notions pertaining to persistent homology that are necessary for our setting. We will follow the presentation from [LMO24, Section 2.1].

Definition 2.10 (Persistence module). A *persistence module* $(V_r, \Phi_{r,s})_{0 < r \leq s}$ over $\mathbb{R}_{>0}$ is a family of \mathbb{F} -vector spaces V_r for some field \mathbb{F} with morphisms $\Phi_{r,s} : V_r \rightarrow V_s$ for each $r \leq s$ such that

- $\Phi_{r,r} = \text{id}_{V_r}$,
- $\Phi_{s,t} \circ \Phi_{r,s} = \Phi_{r,t}$ for each $r \leq s \leq t$.

For conciseness we will denote by V_\bullet the persistence module given by $(V_r, \Phi_{r,s})_{0 < r \leq s}$. The morphisms $\Phi_{\bullet,\bullet}$ are referred to as the *structure maps* of V_\bullet . Note that a persistence module V_\bullet can be regarded as a functor from the poset $(\mathbb{R}_{>0}, \leq)$ to the category of vector spaces.

Definition 2.11 (Interval persistence module). Given an interval I in $\mathbb{R}_{>0}$ and a field \mathbb{F} , the *interval persistence module* induced by I is the persistence module $\mathbb{F}_\bullet[I]$ defined as follows: The vector space at r is \mathbb{F} if r is in I and zero otherwise. Given $r \leq s$, the morphism corresponding to the pair (r, s) is the identity if r, s are both contained in I and zero otherwise.

Definition 2.12 (Barcode and persistence diagram). For a given persistence module V_\bullet , if there is a multiset of intervals $(I_\lambda)_{\lambda \in \Lambda}$ such that V_\bullet is isomorphic to $\bigoplus_{\lambda \in \Lambda} \mathbb{F}_\bullet[I_\lambda]$, then that multiset is referred to as a *(persistence) barcode* associated to the persistence module V_\bullet . Persistence Modules admitting such a multiset of intervals are said to be *interval decomposable*. The *persistence diagram* of V_\bullet is then given as the multiset of points $(b_\lambda, d_\lambda) \in \mathbb{R}^2$, where b_λ is the left endpoint of I_λ and d_λ is its right endpoint.⁴

In applied algebraic topology, many persistence modules arise as follows.

Definition 2.13 (Persistent homology of a filtration). For any $k \geq 0$, applying the k -dimensional *reduced homology* functor (with coefficients in a field \mathbb{F}) to a filtration $\Delta_\bullet = (\Delta_r, \iota_{r,s})_{0 < r \leq s}$ produces the persistence (homology) module $\tilde{H}_k(\Delta_\bullet; \mathbb{F}) = (\tilde{H}_k(\Delta_r; \mathbb{F}), \Phi_{r,s}^k)_{0 < r \leq s}$ where the morphisms $\Phi_{r,s}^k$ are those induced by $\iota_{r,s}$.⁵

In what follows we will drop the field \mathbb{F} from the notation since all of our results hold for an arbitrary choice of \mathbb{F} .

Under suitable assumptions, the persistence modules obtained from filtrations, as described above, are interval decomposable. In particular, the persistence modules obtained from neighborhood filtrations of compact subsets of a Banach space are interval decomposable (so that they admit barcodes).

Theorem 2.14 ([LMO24, Theorem 1]). *Assume \mathcal{X} is a compact subset of a Banach space \mathbf{Y} . Then there is a (unique) persistence barcode associated to the persistence module $\tilde{H}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$. In particular, $\tilde{H}_k(\mathcal{V}_\bullet(\mathcal{X}))$ admits (unique) persistence barcode.*⁶

We will henceforth use $\text{PD}_k(\mathcal{X})$ and $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$ to respectively denote the persistence diagrams of the Vietoris–Rips and Čech filtrations of \mathcal{X} .

⁴Not every persistence module is interval decomposable; see Crawley-Boevey [CB15] for more details.

⁵Note that we are using reduced homology in our definition in order to dispense with the usual infinite length bar at the level of degree zero persistent homology.

⁶In [LMO24, Theorem 1] the authors only contemplate the case of the \mathbf{Y} being equal to $L^\infty(\mathcal{X})$ for some compact metric space \mathcal{X} . However, the proof of their result directly applies to the setting in the statement.

2.1.3. *Stability theorems.* Persistence diagrams are an effective methodology for encoding topological properties of the space \mathcal{X} and its neighborhoods in \mathbf{Y} . Persistence diagrams are stable, as expressed by the following stability theorems. Informally, these state that if the shape and the size of two spaces are similar, then their persistence diagrams are close to each other. To give formal statements, let $\mathbf{d}_b(\cdot, \cdot)$ denote the bottleneck distance between persistence diagrams; see Edelsbrunner and Harer [EH10]. Let \mathbf{d}_H^Y be the Hausdorff distance between two subsets of the same Banach space \mathbf{Y} , and let \mathbf{d}_{GH} be the Gromov–Hausdorff distance between two metric spaces; see Burago, Burago and Ivanov [BBI01, Chapter 7].

Lemma 2.15 (Stability theorem—VR; [CCSG⁺09, Theorem 3.1] and [CDSO14, Theorem 5.2]). *Let \mathcal{X} and \mathcal{X}' be two compact metric spaces. Then,*

$$\mathbf{d}_b(\text{PD}_k(\mathcal{X}), \text{PD}_k(\mathcal{X}')) \leq 2 \mathbf{d}_{GH}(\mathcal{X}, \mathcal{X}').$$

Lemma 2.16 (Stability theorem—Čech; [CDSO14, Theorem 5.6]). *Let $\mathcal{X}, \mathcal{X}'$ be two compact subsets of a Banach space \mathbf{Y} . Then,*

$$\mathbf{d}_b(\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y}), \widetilde{\text{PD}}_k(\mathcal{X}' \subset \mathbf{Y})) \leq \mathbf{d}_H^Y(\mathcal{X}, \mathcal{X}').$$

Note that Lemma 2.16 implies Lemma 2.15. Indeed, this was implicitly used in the proof of [CCSG⁺09, Theorem 3.1]. Notice that the coefficient 2 does not appear in the second stability theorem.

2.1.4. *Homological spectra.* As described in Section 1, the chief goal of our paper is to provide effective bounds for the lifetime of all homology classes that appear along a (geometric) filtration of a metric space. We will formulate and realize this goal in a setting that encompasses, but is more general than, persistence diagrams.

Remark 2.17. Notice that it is not true that the only homology classes that show up across the filtration are those coming from the initial space. One well-known example is that of the circle S^1 (with its geodesic distance) and the Vietoris–Rips filtration. Indeed, as shown by Adamszek and Adams in [AA17], whereas $\mathcal{V}_t(S^1)$ has the homotopy type of S^1 for $t \in (0, \frac{2\pi}{3}]$, its homotopy type is that of S^3 as soon as $t \in (\frac{2\pi}{3}, \frac{4\pi}{5}]$. In fact, they show that $\mathcal{V}_\bullet(S^1)$ eventually attains the homotopy types of all odd-dimensional spheres.

We recall some additional definitions and results from [LMO24].

Definition 2.18. For an integer $k \geq 0$, a given field \mathbb{F} , and a filtration

$$\Delta_\bullet = (\Delta_r, \iota_{r,s})_{0 < r \leq s},$$

let

$$(1) \quad \text{Spec}_k(\Delta_\bullet) := \bigcup_{r>0} \left(\widetilde{\text{H}}_k(\Delta_r; \mathbb{F}) \setminus \{0\} \times \{r\} \right)$$

be the k -th *homological spectrum* of Δ_\bullet (with coefficients in \mathbb{F}).

Now, fix an arbitrary $(\omega, s) \in \text{Spec}_k(\Delta_\bullet)$. Then, let

(2)

$$b_{(\omega,s)} := \inf\{r > 0 : r \leq s \text{ and } \exists \text{ nonzero } \omega_r \in \tilde{H}_k(\Delta_r) \text{ such that } (\iota_{r,s})_*(\omega_r) = \omega\},$$

(3)

$$d_{(\omega,s)} := \sup\{t > 0 : t \geq s \text{ and } \exists \text{ nonzero } \omega_t \in \tilde{H}_k(\Delta_t) \text{ such that } (\iota_{s,t})_*(\omega) = \omega_t\}$$

$$(4) \quad = \sup\{t \geq s : (\iota_{s,t})_*(\omega) \neq 0\}.$$

Whenever Δ_\bullet is a neighborhood filtration $\mathcal{N}_\bullet(\mathcal{X} \subset \mathbf{Y})$, as in [LMO24, Theorem 8], one has that $b_{(\omega,s)} < s \leq d_{(\omega,s)}$.⁷ Let

$$I_{(\omega,s)} := (b_{(\omega,s)}, d_{(\omega,s)}].$$

Informally, the interval $I_{(\omega,s)}$ encodes the maximal region around $s \in \mathbb{R}_{>0}$ inside which the class ω is “alive”.⁸

Definition 2.19. The value $b_{(\omega,s)}$ is referred to as the *birth time* of ω whereas $d_{(\omega,s)}$ is the *death time* of ω . The value $d_{(\omega,s)} - b_{(\omega,s)}$ will be referred to as the *lifespan* of ω .

Remark 2.20 (A caveat). We focus on the special case of neighborhood filtrations. The birth time $b_{(\omega,s)}$ was defined as the infimum of all times $r \leq s$ when a “predecessor” of ω exists in $\tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$, and it is natural to ask whether there exists a homology class supported on

$$\bigcap_{b_{(\omega,s)} < r \leq s} \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) = \overline{\mathcal{N}}_{b_{(\omega,s)}}(\mathcal{X} \subset \mathbf{Y})$$

that is also homologous to ω in $\mathcal{N}_s(\mathcal{X} \subset \mathbf{Y})$. It turns out this is not the case even for the neighborhood filtration of a compact set in \mathbb{R}^2 . Namely, there exists a compact set $\mathcal{X} \subset \mathbb{R}^2$, known as the *Warsaw circle* or *closed topologist’s sine curve*, satisfying the following counter-intuitive property: its first singular homology is zero, but every open neighborhood of it contains a homologically nontrivial circle; see Borsuk [Bor75]. Every two of those circles are homologous to each other (within the union of the two neighborhoods). But these circles do not converge, as we shrink the neighborhood, to a nontrivial homology class of \mathcal{X} , because $\tilde{H}_1(\mathcal{X}) = 0$. The natural way to treat this “limit circle” is to consider Čech homology instead of ordinary singular homology. The Čech homology $\tilde{H}_1(\mathcal{X})$ is nontrivial, and contains the “limit circle”. In general, a predecessor of the class ω naturally lives in the Čech homology $\tilde{H}_k(\overline{\mathcal{N}}_{b_{(\omega,s)}}(\mathcal{X} \subset \mathbf{Y}))$. Nonetheless, we consistently use singular homology throughout the paper in order to avoid overly technical details.

To ease the notational burden, we will often drop the parameter s and will use the more succinct notation b_ω , d_ω , I_ω , etc.

Remark 2.21. A priori, one would expect the collection of all intervals

$$\{I_\omega; \omega \in \text{Spec}_k(\Delta_\bullet)\}$$

⁷In general, the type of intervals (open-open, closed-open, etc) one obtains depends on whether the filtration is defined via open or closed neighborhoods. Note that we’ve defined neighborhood filtrations via *open* neighborhoods.

⁸Note that, since we are using reduced homology, there is no degree 0 class ω such that $d_{(\omega,s)} = \infty$.

to be closely related to the k -th persistence diagram of Δ_\bullet . Whereas Proposition 2.22 below establishes a sense in which this is the case, it is not always true that all such intervals appear in the interval decomposition of $\tilde{H}_k(\Delta_\bullet)$ (whenever it exists). An example showing this discrepancy in the case of Vietoris–Rips filtrations can be found in [LMO24, Example 9.16].

Proposition 2.22 ([LMO20, Proposition 9.2]). *Let $\Delta_\bullet = (\Delta_r, \iota_{r,s})_{0 < r \leq s}$ be a neighborhood filtration and let $k \geq 1$ be an integer. Then, for all $r < s$, the multiplicity of the interval $(r, s]$ in the barcode of $H_k(\Delta_\bullet)$ is equal to*

$$\max \left\{ m \in \mathbb{Z}_{\geq 0} \mid \begin{array}{l} \exists \text{ linearly independent vectors } \omega_1, \dots, \omega_m \in H_k(\Delta_s) \text{ s.t. } I_{(\omega_i, s)} = (r, s] \forall i \\ \text{and no nonzero linear combination of these vectors belongs to } \text{Im}((\iota_{r,s})_*) \end{array} \right\}.$$

Proposition 2.22 indicates that for each interval $I = (r, s]$ in the barcode of Δ_\bullet there is a finite linearly independent collection $\omega_1, \dots, \omega_m \in H_k(\Delta_s)$ satisfying the conditions above such that $I = I_{(\omega_i, s)}$ for all i . One calls any such ω_i a *representative* of the interval I .⁹ See [LMO24, page 42] for an example demonstrating the role of the condition that no nonzero linear combination of these vectors belongs to $\text{Im}((\iota_{r,s})_*)$.

Remark 2.23. Proposition 2.22 implies that if I is an interval in the barcode of a neighborhood filtration, then $I = I_{\omega_i}$ so that, in particular, its length is equal to that of I_{ω_i} , where ω_i is as in the statement. Therefore, and as we will do in the rest of the paper, if we have an upper bound for the length of all intervals I_ω where $\omega \in \text{Spec}_k(\Delta_\bullet)$ then we will automatically have an upper bound for the length of every interval in the barcode of Δ_\bullet .

2.2. Geometry of $L^\infty(\mathcal{X})$ and tight spans. While the discussion in this paper applies to subsets of any Banach space, special attention is paid to the important case of L^∞ spaces. There are two main reasons for that. The first one is that if we start with a compact metric space that is a priori not an isometrically embedded into a Banach space, there is a nice way of placing it inside $L^\infty(\mathcal{X})$, the space of bounded functions on \mathcal{X} with the supremum norm.

Definition 2.24 (Kuratowski embedding). For a compact metric space $(\mathcal{X}, \mathbf{d}_\mathcal{X})$, the map $\kappa : \mathcal{X} \rightarrow L^\infty(\mathcal{X})$, defined as $x \mapsto \mathbf{d}_\mathcal{X}(x, \cdot)$, is a distance-preserving embedding, and it is called the *Kuratowski embedding*.

The second reason is that L^∞ spaces enjoy the *hyperconvexity* property: If several balls have non-empty pairwise intersection then they share a point in common.

Definition 2.25 (Hyperconvex space). A metric space $(\mathcal{E}, \mathbf{d}_\mathcal{E})$ is called *hyperconvex* if for every family $(x_i, r_i)_{i \in I}$ of x_i in \mathcal{E} and $r_i \geq 0$ such that $\mathbf{d}_\mathcal{E}(x_i, x_j) \leq r_i + r_j$ for every i, j in I , there exists a point $x \in \mathcal{E}$ such that $\mathbf{d}_\mathcal{E}(x_i, x) \leq r_i$ for every i in I .

The hyperconvexity of $L^\infty(\mathcal{X})$ implies that the Čech and VR filtrations of a subset of $L^\infty(\mathcal{X})$ coincide (up to a factor of two in the filtration index). This gives a way to study VR-lifespans of homology of \mathcal{X} by immersing it in $L^\infty(\mathcal{X})$ and then using the tools applicable to Čech-lifespans. Note that this idea has been successfully used in the context of VR-persistence, cf. Proposition 2.34.

⁹In [LMO24, Proposition 9.2] the authors consider the case of the VR filtration of a totally bounded metric space. The same proof applies to the more general statement given above.

Hyperconvexity implies certain universal properties of L^∞ spaces, and these properties will be implicitly used below in relation to Urysohn width (see Remark 3.16), and tight spans (a.k.a. injective hulls). We briefly discuss the latter now and we refer the reader to Lang's survey [Lan13] for more information.

Definition 2.26 (Injective metric space). A metric space \mathcal{E} is called *injective* if for every 1-Lipschitz map $\phi : \mathcal{X} \rightarrow \mathcal{E}$ and distance preserving embedding of \mathcal{X} into $\tilde{\mathcal{X}}$, there exists a 1-Lipschitz map $\tilde{\phi} : \tilde{\mathcal{X}} \rightarrow \mathcal{E}$ extending ϕ :

$$\begin{array}{ccc} \mathcal{X} & \xhookrightarrow{\quad} & \tilde{\mathcal{X}} \\ & \searrow \phi & \downarrow \tilde{\phi} \\ & & \mathcal{E} \end{array}$$

It turns out that injectivity coincides with hyperconvexity.

Proposition 2.27. *A metric space is injective if and only if it is hyperconvex.*

The proof of Proposition 2.27 can be found in Aronszajn and Panitchpakdi [AP56] and in [Lan13, Proposition 2.3].

It is known that $L^\infty(\mathcal{X})$ is injective [Lan13]. However, there exists a more efficient injective space containing \mathcal{X} .

Definition 2.28 (Tight span). The *tight span* $E(\mathcal{X})$ of a compact metric space \mathcal{X} is the minimal injective metric space admitting an isometric embedding of \mathcal{X} . Minimality here means that any other injective metric space admitting an isometric embedding of \mathcal{X} contains an isometric copy of $E(\mathcal{X})$.

Tight spans are sometimes called *injective hulls* or *hyperconvex hulls*. The notions of injectivity and hyperconvexity were first proposed by Aronszajn and Panitchpakdi [AP56]. Isbell [Isb64] first identified the notion of tight span (although the author used the term *injective envelope*). Additional contributions were made by Dress [Dre84] and Lang [Lan13]; see Chepoi [Che97] for a historical account.

Proposition 2.29 (Properties of the tight span [Lan13]). *The tight span $E(\mathcal{X})$ of a compact metric space \mathcal{X} exists and satisfies:*

- (1) $E(\mathcal{X})$ is compact.
- (2) $E(\mathcal{X})$ is contractible.
- (3) $E(\mathcal{X})$ is isometric to \mathcal{X} for any metric tree \mathcal{X} .
- (4) $\text{diam}(E(\mathcal{X})) = \text{diam}(\mathcal{X})$.

One particular realization of the tight span $E(\mathcal{X})$ of \mathcal{X} as a subset of $L^\infty(\mathcal{X})$ is given as follows [Lan13, Section 3]:

$$E(\mathcal{X}) := \{f \in \Delta(\mathcal{X}) : \text{if } g \in \Delta(\mathcal{X}) \text{ and } g \leq f, \text{ then } g = f \text{ (i.e., } f \text{ is minimal)}\},$$

where

$$\Delta(\mathcal{X}) := \{f \in L^\infty(\mathcal{X}) : f(x) + f(x') \geq \mathbf{d}_{\mathcal{X}}(x, x') \text{ for all } x, x' \in \mathcal{X}\}.$$

Note that from the realization of the tight span recalled above, for any $f \in E(\mathcal{X})$ and any $x \in \mathcal{X}$, it holds that

$$(5) \quad f(x) = \max_{x' \in \mathcal{X}} (\mathbf{d}_{\mathcal{X}}(x, x') - f(x')) = \|\mathbf{d}_{\mathcal{X}}(x, \cdot) - f\|_\infty.$$

Remark 2.30. Not only can the tight span $E(\mathcal{X})$ be regarded as a subset of $L^\infty(\mathcal{X})$ but also, directly from the fact that it is an injective metric space, there is a 1-Lipschitz retraction of $L^\infty(\mathcal{X})$ to $E(\mathcal{X})$; see [Lan13, Proposition 2.2].

2.3. Filling radius. The filling radius is a key notion in metric geometry introduced by Gromov as a measure of largeness of a given closed manifold [Gro83]. To *fill* a manifold \mathcal{M} of dimension n , treated as a singular n -cycle, means to find an $(n+1)$ -dimensional singular chain \mathcal{D} with boundary \mathcal{M} , i.e. $\partial\mathcal{D} = \mathcal{M}$. In this case we will also be saying that the cycle \mathcal{M} *bounds*. The ambient space in which filling happens, as well as the coefficients of singular homology, should be specified, as we discuss below.

Definition 2.31 (Gromov's filling radius [Gro83]). The filling radius $\rho(\mathcal{M})$ of a closed n -dimensional Riemannian manifold \mathcal{M} is the infimal number $R > 0$ such that \mathcal{M} can be filled inside of the R -neighborhood of its Kuratowski image in $L^\infty(\mathcal{M})$.

Remark 2.32 (Coefficients). This definition makes sense with any homology coefficients. Some common choices include \mathbb{Z} -coefficients if \mathcal{M} is oriented, \mathbb{Z}_2 -coefficients if \mathcal{M} is not oriented, and \mathbb{Q} -coefficients in some contexts where torsion is a problem. However, the usual persistence homology is well-defined only over fields,¹⁰ so in the rest of the paper we implicitly assume that an arbitrary choice of a field is made (e.g., \mathbb{Z}_2), and that all filling radii and all persistence features are considered over this field. All of our results hold for any field and, for this reason, the field is omitted from the notation.

Before Gromov, a different type of filling radius notion was discussed for submanifolds (or more generally, cycles) in \mathbb{R}^N in geometric analysis, especially in the context of the isoperimetric problem [FF60, MS73, BS83]. We summarize both types of filling radii in Definition 2.33, adapting it to the context of persistence.

Definition 2.33 (Relative and absolute filling radii of a homology class). Let $(\mathcal{X}, \mathbf{d}_\mathcal{X})$ be a metric space, and let $\omega \in \tilde{H}_k(\mathcal{X})$ be a nontrivial reduced singular homology class.

- (1) Assume additionally that \mathcal{X} is a subset of a Banach space \mathbf{Y} , so that the metric $\mathbf{d}_\mathcal{X}$ agrees with the one inherited from \mathbf{Y} (that is, the embedding $\mathcal{X} \hookrightarrow \mathbf{Y}$ is distance-preserving). The *relative filling radius* of ω is the infimal number r such that the image of ω under the induced map $\tilde{H}_k(\mathcal{X}) \rightarrow \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$ is trivial. In other words, it is the infimal size of a neighborhood in \mathbf{Y} where some cycle representing ω bounds a $(k+1)$ -chain. We will use the following notation for the relative filling radius: $\rho(\omega; \mathcal{X} \subset \mathbf{Y})$.
- (2) In case when no ambient space \mathbf{Y} is specified, one can take $\mathbf{Y} = L^\infty(\mathcal{X})$, and measure the relative filling radius of the Kuratowski image $\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})$. This way one obtains the *(absolute) filling radius* of ω : $\rho(\omega; \mathcal{X}) = \rho(\kappa_*(\omega); \kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))$; cf. [LMO24, Definition 24].

The word “absolute”, which is usually omitted, is justified by the fact that $\rho(\omega; \mathcal{X})$ equals the infimum of $\rho(\iota_*(\omega); \iota(\mathcal{X}) \subset \mathbf{Y})$ over all distance-preserving embeddings $\iota : \mathcal{X} \hookrightarrow \mathbf{Y}$ [Gro83, page 8]; we will refer to this as the *universal property*. Note that when $\mathcal{X} = \mathcal{M}$ is a closed Riemannian manifold, Gromov's filling radius (Definition 2.31) over \mathbb{Z}_2 is the same as the absolute filling radius of the \mathbb{Z}_2 -fundamental class $[\mathcal{M}]$.

¹⁰See, however, [Pat18].

The relationship between the VR filtration of a compact metric space and the absolute filling radii $\rho(\omega; \mathcal{X})$ is studied in [LMO24, Section 9.3]. Proposition 2.34 applies in the general context of absolute neighborhood retracts (ANRs), which includes Riemannian manifolds, metric graphs and other commonly appearing metric spaces.¹¹

Proposition 2.34 ([LMO24, Propositions 9.28 and 9.46]). *Let \mathcal{X} be a compact ANR metric space. Then, for any integer $k \geq 1$ and any nonzero $\omega \in \tilde{H}_k(\mathcal{X})$ we have:*

- $\rho(\omega; \mathcal{X}) > 0$;
- the interval $(0, 2\rho(\omega; \mathcal{X}))$ appears in the degree- k barcode of $\mathcal{V}_\bullet(\mathcal{X})$.

Additionally, if \mathcal{X} is a (closed and connected) Riemannian manifold, then there are no other intervals with left-endpoint (birth time) equal to zero in the degree- k barcode of $\mathcal{V}_\bullet(\mathcal{X})$.¹² Here both the filling radius and persistent homology can be computed with coefficients in an arbitrary field when \mathcal{M} is orientable, and with coefficients in \mathbb{Z}_2 when \mathcal{M} is not orientable.

Remark 2.35 (Relative filling radius vs. absolute filling radius). To give an idea about the difference between relative filling radius in \mathbf{Y} and absolute filling radius, here we give a toy example of a flying saucer in $\mathbf{Y} = \mathbb{R}^3$ (with Euclidean metric). Let E be the ellipsoid in \mathbb{R}^3 given by $E = \{(x, y, z) \mid \frac{x^2}{100} + \frac{y^2}{100} + \frac{z^2}{1} = 1\}$, and let $[E]$ be its fundamental class in $H_2(E)$. There are two different ways to treat E as a metric space, resulting in the different values of its filling radius.

- (1) One way is to consider the Riemannian metric g on E induced by the Euclidean metric of \mathbb{R}^3 . The Riemannian surface E_g thus obtained has *absolute filling radius* of about 10: $\rho([E]; E_g) \sim 10$ (the exact computation is tricky).
- (2) The other way is to borrow the extrinsic distance function from \mathbb{R}^3 . This does not make E a Riemannian manifold, but rather just a compact metric space embedded in \mathbb{R}^3 in a distance-preserving way. The corresponding *relative filling radius* equals 1: $\rho([E]; E \subset \mathbb{R}^3) = 1$.

In this example we have $\rho([E]; E_g) > \rho([E]; E \subset \mathbb{R}^3)$, which might seem to contradict the note above saying that $\rho(\omega; \mathcal{X}) \leq \rho(\omega; \mathcal{X} \subset \mathbf{Y})$ for distance-preserving embeddings $\mathcal{X} \subset \mathbf{Y}$. There is no contradiction here: even though the embedding $E_g \subset \mathbb{R}^3$ is a Riemannian isometry, it is not distance-preserving. Indeed, the extrinsic distance in \mathbb{R}^3 between two points of E is smaller than the intrinsic distance between them inside E (the length of the shortest path in E). If one computes the filling radius of E with the extrinsic metric of \mathbb{R}^3 , it will be at most 1.

In the following sections, we will see that Gromov's filling radius of a Riemannian manifold isometrically embedded in an ambient space \mathbf{Y} can highly overestimate the lifespan of a topological feature in its Čech filtration and that the relative filling radius is better adapted to this context.

¹¹Recall that an ANR is any metric space \mathcal{X} with the property that whenever it is embedded into another metric space \mathcal{Z} through a homeomorphism $h : \mathcal{X} \rightarrow \mathcal{Z}$, then there is an open neighborhood U of $h(\mathcal{X})$ such that $h(\mathcal{X})$ is a retract of U ; see Borsuk [Bor32] and Hu [Hu65]. It is known that every compact (topologically) finite-dimensional locally contractible metric space is an ANR. Thus, all Riemannian manifolds are ANRs.

¹²In particular, if \mathcal{X} is an n -dimensional Riemannian manifold, then $(0, 2\rho(\mathcal{X}))$ is the unique interval with left endpoint zero in the degree- n barcode.

Remark 2.36 (Comment about estimation of filling invariants). The type of connections between persistent homology and metric geometry that we explore in this paper have the potential of permitting the estimation of quantities such as the filling radii $\rho(\omega; \mathcal{X})$, $\omega \in \tilde{H}_k(\mathcal{X})$, via the polynomial time algorithms that have been developed for computing persistent homology (see Section 2.1). Indeed, such an estimate would be obtained via Proposition 2.34 and the stability of PH (Lemma 2.15) through computing the VR-barcodes of a carefully chosen ϵ -net for \mathcal{X} , for some $\epsilon > 0$.

Filling radii & lifespans. The notion of *lifespan* of a homology class (Definition 2.19) is directly related to the filling radius as follows. Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Let (b_ω, d_ω) be the homological birth and death times of a degree- k homology class ω present at time s in the neighborhood filtration $\mathcal{N}_\bullet(\mathcal{X} \subset \mathbf{Y})$. Directly from the definition of the relative filling radius, we obtain

$$d_\omega - s = \rho(\omega; \mathcal{N}_s(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}).$$

A bit more generally, if $\omega_r \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$ is a predecessor of ω (that is, it is mapped to ω by the map in homology induced by the inclusion), then

$$d_\omega - r = \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}).$$

Letting $r \rightarrow b_\omega$, we obtain a formula for the lifespan of ω in terms of the filling radii of the predecessors of ω :

$$\begin{aligned} (\star) \quad d_\omega - b_\omega &= \sup_{b_\omega < r \leq s} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \\ &= \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}). \end{aligned}$$

Note that both sides of this formula depend on homology coefficients lying in the same field, which can be arbitrary.

Remark 2.37. As clarified in Remark 2.20, this lifespan cannot be written in terms of the filling radius of a homology class of $\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \mathbf{Y})$. We cannot expect (a predecessor of) ω to be present in the homology of $\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \mathbf{Y})$, unless we work with Čech homology, and redefine the filling radius correspondingly.

Formula (\star) was given for the neighborhood filtration. In view of the functorial nerve theorem (Theorem 2.6, Corollary 2.7), the formula specializes to two important cases:

- For any $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$,

$$d_\omega - b_\omega = \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}).$$

- For any $\omega \in \text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$,

$$d_\omega - b_\omega = 2 \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})).$$

Remark 2.38. For the particular case of VR-lifespans, i.e. when $\mathbf{Y} = L^\infty(\mathcal{X})$, it is important to understand the behavior of the filling radii $\rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X}))$. We make a comment that this cumbersome notation is a bit redundant; it turns out that

$$\rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})) = \rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))).$$

This is not obvious by default, since the absolute filling radius in the right-hand side should be computed in the space $L^\infty(\mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})))$. Nevertheless, this equality holds true, as it is explained in Section 3.3 (see Lemma 3.15).

Auxiliary definitions. Before concluding the background section, we define two versions of the radius of a set \mathcal{X} , which will be used in the remainder of the paper. In the first one, there is no reference to an ambient space, and the center is in the set \mathcal{X} . In the second one (circumradius), the radius of \mathcal{X} is computed in an ambient space, and the center may not be in \mathcal{X} .

Definition 2.39 (Radius). Let $(\mathcal{X}, \mathbf{d}_{\mathcal{X}})$ be a compact metric space.

(1) The *radius* of \mathcal{X} is

$$\text{rad}(\mathcal{X}) := \inf_{x_0 \in \mathcal{X}} \sup_{x \in \mathcal{X}} \mathbf{d}_{\mathcal{X}}(x, x_0).$$

(2) Assume additionally that \mathcal{X} is a compact subset of a Banach space $(\mathbf{Y}, \|\cdot\|)$.¹³ The *circumradius* of \mathcal{X} in \mathbf{Y} is

$$\text{rad}(\mathcal{X} \subset \mathbf{Y}) := \inf_{y \in \mathbf{Y}} \sup_{x \in \mathcal{X}} \|x - y\|.$$

3. BOUNDING LIFESPANS VIA WIDTHS

In this section, we recall and establish upper bounds for the filling radius which can be used to estimate the lifespans in both Čech and Vietoris–Rips settings. The basic idea behind many constructions is simple and can be illustrated as follows. Let \mathcal{X} be a subset of a Banach space \mathbf{Y} , and let $k > \ell$ be positive integers. A degree- k homology class of \mathcal{X} can be “killed” by deforming \mathcal{X} to an ℓ -dimensional complex inside \mathbf{Y} (a “core”), and if every point moves by some controlled distance, then we obtain an estimate for the filling radius of the degree- k homology.

3.1. Background on widths. Here we recall three classical approaches to measuring approximate dimension. Informally, the k -width of a space \mathcal{X} measures the extent to which \mathcal{X} fails to be k -dimensional.

The Urysohn width was historically the first one to be introduced. Definition 3.1 is equivalent to the one given by Urysohn around 1923 in the context of dimension theory; it was posthumously published by Alexandrov [Ale26].

Definition 3.1 (Urysohn width). Let \mathcal{X} be a compact metric space. For an integer $k \geq 0$, the *Urysohn k -width* of \mathcal{X} is defined as

$$\text{UW}_k(\mathcal{X}) = \inf_f \sup_p \text{diam}(f^{-1}(p)),$$

where $f : \mathcal{X} \rightarrow \Delta^k$ is any continuous map to any finite k -dimensional simplicial complex.

By definition, widths are monotone in the dimension parameter:

$$\text{UW}_0(\mathcal{X}) \geq \text{UW}_1(\mathcal{X}) \geq \text{UW}_2(\mathcal{X}) \geq \dots,$$

and $\text{UW}_0(\mathcal{X}) := \text{diam}(\mathcal{X})$ if \mathcal{X} is connected. The n -width of an n -dimensional manifold is zero, and all preceding k -widths are positive for $0 \leq k < n$.

The crucial connection between the Urysohn width and the filling radius is due to Gromov.

¹³So that $\mathbf{d}_{\mathcal{X}}(x, x') = \|x - y\|$ for all $x, x' \in \mathcal{X}$.

Theorem 3.2 ([Gro83]). *For any closed n -dimensional Riemannian manifold \mathcal{M} ,*

$$\rho(\mathcal{M}) \leq \frac{1}{2} \text{UW}_{n-1}(\mathcal{M}).$$

The homology coefficients in the definition of $\rho(\mathcal{M})$ do not matter here; everywhere below we assume them to lie in an arbitrary field, which is fixed and omitted from the notation. The importance of this result for bounding lifespans becomes immediate once one notices that the proof of this inequality can be easily generalized for any homological feature, and not just the fundamental class.

Theorem 3.3. *For any compact metric space \mathcal{X} , any integer $k \geq 1$, and any homology class $\omega \in \widetilde{H}_k(\mathcal{X})$,*

$$\rho(\omega; \mathcal{X}) \leq \frac{1}{2} \text{UW}_{k-1}(\mathcal{X}).$$

Proof sketch, following [Gro20]. Let $f : \mathcal{X} \rightarrow \Delta^{k-1}$ be a map for which $\delta = \sup_p \text{diam}(f^{-1}(p))$ is just a tiny amount bigger than $\text{UW}_{k-1}(\mathcal{X})$. Consider the cylinder $\mathcal{X} \times [0, \delta/2]$ and glue its end $\mathcal{X} \times \{\delta/2\}$ to a copy of Δ^{k-1} along the map f ; that means, pinch every fiber of f inside $\mathcal{X} \times \{\delta/2\}$ to a point. The resulting space \mathcal{C}_f can be endowed with a metric that restricts on $\mathcal{X} \times \{0\}$ to the original metric of \mathcal{X} , and makes the length of every interval $\{x\} \times [0, \delta/2]$ equal to $\delta/2$. Now, any k -cycle S in $\mathcal{X} \times \{0\}$ representing ω becomes null-homologous in \mathcal{C}_f (informally, one can just slide it towards the pinched end of the cylinder, where it degenerates to a $(k-1)$ -dimensional set). Next, we embed \mathcal{C}_f to $L^\infty(\mathcal{C}_f)$ in a distance-preserving way, and post-compose it with the 1-Lipschitz restriction $L^\infty(\mathcal{C}_f) \rightarrow L^\infty(\mathcal{X})$, corresponding to the inclusion $\mathcal{X} \times \{0\} \subset \mathcal{C}_f$. The $(k+1)$ -chain that we built in \mathcal{C}_f to fill S pushes forward to $L^\infty(\mathcal{X})$, and there it lies within distance $\delta/2$ of the Kuratowski image of \mathcal{X} . \square

It is natural to look for a relative version of Theorem 3.3. We provide such an inequality below (Theorem 3.5), additionally replacing its right-hand side by the quantity that is comparable to the Urysohn width but easier to compute; the corresponding metric invariant implicitly appeared in the work of Alexandrov [Ale33] on dimension theory.

Definition 3.4 (Alexandrov width). Let \mathcal{X} be a compact subset of a Banach space $(\mathbf{Y}, \|\cdot\|)$. For an integer $k \geq 0$, the *Alexandrov k -width* of \mathcal{X} (relative to \mathbf{Y}) is defined as

$$\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) := \inf_f \sup_{x \in \mathcal{X}} \|x - f(x)\|,$$

where $f : \mathcal{X} \rightarrow \mathbf{Y}$ is any continuous map whose image is a finite simplicial complex of dimension at most k .

The Alexandrov width enjoys similar properties with the Urysohn width: it is monotonically decreasing in k until it reaches zero when k becomes equal to the dimension of \mathcal{X} . If \mathcal{X} is connected, $\text{AW}_0(\mathcal{X} \subset \mathbf{Y})$ equals the circumradius of \mathcal{X} in \mathbf{Y} , that is, $\text{AW}_0(\mathcal{X} \subset \mathbf{Y}) = \text{rad}(\mathcal{X} \subset \mathbf{Y})$ (see Definition 2.39).

It is easy to see directly from the definitions that $\text{UW}_k(\mathcal{X}) \leq 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y})$. Combining this with Theorem 3.3, we obtain the estimate $\rho(\omega; \mathcal{X}) \leq \text{AW}_{k-1}(\mathcal{X} \subset \mathbf{Y})$, which we will improve by replacing the absolute filling radius by the relative one.

Theorem 3.5. *For any compact set \mathcal{X} sitting in a Banach space \mathbf{Y} , any integer $k \geq 1$, and any homology class $\omega \in H_k(\mathcal{X})$,*

$$\rho(\omega; \mathcal{X} \subset \mathbf{Y}) \leq \text{AW}_{k-1}(\mathcal{X} \subset \mathbf{Y}).$$

This is the relative counterpart of Theorem 3.3. It will follow from a stronger estimate below (Theorem 3.10), but the intuition behind it is simple as described at the preamble of this section: One can kill higher homology of \mathcal{X} by deforming it to a low-dimensional complex in \mathbf{Y} , and if every point moves by some controlled distance, then we obtain an estimate for the filling radius.

The virtues of Theorem 3.5 are twofold. First, the Alexandrov width seems to be easier to estimate than the Urysohn width. Basically, the Urysohn width considers all maps to a k -dimensional space, whereas the Alexandrov width only considers a special class of those maps with the images lying in the same ambient space \mathbf{Y} ; this is also the reason why $\text{UW}_k(\mathcal{X}) \leq 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y})$. Second, in Appendix A we explain that $\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{UW}_k(\mathcal{X})$. Hence, together with Theorem 3.5, this implies one can bound the relative filling radius with the Urysohn width (if \mathcal{X} is compact), too, and we do not lose much when substituting the widths, since

$$\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{UW}_k(\mathcal{X}) \leq 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y}).$$

When \mathbf{Y} is hyperconvex (for example, $L^\infty(\mathcal{X})$), one can claim more (see Remark A.2):

$$\text{UW}_k(\mathcal{X}) = 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y}).$$

While Urysohn's and Alexandrov's notions measure *nonlinear* width, a simpler concept of *linear* width appeared in the work of Kolmogorov [Kol36] in the context of approximation theory. It is commonly used in infinite-dimensional settings.

Definition 3.6 (Kolmogorov width). Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . For an integer $k \geq 0$, the *Kolmogorov k -width* of \mathcal{X} (relative to \mathbf{Y}) is defined as

$$\text{KW}_k(\mathcal{X} \subset \mathbf{Y}) := \inf\{r \mid \mathcal{X} \subset \mathcal{N}_r(\mathcal{P} \subset \mathbf{Y}) \text{ for some affine } k\text{-plane } \mathcal{P} \subset \mathbf{Y}\}.$$

Again, these quantities monotonically decrease in k . It is immediate from definitions that $\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{KW}_k(\mathcal{X} \subset \mathbf{Y})$. From Theorem 3.5 it follows then the relative filling radius of degree- k homology can be bounded as $\rho(\cdot; \mathcal{X} \subset \mathbf{Y}) \leq \text{KW}_{k-1}(\mathcal{X} \subset \mathbf{Y})$. The following finer estimate is another corollary of Theorem 3.10.

Theorem 3.7. *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Then for any homology class $\omega \in \tilde{H}_k(\mathcal{X})$, $k \geq 0$,*

$$\rho(\omega; \mathcal{X} \subset \mathbf{Y}) \leq \text{KW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Remark 3.8 (Filling radius and k -widths). We give a toy example to illustrate the relation between the k -width and filling radius.

For $m \leq N$, consider the $(m-1)$ -dimensional ellipsoid $E \subset \mathbb{R}^N$ (with the extrinsic metric) given by

$$E := \left\{ \mathbf{x} \in \mathbb{R}^N \mid \sum_{i=1}^m \frac{x_i^2}{a_i^2} = 1, \text{ and } x_j = 0 \text{ for } j > m \right\},$$

where $a_1 > a_2 > \dots > a_m > 0$. The widths can be roughly estimated as $\text{UW}_k(E) \sim \text{AW}_k(E \subset \mathbb{R}^N) \sim a_{k+1}$ for $0 \leq k < m$. On the other hand, the relative filling radius of E in \mathbb{R}^N equals the length of the shortest axis: $\rho([E]; E \subset \mathbb{R}^N) = a_m$. Here, the dimensions m, N and the sequence $\{a_i\}$ are arbitrary. Theorem 3.5 tells us

that $\rho([E]; E \subset \mathbb{R}^N) \leq \text{AW}_{m-2}(E \subset \mathbb{R}^N)$, but we see that the discrepancy between $\rho([E]; E \subset \mathbb{R}^N) = a_m$ and $\text{AW}_{m-2}(E \subset \mathbb{R}^N) \sim a_{m-1}$ can be arbitrarily large. In other words, the Alexandrov (and Urysohn) width can highly overestimate the filling radius. However, the estimate of Theorem 3.7 in this example is sharp: $\rho([E]; E \subset \mathbb{R}^N) = \text{KW}_{m-1}(E \subset \mathbb{R}^N) = a_m$.

For more information on widths, we refer the reader to Balitskiy [Bal21].

3.2. Bounds for Čech lifespans via treewidth. Recall (from formula (\star) in Section 2.3) that $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$ -lifespans can be written in terms of relative filling radii as follows:

$$d_\omega - b_\omega = \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}),$$

where ω_r maps to ω by the homology map induced by the inclusion (see Section 2.3). To estimate lifespans, we need upper bounds for filling radii, such as in Theorem 3.5 and Theorem 3.7. Both estimates can be simultaneously strengthened using the following new width invariant. Recall that throughout the paper, we use homology with coefficients in a fixed field, which is omitted from the notation.

Definition 3.9 (Treewidth). Let \mathcal{X} be a compact subset of a Banach space $(\mathbf{Y}, \|\cdot\|)$. For an integer $k \geq 0$, let us define the k^{th} *treewidth* of \mathcal{X} as

$$\text{TW}_k(\mathcal{X} \subset \mathbf{Y}) := \inf_f \sup_{x \in \mathcal{X}} \|x - f(x)\|,$$

where $f : \mathcal{X} \rightarrow \mathbf{Y}$ is any continuous map whose image $f(\mathcal{X})$ is a finite simplicial complex of dimension at most k with trivial k^{th} reduced homology.

The treewidth satisfies evident properties:

$$\text{rad}(\mathcal{X} \subset \mathbf{Y}) = \text{TW}_0(\mathcal{X} \subset \mathbf{Y}) \geq \text{TW}_1(\mathcal{X} \subset \mathbf{Y}) \geq \text{TW}_2(\mathcal{X} \subset \mathbf{Y}) \geq \dots$$

An equivalent definition of the k^{th} treewidth (perhaps, better motivating our choice of the word *treewidth*) is obtained if one only considers maps f such that the image $f(\mathcal{X})$ is a subcomplex of a finite k -dimensional *contractible* complex inside \mathbf{Y} . Indeed, to any k -complex with trivial $\tilde{H}_k(\cdot)$ one can glue several cells of dimension $\leq k$ to kill all its homology in lower degrees as well as its fundamental group. In the opposite direction, any subset of a k -dimensional contractible complex has trivial $\tilde{H}_k(\cdot)$.

The motivation behind the definition of treewidth is to give a common generalization of Theorem 3.5 and Theorem 3.7. Both follow from the following result combined with the evident bounds $\text{TW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{AW}_{k-1}(\mathcal{X} \subset \mathbf{Y})$ and $\text{TW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{KW}_k(\mathcal{X} \subset \mathbf{Y})$.

Theorem 3.10. *For any compact set \mathcal{X} sitting in a Banach space \mathbf{Y} , any integer $k \geq 0$, and any homology class $\omega \in \tilde{H}_k(\mathcal{X})$,*

$$\rho(\omega; \mathcal{X} \subset \mathbf{Y}) \leq \text{TW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Proof. Fix any number $\delta > \text{TW}_k(\mathcal{X} \subset \mathbf{Y})$, and pick a witness map $f : \mathcal{X} \rightarrow \mathbf{Y}$, whose image $f(\mathcal{X}) = \mathcal{T}$ has dimension at most k and trivial $\tilde{H}_k(\mathcal{T})$, such that $\|x - f(x)\| \leq \delta$ for all $x \in \mathcal{X}$. Pick a cycle $S \subset \mathcal{X}$ representing ω and continuously deform it via the linear homotopy $h : S \times [0, 1] \rightarrow \mathbf{Y}$ given by

$$h(x, t) = (1 - t)x + tf(x).$$

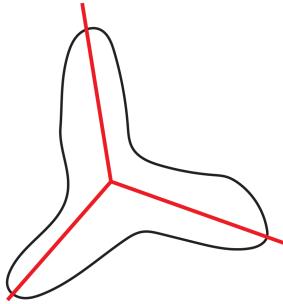


FIGURE 1. A shape with small 1-treewidth but considerable Alexandrov 0-width and Kolmogorov 1-width

Every point $x \in S$ moves by distance at most δ , so the continuous deformation always stays in $\mathcal{N}_\delta(S \subset \mathbf{Y})$. At time $t = 1$, the deformed cycle $h(S, 1)$ lies in \mathcal{T} and bounds a $(k + 1)$ chain within $h(S, 1)$, because \mathcal{T} is of dimension k with trivial $\tilde{H}_k(\mathcal{T})$. Therefore, S is null-homologous in $\mathcal{N}_\delta(S \subset \mathbf{Y})$ (basically, S bounds the $(k + 1)$ -dimensional trace of the homotopy h). \square

A simple example in which Theorem 3.10 gives a stronger bound than Theorem 3.5 and Theorem 3.7 is depicted in Figure 1; the fundamental class of this circle has a small filling radius, which can be efficiently estimated by the 1-treewidth (approximating the shape by a tree, red in the figure), while approximations by straight lines or points give significantly worse bounds.

We proceed with a strengthening applicable to noncompact sets of the form $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$. Notice that the width in the right-hand side of the following estimate is computed on a compact set, since in Banach spaces the closure of the convex hull of a compact set is compact; see e.g. Lax [Lax02, Section 13, Exercise 9].

Theorem 3.11. *For any compact set \mathcal{X} sitting in a Banach space \mathbf{Y} , any integer $k \geq 0$, any positive number r , and any homology class $\omega \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$,*

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \leq \text{TW}_k(\overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

Proof. Let our homology class $\omega \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$ be represented by a singular cycle S_0 supported in $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$. Since S_0 is a compact subset of \mathbf{Y} , it can be covered by finitely many open balls $B_r(x), x \in \mathcal{X}' \subset \mathcal{X}, |\mathcal{X}'| < \infty$ (recall that $B_r(x)$ denotes the open ball of radius r centered at x). Consider the union $U = \bigcup_{x \in \mathcal{X}'} B_r(x)$

and choose any partition of unity subordinate to the cover of U by $B_r(x), x \in \mathcal{X}'$. This partition of unity gives rise to a map from U to the nerve L of that cover. By the nerve theorem, this map $\phi : U \rightarrow L$ is a homotopy equivalence, with an evident homotopy inverse $\psi : L \rightarrow U$, which sends every vertex of L to the center of the corresponding ball in the cover $U = \bigcup_{x \in \mathcal{X}'} B_r(x)$, and extends affinely on the rest of

L . Therefore, $\omega|_U = \psi_* \circ \phi_*(\omega|_U) \in \tilde{H}_k(U)$, and the class ω can be represented by a cycle S supported in $\psi(L) \subset \overline{\text{conv}}(\mathcal{X})$. A similar argument shows that if S bounds in $\mathcal{N}_R(\mathcal{X} \subset \mathbf{Y})$, for some $R > r$, then the filling chain can also be taken with the support in $\overline{\text{conv}}(\mathcal{X})$. Therefore,

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \leq \rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

Now we are in position to apply Theorem 3.10 to the compact set $\overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X}))$:

$$\begin{aligned} \rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}) &= \rho(\omega; \overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}) \\ &\leq \text{TW}_k(\overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}). \end{aligned}$$

□

Corollary 3.12 (Čech lifespans via treewidth). *Let \mathcal{X} be compact subset of a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 0$. Then,*

$$d_\omega - b_\omega \leq \text{TW}_k(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

In particular,

$$d_\omega - b_\omega \leq \text{TW}_k(\overline{\text{conv}}(\mathcal{X}) \subset \mathbf{Y}).$$

Remark 3.13. Recall the monotonicity of the treewidth: $\text{TW}_k(\cdot) \leq \text{TW}_0(\cdot) = \text{rad}(\cdot)$. Therefore, as a trivial consequence of Corollary 3.12, we can upper-bound lifespans in all dimensions by $\text{TW}_0(\overline{\text{conv}}(\mathcal{X}) \subset \mathbf{Y}) = \text{rad}(\mathcal{X} \subset \mathbf{Y})$.

Proof of Corollary 3.12. We use the aforementioned formula

$$d_\omega - b_\omega = \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}),$$

and bound each filling radius by the treewidth of $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$:

$$\rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \leq \text{TW}_k(\overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

To conclude, we need a continuity property:

$$\lim_{r \rightarrow b_\omega + 0} \text{TW}_k(\overline{\mathcal{N}}_r(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}) = \text{TW}_k(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

Its proof is explained in the appendix; see Lemma A.3. □

Combining this with the trivial inequality $\text{TW}_k(\cdot) \leq \text{AW}_{k-1}(\cdot)$, we obtain the following bound.

Corollary 3.14 (Čech lifespans via Alexandrov width). *Let \mathcal{X} be compact subset of a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 1$. Then,*

$$d_\omega - b_\omega \leq \text{AW}_{k-1}(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset \overline{\text{conv}}(\mathcal{X})) \subset \mathbf{Y}).$$

In particular,

$$d_\omega - b_\omega \leq \text{AW}_{k-1}(\overline{\text{conv}}(\mathcal{X}) \subset \mathbf{Y}).$$

We would like to point out that while k -widths are effective for bounding lifespans of homology classes of degree $> k$, they do not say much about the lifespans of homology classes in lower degrees. In Figure 2, we give a simple 2-dimensional example S . Here, $\text{AW}_1(S \subset \mathbb{R}^3)$ is small, and hence, the homology class in $\tilde{H}_2(S)$ has short lifespans by Corollary 3.14. However, one can easily see that homology classes $\alpha \in \tilde{H}_1(S)$ can still have very large lifespans. One can generalize this example to any dimension and codimension. In Section 5 we take a different perspective, and give global bounds for lifespans and death times in all degrees simultaneously.

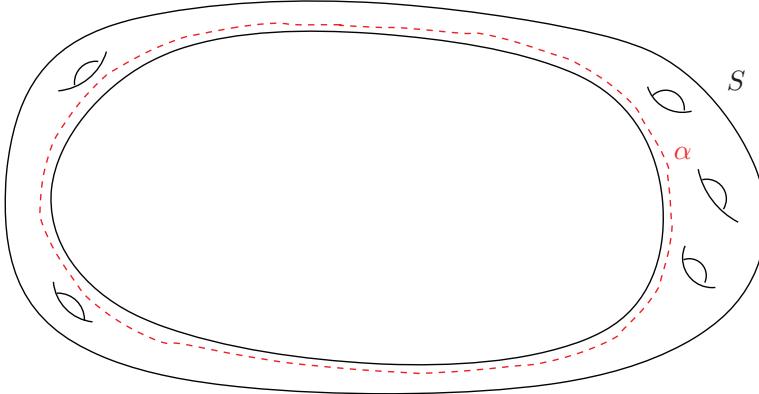


FIGURE 2. For the surface S , while $\text{AW}_1(S \subset \mathbb{R}^3)$ is small, the lifespan of red curve $\alpha \in \tilde{H}_1(S)$ is large

3.3. Bounds for Vietoris–Rips lifespans via Urysohn width. Recall the relationship between $\text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$ -lifespans and filling radii implied by formula (\star) of Section 2.3:

$$d_\omega - b_\omega = 2 \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})),$$

where ω_r maps to ω by the homology level map induced by the inclusion, and $\kappa : \mathcal{X} \rightarrow L^\infty(\mathcal{X})$ is the Kuratowski embedding. Recall also (Remark 2.38) that the right-hand side can be rewritten a bit shorter:

$$\rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})) = \rho(\omega_r; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))).$$

This follows from Lemma 3.15.

Lemma 3.15. *Let \mathcal{Z} be a subset of $L^\infty(\mathcal{X})$ containing the Kuratowski image $\kappa(\mathcal{X})$. Then, for any homology class $\omega \in \tilde{H}_k(\mathcal{Z})$, the following is true:*

$$\rho(\omega; \mathcal{Z} \subset L^\infty(\mathcal{X})) = \rho(\omega; \mathcal{Z}).$$

Proof. The inequality “ \geq ” follows from the universal property of the absolute filling radius (see the comment after Definition 2.33). To show the inequality “ \leq ”, we first recall that

$$\rho(\omega; \mathcal{Z}) = \rho(\kappa'_*(\omega); \kappa'(\mathcal{Z}) \subset L^\infty(\mathcal{Z})),$$

where $\kappa' : \mathcal{Z} \rightarrow L^\infty(\mathcal{Z})$ is the Kuratowski map. Second, we make use of the injectivity (Definition 2.26) of $L^\infty(\mathcal{X})$ in the following way. The inclusion $\mathcal{Z} \subset L^\infty(\mathcal{X})$, via the Kuratowski identification $\mathcal{Z} \xrightarrow{\sim} \kappa'(\mathcal{Z})$, gives rise to a distance preserving (hence, 1-Lipschitz) map $\kappa'(\mathcal{Z}) \rightarrow L^\infty(\mathcal{X})$. This map can be extended to a 1-Lipschitz map $p : L^\infty(\mathcal{Z}) \rightarrow L^\infty(\mathcal{X})$, by invoking the injectivity of $L^\infty(\mathcal{X})$. This map pushes forward the cycle $\kappa'_*(\omega)$ to ω . Moreover, it pushes forward any chain filling $\kappa'_*(\omega)$ inside $\mathcal{N}_r(\kappa'(\mathcal{Z})) \subset L^\infty(\mathcal{Z})$, to a chain filling ω inside $\mathcal{N}_r(\mathcal{Z} \subset L^\infty(\mathcal{X}))$. \square

To estimate VR-lifespans, we need to know how to upper-bound quantities like $\rho(\cdot; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})))$, for which Theorem 3.3 is not applicable: the set $\mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))$ is not compact (even after taking the closure). One way to deal with noncompactness is to intersect $\mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))$ with the closure of

the convex hull of \mathcal{X} , like we did in Theorem 3.11. That theorem, together with the inequality $\text{TW}_k(\cdot) \leq \text{AW}_{k-1}(\cdot)$, will give us a bound on $\rho(\cdot; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})))$ in terms of the Alexandrov width of $\overline{\mathcal{N}}_r(\kappa(\mathcal{X}) \subset \overline{\text{conv}}(\kappa(\mathcal{X})))$. It turns out the latter can be interpreted in terms of the Urysohn width.

Remark 3.16 (L^∞ -interpretation of Urysohn width). In the inequality $\rho(\omega; \mathcal{X}) \leq \frac{1}{2} \text{UW}_{k-1}(\mathcal{X})$, the left-hand side is defined via the Kuratowski embedding in $L^\infty(\mathcal{X})$, and it is instructive to interpret the right-hand side as well, in a way compatible with the picture of $\kappa(\mathcal{X})$ sitting in $L^\infty(\mathcal{X})$. This interpretation goes back to Tikhomirov [Tik76]. Some references in English are [Gro83, Appendix 1, Proposition (D)] by Gromov and [Bal21, Theorem 2.1.9 and Definition 2.3.1] by Balitskiy. Here we provide a short summary. The width $\text{UW}_{k-1}(\mathcal{X})$ can be equivalently defined as the infimal number $\delta > 0$ such that there is map $f : \mathcal{X} \rightarrow L^\infty(\mathcal{X})$ whose image is at most $(k-1)$ -dimensional complex in $L^\infty(\mathcal{X})$, and such that $\|\kappa(x) - f(x)\|_\infty \leq \delta/2$ for all $x \in \mathcal{X}$. In a nutshell, Theorem 3.3 says that one can try killing higher homology of \mathcal{X} by deforming $\kappa(\mathcal{X})$ to a low-dimensional complex inside $L^\infty(\mathcal{X})$, and if every point moves by some controlled distance, then we get an estimate for the filling radius. This interpretation can be shortly reformulated as follows:

$$\text{UW}_{k-1}(\mathcal{X}) = 2 \text{AW}_{k-1}(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})).$$

Given that interpretation, the following estimate is immediate:

$$\rho(\omega; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X}))) \leq \frac{1}{2} \text{UW}_{k-1}(\overline{\mathcal{N}}_r(\kappa(\mathcal{X}) \subset \overline{\text{conv}}(\kappa(\mathcal{X})))).$$

We will go further and give an even better estimate for estimate $\rho(\cdot; \mathcal{N}_r(\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})))$ by replacing the convex hull with the tight span (Definition 2.28). The following result subsumes Theorem 3.3 when $r = 0$. The specific realization of $\text{E}(\mathcal{X})$ in $L^\infty(\mathcal{X})$ that was described in Section 2.2 has an isometric copy of \mathcal{X} inside, and the corresponding embedding of \mathcal{X} in $L^\infty(\mathcal{X})$ agrees with the Kuratowski embedding.

To simplify notation a bit, in the rest of this section, we omit κ and simply write $\mathcal{X} \subset \text{E}(\mathcal{X}) \subset L^\infty(\mathcal{X})$. Notice that the set $\overline{\mathcal{N}}_r(\mathcal{X} \subset \text{E}(\mathcal{X}))$ is compact, as a closed subset of the compact metric space $\text{E}(\mathcal{X})$.

Theorem 3.17. *For any compact metric space \mathcal{X} , any integer $k \geq 1$, any positive number r , and any homology class $\omega \in H_k(\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})))$,*

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))) \leq \frac{1}{2} \text{UW}_{k-1}(\overline{\mathcal{N}}_r(\mathcal{X} \subset \text{E}(\mathcal{X}))).$$

Proof. We have a chain of distance-preserving embeddings $\mathcal{X} \subset \text{E}(\mathcal{X}) \subset L^\infty(\mathcal{X})$ as well as $\mathcal{X} \subset \mathcal{N}_r(\mathcal{X} \subset \text{E}(\mathcal{X})) \subset \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$. All distances below are measured using the L^∞ -norm. We need to bound $\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})))$.

Recall the following property of the tight span (Remark 2.30): there exists a 1-Lipschitz retraction $\pi : L^\infty(\mathcal{X}) \rightarrow \text{E}(\mathcal{X})$. Our first observation is that π maps $\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$ to $\mathcal{N}_r(\mathcal{X} \subset \text{E}(\mathcal{X}))$. Indeed, take any point $y \in \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$. Since $\|y - x\| < r$ for some $x \in \mathcal{X}$, and π is 1-Lipschitz retraction, it follows that $\|\pi(y) - x\| = \|\pi(y) - \pi(x)\| \leq \|y - x\| < r$. Therefore, $\pi(\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))) = \mathcal{N}_r(\mathcal{X} \subset \text{E}(\mathcal{X}))$. Denote the reverse inclusion

$$\iota : \mathcal{N}_r(\mathcal{X} \subset \text{E}(\mathcal{X})) \hookrightarrow \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})).$$

Since both y and $\pi(y)$ lie in the open radius r ball centered at x , the whole straight line segment between them lies in that ball and in $\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$. Therefore, a cycle representing ω can be continuously deformed to a cycle representing $\iota_* \circ \pi_*(\omega)$ by letting each y in the support of ω slide along the straight line segment towards $\pi(y)$. This homotopy takes place entirely in $\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$, and therefore, $\omega = \iota_* \circ \pi_*(\omega)$.

It should be obvious now that

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})) \leq \rho(\pi_*(\omega); \mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X})) \subset L^\infty(\mathcal{X})).$$

Indeed, any cycle S representing $\pi_*(\omega) \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X})))$ lies entirely in $\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$ and there represents the homology class $\iota_* \circ \pi_*(\omega) = \omega \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})))$. If S bounds in the neighborhood of $\mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X}))$ of certain radius, then it bounds even in a smaller neighborhood of $\mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X}))$.

Using the observation at the beginning of this subsection, we deduce that

$$\rho(\pi_*(\omega); \mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X})) \subset L^\infty(\mathcal{X})) = \rho(\pi_*(\omega); \mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X}))).$$

The next step follows trivially from the definition of the filling radius:

$$\rho(\pi_*(\omega); \mathcal{N}_r(\mathcal{X} \subset E(\mathcal{X}))) = \rho(\pi_*(\omega); \overline{\mathcal{N}}_r(\mathcal{X} \subset E(\mathcal{X}))).$$

The final step is to apply Theorem 3.3 to the compact set $\overline{\mathcal{N}}_r(\mathcal{X} \subset E(\mathcal{X}))$, to conclude that

$$\rho(\pi_*(\omega); \overline{\mathcal{N}}_r(\mathcal{X} \subset E(\mathcal{X}))) \leq \frac{1}{2} \text{UW}_{k-1}(\overline{\mathcal{N}}_r(\mathcal{X} \subset E(\mathcal{X}))).$$

Assembling all these inequalities together, we obtain the result. \square

Corollary 3.18 (VR lifespans via Urysohn width). *Let \mathcal{X} be a compact metric space, and let $\omega \in \text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$, $k \geq 1$. Then,*

$$d_\omega - b_\omega \leq \text{UW}_{k-1}(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset E(\mathcal{X}))).$$

In particular,

$$d_\omega - b_\omega \leq \text{UW}_{k-1}(E(\mathcal{X})).$$

Remark 3.19. It follows that lifespans in dimensions 1 and higher can be upper-bounded by $\text{UW}_0(E(\mathcal{X})) = \text{diam}(\mathcal{X})$. It is easy to see that lifespans in dimension 0 are also bounded above by $\text{diam}(\mathcal{X})$.

Proof of Corollary 3.18. Recall that

$$d_\omega - b_\omega = 2 \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset L^\infty(\mathcal{X})) \subset L^\infty(\mathcal{X})),$$

where ω_r maps to ω by the homology map induced by the inclusion. Upper-bounds for these filling radii are given by Theorem 3.17. To conclude, it remains to use the continuity property of the Urysohn width [Bal21, Theorem 2.4.1]:

$$\lim_{r \rightarrow b_\omega + 0} \text{UW}_{k-1}(\overline{\mathcal{N}}_r(\mathcal{X} \subset E(\mathcal{X}))) = \text{UW}_{k-1}(\overline{\mathcal{N}}_{b_\omega}(\mathcal{X} \subset E(\mathcal{X}))).$$

\square

3.4. Robust bounds for Čech lifespans via treewidth. The estimates provided in the preceding sections have a drawback: there is no a priori method to directly relate the width of $\mathcal{X} \subset \mathbf{Y}$ to the width of its r -neighborhood $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$. This poses a challenge, as the bounds in Corollaries 3.12, 3.14, and 3.18 rely on the width of a specific neighborhood of \mathcal{X} . However, by further refining the class of potential “cores”, we can ensure the derivation of more effective bounds. For example, this is possible if the core is an affine k -dimensional subspace (see Section 3.5).

Here, we introduce a modified version of treewidth that enables bounding lifespans directly in terms of \mathcal{X} itself, rather than its neighborhood. To this end, we introduce an auxiliary definition to help constrain the geometric complexity of potential “cores”.

Definition 3.20. Let $(\mathbf{Y}, \|\cdot\|)$ be a Banach space, and let $\mathcal{T} \subset \mathbf{Y}$ be a simplicial complex. Suppose there is a retraction map $f : \mathbf{Y} \rightarrow \mathcal{T}$. We say that f is *C-robust* if for any finite set $\mathcal{X} \subset \mathbf{Y}$, and any $x \in \text{conv}(\mathcal{X})$,

$$\|f(x) - x\| \leq \max_{x' \in \mathcal{X}} \|f(x') - x'\| + C \cdot \text{rad}(\mathcal{X} \subset \mathbf{Y}).$$

Remark 3.21. Recall that, given $\mathcal{T} \subset \mathbf{Y}$, a retraction $f : \mathbf{Y} \rightarrow \mathcal{T}$ is a continuous map fixing every point of \mathcal{T} . Since \mathbf{Y} is contractible, the existence of a retraction implies that \mathcal{T} is contractible.

Examples.

- (1) If \mathcal{T} is a closed affine subspace of \mathbf{Y} , and f is any linear projection on \mathcal{T} , then f is 0-robust.
- (2) If \mathcal{T} is nonempty closed convex subset of \mathbf{Y} , and f is a continuous nearest-point projection (metric projection) on \mathcal{T} , then f is 0-robust.¹⁴
- (3) If the assignment $x \mapsto \|f(x) - x\|$ is *C*-Lipschitz, then f is *C*-robust.
- (4) If $\mathbf{Y} = \mathbb{R}^2$ is the Euclidean plane, and \mathcal{T} is the union of $m \geq 3$ equiangular rays (or intervals) with a common endpoint, then there is a natural retraction to \mathcal{T} which is $\frac{1}{\tan(\pi/m)}$ -robust.

Definition 3.22 (Robust treewidth). Let \mathcal{X} be a subset of a Banach space $(\mathbf{Y}, \|\cdot\|)$. For an integer $k \geq 0$, and a real number $C \geq 0$, we say that the *C-robust k-dimensional treewidth* of \mathcal{X} is at most δ , if there is a *C*-robust retraction $f : \mathbf{Y} \rightarrow \mathcal{T}$ to a finite simplicial complex $\mathcal{T} \subset \mathbf{Y}$ of dimension at most k , displacing every point of \mathcal{X} by distance at most δ :

$$\|f(x) - x\| \leq \delta \quad \forall x \in \mathcal{X}.$$

The infimum of the numbers δ satisfying this condition will be denoted by $\text{TW}_k^C(\mathcal{X} \subset \mathbf{Y})$.

It is immediate from the definition that $\text{TW}_k^{C_1}(\mathcal{X} \subset \mathbf{Y}) \geq \text{TW}_k^{C_2}(\mathcal{X} \subset \mathbf{Y})$ for any $C_1 \leq C_2$, and that $\text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}) \geq \text{TW}_k(\mathcal{X} \subset \mathbf{Y})$ for any C .

Theorem 3.23. For any compact set \mathcal{X} sitting in a Banach space \mathbf{Y} , any integer $k \geq 0$, any real numbers $C \geq 0$ and $r > 0$, and any homology class $\omega \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$,

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})) \leq \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}) + Cr.$$

¹⁴It is known that a nearest-point projection exists, if \mathbf{Y} is reflexive, but it does not have to be continuous in general; see for example the discussion after [Meg12, Corollary 5.1.19].

Proof. Suppose we are given a C -robust retraction $f : \mathbf{Y} \rightarrow \mathcal{T}$ to a k -dimensional simplicial complex $\mathcal{T} \subset \mathbf{Y}$, such that $\|x - f(x)\| \leq \delta$ for all $x \in \mathcal{X}$.

Let our homology class $\omega \in \tilde{H}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$ be represented by a singular cycle S_0 supported in $\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y})$. Since S_0 is compact as a subset of \mathbf{Y} , it can be covered by finitely many open balls $B_r(x), x \in \mathcal{X}' \subset \mathcal{X}, |\mathcal{X}'| < \infty$ (here $B_r(x)$ denotes the open ball of radius r centered at x). Consider the union $U = \bigcup_{x \in \mathcal{X}'} B_r(x)$ and choose any partition of unity subordinate to the cover of U by $B_r(x), x \in \mathcal{X}'$. This partition of unity induces a map from U to the nerve L of that cover. By the nerve theorem this map $\phi : U \rightarrow L$ is a homotopy equivalence, with an evident homotopy inverse $\psi : L \rightarrow U$, which sends every vertex of L to the center of the corresponding ball in the cover $U = \bigcup_{x \in \mathcal{X}'} B_r(x)$, and extends affinely on the rest of L . Therefore, $\omega|_U = \psi_* \circ \phi_*(\omega|_U) \in \tilde{H}_k(U)$, and the class ω can be represented by a cycle S supported in $\psi(L)$; in other words, S is obtained by gluing flat k -simplices with vertices in \mathcal{X}' , and with the circumradius of each simplex at most r .

Now we continuously deform S using the linear homotopy $h : S \times [0, 1] \rightarrow \mathbf{Y}$ given by

$$h(x, t) = (1 - t)x + t f(x).$$

By the robustness property of f , every point $x \in S$ is moved by a distance at most $\delta + Cr$, so the continuous deformation always stays in $\mathcal{N}_{\delta+Cr}(S \subset \mathbf{Y})$. At time $t = 1$, the deformed cycle $h(S, 1)$ lies in \mathcal{T} and bounds a $(k+1)$ chain within $h(S, 1)$, because \mathcal{T} is of dimension k with trivial $\tilde{H}_k(\mathcal{T})$. Therefore, S is nullhomologous in $\mathcal{N}_{\delta+Cr}(S \subset \mathbf{Y})$ (basically, S bounds the $(k+1)$ -dimensional trace of the homotopy h). \square

Corollary 3.24 (Čech lifespans via treewidth). *Let \mathcal{X} be compact subset in a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, with birth time $b_\omega \geq 1$. Then,*

$$\frac{d_\omega}{b_\omega} \leq C + 1 + \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}).$$

Proof. We again employ formula (\star) , now bounding filling radii by Theorem 3.23:

$$\begin{aligned} d_\omega - b_\omega &= \lim_{r \rightarrow b_\omega + 0} \rho(\omega_r; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \\ &\leq \lim_{r \rightarrow b_\omega + 0} \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}) + Cr \\ &= \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y}) + Cb_\omega. \end{aligned}$$

For $b_\omega \geq 1$, if we divide the entire inequality by b_ω , we obtain the result. \square

Remark 3.25 (Comments about Corollary 3.24). We make the following remarks.

- Notice that on the right-hand side, there is no reference to the birth time b_ω . Therefore, C -robust treewidth yields strong control on *late-born* homology classes (i.e. those satisfying $b_\omega \geq 1$).
- If the assumption $b_\omega \geq 1$ is changed to $b_\omega \geq \frac{1}{\alpha}$ for some $\alpha > 0$ then the conclusion becomes $\frac{d_\omega}{b_\omega} \leq 1 + C + \alpha \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y})$.
- The corollary implies a robust estimate on the *multiplicative persistence* of intervals in the persistence barcode of the Čech filtration $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y})$. See Bobrowski, Kahle and Skraba [BKS17] and Adams and Coskunuzer [AC22] for studies of multiplicative persistence.

- Additionally, the proof above gives the following estimates for (b_ω, d_ω) :
 - If $b_\omega = 0$, then $d_\omega \leq \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y})$.
 - If $b_\omega \geq \text{TW}_k^C(\mathcal{X} \subset \mathbf{Y})$, then $\frac{d_\omega}{b_\omega} \leq C + 2$.

3.5. Robust bounds for Čech lifespans via Kolmogorov width and PCA_∞ . For a given *finite* set \mathcal{X} in \mathbb{R}^N , standard Principal Component Analysis or PCA (see [Has09, §14.5]) approximates \mathcal{X} by affine k -dimensional subspaces through minimizing over all such subspaces \mathcal{P} the sum of squared distances of points in \mathcal{X} to \mathcal{P} . The infimum obtained through this procedure is called the $(k+1)^{\text{th}}$ -variance of \mathcal{X} .

Since its introduction by Pearson in 1901 [Pea01], PCA has become one of the most widely used data analysis techniques, owing to its interpretability and the availability of highly efficient computational algorithms. In this work, we investigate a specific ℓ^∞ -variant of classical PCA, designed for compact subsets of Banach spaces, and examine its connections to persistent homology by analyzing how it constrains the lifespans of homology classes.

Definition 3.26 (PCA_∞). Let \mathcal{X} be a compact subset in a Banach space $(\mathbf{Y}, \|\cdot\|)$. For $0 \leq k < \dim(\mathbf{Y})$, let \mathcal{A}_k be the space of affine k -dimensional subspaces in \mathbf{Y} . Then, we define the $(k+1)^{\text{th}}$ ℓ^∞ -variance of \mathcal{X} as

$$\nu_{k+1}(\mathcal{X} \subset \mathbf{Y}) := \inf_{\mathcal{P} \in \mathcal{A}_k} \sup_{x \in \mathcal{X}} \mathbf{d}_\mathbf{Y}(x, \mathcal{P}),$$

where $\mathbf{d}_\mathbf{Y}(x, \mathcal{P})$ is the distance from x to the k -subspace \mathcal{P} , i.e., $\mathbf{d}_\mathbf{Y}(x, \mathcal{P}) := \inf_{p \in \mathcal{P}} \|x - p\|$.

Remark 3.27. The term ℓ^∞ -variance is just a different name for Kolmogorov widths (cf. Definition 3.6):

$$\nu_{k+1}(\mathcal{X} \subset \mathbf{Y}) = \text{KW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Theorem 3.28. Let \mathcal{X} be a subset of a Banach space \mathbf{Y} . Then, for any $r > 0$ and any homology class $\omega \in \widetilde{\text{H}}_k(\mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 0$,

$$\rho(\omega; \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) \subset \mathbf{Y}) \leq \nu_{k+1}(\mathcal{X} \subset \mathbf{Y}) = \text{KW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Proof. Apply Theorem 3.23 with \mathcal{T} being a k -dimensional affine subspace, and f being a 0-robust nearest-point projection. \square

Čech Lifespans of homology classes in degree $\geq k$ cannot exceed the $(k+1)^{\text{th}}$ variance, $\nu_{k+1}(\mathcal{X})$.

Corollary 3.29 (Čech lifespans via ℓ^∞ -variance). Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} and let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 0$. Then,

$$d_\omega - b_\omega \leq \nu_{k+1}(\mathcal{X} \subset \mathbf{Y}) = \text{KW}_k(\mathcal{X} \subset \mathbf{Y}).$$

Proof. Combine formula (\star) with Theorem 3.28. \square

Remark 3.30. Notice that Corollary 3.29 gives strong estimates for point clouds in \mathbb{R}^N . For example, if the variance $\nu_{k_0+1}(\mathcal{X})$ is small, then, in applications where only long-lived homological features are relevant, the corollary implies that one can ignore homology classes in $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$ for $k \geq k_0$. We should note that, by its very definition, PCA_∞ is more sensitive than the original PCA against outliers. We discuss this difference and pose a statistical question about the relationship between the original PCA and persistent homology in Section 6.

Remark 3.31. Notice that all lifespan bounds in Corollaries 3.12, 3.14, 3.18, 3.24, and 3.29 are stated in terms of homological lifespans ($d_\omega - b_\omega$). Remark 2.23 implies that all these bounds therefore apply to the length of every interval in either $\text{PD}_k(\mathcal{X})$ or $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$.

4. BOUNDING LIFESPANS BY SPREAD

This section is concerned with global upper bounds for Čech lifespans of homology classes in arbitrary degrees dimension. In terms of the conceptual description laid out on page 8441, to formulate these results we will choose cores to be übercontractible spaces, as defined below.

Our results in this section are somewhat analogous to the following result concerning VR lifespans.

Theorem 4.1 (VR-lifespans via spread; [LMO24, Proposition 9.19 and Remark 9.18]). *Let \mathcal{X} be a compact metric space and let $\omega \in \text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$. Then $d_\omega - b_\omega \leq \text{spread}(\mathcal{X})$.*

Here Katz's notion of *spread* is used [Kat83]; see also [Wil92, Definition 4]. By definition, $\text{spread}(\mathcal{X})$ is the infimum of the real numbers $\delta \geq 0$ for which there is a finite subset $A = \{a_1, \dots, a_m\} \subset \mathcal{X}$ with $\text{diam}(A) \leq \delta$ and $\mathbf{d}_H^{\mathcal{X}}(\mathcal{X}, A) \leq \delta$.

In the rest of this section, we work in a Banach space \mathbf{Y} and, for each compact set $\mathcal{X} \subset \mathbf{Y}$, we consider a certain variant of the notion spread for which we prove an analogue of Theorem 4.1.

Definition 4.2 (Übercontractibility). A set \mathcal{T} , sitting in a Banach space \mathbf{Y} , will be said to be δ -übercontractible if its neighborhood $\mathcal{N}_r(\mathcal{T} \subset \mathbf{Y})$ is contractible for every $r \geq \delta$.

Remark 4.3 (Acyclicity). For our purposes a slightly weaker condition would suffice: all neighborhoods $\mathcal{N}_r(\mathcal{T} \subset \mathbf{Y})$, $r \geq \delta$, are acyclic (have trivial homology groups) rather than contractible (have trivial homotopy groups). The acyclicity condition is equivalent to having almost trivial Čech persistence diagram (trivial beyond death time δ), and is implied by übercontractibility as defined above.

Definition 4.4 (Überspread). For a compact subset \mathcal{X} of a Banach space \mathbf{Y} , we define its überspread as

$$\ddot{\text{u}}\text{-spread}(\mathcal{X} \subset \mathbf{Y}) := \inf \left\{ \delta \geq 0 \mid \begin{array}{l} \exists \mathcal{T} \subset \mathbf{Y} \text{ a } \delta\text{-übercontractible simplicial complex} \\ \text{such that } \mathbf{d}_H^{\mathbf{Y}}(\mathcal{X}, \mathcal{T}) \leq \delta \end{array} \right\}.$$

Note that for any point $p \in \mathbf{Y}$ the singleton set $\{p\}$ is δ -übercontractible for all $\delta \geq 0$. In particular, this implies that

$$\ddot{\text{u}}\text{-spread}(\mathcal{X} \subset \mathbf{Y}) \leq \text{rad}(\mathcal{X} \subset \mathbf{Y}).$$

For additional constructive examples of übercontractible sets, we refer the reader to Appendix B, where we show that the cut-locus of a convex set is typically 0-übercontractible.

Theorem 4.5 (Čech lifespans via überspread). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Let $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$ for any $k \geq 0$. Then,*

$$d_\omega - b_\omega \leq 2 \ddot{\text{u}}\text{-spread}(\mathcal{X} \subset \mathbf{Y}).$$

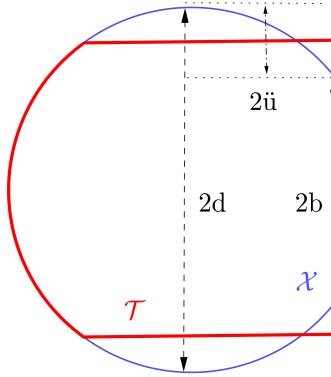


FIGURE 3. The bound $d - b \leq 2\ddot{u}$ in Theorem 4.5 is sharp; here $\mathbf{Y} = \mathbb{R}^2$ and $\ddot{u} = \ddot{u}\text{-spread}(\mathcal{X} \subset \mathbb{R}^2)$

Proof. Via the functorial nerve theorem (Theorem 2.6), we will argue at the level of neighborhood filtrations, as opposed to simplicial filtrations. Let \mathcal{T} be a δ -übercontractible simplicial complex in \mathbf{Y} which is Hausdorff distance at most δ away from \mathcal{X} , where δ is just a tiny bit greater than $\ddot{u}\text{-spread}(\mathcal{X} \subset \mathbf{Y})$. Then, for all $r > 0$ we have

$$\begin{array}{ccc} \mathcal{N}_r(\mathcal{X} \subset \mathbf{Y}) & \xrightarrow{i_{r,r+2\delta}} & \mathcal{N}_{r+2\delta}(\mathcal{X} \subset \mathbf{Y}) \\ & \searrow p_r & \swarrow q_{r+\delta} \\ & \mathcal{N}_{r+\delta}(\mathcal{T} \subset \mathbf{Y}) & \end{array}$$

where $i_{r,r+2\delta}$, p_r , and $q_{r+\delta}$ are the obvious inclusion maps. Applying the homology functor to the above diagram gives, for all $r > 0$, that

$$\tilde{H}_k(i_{r,r+2\delta}) = \tilde{H}_k(p_r) \circ \tilde{H}_k(q_{r+\delta}) = 0$$

since $\mathcal{N}_{r+\delta}(\mathcal{T} \subset \mathbf{Y})$ is contractible. Then this means that $d_\omega - b_\omega \leq 2\delta$. The proof follows. \square

Remark 4.6 (Katz spread vs. \ddot{u} -spread). Let us compare Theorem 4.5 with [LMO24, Proposition 9.19]. Suppose we are given a compact metric space \mathcal{X} and a finite subset $A = \{a_1, \dots, a_m\} \subset \mathcal{X}$ with $\text{diam}(A) \leq \delta$ and $\mathbf{d}_H^{\mathcal{X}}(\mathcal{X}, A) \leq \delta$. In other words, the *Katz spread* of \mathcal{X} does not exceed δ :

$$\text{spread}(\mathcal{X}) \leq \delta.$$

Embed \mathcal{X} in $L^\infty(\mathcal{X})$ via the Kuratowski map κ , and note that the set $\kappa(A) \subset L^\infty(\mathcal{X})$ is $\delta/2$ -übercontractible. Hence, we can apply Theorem 4.5 and conclude that Čech lifespans of $\kappa(\mathcal{X}) \subset L^\infty(\mathcal{X})$ do not exceed 2δ . Therefore, VR lifespans of \mathcal{X} do not exceed 4δ . This is weaker than the conclusion of [LMO24, Proposition 9.19], which tells us that in this situation VR lifespans of \mathcal{X} do not exceed $\text{spread}(\mathcal{X}) \leq \delta$.

In general, the estimate in Theorem 4.5 cannot be improved (see Figure 3). However, the factor of 2 in that estimate can sometimes be removed if the comparison set \mathcal{T} (i.e. the “core”) admits a 1-Lipschitz nearest-point projection from \mathbf{Y} .

Theorem 4.7. *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Assume that $\text{conv}(\mathcal{X})$ admits a 1-Lipschitz nearest-point projection to a closed convex set $\mathcal{T} \subset \mathbf{Y}$.¹⁵ Then any $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, $k \geq 0$, has lifespan at most $\mathbf{d}_H^\mathbf{Y}(\mathcal{X}, \mathcal{T})$.*

Remark 4.8. Theorem 4.7 cannot be derived from the stability results in Lemma 2.15 and Lemma 2.16, as those results apply exclusively to homology classes associated with points in the persistence diagram $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$. In contrast, Theorem 4.7 applies to all elements of $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$, which, as expressed by Remark 2.21 and Proposition 2.22, generally includes strictly more homology classes than those represented by points in the persistence diagram $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$. Nevertheless, the proof of Theorem 4.7 shares ideas with (certain simplicial-level) proofs of the stability results.

Proof of Theorem 4.7. Let $\delta = \mathbf{d}_H^\mathbf{Y}(\mathcal{X}, \mathcal{T})$. By assumption, there exists a 1-Lipschitz nearest-point projection $\pi : \text{conv}(\mathcal{X}) \rightarrow \mathcal{T}$; there is also a (possibly discontinuous) map $j : \mathcal{T} \rightarrow \mathcal{X}$ such that

$$\|x - \pi(x)\| \leq \delta \text{ and } \|a - j(a)\| \leq \delta \quad \forall x \in \mathcal{X} \text{ and } a \in \mathcal{T}.$$

It is easy to see that π and j induce, for each $\epsilon > 0$, simplicial maps

$$\pi_\epsilon : \mathcal{C}_\epsilon(\mathcal{X} \subset \mathbf{Y}) \rightarrow \mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y}) \text{ and } j_\epsilon : \mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y}) \rightarrow \mathcal{C}_{\epsilon+\delta}(\mathcal{X} \subset \mathbf{Y}).$$

To see this in the case of π we proceed as follows. Let $\sigma = \{x_0, \dots, x_n\} \in \mathcal{C}_\epsilon(\mathcal{X} \subset \mathbf{Y})$. Let $y_\sigma \in \text{conv}(\mathcal{X})$ satisfy $\|x_i - y_\sigma\| \leq \epsilon$ for all i . Then, since $\pi : \text{conv}(\mathcal{X}) \rightarrow \mathcal{T}$ is 1-Lipschitz, we have

$$\|\pi(x_i) - \pi(y_\sigma)\| \leq \|x_i - y_\sigma\| \leq \epsilon$$

for all i . Hence, $\pi_\epsilon(\sigma) \in \mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y})$.

In the case of j , let $\sigma = \{x_0, \dots, x_n\} \in \mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y})$. Let $y_\sigma \in \mathbf{Y}$ satisfy $\|x_i - y_\sigma\| \leq \epsilon$ for all i . Then,

$$\|j(x_i) - y_\sigma\| \leq \|j(x_i) - x_i\| + \|x_i - y_\sigma\| \leq \delta + \epsilon$$

for all i . Hence, $j_\epsilon(\sigma) \in \mathcal{C}_{\epsilon+\delta}(\mathcal{X} \subset \mathbf{Y})$.

We then have the following (not-necessarily commutative) diagram:

$$\begin{array}{ccc} \mathcal{C}_\epsilon(\mathcal{X} \subset \mathbf{Y}) & \xrightarrow{i_{\epsilon, \epsilon+\delta}} & \mathcal{C}_{\epsilon+\delta}(\mathcal{X} \subset \mathbf{Y}) \\ \searrow \pi_\epsilon & & \swarrow j_\epsilon \\ & \mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y}) & \end{array}$$

where $i_{\epsilon, \epsilon+\delta}$ is the obvious inclusion map.

Claim. For each $\epsilon > 0$ the maps $j_\epsilon \circ \pi_\epsilon$ and $i_{\epsilon, \epsilon+\delta}$ are contiguous.

To see the claim, let $\sigma = \{x_0, \dots, x_n\}$ be any simplex in $\mathcal{C}_\epsilon(\mathcal{X} \subset \mathbf{Y})$. We will prove that $\tau := \sigma \cup j \circ \pi(\sigma)$ is a simplex in $\mathcal{C}_{\epsilon+\delta}(\mathcal{X} \subset \mathbf{Y})$. This requires us to find a point $y_\tau \in \mathbf{Y}$ such that $\|v - y_\tau\| \leq \epsilon + \delta$ for all $v \in \tau$.

Let $y_\sigma \in \mathbb{R}^d$ be such that $\|x_i - y_\sigma\| \leq \epsilon$ for all i . Notice that then we can write

$$\tau = \{x_0, \dots, x_n\} \cup \{j \circ \pi(x_0), \dots, j \circ \pi(x_n)\}.$$

Then, we let $y_\tau := \pi(y_\sigma)$ and calculate

$$\max_{v \in \tau} \|v - y_\tau\| = \max \left(\max_i \|x_i - \pi(y_\sigma)\|, \max_i \|j \circ \pi(x_i) - \pi(y_\sigma)\| \right).$$

¹⁵For example, this is true whenever \mathbf{Y} is Euclidean.

For the first argument of the maximum above,

$$\begin{aligned}\|x_i - \pi(y_\sigma)\| &\leq \|x_i - \pi(x_i)\| + \|\pi(x_i) - \pi(y_\sigma)\| \\ &\leq \|x_i - \pi(x_i)\| + \|x_i - y_\sigma\| \\ &\leq \delta + \epsilon.\end{aligned}$$

For the second argument,

$$\begin{aligned}\|j \circ \pi(x_i) - \pi(y_\sigma)\| &\leq \|j \circ \pi(x_i) - \pi(x_i)\| + \|\pi(x_i) - \pi(y_\sigma)\| \\ &\leq \|j \circ \pi(x_i) - \pi(x_i)\| + \|x_i - x_\sigma\| \\ &\leq \delta + \epsilon,\end{aligned}$$

which establishes the claim.

Now, going back to the proof of the theorem: apply the homology functor \tilde{H}_k to the diagram above and obtain the following *commutative* diagram:

$$\begin{array}{ccc}\tilde{H}_k(\mathcal{C}_\epsilon(\mathcal{X} \subset \mathbf{Y})) & \xrightarrow{\tilde{H}_k(i_{\epsilon, \epsilon+\delta})} & \tilde{H}_k(\mathcal{C}_{\epsilon+\delta}(\mathcal{X} \subset \mathbf{Y})) \\ & \searrow \tilde{H}_k(\pi_\epsilon) & \nearrow \tilde{H}_k(j_\epsilon) \\ & \tilde{H}_k(\mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y})) & \end{array}$$

Now, $\tilde{H}_k(\mathcal{C}_\epsilon(\mathcal{T} \subset \mathbf{Y})) = 0$ since \mathcal{T} is convex. Then,

$$\tilde{H}_k(i_{\epsilon, \epsilon+\delta}) = \tilde{H}_k(j_\epsilon) \circ \tilde{H}_k(\pi_\epsilon) = 0,$$

which completes the proof. \square

5. BOUNDING EXTINCTION TIMES

In Section 3, we discussed bounds for the lifespans of individual homology classes in degree $\geq k$ via different notions of k -width from metric geometry. However, as Figure 2 suggests, these k -widths do not give any bound for the lower homology classes in lower degrees. In other words, if $\text{AW}_k(\mathcal{M} \subset \mathbb{R}^N)$ (or another k -width) is small, the filling radius (or lifespan) of a homology class $\alpha \in \tilde{H}_j(\mathcal{M})$ for $j < k$ can still be very large. It is easy to see that the example in Figure 2 can be generalized to any dimension and codimension. In particular, k -widths do not say much about the size of topological features in dimensions lower than k .

In this section, we use a different approach and propose a global bound to the lifespans of homology classes in all degrees at once. Our goal is to bound the following quantity.

Definition 5.1 (Extinction time). Let \mathcal{X} be a compact metric space. Then,

$$\xi(\mathcal{X}) := \sup_{k \geq 0} \sup \{d_\omega \mid \omega \in \text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))\}$$

is called the *VR-extinction time* of \mathcal{X} . We similarly, define $\check{\xi}(\mathcal{X} \subset \mathbf{Y})$, the $\check{\text{Cech}}$ -extinction time for the $\check{\text{Cech}}$ filtration of a subset \mathcal{X} of a Banach space \mathbf{Y} . We will use the terms VR-extinction (respectively $\check{\text{Cech}}$ -extinction) for short.

Notice that via the VR-extinction $\xi(\mathcal{X})$, we are not only bounding the lifespans of homological features with birth time = 0, but we are bounding the death time of *all* homological features (in degree 1 and higher) appearing throughout the filtration. Also, note that, as pointed out in [LMO24, Remark 9.8], the radius of \mathcal{X}

(Definition 2.39) automatically gives an upper bound for the VR-extinction time, that is, $\xi(\mathcal{X}) \leq \text{rad}(\mathcal{X})$. This is easy to see since the VR complex $\mathcal{V}_r(\mathcal{X})$ is a simplicial cone for $r \geq \text{rad}(\mathcal{X})$. A similar statement is clearly true for Čech-extinction, i.e. $\check{\xi}(\mathcal{X} \subset \mathbf{Y}) \leq \text{rad}(\mathcal{X} \subset \mathbf{Y})$. See Remark 5.2 for comments on the usefulness of this type of bounds.

Remark 5.2 (Utility of extinction time bounds). Bounds on extinction times can be useful in practice for reducing the computational effort incurred by algorithms designed for calculating PH, as we now explain. By Proposition 2.22, any bound on the extinction times of a compact metric space provides an upper bound on the right endpoints of every interval in the barcode of the VR-filtration for that space; see Remark 2.23. This is particularly useful due to the inherent structure of algorithmic procedures for computing persistence diagrams of simplicial filtrations whose complexity increases with the total number of simplices; see [EH10, Chapter VII]. For instance, both the software packages Ripser [Bau21] and Eirene [HG16] use $\text{rad}(X)$ as a cut-off value for the filtration parameter. Any computationally feasible approximation to the extinction time bounds below could similarly boost efficiency in practical applications.

In the following, we aim to give much finer estimates for extinction times in both the VR and Čech settings.

Remark 5.3 (Motivating example). Here, we give a toy example to motivate the notion of extinction defined above. Let E be the $(N-1)$ -dimensional ellipsoid in \mathbb{R}^N given by

$$E := \left\{ \mathbf{x} \in \mathbb{R}^N \left| \sum_{i=1}^N \frac{x_i^2}{a_i^2} = 1 \right. \right\},$$

where $a_1 > a_2 > \dots > a_N > 0$. While E has trivial homology groups in low degrees, one can easily add some topology to E by adding k -handles for $0 < k < N-1$ as follows. Fix $0 < k < N-1$. Let

$$\bar{E} := \left\{ \mathbf{x} \in \mathbb{R}^N \left| \sum_{i=1}^N \frac{x_i^2}{a_i^2} \leq 1 \right. \right\}$$

be the solid ellipsoid. Given an integer $m \geq 1$, we will use notation $\bar{E}^m := \{\mathbf{x} \in \bar{E} \mid x_i = 0 \text{ for } i > m\}$. Let $\{D_1, D_2, \dots, D_{\ell_k}\}$ be ℓ_k disjoint small disks in \bar{E}^{N-k} . Then, consider the following surgery operation. For $1 \leq j \leq \ell_k$, let $\Omega_j := \bar{E} \cap (D_j \times \mathbb{R}^k)$. Let $S_j := E \cap \Omega_j$, and $T_j := \text{cl}(\partial\Omega_j \setminus S_j)$. Then, swapping the S_j and the T_j will give a new closed manifold

$$\hat{E} := \left(E \setminus \bigcup_{j=1}^{\ell_k} S_j \right) \cup \bigcup_{j=1}^{\ell_k} T_j \subset \mathbb{R}^N.$$

See Figure 4 for an illustration. While $\tilde{H}_k(E)$ is trivial for $0 < k < N-1$, $\text{rank}(\tilde{H}_k(\hat{E})) \geq \ell_k$ because of the homology classes generated by the k -handles $\{T_j\}$. By choosing $\{D_j\}$ such that $\{\Omega_j\}$ are all pairwise disjoint, one can obtain \hat{E} with nontrivial homology in all desired degrees k .

Notice that, through Kolmogorov widths, Corollary 3.29 yields that

$$d_\omega - b_\omega \leq \text{KW}_k(\hat{E} \subset \mathbb{R}^N) = a_{k+1}$$

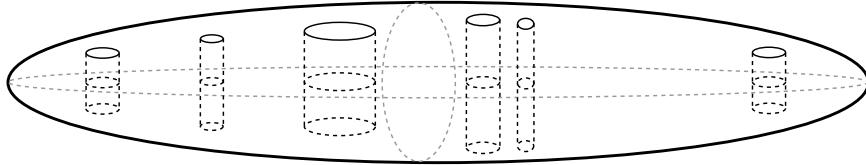


FIGURE 4. The space \hat{E} from Remark 5.3 for the case $N = 3$ and $k = 1$

for any class $\omega \in \text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbb{R}^N))$. By suitably choosing the numbers $\{a_i\}$, all k -widths for $k \leq \dim(\hat{E}) - 1$ can be made arbitrarily large. However, it is not hard to see that, in this particular example, the death times of all these homology classes do not exceed the smallest axis length a_N . This example therefore shows that k -widths can highly overestimate lifespans. In other words, we have that the Čech extinction satisfies $\check{\xi}(\hat{E} \subset \mathbb{R}^N) \leq a_N$. In the following sections, we estimate extinction radii by comparing a given space to a nearby topologically trivial space.

5.1. Bounding Čech-extinction via convex hulls.

Definition 5.4 (Convexity deficiency). For a compact set \mathcal{X} sitting in a Banach space $(\mathbf{Y}, \|\cdot\|)$, we define the *convexity deficiency* of \mathcal{X} as

$$\text{cdef}(\mathcal{X} \subset \mathbf{Y}) := \mathbf{d}_H^{\mathbf{Y}}(\mathcal{X}, \text{conv}(\mathcal{X})) = \sup_{y \in \text{conv}(\mathcal{X})} \inf_{x \in \mathcal{X}} \|y - x\|.$$

Theorem 5.5 (Bounding Čech extinction). *Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Then,*

$$\check{\xi}(\mathcal{X} \subset \mathbf{Y}) \leq \text{cdef}(\mathcal{X} \subset \mathbf{Y}).$$

Proof. For any $r \geq 0$, and any points $x_1, \dots, x_m \in \mathcal{X}$, the intersection of balls $\bigcap_{i=1}^m B_r(x_i)$ is nonempty in \mathbf{Y} if and only if the intersection $\bigcap_{i=1}^m (B_r(x_i) \cap \text{conv}(\mathcal{X}))$ is nonempty. Therefore, the filtered Čech complexes $\mathcal{C}(\mathcal{X} \subset \mathbf{Y})$ and $\mathcal{C}(\mathcal{X} \subset \text{conv}(\mathcal{X}))$ are identical. However, the homotopy type of $\mathcal{C}_r(\mathcal{X} \subset \text{conv}(\mathcal{X}))$ is trivial whenever $r \geq \text{cdef}(\mathcal{X} \subset \mathbf{Y})$, so no topological feature persists beyond time $\text{cdef}(\mathcal{X} \subset \mathbf{Y})$. \square

Remark 5.6 (Stability vs. extinction). The stability theorems (Section 2.1.3) do not imply the previous result for the following reasons:

- (1) In contrast with the stability theorems Lemma 2.15 and Lemma 2.16, which apply *only* to homology classes associated to points in the persistence diagram $\text{PD}_k(\mathcal{X} \subset \mathbf{Y})$, Theorem 5.5 applies to *all elements* of $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$. Recall, from Remark 2.21 and Proposition 2.22, that in general $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$ contains strictly more homology classes than those which can be associated with points in the persistence diagram $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$.
- (2) Strictly speaking, one should not expect that, in general, stability holds for the lifetime or extinction of arbitrary classes in $\text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$ or $\text{Spec}_k(\mathcal{C}_\bullet(\mathcal{X} \subset \mathbf{Y}))$. For example, the filling radius of an m -dimensional manifold \mathcal{M} coincides with both the lifetime and extinction of its fundamental class $[\mathcal{M}] \in \text{Spec}_m(\mathcal{V}_\bullet(\mathcal{M}))$. However, the filling radius of \mathcal{M} is *not* stable under

the Gromov–Hausdorff distance as explained in [LMO24, Sections 9.4 and 9.5].

(3) Even if we restricted ourselves to those homology classes associated with points in the persistence diagram, the stability theorems would only yield bounds on lifespans but not on death times (i.e. on extinction). Indeed, Lemma 2.16 implies that

$$\mathbf{d}_b(\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y}), \widetilde{\text{PD}}_k(\text{conv}(\mathcal{X}) \subset \mathbf{Y})) \leq \mathbf{d}_{\text{H}}^{\mathbf{Y}}(\mathcal{X}, \text{conv}(\mathcal{X})) = \text{cdef}(\mathcal{X} \subset \mathbf{Y}).$$

Since $\widetilde{\text{PD}}_k(\text{conv}(\mathcal{X}) \subset \mathbf{Y}) = \emptyset$, the stability result and the definition of the bottleneck distance [EH10] imply that for any point $(b, d) \in \widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$,

$$d - b \leq \text{cdef}(\mathcal{X} \subset \mathbf{Y}).$$

However, both b and d can be arbitrarily large without violating this inequality. Meanwhile, Theorem 5.5 bounds from above the second coordinate of *every* element of $\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$: for any $(b, d) \in \widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$,

$$d \leq \text{cdef}(\mathcal{X} \subset \mathbf{Y}).$$

Remark 5.7. Example in Figure 3 shows that it is not true that the extinction $\xi(\mathcal{X} \subset \mathbf{Y})$ is bounded from above by the überspread ü-spread($\mathcal{X} \subset \mathbf{Y}$).

While Remark 5.7 shows that the distance to the nearest übercontractible space (überspread) fails to bound extinction, it might be still true that distance to a specific cleverly chosen übercontractible space might give an estimate on extinction. One natural choice for the role of an übercontractible space \mathcal{T} approximating $\mathcal{X} \subset \mathbf{Y}$ is the convex hull of \mathcal{X} ; this is exactly what has just been discussed; with this choice, there is an extinction bound (Theorem 5.5), but it is tempting to improve it by choosing a finer \mathcal{T} . Another natural choice of \mathcal{T} is given by the cut-locus of the boundary of the convex hull of \mathcal{X} (see the discussion in Appendix B). Unfortunately, the extinction time is not bounded from above by the distance from \mathcal{X} to this \mathcal{T} (chosen as the cut-locus of $\partial \text{conv}(\mathcal{X})$, assuming $\text{conv}(\mathcal{X})$ full-dimensional). We omit the discussion of examples in view of the negative nature of the result, but they can be obtained as subsets of a square (one can take a square and cut out a large off-centered disk).

Question 5.8 (Čech extinction and cut-locus of $\text{conv}(\mathcal{X})$). Let \mathcal{X} be a compact subset of a Banach space \mathbf{Y} . Let $\mathcal{C}^{\partial \text{conv}(\mathcal{X})}$ be the cut-locus of the convex hull of \mathcal{X} (see Appendix B). Is it possible to upper-bound the extinction $\check{\xi}(\mathcal{X} \subset \mathbf{Y})$ in terms of the Hausdorff distance to $\mathcal{C}^{\partial \text{conv}(\mathcal{X})}$? That is, is it true that

$$\check{\xi}(\mathcal{X} \subset \mathbf{Y}) \leq C_{\mathbf{Y}} \cdot \mathbf{d}_{\text{H}}^{\mathbf{Y}}(\mathcal{X}, \mathcal{C}^{\partial \text{conv}(\mathcal{X})})$$

for some $C_{\mathbf{Y}} > 0$?

5.2. Bounding VR-extinction via the tight span. To derive a bound on VR-extinction times, Theorem 5.5 can be applied to the Kuratowski (distance-preserving) embedding of \mathcal{X} into $L^\infty(\mathcal{X})$. However, this approach yields a suboptimal result. Here, we present a refinement.

Definition 5.9 is analogous to Definition 5.4 in that the convex hull of $\mathcal{X} \subset \mathbf{Y}$ is supplanted by $\text{E}(\mathcal{X})$ (see Definition 2.28). Recall that \mathcal{X} naturally embeds in its tight span, so we can assume that $\mathcal{X} \subset \text{E}(\mathcal{X})$.

Definition 5.9 (Hyperconvexity deficiency). The *hyperconvexity deficiency* of a compact metric space \mathcal{X} is defined as the number

$$\text{hcdef}(\mathcal{X}) := \mathbf{d}_H^{\text{E}(\mathcal{X})}(\mathcal{X}, \text{E}(\mathcal{X})) = \sup_{f \in \text{E}(\mathcal{X})} \inf_{x \in \mathcal{X}} \|f - \mathbf{d}_{\mathcal{X}}(x, \cdot)\|_{\infty}.$$

Corollary 5.10 to Corollary 2.7 is analogous to Theorem 5.5.

Corollary 5.10 (Bounding VR extinction). *Let \mathcal{X} be a compact metric space. Then,*

$$\xi(\mathcal{X}) \leq 2 \text{hcdef}(\mathcal{X}).$$

Furthermore, this bound is tight (see Remark 5.11).

Proof. The claim follows from Corollary 2.7 and Remark 2.8 together with the facts that (1) $\mathcal{N}_t(\mathcal{X} \subset \text{E}(\mathcal{X})) = \text{E}(\mathcal{X})$ for all $t \geq \text{hcdef}(\mathcal{X})$ and (2) $\text{E}(\mathcal{X})$ is contractible. \square

Remark 5.11 (Tightness of the bound). Let \mathcal{X} be the unit ℓ^{∞} sphere in \mathbb{R}^2 . Then, in that case, by results of Kılıç and Koçak [KK16], $\text{E}(\mathcal{X})$ is isometric to $([-1, 1] \times [-1, 1], \ell^{\infty})$ and we compute that $\text{hcdef}(\mathcal{X}) = 1$. Since $\mathcal{N}_t(\mathcal{X} \subset \text{E}(\mathcal{X})) \simeq \mathcal{X} \simeq S^1$ for every $0 < t < 1$, and for $t \geq 1$ we have $\mathcal{N}_t(\mathcal{X} \subset \text{E}(\mathcal{X})) = \text{E}(\mathcal{X})$, which is contractible, by Corollary 2.7 we have

- $\mathcal{V}_r(\mathcal{X}) \simeq S^1$ for every $0 < r < 2$, and
- $\mathcal{V}_r(\mathcal{X})$ is contractible for $r \geq 2$.

Hence, for $\omega = [S^1]$, $d_{\omega} = 2$; see also [LMO24, Corollary 7.13].

Recall from [LMO24, Remark 9.8] that $\mathcal{V}_r(\mathcal{X})$ becomes contractible as soon as $r \geq \text{rad}(\mathcal{X})$. Proposition 5.12 proves that the bound in Corollary 5.10 is never worse than this bound.

Proposition 5.12. *The inequality*

$$2 \text{hcdef}(\mathcal{X}) \leq \text{rad}(\mathcal{X})$$

holds for every compact metric space \mathcal{X} .

Remark 5.13 (Comparison of $\text{rad}(\mathcal{X})$ and $\text{hcdef}(\mathcal{X})$). The upper bound given by Corollary 5.10 can be much smaller than $\text{rad}(\mathcal{X})$. Indeed, Let \mathcal{X} be any metric tree, then, in that case (by item (3) of Proposition 2.29) $\text{E}(\mathcal{X}) = \mathcal{X}$ so that $\text{hcdef}(\mathcal{X}) = 0$. However, \mathcal{X} can be chosen so that $\text{rad}(\mathcal{X})$ (and also its spread) are arbitrarily large.

Proof of Proposition 5.12. Assume that $\delta > \text{rad}(\mathcal{X})$ and let $x_0 \in \mathcal{X}$ be a point such that $\mathbf{d}_{\mathcal{X}}(x_0, x) < \delta$ for all $x \in \mathcal{X}$. Pick any $f \in \text{E}(\mathcal{X})$ and recall that, according to equation (5), we have that

$$f(x_0) = \max_{x' \in \mathcal{X}} (\mathbf{d}_{\mathcal{X}}(x_0, x') - f(x')).$$

Let $x'_0 \in \mathcal{X}$ be such that $f(x_0) = \mathbf{d}_{\mathcal{X}}(x_0, x'_0) - f(x'_0)$. Notice that then, by equation (5), we have both

$$\|f - \mathbf{d}_{\mathcal{X}}(x_0, \cdot)\|_{\infty} = f(x_0) \text{ and } \|f - \mathbf{d}_{\mathcal{X}}(x'_0, \cdot)\|_{\infty} = f(x'_0).$$

Adding these two expressions together we obtain that

$$\|f - \mathbf{d}_{\mathcal{X}}(x_0, \cdot)\|_{\infty} + \|f - \mathbf{d}_{\mathcal{X}}(x'_0, \cdot)\|_{\infty} = f(x_0) + f(x'_0) = \mathbf{d}_{\mathcal{X}}(x_0, x'_0) < \delta.$$

From this, we conclude that

$$\inf_{x \in \mathcal{X}} \|f - \mathbf{d}_{\mathcal{X}}(x, \cdot)\|_{\infty} \leq \min(\|f - \mathbf{d}_{\mathcal{X}}(x_0, \cdot)\|_{\infty}, \|f - \mathbf{d}_{\mathcal{X}}(x'_0, \cdot)\|_{\infty}) < \frac{\delta}{2}.$$

Since $f \in E(\mathcal{X})$ was arbitrary, this proves that $\text{hcdef}(\mathcal{X}) < \frac{\delta}{2}$ from which the claim follows. \square

6. FINAL REMARKS

Here we provide some remarks that could suggest further exploration.

Widths and lifespans. Although our results providing bounds on lifespans via widths are primarily theoretical, they offer practical value both in terms of improving the interpretability of PH features and in applications. Computing the exact Alexandrov, Urysohn, or Kolmogorov k -width for a given set \mathcal{X} is often computationally challenging. However, by definition, these widths arise as infima of certain measurements over k -dimensional spaces (the “cores”). While identifying the optimal k -dimensional space could be highly complex, any meaningful and well-chosen k -dimensional space can yield relevant measurements that serve as upper bounds for these widths and, consequently, for the lifespans. Thus, even if calculating the optimal bound is infeasible, our results can be effectively leveraged to provide rough yet meaningful upper bounds for the lifespans of significant topological features. See Remark 5.2 for other considerations related to potential uses of our bounds.

We highlight that our framework enables a bidirectional exchange of concepts, integrating ideas from metric geometry into applied algebraic topology and supporting their application in the reverse direction. Specifically, since lifespans are bounded above by various notions of width introduced in Section 3, they consequently provide lower bounds for these quantities. In other words, for a given metric space \mathcal{X} , the maximum lifespan over classes in $\text{Spec}_k(\mathcal{V}_\bullet(\mathcal{X}))$ (with zero birth time) serves as a lower bound for the corresponding width (e.g., $\text{UW}_{k-1}(\mathcal{X})$). This relationship offers a practical approach for estimating widths by leveraging topological persistence. We exemplify this now.

Example 6.1 (The Urysohn width of the n -torus). Let $n \geq 1$ be any integer and $a_1 \geq a_2 \geq \dots \geq a_n > 0$. Consider the n -torus $T^n := a_1 S^1 \times a_2 S^1 \times \dots \times a_n S^1$ endowed with the ℓ^∞ product metric, and where each $a_i S^1$ factor has the geodesic metric (with diameter πa_i). Then, we claim that for any $k \in \{1, \dots, n\}$ we have

$$\frac{2\pi}{3} a_k \leq \text{UW}_{k-1}(T^n) \leq \pi a_k.$$

The upper bound can be trivially obtained by considering the projection onto the first $k-1$ factors of T^n . The lower bound can be obtained through an argument via PH as follows:

- (1) As proved by Adamaszek and Adams in [AA17], $I := (0, \frac{2\pi}{3}]$ is the only bar in the VR-barcode of S^1 whose left endpoint is zero. Furthermore, this interval appears in degree-1.
- (2) By the Künneth formula for the VR-barcodes of ℓ^∞ -products of compact metric spaces (see e.g. [LMO24, Theorem 6.1 and Example 6.4]), the only intervals in the degree- k VR-barcode of T^n with zero left endpoint must arise from intersecting exactly k of the intervals $a_1 I, a_2 I, \dots, a_n I$ each corresponding to the degree-1 VR-barcode of one of the n S^1 -factors of T^n .

These intersections are precisely of the form $(0, d(L)]$ where $L \subset \{1, \dots, n\}$ s.t. $|L| = k$ and

$$d(L) := \frac{2\pi}{3} \min_{\ell \in L} a_\ell.$$

Then, the maximum of $d(L)$ over all such subsets L equals $\frac{2\pi}{3} a_k$.

(3) By Theorem 3.3 and Proposition 2.34 we now conclude that $\frac{2\pi}{3} a_k \leq \text{UW}_{k-1}(T^n)$.

Compare with [Gro88, (E₁)] and see also [Gro83, page 8].

Standard PCA vs. lifespans. In Section 3.5, we give bounds for Čech lifespans via the variances $\{\nu_k\}$ induced by PCA_∞ (Corollary 3.29). As one can easily notice, in order to have such a rigorous bound, we modified the usual PCA definition and considered an ℓ^∞ -variant. However, this makes the PCA_∞ structure highly sensitive to outliers in comparison with the original PCA. On the other hand, standard PCA is a mainstream, highly effective dimension reduction tool for real-life applications with several very efficient computational techniques available. While our results do not say anything about the relation between the original PCA and PH, an experimental result relating the lifespans of bars in the persistence diagram of the Čech filtration $\mathcal{C}_\bullet(\mathcal{X} \subset \mathbb{R}^N)$ with variances $\tilde{\nu}_{k+1}(\mathcal{X})$ of original PCA would be very interesting, for a given finite set of points $\mathcal{X} \subset \mathbb{R}^N$. It would be particularly interesting and useful for real-life applications to carry out a statistical comparative analysis for random finite subsets \mathcal{X} in \mathbb{R}^N (and for $k \geq 1$).

Principal curves and surfaces. One can notice that when describing the width-based arguments in Section 3, we first introduce a k -dimensional optimal core Λ_k for a given set \mathcal{X} , then the k -width $\text{W}_k(\mathcal{X})$ is defined as some kind of “distance” from \mathcal{X} to k -core Λ_k . Hence, when $\text{W}_k(\mathcal{X})$ is small, in metric geometry, \mathcal{X} is regarded as “essentially k -dimensional”. There is a similar notion in statistics called *Principal Curves and Surfaces*; see Hastie and Stuetzle [HS89], Delicado [Del01], and Ozertem and Erdogmus [OE11]. While principal curves and surfaces are defined as 1- and 2-dimensional objects, one can easily generalize the idea to any dimension k , e.g. principal k -manifolds. In our setting, for a compact subset \mathcal{X} in a Banach space \mathbf{Y} , principal curves and surfaces can be considered as k -dimensional objects Σ_k which minimize the ℓ^2 -distance from \mathcal{X} to Σ_k for $k = 1, 2$ with some normalization condition on Σ_k . In this sense, principal curves and surfaces can be viewed as nonlinear generalizations of principal component analysis (PCA).

Similarly, for a given set \mathcal{X} , our k -cores and principal k -manifolds can be regarded as analogous constructs which extend the underlying idea to a broader geometric framework. In Section 3.5, we defined PCA_∞ as ℓ^∞ -version of the original ℓ^2 PCA. Similarly, our k -cores Λ_k minimize the ℓ^∞ -distance between \mathcal{X} and Λ_k while principal k -manifolds Σ_k minimize the ℓ^2 -distance between \mathcal{X} and Σ_k . By using this analogy, as principal curves and surfaces are suggested as *dimension reduction* method, one can consider our k -cores represent the essential structure of \mathcal{X} when $\text{W}_k(\mathcal{X})$ is small. Furthermore, just like the discussion in the preceding paragraph (PCA vs PCA_∞), it would be interesting to carry out a statistical study of the relationship between the Čech lifespans and ℓ^2 -distance to principal k -manifolds.

APPENDIX A. AUXILIARY PROPERTIES OF WIDTHS

It was mentioned in Section 3.1 that

$$\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{UW}_k(\mathcal{X}) \leq 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y}).$$

The right-hand side inequality is trivial, whereas the inequality on the left requires an explanation.

Lemma A.1 (Alexandrov [Ale33]). *For any compact set \mathcal{X} in a Banach space \mathbf{Y} ,*

$$\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) \leq \text{UW}_k(\mathcal{X}).$$

Proof. Suppose $\text{UW}_k(\mathcal{X}) < \delta$, and let us show that $\text{AW}_k(\mathcal{X} \subset \mathbf{Y}) < \delta$. There is a continuous map $f : \mathcal{X} \rightarrow \Delta^k$ to a finite k -dimensional complex with fibers of diameter $< \delta$. Subdivide Δ^k very finely, so that the preimage of any open star of Δ^k under the map f has diameter $< \delta$.¹⁶ For each vertex $v \in \Delta^k$, pick a point c_v in the preimage of the open star S_v (unless this preimage is empty; in this case, we can safely remove S_v from Δ^k). The ball $B_\delta(c_v) \subset \mathbf{Y}$ of radius δ centered at c_v covers $f^{-1}(S_v)$. Consider an auxiliary map $\gamma : \Delta^k \rightarrow \mathbf{Y}$ defined by sending any vertex $v \in \Delta^k$ to $c_v \in \mathbf{Y}$, and then extending linearly on Δ^k . Compose this map with f , and consider $\gamma \circ f : \mathcal{X} \rightarrow \mathbf{Y}$. Its image is a simplicial complex of dimension at most k , and to complete the proof it suffices to show that every point $x \in \mathcal{X}$ is moved by distance $\|x - \gamma(f(x))\| < \delta$. Let $f(x)$ lie in the relative interior of a simplex of Δ^k with the vertices v_0, \dots, v_m . Then $x \in f^{-1}(S_{v_i}) \subset B_\delta(c_{v_i})$, and $\|x - c_{v_i}\| < \delta$, where $0 \leq i \leq m$. By construction of γ , the point $\gamma(f(x))$ lies in the convex hull of the points c_{v_i} , $0 \leq i \leq m$. Therefore, $\|x - \gamma(f(x))\|$ does not exceed the maximum of $\|x - c_{v_i}\|$ over $0 \leq i \leq m$, and this maximum is less than δ . \square

Remark A.2. Depending on the geometry of \mathbf{Y} , this inequality may be slightly improved if for each vertex $f^{-1}(S_v)$ we cover $f^{-1}(S_v)$ by a ball of the smallest possible radius. For example, if $\mathbf{Y} = \mathbb{R}^N$ is Euclidean, then it follows from Jung's theorem [Jun01] that this radius can be taken to be $\delta \sqrt{\frac{N}{2(N+1)}}$ (instead of δ). The rest of the proof runs without changes, and the final result is $\text{AW}_k(\mathcal{X} \subset \mathbb{R}^N) \leq \sqrt{\frac{N}{2(N+1)}} \text{UW}_k(\mathcal{X}) < \frac{1}{\sqrt{2}} \text{UW}_k(\mathcal{X})$. Another extreme example is when \mathbf{Y} is hyperconvex (for example, $L^\infty(\mathcal{X})$), and any bounded set \mathcal{A} can be covered by a ball of radius $\frac{1}{2} \text{diam}(\mathcal{A})$. In this case, our estimates actually imply that $\text{UW}_k(\mathcal{X}) = 2 \text{AW}_k(\mathcal{X} \subset \mathbf{Y})$.

The following property was used in Section 3.2.

Lemma A.3. *The treewidth enjoys the following continuity property. Let $\mathcal{X}_1 \supset \mathcal{X}_2 \supset \dots$ be a nested sequence of compact sets in a Banach space \mathbf{Y} , and let k be a nonnegative integer. Then*

$$\lim_{i \rightarrow \infty} \text{TW}_k(\mathcal{X}_i \subset \mathbf{Y}) = \text{TW}_k\left(\bigcap_i \mathcal{X}_i \subset \mathbf{Y}\right).$$

Proof. Denote $\mathcal{X} = \bigcap \mathcal{X}_i$, $w = \text{TW}_k(\mathcal{X} \subset \mathbf{Y})$, and let $f : \mathcal{X} \rightarrow \mathbf{Y}$ be a witness map: a continuous map whose image lies in a finite contractible k -dimensional simplicial complex $\Delta \subset \mathbf{Y}$, and such that $\|x - f(x)\| < w + \epsilon/3$ for all $x \in \mathcal{X}$ and some

¹⁶Recall that the open star S_v of a vertex $v \in \Delta^k$ is the union of the relative interiors of all simplices of Δ^k that contain v .

arbitrarily chosen $\epsilon > 0$. For each point $y \in \Delta$, the fiber $f^{-1}(y)$ lies in the open ball $B_{w+\epsilon/3}(y)$ (note that $f^{-1}(y)$ can be empty). Since f is continuous on the compact set \mathcal{X} , it is uniformly continuous, hence there is a tiny radius $\rho > 0$ such that the “thickened fiber” $f^{-1}(B_\rho(y))$ lies in the open ball $B_{w+2\epsilon/3}(y)$, for each $y \in \Delta$. Since Δ is a finite complex of dimension k , there is a finite closed cover $\bigcup D_j = \Delta$ such that its multiplicity is $k+1$, its nerve is homotopy equivalent to Δ , each D_j fits in a ball $B_\rho(y_j)$, and $C_j := f^{-1}(D_j)$ fits in the ball $B_{w+2\epsilon/3}(y_j)$.

We claim that the sets in the closed cover $\bigcup C_j = \mathcal{X}$ can be inflated a little bit while preserving their intersection pattern. For each collection of indices $J = \{j_1, \dots, j_n\}$ such that the intersection $\bigcap_{j \in J} C_j$ is empty, it follows from the compactness of \mathcal{X} that the number

$$\delta_J = \min_{x \in \mathcal{X}} \max_{j \in J} \text{dist}(x, C_j)$$

is attained and positive. Take a positive δ smaller than $\epsilon/3$ and also smaller than each δ_J over all collections J such that $\bigcap_{j \in J} C_j = \emptyset$. Consider the open cover $\{U_j\}$ of \mathcal{X} , where $U_j = \mathcal{N}_\delta(C_j \subset \mathbf{Y})$. It has the same nerve as $\{C_j\}$, and each U_j is contained in $B_{w+\epsilon}(y_j)$. We have

$$\mathcal{X} \subset \bigcup_j U_j = \mathcal{N}_\delta(\mathcal{X} \subset \mathbf{Y}).$$

Now pick a partition of unity $\{\psi_j\}$ subordinate to $\{U_j\}$. Use it to map $\bigcup_j U_j$ to its nerve. Namely, a point $x \in \bigcup_j U_j$ is mapped to $\sum_j \psi_j(x) y_j$. This gives a (possibly nonsurjective) map $\bigcup_j U_j \rightarrow \Delta$. Since $U_j \subset B_{w+\epsilon}(y_j)$, every point under this map is shifted by distance less than $w + \epsilon$.

Notice that the open set $\bigcup_j U_j$ contains all \mathcal{X}_i for i large enough, as it can be easily deduced from compactness. Therefore, we have a map from \mathcal{X}_i , for all i large enough, to a contractible k -dimensional simplicial complex inside \mathbf{Y} , and every point is shifted by a distance less than $w + \epsilon$. Hence,

$$\text{TW}_k(\mathcal{X}_i \subset \mathbf{Y}) - \text{TW}_k(\mathcal{X} \subset \mathbf{Y}) < \epsilon.$$

Since ϵ was arbitrary, the result follows. \square

APPENDIX B. ÜBERCONTRACTIBLE SETS

Here we only work in $\mathbf{Y} = \mathbb{R}^N$ with Euclidean metric, and we are interested in 0-übercontractible sets, which we call übercontractible. A rich source of übercontractible sets comes from the *cut-locus* construction, usually considered in a more general Riemannian setting; see, e.g., Wolter [Wol85]. A very close concept is that of *skeleta* (see, e.g., Fremlin [Fre97]), different from our definition only in that we take the closure. In computational geometry, low-dimensional skeleton constructions are also called *medial axes*; see, e.g., the survey by Saha, Borgefors and Sanniti di Baja [SBdB16]. We only formulate the definition of the cut-locus for convex hypersurfaces in $\mathbf{Y} = \mathbb{R}^N$.

Definition B.1. Let $K \subset \mathbf{Y}$ be a convex body (that is, a compact convex set with a nonempty interior). For each $x \in K$, consider the largest closed ball centered at x and contained in K ; let $r^{\partial K}(x) \geq 0$ be its radius, and $B^{\partial K}(x) \subset \partial K$ be the set of the points where this ball touches the boundary of K . The *cut-locus* of ∂K is

the closure of the set of all such $x \in K$ for which the cardinality of $B^{\partial K}(x)$ is at least 2:

$$C^{\partial K} = \text{cl}(\{x \in K : |B^{\partial K}(x)| > 1\}).$$

There is a canonical retraction $c^{\partial K} : K \rightarrow C^{\partial K}$ defined as follows. For $x \in C^{\partial K}$, set $c^{\partial K}(x) = x$. For each point $x \in K \setminus C^{\partial K}$, let $b^{\partial K}(x)$ be the only element of $B^{\partial K}(x)$, and consider the largest closed ball contained in K and touching ∂K at $b^{\partial K}(x)$. The center of this ball belongs to the cut-locus and will be denoted $c^{\partial K}(x)$.

Theorem B.2. *The cut-locus $C^{\partial K}$ of the boundary of a convex polytope K is übercontractible.*

The key lemma in the proof comes from the Morse theory for the distance functional. It was applied by Bárány, Holmsen and Karasev in [BHK15, Section 3] to give a sufficient condition for a set to be contractible, but the proof there, in fact, guarantees that the set is übercontractible.

Lemma B.3 (Cf. [BHK15, Theorem 2]). *Let $\mathcal{T} \subset \mathbf{Y}$ be a union of finitely many compact convex sets. For each $x \notin \mathcal{T}$ consider the largest ball centered at x whose interior does not meet \mathcal{T} . Let $B^{\mathcal{T}}(x)$ be the set of the points where this ball touches \mathcal{T} . Suppose that for each $x \notin \mathcal{T}$ we have $x \notin \text{conv}(B^{\mathcal{T}}(x))$. Then \mathcal{T} is übercontractible.*

Proof of Theorem B.2. Suppose K is a convex polytope. It can be shown that $C^{\partial K}$ is a polyhedral complex of codimension 1 in Y . Let F be a facet of K , and let n_F be the corresponding inner normal vector. The (relative) boundary of F lies in $C^{\partial K}$. The map $c^{\partial K}$, restricted to F , sends each point $x \in F$ along n_F until it hits $C^{\partial K}$. Together F and the polyhedral surface $c^{\partial K}(F)$ bound a convex polytope K_F .

Consider a point $x \in K \setminus C^{\partial K}$. It then lies in K_F , for some facet F . It is sufficient to verify the assumption of Lemma B.3: one needs to check that $x \notin \text{conv}(B^{\mathcal{T}}(x))$, where $\mathcal{T} = C^{\partial K}$. Indeed, for every $y \in B^{\mathcal{T}}(x)$, the vector $y - x$ forms an acute angle with n_F , and therefore, the entire set $B^{\mathcal{T}}(x)$ lies in the open halfspace $\{z \in Y : \langle z - x, n_F \rangle > 0\}$. Therefore, $x \notin \text{conv}(B^{\mathcal{T}}(x))$. \square

We conclude this section by speculating how Theorem B.2 can be proven for convex bodies other than polytopes. Unfortunately, the case of a general convex body K cannot be proven by approximating K with polytopes K_i , because in order to guarantee C^0 -convergence $C^{\partial K_i} \rightarrow C^{\partial K}$, we cannot get away just with polytopes. We sketch an argument that works for fairly general convex bodies, modulo some technicalities. We take a different approach and assume the following “tameness” assumption: $C^{\partial K}$ consists of finitely many compact convex sets. It allows us to apply Lemma B.3, but it is highly likely that the lemma can also be stated and proven in greater generality. Not only the class of tame convex bodies includes polytopes, but it also seems to be C^2 -dense among all C^2 -smooth convex bodies (we do not discuss this in detail since this digresses too far from the main topic of the paper).

Sketch of the proof of Theorem B.2 for tame convex bodies. The following two properties will be shown to imply the übercontractibility of $C^{\partial K}$.

- (1) The cut-locus of ∂K consists of finitely many compact convex sets (so we can use Lemma B.3).
- (2) The boundary of K is C^2 -smooth and strongly convex in the sense that its second fundamental form is positive definite.

The tameness assumption guarantees the first property. We argue that we can inflate K slightly to get a new convex body K' that satisfies the second property, while preserving cut-locus $C^{\partial K'} = C^{\partial K} = \mathcal{T}$. This amounts to choosing $r^{\partial K'}(x)$ for $x \in \mathcal{T}$ carefully, so that on every face of \mathcal{T} , $r^{\partial K'}$ is C^1 -close to $r^{\partial K}$, and $r^{\partial K'}$ is strongly convex. The details are omitted.

In the rest of the proof we assume that K satisfies the two properties above, and we will verify the assumption of Lemma B.3 for $\mathcal{T} = C^{\partial K}$ to show that it is übercontractible. Pick points $x \in K \setminus \mathcal{T}$, and $y \in B^{\mathcal{T}}(x)$. Our goal is to prove that $\langle y - x, n \rangle > 0$, where n is the inner normal to ∂K at $b^{\partial K}(x)$. This will imply that $x \notin \text{conv}(B^{\mathcal{T}}(x)) \subset \{z \in \mathbf{Y} : \langle z - x, n \rangle > 0\}$.

Suppose, for the contrary, that $\langle y - x, n \rangle \leq 0$. Parametrize the straight line segment $[x, y]$ linearly as $x(t), t \in [0, 1]$, $x(0) = x$, $x(1) = y$. For each $t \in [0, 1]$, let $u(t) = b^{\partial K}(x(t))$ be the only element $B^{\partial K}(x(t))$. Consider also $v(t) = c^{\partial K}(x(t))$, for $t \in [0, 1]$, and notice that $v(1) = y$. The key idea is to look at the continuous family of straight line segments $[x(t), v(t)]$, and investigate how they intersect the open ball O centered at x of radius $r^{\partial K}(x)$. We make the following observations.

- The open ball O does not intersect \mathcal{T} , and $v(t) \in \mathcal{T}$. So for each $t \in [0, 1)$, the segment $[x(t), v(t)]$ starts inside O and ends outside of O .
- The strong convexity of ∂K implies that the angle between the vectors $y - x$ and $v(t) - x(t)$ is strictly increasing for $t \in [0, 1)$. To see that, one needs to differentiate in t the inner normal to ∂K at $u(t)$ (which is collinear with $v(t) - x(t)$).
- Initially, the angle between the vectors $y - x$ and $v(0) - x(0)$ is nonacute (since by assumption $\langle y - x, n \rangle \leq 0$). Therefore, for t close to 1 the angle between $y - x$ and $v(t) - x(t)$ is obtuse, and tends to a limit that is obtuse. It follows that the length of the part of $[x(t), v(t)]$ that lies in O is bounded away from 0.
- But, as $t \rightarrow 1$, the segment $[x(t), v(t)]$ degenerates to the point y , so its length must approach zero. This contradiction concludes the proof.

□

APPENDIX C. NOTATION TABLE

TABLE 1. Notation and main symbols

Notation	Definition	Place
$\text{conv}(\mathcal{X})$	Convex hull of \mathcal{X} in a Banach space	
$\overline{\text{conv}}(\mathcal{X})$	Closure of convex hull of \mathcal{X} in a Banach space	
$H_k(\cdot)$	Singular homology in degree k	
$\tilde{H}_k(\cdot)$	Reduced singular homology in degree k	
$\mathcal{N}_r(\mathcal{X} \subset \mathcal{Z})$	Open r -neighborhood of \mathcal{X} in \mathcal{Z}	Section 2.1
$\overline{\mathcal{N}}_r(\mathcal{X} \subset \mathcal{Z})$	Closed r -neighborhood of \mathcal{X} in \mathcal{Z}	Section 2.1
$\mathcal{V}_r(\mathcal{X})$	Vietoris–Rips complex of \mathcal{X} for distance r	Definition 2.1
$\mathcal{C}_r(\mathcal{X} \subset \mathcal{Z})$	Čech complex of $\mathcal{X} \subset \mathcal{Z}$ for distance r	Definition 2.2
$\text{Spec}_k(\Delta_\bullet)$	Homological spectrum of the filtration Δ_\bullet	Definition 2.18
$\text{PD}_k(\mathcal{X})$	k^{th} -persistence diagram of \mathcal{X} with VR filtration	Section 2.1.2
$\widetilde{\text{PD}}_k(\mathcal{X} \subset \mathbf{Y})$	k^{th} -persistence diagram of $\mathcal{X} \subset \mathbf{Y}$ with Čech filtration	Section 2.1.2
b_ω, d_ω	Birth and death time of the homology class ω	Definition 2.19
$d_\omega - b_\omega$	Persistence (lifespan) of the homology class ω	Definition 2.19
$d_b(\cdot, \cdot)$	Bottleneck distance between persistence diagrams	Section 2.1.3
$d_H^Y(\cdot, \cdot)$	Hausdorff distance between two subsets in \mathbf{Y}	Section 2.1
$d_{\text{CH}}(\cdot, \cdot)$	Gromov–Hausdorff distance between two metric spaces	Section 2.1
$L^\infty(\mathcal{X})$	Space of bounded functions on \mathcal{X} with sup norm	Section 2.2
$E(\mathcal{X})$	Tight span of \mathcal{X}	Definition 2.28
$\rho(\mathcal{M})$	Gromov’s filling radius of \mathcal{M}	Definition 2.31
$\rho(\omega; \mathcal{X} \subset \mathbf{Y})$	The relative filling radius of $\omega \in \widetilde{H}_k(\mathcal{X})$ relative to \mathbf{Y}	Definition 2.33
$\rho(\omega; \mathcal{X})$	The absolute filling radius of $\omega \in \widetilde{H}_k(\mathcal{X})$	Definition 2.33
$\text{rad}(\mathcal{X})$	The radius of \mathcal{X}	Definition 2.39
$\text{rad}(\mathcal{X} \subset \mathbf{Y})$	The circumradius of \mathcal{X} relative to \mathbf{Y}	Definition 2.39
$\text{UW}_k(\mathcal{X})$	Urysohn k -width of \mathcal{X}	Definition 3.1
$\text{AW}_k(\mathcal{X} \subset \mathbf{Y})$	Alexandrov k -width of \mathcal{X} relative to \mathbf{Y}	Definition 3.4
$\text{TW}_k(\mathcal{X} \subset \mathbf{Y})$	k^{th} treewidth of \mathcal{X} relative to \mathbf{Y}	Definition 3.9
$\text{TW}_k^C(\mathcal{X} \subset \mathbf{Y})$	C -robust k -dimensional treewidth of \mathcal{X}	Definition 3.22
$\text{KW}_k(\mathcal{X} \subset \mathbf{Y})$	Kolmogorov k -width of \mathcal{X} relative to \mathbf{Y}	Definition 3.6
$\nu_k(\mathcal{X})$	k^{th} variance of \mathcal{X} (PCA $_\infty$)	Definition 3.26
$\ddot{\nu}$ -spread($\mathcal{X} \subset \mathbf{Y}$)	Überspread of \mathcal{X} relative to \mathbf{Y}	Definition 4.4
$\xi(\mathcal{X})$	VR-extinction time of \mathcal{X}	Definition 5.1
$\check{\xi}(\mathcal{X} \subset \mathbf{Y})$	Čech-extinction time of \mathcal{X} relative to \mathbf{Y}	Definition 5.1
$\text{cdef}(\mathcal{X} \subset \mathbf{Y})$	Convexity deficiency of \mathcal{X} relative to \mathbf{Y}	Definition 5.4
$\text{hcdef}(\mathcal{X})$	Hyperconvexity deficiency of \mathcal{X}	Definition 5.9

ACKNOWLEDGMENT

A major part of this work was done while the first-named author was a postdoc at the Institute for Advanced Study, Princeton.

REFERENCES

- [AA17] Michał Adamaszek and Henry Adams, *The Vietoris–Rips complexes of a circle*, Pacific J. Math. **290** (2017), no. 1, 1–40, DOI 10.2140/pjm.2017.290.1. MR3673078
- [AAF19] Mehmet E. Aktas, Esra Akbas, and Ahmed El Fatmaoui, *Persistence homology of networks: methods and applications*, Appl. Network Sci. **4** (2019), no. 1, 1–28.
- [AAR19] Michał Adamaszek, Henry Adams, and Samadwara Reddy, *On Vietoris–Rips complexes of ellipses*, J. Topol. Anal. **11** (2019), no. 3, 661–690, DOI 10.1142/S1793525319500274. MR3999516

- [ABC⁺22] Henry Adams, Johnathan Bush, Nate Clause, Florian Frick, Mario Gómez, Michael Harrison, R. Amzi Jeffs, Evgeniya Lagoda, Sunhyuk Lim, Facundo Mémoli, et al., *Gromov–Hausdorff distances, Borsuk–Ulam theorems, and Vietoris–Rips complexes*, Preprint, [arXiv:2301.00246](https://arxiv.org/abs/2301.00246), 2022.
- [AC22] Henry Adams and Baris Coskunuzer, *Geometric approaches to persistent homology*, SIAM J. Appl. Algebra Geom. **6** (2022), no. 4, 685–710, DOI 10.1137/21M1422914. MR4522864
- [AFV23] Henry Adams, Florian Frick, and Žiga Virk, *Vietoris thickenings and complexes have isomorphic homotopy groups*, J. Appl. Comput. Topol. **7** (2023), no. 2, 221–241, DOI 10.1007/s41468-022-00106-5. MR4588139
- [Ale26] Paul Alexandroff, *Notes supplémentaires au “Mémoire sur les multiplicités Cantoriennes”, rédigées d’après les papiers posthumes de Paul Urysohn*, Fund. Math. **1** (1926), no. 8, 352–359.
- [Ale28] Paul Alexandroff, *Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung* (German), Math. Ann. **98** (1928), no. 1, 617–635, DOI 10.1007/BF01451612. MR1512423
- [Ale33] Paul Alexandroff, *Über die Urysohnschen Konstanten*, Fund. Math. **20** (1933), no. 1, 140–150.
- [ALS13] Dominique Attali, André Lieutier, and David Salinas, *Vietoris–Rips complexes also provide topologically correct reconstructions of sampled shapes*, Comput. Geom. **46** (2013), no. 4, 448–465, DOI 10.1016/j.comgeo.2012.02.009. MR3003901
- [AMMW24] Henry Adams, Facundo Mémoli, Michael Moy, and Qingsong Wang, *The persistent topology of optimal transport based metric thickenings*, Algebr. Geom. Topol. **24** (2024), no. 1, 393–447, DOI 10.2140/agt.2024.24.393. MR4721372
- [AP56] N. Aronszajn and P. Panitchpakdi, *Extension of uniformly continuous transformations and hyperconvex metric spaces*, Pacific J. Math. **6** (1956), 405–439. MR84762
- [AQO⁺20] Erik J. Amézquita, Michelle Y. Quigley, Tim Ophelders, Elizabeth Munch, and Daniel H. Chitwood, *The shape of things to come: TDA and biology, from molecules to organisms*, Dev. Dyn. **249** (2020), no. 7, 816–833.
- [Bal21] Alexey Balitskiy, *Bounds on Urysohn Width*, ProQuest LLC, Ann Arbor, MI, 2021. Thesis (Ph.D.)–Massachusetts Institute of Technology. MR4464413
- [Bar94] S. A. Barannikov, *The framed Morse complex and its invariants*, Singularities and bifurcations, Adv. Soviet Math., vol. 21, Amer. Math. Soc., Providence, RI, 1994, pp. 93–115. MR1310596
- [Bau21] Ulrich Bauer, *Ripser: efficient computation of Vietoris–Rips persistence barcodes*, J. Appl. Comput. Topol. **5** (2021), no. 3, 391–423, DOI 10.1007/s41468-021-00071-5. MR4298669
- [BBI01] Dmitri Burago, Yuri Burago, and Sergei Ivanov, *A course in metric geometry*, Graduate Studies in Mathematics, vol. 33, American Mathematical Society, Providence, RI, 2001, DOI 10.1090/gsm/033. MR1835418
- [BH10] Michael Brunnbauer and Bernhard Hanke, *Large and small group homology*, J. Topol. **3** (2010), no. 2, 463–486, DOI 10.1112/jtopol/jtq014. MR2651368
- [BHK15] Imre Bárány, Andreas F. Holmsen, and Roman Karasev, *Topology of geometric joins*, Discrete Comput. Geom. **53** (2015), no. 2, 402–413, DOI 10.1007/s00454-015-9665-2. MR3316229
- [BHPG⁺21] Michael Bleher, Lukas Hahn, Juan Ángel Patino-Galindo, Mathieu Carriere, Ulrich Bauer, Raúl Rabadán, and Andreas Ott, *Topology identifies emerging adaptive mutations in SARS-CoV-2*, medRxiv, 2021, pp. 2021–06.
- [BHPW20] Peter Bubenik, Michael Hull, Dhruv Patel, and Benjamin Whittle, *Persistent homology detects curvature*, Inverse Problems **36** (2020), no. 2, 025008, 23, DOI 10.1088/1361-6420/ab4ac0. MR4063201
- [BKRR23] Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle, *A unified view on the functorial nerve theorem and its variations*, Expo. Math. **41** (2023), no. 4, Paper No. 125503, 52, DOI 10.1016/j.exmath.2023.04.005. MR4652781
- [BKS17] Omer Bobrowski, Matthew Kahle, and Primoz Skraba, *Maximally persistent cycles in random geometric complexes*, Ann. Appl. Probab. **27** (2017), no. 4, 2032–2060, DOI 10.1214/16-AAP1232. MR3693519

- [BM24] Peter Bubenik and Nikola Milićević, *Homotopy, homology, and persistent homology using closure spaces*, J. Appl. Comput. Topol. **8** (2024), no. 3, 579–641, DOI 10.1007/s41468-024-00183-8. MR4799024
- [Bor32] Karol Borsuk, *Über eine Klasse von lokal zusammenhängenden Räumen*, Fund. Math. **19** (1932), no. 1, 220–242.
- [Bor48] Karol Borsuk, *On the imbedding of systems of compacta in simplicial complexes*, Fund. Math. **35** (1948), 217–234, DOI 10.4064/fm-35-1-217-234. MR28019
- [Bor75] Karol Borsuk, *Theory of shape*, Monografie Matematyczne [Mathematical Monographs], Tom 59, PWN—Polish Scientific Publishers, Warsaw, 1975. MR418088
- [BS83] Enrico Bombieri and Leon Simon, *On the Gehring link problem*, Seminar on minimal submanifolds, Ann. of Math. Stud., vol. 103, Princeton Univ. Press, Princeton, NJ, 1983, pp. 271–274. MR795242
- [Cai94] Mingliang Cai, *On Gromov’s large Riemannian manifolds*, Geom. Dedicata **50** (1994), no. 1, 37–45, DOI 10.1007/BF01263649. MR1280793
- [Car09] Gunnar Carlsson, *Topology and data*, Bull. Amer. Math. Soc. (N.S.) **46** (2009), no. 2, 255–308, DOI 10.1090/S0273-0979-09-01249-X. MR2476414
- [CB15] William Crawley-Boevey, *Decomposition of pointwise finite-dimensional persistence modules*, J. Algebra Appl. **14** (2015), no. 5, 1550066, 8, DOI 10.1142/S0219498815500668. MR3323327
- [CCI⁺20] Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda, *Perslay: a neural network layer for persistence diagrams and new graph topological signatures*, International Conference on Artificial Intelligence and Statistics, pp. 2786–2796, PMLR, 2020.
- [CCR13] Joseph Minhow Chan, Gunnar Carlsson, and Raul Rabadan, *Topology of viral evolution*, Proc. Natl. Acad. Sci. USA **110** (2013), no. 46, 18566–18571, DOI 10.1073/pnas.1313480110. MR3153945
- [CCSG⁺09] Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Mémoli, and Steve Y. Oudot, *Gromov–Hausdorff stable signatures for shapes using persistence*, Computer Graphics Forum, vol. 28, pp. 1393–1403, Wiley Online Library, 2009.
- [CDSGO16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot, *The structure and stability of persistence modules*, SpringerBriefs in Mathematics, Springer, [Cham], 2016, DOI 10.1007/978-3-319-42545-0. MR3524869
- [CDSO14] Frédéric Chazal, Vin de Silva, and Steve Oudot, *Persistence stability for geometric complexes*, Geom. Dedicata **173** (2014), 193–214, DOI 10.1007/s10711-013-9937-z. MR3275299
- [Che97] Victor Chepoi, *A T_X -approach to some results on cuts and metrics*, Adv. in Appl. Math. **19** (1997), no. 4, 453–470, DOI 10.1006/aama.1997.0549. MR1479014
- [CI08] Carina Curto and Vladimir Itskov, *Cell groups reveal structure of stimulus space*, PLoS Comput. Biol. **4** (2008), no. 10, e1000205, 13, DOI 10.1371/journal.pcbi.1000205. MR2457124
- [CVJ21] Gunnar Carlsson and Mikael Vejdemo-Johansson, *Topological data analysis with applications*, Cambridge University Press, Cambridge, 2022, DOI 10.1017/9781108975704. MR4346385
- [CW17] Zixuan Cang and Guo-Wei Wei, *Topologynet: topology based deep convolutional and multi-task neural networks for biomolecular property predictions*, PLoS Comput. Biol. **13** (2017), no. 7, e1005690.
- [Deh55] René Deheuvels, *Topologie d’une fonctionnelle* (French), Ann. of Math. (2) **61** (1955), 13–72, DOI 10.2307/1969619. MR69488
- [Del01] Pedro Delicado, *Another look at principal curves and surfaces*, J. Multivariate Anal. **77** (2001), no. 1, 84–116, DOI 10.1006/jmva.2000.1917. MR1838715
- [DMFC12] Yuri Dabaghian, Facundo Mémoli, Loren Frank, and Gunnar Carlsson, *A topological paradigm for hippocampal spatial map formation using persistent homology*, PLoS Comput. Biol. (2012).
- [Dre84] Andreas W. M. Dress, *Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups: a note on combinatorial properties of metric spaces*, Adv. in Math. **53** (1984), no. 3, 321–402, DOI 10.1016/0001-8708(84)90029-X. MR753872

[EH10] Herbert Edelsbrunner and John L. Harer, *Computational topology*, American Mathematical Society, Providence, RI, 2010. An introduction, DOI 10.1090/mhk/069. MR2572029

[ELZ02] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian, *Topological persistence and simplification*, Discrete Comput. Geom. **28** (2002), no. 4, 511–533, DOI 10.1007/s00454-002-2885-2. Discrete and computational geometry and graph drawing (Columbia, SC, 2001). MR1949898

[FF60] Herbert Federer and Wendell H. Fleming, *Normal and integral currents*, Ann. of Math. (2) **72** (1960), 458–520, DOI 10.2307/1970227. MR123260

[Fre97] D. H. Fremlin, *Skeletons and central sets*, Proc. London Math. Soc. (3) **74** (1997), no. 3, 701–720, DOI 10.1112/S0024611597000233. MR1434446

[Fro90] Patrizio Frosini, *A distance for similarity classes of submanifolds of a Euclidean space*, Bull. Austral. Math. Soc. **42** (1990), no. 3, 407–416, DOI 10.1017/S0004972700028574. MR1083277

[GHI⁺15] Marcio Gameiro, Yasuaki Hiraoka, Shunsuke Izumi, Miroslav Kramar, Konstantin Mischaikow, and Vidit Nanda, *A topological measurement of protein compressibility*, Jpn. J. Ind. Appl. Math. **32** (2015), no. 1, 1–17, DOI 10.1007/s13160-014-0153-5. MR3318898

[Ghr18] Robert Ghrist, *Homological algebra and data*, The mathematics of data, IAS/Park City Math. Ser., vol. 25, Amer. Math. Soc., Providence, RI, 2018, pp. 273–325. MR3839171

[Gil24] Patrick Gillespie, *Vietoris thickenings and complexes are weakly homotopy equivalent*, J. Appl. Comput. Topol. **8** (2024), no. 1, 35–53, DOI 10.1007/s41468-023-00135-8. MR4707456

[GK18] Marian Gidea and Yuri Katz, *Topological data analysis of financial time series: landscapes of crashes*, Phys. A **491** (2018), 820–834, DOI 10.1016/j.physa.2017.09.028. MR3721543

[Gro83] Mikhael Gromov, *Filling Riemannian manifolds*, J. Differential Geom. **18** (1983), no. 1, 1–147. MR697984

[Gro86] M. Gromov, *Large Riemannian manifolds*, Curvature and topology of Riemannian manifolds (Katata, 1985), Lecture Notes in Math., vol. 1201, Springer, Berlin, 1986, pp. 108–121, DOI 10.1007/BFb0075649. MR859578

[Gro88] M. Gromov, *Width and related invariants of Riemannian manifolds* (English, with French summary), Astérisque **163–164** (1988), 6, 93–109, 282 (1989). On the geometry of differentiable manifolds (Rome, 1986). MR999972

[Gro20] Misha Gromov, *No metrics with positive scalar curvatures on aspherical 5-manifolds*, Preprint, [arXiv:2009.05332](https://arxiv.org/abs/2009.05332), 2020.

[Gut11] Larry Guth, *Volumes of balls in large Riemannian manifolds*, Ann. of Math. (2) **173** (2011), no. 1, 51–76, DOI 10.4007/annals.2011.173.1.2. MR2753599

[Gut17] Larry Guth, *Volumes of balls in Riemannian manifolds and Uryson width*, J. Topol. Anal. **9** (2017), no. 2, 195–219, DOI 10.1142/S1793525317500029. MR3622232

[Has09] Trevor Hastie, *The elements of statistical learning: data mining, inference, and prediction*, 2009.

[Hau95] Jean-Claude Hausmann, *On the Vietoris-Rips complexes and a cohomology theory for metric spaces*, Prospects in topology (Princeton, NJ, 1994), Ann. of Math. Stud., vol. 138, Princeton Univ. Press, Princeton, NJ, 1995, pp. 175–188. MR1368659

[HG16] Gregory Henselman and Robert Ghrist, *Matroid filtrations and computational persistent homology*, Preprint, [arXiv:1606.00199](https://arxiv.org/abs/1606.00199), 2016.

[HGR⁺20] Christoph Hofer, Florian Graf, Bastian Rieck, Marc Niethammer, and Roland Kwitt, *Graph filtration learning*, International Conference on Machine Learning, pp. 4314–4323, PMLR, 2020.

[HJ23] Michael Harrison and R. Amzi Jeffs, *Quantitative upper bounds on the Gromov-Hausdorff distance between spheres*, Preprint, [arXiv:2309.11237](https://arxiv.org/abs/2309.11237), 2023.

[HKN19] Christoph D. Hofer, Roland Kwitt, and Marc Niethammer, *Learning representations of persistence barcodes*, J. Mach. Learn. Res. **20** (2019), Paper No. 126, 45. MR4002880

- [HKNU17] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl, *Deep learning with topological signatures*, Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, pp. 1633–1643.
- [HMMN14] Shaun Harker, Konstantin Mischaikow, Marian Mrozek, and Vidit Nanda, *Discrete Morse theoretic algorithms for computing homology of complexes and maps*, Found. Comput. Math. **14** (2014), no. 1, 151–184, DOI 10.1007/s10208-013-9145-0. MR3160710
- [HNH⁺16] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G. Escolar, Kaname Matsue, and Yasumasa Nishiura, *Hierarchical structures of amorphous solids characterized by persistent homology*, Proc. Natl. Acad. Sci. USA **113** (2016), no. 26, 7035–7040.
- [HS89] Trevor Hastie and Werner Stuetzle, *Principal curves*, J. Amer. Statist. Assoc. **84** (1989), no. 406, 502–516. MR1010339
- [Hu65] Sze-tsen Hu, *Theory of retracts*, Wayne State University Press, Detroit, MI, 1965. MR181977
- [Isb64] J. R. Isbell, *Six theorems about injective metric spaces*, Comment. Math. Helv. **39** (1964), 65–76, DOI 10.1007/BF02566944. MR182949
- [JJ23] Parvaneh Joharinad and Jürgen Jost, *Mathematical principles of topological and geometric data analysis*, Mathematics of Data, vol. 2, Springer, Cham, [2023] ©2023, DOI 10.1007/978-3-031-33440-5. MR4648028
- [Jun01] Heinrich Jung, *Ueber die kleinste Kugel, die eine räumliche Figur einschliesst* (German), J. Reine Angew. Math. **123** (1901), 241–257, DOI 10.1515/crll.1901.123.241. MR1580570
- [Kat83] Mikhail Katz, *The filling radius of two-point homogeneous spaces*, J. Differential Geom. **18** (1983), no. 3, 505–511. MR723814
- [Kat89] Mikhail Katz, *Diameter-extremal subsets of spheres*, Discrete Comput. Geom. **4** (1989), no. 2, 117–137, DOI 10.1007/BF02187719. MR973541
- [Kat90] Mikhail Katz, *The filling radius of homogeneous manifolds*, Séminaire de Théorie Spectrale et Géométrie, No. 9, Année 1990–1991, Sémin. Théor. Spectr. Géom., vol. 9, Univ. Grenoble I, Saint-Martin-d’Hères, 1991, pp. 103–109, DOI 10.5802/tsg.88. MR1715933
- [Kat91] Mikhail Katz, *On neighborhoods of the Kuratowski imbedding beyond the first extremum of the diameter functional*, Fund. Math. **137** (1991), no. 3, 161–175, DOI 10.4064/fm-137-3-161-175. MR1110030
- [KK16] Mehmet Kılıç and Şahin Koçak, *Tight span of subsets of the plane with the maximum metric*, Adv. Math. **301** (2016), 693–710, DOI 10.1016/j.aim.2016.05.026. MR3539386
- [Kol36] A. Kolmogoroff, *Über die beste Annäherung von Funktionen einer gegebenen Funktionenklasse* (German), Ann. of Math. (2) **37** (1936), no. 1, 107–110, DOI 10.2307/1968691. MR1503273
- [Lan13] Urs Lang, *Injective hulls of certain discrete metric spaces and groups*, J. Topol. Anal. **5** (2013), no. 3, 297–331, DOI 10.1142/S1793525313500118. MR3096307
- [Lat01] Janko Latschev, *Vietoris–Rips complexes of metric spaces near a closed Riemannian manifold*, Arch. Math. (Basel) **77** (2001), no. 6, 522–528, DOI 10.1007/PL00000526. MR1879057
- [Lax02] Peter D. Lax, *Functional analysis*, Pure and Applied Mathematics (New York), Wiley-Interscience [John Wiley & Sons], New York, 2002. MR1892228
- [LBD⁺18] Yongjin Lee, Senja D. Barthel, Paweł Dłotko, Seyed Mohamad Moosavi, Kathryn Hess, and Berend Smit, *High-throughput screening approach for nanoporous materials genome using topological data analysis: application to zeolites*, J. Chem. Theory Comput. **14** (2018), no. 8, 4427–4437.
- [LLNR22] Yevgeny Liokumovich, Boris Lishak, Alexander Nabutovsky, and Regina Rotman, *Filling metric spaces*, Duke Math. J. **171** (2022), no. 3, 595–632, DOI 10.1215/00127094-2021-0039. MR4382977
- [LMO20] Sunhyuk Lim, Facundo Memoli, and Osman Berat Okutan, *Vietoris–Rips persistent homology, injective metric spaces, and the filling radius*. arXiv preprint arXiv:2001.07588, 2020.

- [LMO24] Sunhyuk Lim, Facundo Mémoli, and Osman Berat Okutan, *Vietoris-Rips persistent homology, injective metric spaces, and the filling radius*, Algebr. Geom. Topol. **24** (2024), no. 2, 1019–1100, DOI 10.2140/agt.2024.24.1019. MR4735051
- [LMS23] Sunhyuk Lim, Facundo Mémoli, and Zane Smith, *The Gromov-Hausdorff distance between spheres*, Geom. Topol. **27** (2023), no. 9, 3733–3800, DOI 10.2140/gt.2023.27.3733. MR4674839
- [Meg12] Robert E. Megginson, *An introduction to Banach space theory*, Graduate Texts in Mathematics, vol. 183, Springer-Verlag, New York, 1998, DOI 10.1007/978-1-4612-0603-3. MR1650235
- [MN13] Konstantin Mischaikow and Vidit Nanda, *Morse theory for filtrations and efficient computation of persistent homology*, Discrete Comput. Geom. **50** (2013), no. 2, 330–353, DOI 10.1007/s00454-013-9529-6. MR3090522
- [Mor30] Marston Morse, *The foundations of a theory of the calculus of variations in the large in m -space. II*, Trans. Amer. Math. Soc. **32** (1930), no. 4, 599–631, DOI 10.2307/1989343. MR1501555
- [MS73] J. H. Michael and L. M. Simon, *Sobolev and mean-value inequalities on generalized submanifolds of R^n* , Comm. Pure Appl. Math. **26** (1973), 361–379, DOI 10.1002/cpa.3160260305. MR344978
- [MZ24] Facundo Mémoli and Ling Zhou, *Ephemeral persistence features and the stability of filtered chain complexes*, J. Comput. Geom. **15** (2024), no. 2, 258–328. MR4878036
- [NLC11] Monica Nicolau, Arnold J. Levine, and Gunnar Carlsson, *Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival*, Proc. Natl. Acad. Sci. USA **108** (2011), no. 17, 7265–7270.
- [NRS21] Alexander Nabutovsky, Regina Rotman, and Stéphane Sabourau, *Sweepouts of closed Riemannian manifolds*, Geom. Funct. Anal. **31** (2021), no. 3, 721–766, DOI 10.1007/s00039-021-00575-3. MR4311582
- [OE11] Umut Ozertem and Deniz Erdogmus, *Locally defined principal curves and surfaces*, J. Mach. Learn. Res. **12** (2011), 1249–1286. MR2804600
- [ONH22] Ippei Obayashi, Takenobu Nakamura, and Yasuaki Hiraoka, *Persistent homology analysis for materials research and persistent homology software: HomCloud*, J. Phys. Soc. Japan **91** (2022), no. 9, 091013.
- [Pat18] Amit Patel, *Generalized persistence diagrams*, J. Appl. Comput. Topol. **1** (2018), no. 3–4, 397–419, DOI 10.1007/s41468-018-0012-6. MR3975559
- [Pea01] Karl Pearson, *LIII. On lines and planes of closest fit to systems of points in space*, Lond. Edinb. Dublin Philos. Mag. J. Sci. **2** (1901), no. 11, 559–572.
- [PRSZ20] Leonid Polterovich, Daniel Rosen, Karina Samvelyan, and Jun Zhang, *Topological persistence in geometry and analysis*, University Lecture Series, vol. 74, American Mathematical Society, Providence, RI, [2020] ©2020. MR4249570
- [RB19] Raúl Rabadán and Andrew J. Blumberg, *Topological data analysis for genomics and evolution: topology in biology*, Cambridge University Press, 2019.
- [Rie20] Antonio Rieser, *Vietoris-Rips homology theory for semi-uniform spaces*, Preprint, arXiv:2008.05739, 2020.
- [RM24] Saúl Rodríguez Martín, *Gromov-Hausdorff distances from simply connected geodesic spaces to the circle*, Proc. Amer. Math. Soc. Ser. B **11** (2024), 624–637, DOI 10.1090/bproc/243. MR4840249
- [RNS⁺17] Michael W. Reimann, Max Nolte, Martina Scolamiero, Katharine Turner, Rodrigo Perin, Giuseppe Chindemi, Paweł Dłotko, Ran Levi, Kathryn Hess, and Henry Markram, *Cliques of neurons bound into cavities provide a missing link between structure and function*, Front. Comput. Neurosci. **11** (2017), 266051.
- [Rob99] V. Robins, *Towards computing homology from finite approximations*, Proceedings of the 14th Summer Conference on General Topology and its Applications (Brookville, NY, 1999), Topology Proc. **24** (1999), no. Summer, 503–532 (2001). MR1876386
- [RQD23] Simon Rudkin, Wanling Qiu, and Paweł Dłotko, *Uncertainty, volatility and the persistence norms of financial time series*, Expert Syst. Appl. **223** (2023), 119894.
- [RSDFS16] Vanessa Robins, Mohammad Saadatfar, Olaf Delgado-Friedrichs, and Adrian P. Sheppard, *Percolating length scales from topological persistence analysis of micro-CT images of porous materials*, Water Resour. Res. **52** (2016), no. 1, 315–329.

- [Sab20] Stéphane Sabourau, *One-cycle sweepout estimates of essential surfaces in closed Riemannian manifolds*, Amer. J. Math. **142** (2020), no. 4, 1051–1082, DOI 10.1353/ajm.2020.0031. MR4124115
- [SBdB16] Punam K. Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja, *A survey on skeletonization algorithms and their applications*, Pattern Recognit. Lett. **76** (2016), 3–12.
- [SDB⁺22] Søren S. Sørensen, Tao Du, Christophe A. N. Biscio, Lisbeth Fajstrup, and Morten M. Smedskjaer, *Persistent homology: a tool to understand medium-range order glass structure*, J. Non-Crystalline Solids X **16** (2022), 100123.
- [SL22] Yara Skaf and Reinhard Laubenbacher, *Topological data analysis in biomedicine: a review*, J. Biomed. Inform. (2022), 104082.
- [SMI⁺08] Gurjeet Singh, Facundo Memoli, Tigran Ishkhanov, Guillermo Sapiro, Gunnar Carlsson, and Dario L. Ringach, *Topological analysis of population activity in visual cortex*, J. Vision **8** (2008), no. 8, 11–11.
- [Tik76] V. M. Tikhomirov, *Nekotorye voprosy teorii priblizhenii* (Russian), Izdat. Moskov. Univ., Moscow, 1976. MR487161
- [Tur19] Katharine Turner, *Rips filtrations for quasimetric spaces and asymmetric functions with stability results*, Algebr. Geom. Topol. **19** (2019), no. 3, 1135–1170, DOI 10.2140/agt.2019.19.1135. MR3954277
- [UZ16] Michael Usher and Jun Zhang, *Persistent homology and Floer-Novikov theory*, Geom. Topol. **20** (2016), no. 6, 3333–3430, DOI 10.2140/gt.2016.20.3333. MR3590354
- [Vie27] L. Vietoris, *Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen* (German), Math. Ann. **97** (1927), no. 1, 454–472, DOI 10.1007/BF01447877. MR1512371
- [Vir20] Žiga Virk, *1-dimensional intrinsic persistence of geodesic spaces*, J. Topol. Anal. **12** (2020), no. 1, 169–207, DOI 10.1142/S1793525319500444. MR4080099
- [Vir21] Žiga Virk, *Rips complexes as nerves and a functorial Dowker-nerve diagram*, Mediterr. J. Math. **18** (2021), no. 2, Paper No. 58, 24, DOI 10.1007/s00009-021-01699-4. MR4218370
- [Vir22] Žiga Virk, *Footprints of geodesics in persistent homology*, Mediterr. J. Math. **19** (2022), no. 4, Paper No. 160, 29, DOI 10.1007/s00009-022-02089-0. MR4443111
- [Wil92] Frederick H. Wilhelm Jr., *On the filling radius of positively curved manifolds*, Invent. Math. **107** (1992), no. 3, 653–668, DOI 10.1007/BF01231906. MR1150606
- [Wol85] Franz-Erich Wolter, *Cut loci in bordered and unbordered Riemannian manifolds*, Ph.D. Thesis, Technische Universität Berlin, 1985.
- [Zar22] Matthew C. B. Zaremsky, *Bestvina-Brady discrete Morse theory and Vietoris-Rips complexes*, Amer. J. Math. **144** (2022), no. 5, 1177–1200. MR4494179
- [ZYCW20] Qi Zhao, Ze Ye, Chao Chen, and Yusu Wang, *Persistence enhanced graph neural network*, International Conference on Artificial Intelligence and Statistics, pp. 2896–2906, PMLR, 2020.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF LUXEMBOURG, LUXEMBOURG
Email address: `alexey.balitskiy@uni.lu`

DEPARTMENT OF MATHEMATICAL SCIENCES, UT DALLAS, DALLAS, TEXAS
Email address: `coskunuz@utdallas.edu`

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, NEW BRUNSWICK, NEW JERSEY
Email address: `facundo.memoli@rutgers.edu`