

Persistence stability for geometric complexes

Frédéric Chazal · Vin de Silva · Steve Oudot

Received: 10 February 2013 / Accepted: 19 November 2013
© Springer Science+Business Media Dordrecht 2013

Abstract In this paper we study the properties of the homology of different geometric filtered complexes (such as Vietoris–Rips, Čech and witness complexes) built on top of totally bounded metric spaces. Using recent developments in the theory of topological persistence, we provide simple and natural proofs of the stability of the persistent homology of such complexes with respect to the Gromov–Hausdorff distance. We also exhibit a few noteworthy properties of the homology of the Rips and Čech complexes built on top of compact spaces.

Keywords Persistent homology · Vietoris–Rips complex · Čech complex · Gromov–Hausdorff distance

Mathematics Subject Classification (2000) 51F99 · 55N99

1 Introduction

The inference of topological properties of metric spaces from approximations is a problem that has attracted special attention in computational topology in recent years. Given a metric

The authors gratefully acknowledge the following funding sources for this work: Digiteo project C3TTA (including the Digiteo chair held by the second author); European project CG-Learning (EC contract No. 255827); ANR project GIGA (ANR-09-BLAN-0331-01); DARPA project Sensor Topology and Minimal Planning ‘SToMP’ (HR0011-07-1-0002). The second author is a 2013 Simons Fellow, and is supported in part by the Institute for Mathematics and its Applications with funds provided by the National Science Foundation.

F. Chazal · S. Oudot
INRIA Saclay, Ile-de-France, Palaiseau, France
e-mail: frederic.chazal@inria.fr

S. Oudot
e-mail: steve.oudot@inria.fr

V. de Silva (✉)
Pomona College, Claremont, CA, USA
e-mail: vin.desilva@pomona.edu

space (Y, d_Y) approximating an unknown metric space (X, d_X) , the aim is to build a simplicial complex on the vertex set Y whose homology or homotopy type is the same as X . Note that, although Y is finite in many applications, finiteness is not a requirement a priori.

Among the many geometric complexes available to us, the Vietoris–Rips complex (or simply ‘Rips complex’) is particularly useful, being easy to compute and having good approximation properties. We recall the definition. Let (X, d_X) be a metric space and α a real parameter (the ‘scale’). Then $\text{Rips}(X, \alpha)$ is the simplicial complex on X whose simplices are the finite subsets of X with diameter at most α :

$$\sigma = [x_0, x_1, \dots, x_k] \in \text{Rips}(X, \alpha) \iff d_X(x_i, x_j) \leq \alpha \quad \text{for all } i, j$$

When (X, d_X) is a closed Riemannian manifold, Hausmann [15] proved that if $\alpha > 0$ is sufficiently small then the geometric realisation of $\text{Rips}(X, \alpha)$ is homotopy equivalent to X . This result was later generalised by Latschev [16], who proved that if (Y, d_Y) is sufficiently close to (X, d_X) in the Gromov–Hausdorff distance, then there exists $\alpha > 0$ such that $\text{Rips}(Y, \alpha)$ is homotopy equivalent to X . Recently, Attali et al. [1] adapted these results to a class of sufficiently regular compact subsets of Euclidean spaces. For larger classes of compact subsets of Riemannian manifolds, the homology and homotopy of such sets are known to be encoded in nested pairs of Vietoris–Rips complexes [6], yet it remains still open whether or not a single Rips complex can carry this topological information.

These approaches make it possible to recover the topology of a metric space (X, d_X) from a sufficiently close approximation (Y, d_Y) , provided that the parameter a is chosen correctly. Unfortunately, this choice very much depends on the geometry of X and can be difficult (if even possible) to determine in practical applications. One way round this issue is to use topological persistence [11, 18], which encodes the homology of the entire nested family $\mathbb{R}\text{ips}(X) = (\text{Rips}(X, \alpha))_{\alpha \in \mathbb{R}}$ in a single invariant, the persistence diagram. Relevant scales α can then be selected by the user, and the diagram provides an explicit relationship between the choice of a scale α and the homology of the corresponding Vietoris–Rips complex.

The stability of this construction was established by Chazal et al. [4], who proved that

$$d_b(\text{dgm}(\mathbb{R}\text{ips}(X)), \text{dgm}(\mathbb{R}\text{ips}(Y))) \leq 2 d_{\text{GH}}(X, Y), \tag{1.1}$$

for finite metric spaces (X, d_X) and (Y, d_Y) . Here d_b and d_{GH} denote the bottleneck [10, Chap. 8] and Gromov–Hausdorff distances, respectively. The bound turns out to be tight, which motivates the use of persistence diagrams as discriminative signatures to compare geometric shapes represented as finite metric spaces.

In this paper we show that the same inequality holds for all totally bounded metric spaces, and can in fact be extended to a larger class of filtered geometric complexes on such spaces. This includes a new family of examples called Dowker complexes. Our analysis adopts a new perspective, guided by recent developments in the theory of topological and algebraic persistence [7] which result in simple and natural proofs. Our contributions are the following:

- Extending the concept of simplicial map between complexes to the one of ε -simplicial multivalued map between filtered complexes, we show that such maps induce canonical ε -interleavings between the persistent homology modules of these complexes—Sect. 3.
- Applying this result to correspondences between metric spaces, we prove the ε -interleaving of the persistent homology modules of certain families of filtered geometric complexes (including Rips filtrations) built on top of ε -close metric spaces—Sect. 4.
- We prove the tameness of the persistent homology modules of the above filtered complexes when the vertex sets are totally bounded. Combined with the previous results, this result shows that inequality 1.1 can be generalised as claimed above —Sect. 5.1.

In addition to this consistent set of results, the Sect. 5.2 presents a few noteworthy properties of the homology groups of Rips and Čech complexes of totally bounded metric spaces. We finish in Sect. 6 with some results on the persistence diagrams of path-metric and δ -hyperbolic metric spaces.

Remark 1.1 Why total boundedness? Recall that a metric space is totally bounded if for every $\varepsilon > 0$ it admits a finite ε -sample. In other words, such a space is approximable at every resolution by a finite metric space. This explains the good behaviour of these spaces with respect to persistent homology; it is a manifestation of the good behaviour of persistent homology with respect to approximations.

2 Persistence modules and persistence diagrams

We adopt the approach and the notation of [7]. In this section we recall the definitions and results that we need. For a detailed presentation the reader is referred to [7].

A persistence module \mathbb{V} over the real numbers \mathbf{R} is an indexed family of vector spaces¹ $(V_a)_{a \in \mathbf{R}}$ together with a doubly-indexed family of linear maps $(v_a^b : V_a \rightarrow V_b \mid a \leq b)$ which satisfy the composition law $v_b^c \circ v_a^b = v_a^c$ whenever $a \leq b \leq c$, and where v_a^a is the identity map on V_a .

Example 2.1 (homology of a filtered complex) This is the standard example, which we use throughout this paper. Let \mathbb{S} be a filtered simplicial complex: that is, a family $(\mathbb{S}_a)_{a \in \mathbf{R}}$ of subcomplexes of some fixed simplicial complex $\bar{\mathbb{S}}$, such that $\mathbb{S}_a \subseteq \mathbb{S}_b$ whenever $a \leq b$. Let $V_a = H(\mathbb{S}_a)$ be the homology group² of \mathbb{S}_a , and let $v_a^b : H(\mathbb{S}_a) \rightarrow H(\mathbb{S}_b)$ be the linear map induced by the inclusion $\mathbb{S}_a \hookrightarrow \mathbb{S}_b$. Since, for any $a \leq b \leq c$, the inclusion $\mathbb{S}_a \hookrightarrow \mathbb{S}_c$ is the composition of the inclusions $\mathbb{S}_a \hookrightarrow \mathbb{S}_b$ and $\mathbb{S}_b \hookrightarrow \mathbb{S}_c$, it follows by functoriality that the linear maps satisfy $v_a^c = v_b^c \circ v_a^b$ and the family $(H(\mathbb{S}_a))_{a \in \mathbf{R}}$ is a persistence module.

Let \mathbb{U}, \mathbb{V} be persistence modules over \mathbf{R} , and let ε be any real number. A **homomorphism of degree ε** is a collection Φ of linear maps

$$(\phi_a : U_a \rightarrow V_{a+\varepsilon})_{a \in \mathbf{R}}$$

such that $v_{a+\varepsilon}^{b+\varepsilon} \circ \phi_a = \phi_b \circ v_a^b$ for all $a \leq b$. We write

$$\text{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}) = \{\text{homomorphisms } \mathbb{U} \rightarrow \mathbb{V} \text{ of degree } \varepsilon\}.$$

Composition is defined in the obvious way. For $\varepsilon \geq 0$, the most important degree- ε endomorphism is the shift map

$$1_{\mathbb{V}}^\varepsilon \in \text{Hom}^\varepsilon(\mathbb{V}, \mathbb{V})$$

defined to be the collection of maps $(v_a^{a+\varepsilon})$ from the persistence structure on \mathbb{V} . If Φ is a homomorphism $\mathbb{U} \rightarrow \mathbb{V}$ of any degree, then by definition $\Phi 1_{\mathbb{U}}^\varepsilon = 1_{\mathbb{V}}^\varepsilon \Phi$ for all $\varepsilon \geq 0$.

Example 2.2 (continuing Example 2.1) Given $\varepsilon > 0$, if $f : \bar{\mathbb{S}} \rightarrow \bar{\mathbb{S}}'$ is a simplicial map such that f maps \mathbb{S}_a to $\mathbb{S}'_{a+\varepsilon}$ for any $a \in \mathbf{R}$, then f induces a homomorphism of degree ε between the persistence modules $H(\mathbb{S})$ and $H(\mathbb{S}')$.

¹ All vector spaces are taken to be over an arbitrary field \mathbf{k} , fixed throughout this paper.

² We use simplicial homology with coefficients in the field \mathbf{k} .

Two persistence modules \mathbb{U}, \mathbb{V} are said to be ε -interleaved if there are maps

$$\Phi \in \text{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}), \quad \Psi \in \text{Hom}^\varepsilon(\mathbb{V}, \mathbb{U})$$

such that $\Psi\Phi = 1_{\mathbb{U}}^{2\varepsilon}$ and $\Phi\Psi = 1_{\mathbb{V}}^{2\varepsilon}$.

Following [4, 7] we say that a persistence module \mathbb{V} is **q-tame** if

$$\text{rank}(v_a^b) < \infty \quad \text{whenever } a < b.$$

This regularity condition ensures that persistence modules behave well:

Theorem 2.3 [7] *If \mathbb{U} is a q-tame module then it has a well-defined persistence diagram $\text{dgm}(\mathbb{U})$. If \mathbb{U}, \mathbb{V} are q-tame persistence modules that are ε -interleaved then there exists an ε -matching between the multisets $\text{dgm}(\mathbb{U}), \text{dgm}(\mathbb{V})$. Thus, the bottleneck distance between the diagrams satisfies the bound $d_b(\text{dgm}(\mathbb{U}), \text{dgm}(\mathbb{V})) \leq \varepsilon$.*

3 Multivalued maps

The notion of a simplicial map between simplicial complexes extends to the notion of an ε -simplicial map between filtered simplicial complexes in the following way:

Definition 3.1 Let \mathbb{S} and \mathbb{T} be two filtered simplicial complexes with vertex sets X and Y respectively. A map $f : X \rightarrow Y$ is ε -simplicial from \mathbb{S} to \mathbb{T} if it induces a simplicial map $\mathbb{S}_a \rightarrow \mathbb{T}_{a+\varepsilon}$ for every $a \in \mathbf{R}$. Equivalently, f is ε -simplicial if and only if for any $a \in \mathbf{R}$ and any simplex $\sigma \in \mathbb{S}_a$, $f(\sigma)$ is a simplex of $\mathbb{T}_{a+\varepsilon}$.

We wish to extend this concept to multivalued maps. Here are the basic notions.

A **multivalued map** $C : X \rightrightarrows Y$ from a set X to a set Y is a subset of $X \times Y$, also denoted C , that projects surjectively onto X through the canonical projection $\pi_X : X \times Y \rightarrow X$. The **image** $C(\sigma)$ of a subset σ of X is the canonical projection onto Y of the preimage of σ through π_X . A (single-valued) map f from X to Y is **subordinate** to C if we have $(x, f(x)) \in C$ for every $x \in X$; then we write $f : X \xrightarrow{C} Y$. The **composite** of two multivalued maps $C : X \rightrightarrows Y$ and $D : Y \rightrightarrows Z$ is the multivalued map $D \circ C : X \rightrightarrows Z$, defined by:

$$(x, z) \in D \circ C \iff \text{there exists } y \in Y \text{ such that } (x, y) \in C \text{ and } (y, z) \in D$$

The **transpose** of C , denoted C^T , is the image of C through the symmetry map $(x, y) \mapsto (y, x)$. Although C^T is well-defined as a subset of $Y \times X$, it is not always a multivalued map because it may not project surjectively onto Y .

Definition 3.2 Let \mathbb{S} and \mathbb{T} be two filtered simplicial complexes with vertex sets X and Y respectively. A multivalued map $C : X \rightrightarrows Y$ is ε -**simplicial** from \mathbb{S} to \mathbb{T} if for any $a \in \mathbf{R}$ and any simplex $\sigma \in \mathbb{S}_a$, every finite subset of $C(\sigma)$ is a simplex of $\mathbb{T}_{a+\varepsilon}$.

Proposition 3.3 *Let $C : X \rightrightarrows Y$ be an ε -simplicial multivalued map from \mathbb{S} to \mathbb{T} . Then C induces a canonical linear map $H(C) \in \text{Hom}^\varepsilon(H(\mathbb{S}), H(\mathbb{T}))$, equal to $H(f)$ for any f subordinate to C .*

Proof Any choice of f induces a simplicial map $\mathbb{S}_a \rightarrow \mathbb{T}_{a+\varepsilon}$ at each $a \in \mathbf{R}$, and these maps commute with the inclusions $\mathbb{S}_a \hookrightarrow \mathbb{S}_b, \mathbb{T}_{a+\varepsilon} \hookrightarrow \mathbb{T}_{b+\varepsilon}$ for all $a \leq b$. Thus f induces $H(f) \in \text{Hom}^\varepsilon(H(\mathbb{S}), H(\mathbb{T}))$. Any two subordinate maps $f_1, f_2 : X \xrightarrow{C} Y$ induce simplicial maps

$\mathbb{S}_a \rightarrow \mathbb{T}_{a+\varepsilon}$ which are contiguous. Indeed, for any $\sigma \in \mathbb{S}_a$ the two simplices $f_1(\sigma), f_2(\sigma)$ span a simplex of $\mathbb{T}_{a+\varepsilon}$ since their vertices comprise a finite subset of $C(\sigma)$. It follows [17, Theorems 12.4 and 12.5] that $H(f_1) = H(f_2)$. Thus the map $H(C)$ is uniquely defined. \square

Another immediate consequence is that the induced homomorphism is invariant under taking subsets of C that are also multivalued maps:

Proposition 3.4 *If $C' \subseteq C : X \rightrightarrows Y$ and C is ε -simplicial from \mathbb{S} to \mathbb{T} , then C' is ε -simplicial from \mathbb{S} to \mathbb{T} and $H(C') = H(C)$.*

Proof Since C' is a multivalued map contained in C , it is also ε -simplicial, and any map $f : X \rightarrow Y$ that is subordinate to C' is also subordinate to C , so we have $H(C') = H(C)$. \square

Finally, induced homomorphisms compose in the natural way:

Proposition 3.5 *Let $\mathbb{S}, \mathbb{T}, \mathbb{U}$ be filtered complexes with vertex sets X, Y, Z respectively. If*

$$C : X \rightrightarrows Y \text{ is a } \varepsilon\text{-simplicial multivalued map from } \mathbb{S} \text{ to } \mathbb{T},$$

$$D : Y \rightrightarrows Z \text{ is a } \delta\text{-simplicial multivalued map from } \mathbb{T} \text{ to } \mathbb{U},$$

then the composite $D \circ C : X \rightrightarrows Z$ is a $(\varepsilon + \delta)$ -simplicial multivalued map from \mathbb{S} to \mathbb{U} , and $H(D \circ C) = H(D) \circ H(C)$.

Proof $D \circ C$ is $(\varepsilon + \delta)$ -simplicial as an immediate consequence of the definition of ε -simplicial multivalued map. Let $f : X \xrightarrow{C} Y$ be subordinate to C , and let $g : Y \xrightarrow{D} Z$ be subordinate to D . The composite $g \circ f : X \xrightarrow{D \circ C} Z$ is subordinate to $D \circ C$, therefore $H(D \circ C) = H(D) \circ H(C)$. \square

4 Correspondences

4.1 Interleaving persistence modules of filtered complexes through correspondences

Definition 4.1 A multivalued map $C : X \rightrightarrows Y$ is a **correspondence** if the canonical projection $C \rightarrow Y$ is surjective, or equivalently, if C^T is also a multivalued map.

We immediately deduce, if C is a correspondence, that the identity maps $1_X = \{(x, x) : x \in X\}$ and $1_Y = \{(y, y) : y \in Y\}$ satisfy

$$1_X \subseteq C^T \circ C, \quad 1_Y \subseteq C \circ C^T.$$

From this property and Propositions 3.4 and 3.5, we deduce the following:

Proposition 4.2 *Let \mathbb{S}, \mathbb{T} be filtered complexes with vertex sets X, Y respectively. If $C : X \rightrightarrows Y$ is a correspondence such that C and C^T are both ε -simplicial, then together they induce a canonical ε -interleaving between $H(\mathbb{S})$ and $H(\mathbb{T})$, the interleaving homomorphisms being $H(C)$ and $H(C^T)$.*

4.2 Applications to filtered complexes on metric spaces

When (X, d_X) and (Y, d_Y) are metric spaces, the distortion of a correspondence $C : X \rightrightarrows Y$ is defined as follows:

$$\text{dis}(C) = \sup\{|d_X(x, x') - d_Y(y, y')| : (x, y), (x', y') \in C\}$$

The Gromov–Hausdorff distance [3, Theorem 7.3.25] between (X, d_X) and (Y, d_Y) is then defined by taking the infimum of the distortions among all the correspondences between X and Y :

$$d_{\text{GH}}(X, Y) = \frac{1}{2} \inf\{\text{dis}(C) : C \text{ is a correspondence } X \rightrightarrows Y\}$$

Although d_{GH} is not necessarily finite, it is a distance on the set of isometry classes of compact metric spaces: (i) it is zero if and only if the spaces are isometric; (ii) a correspondence and its transpose have the same distortion, so d_{GH} is symmetric; and (iii) the composite of two correspondences C, C' is a correspondence with distortion at most $\text{dis}(C') + \text{dis}(C)$, so d_{GH} satisfies the triangle inequality.

The theme of the next few examples is that low-distortion correspondences give rise to ε -simplicial maps on filtered complexes.

4.2.1 The Vietoris–Rips complex

Let (X, d_X) be a metric space. For $a \in \mathbf{R}$ we define a simplicial complex $\text{Rips}(X, a)$ on the vertex set X by the following condition:

$$\{x_0, x_1, \dots, x_k\} \in \text{Rips}(X, a) \Leftrightarrow d_X(x_i, x_j) \leq a \text{ for all } i, j$$

For $a \leq 0$, note that $\text{Rips}(X, a)$ consists of the vertex set X alone. There is a natural inclusion $\text{Rips}(X, a) \subseteq \text{Rips}(X, b)$ whenever $a \leq b$. Thus, the simplicial complexes $\text{Rips}(X, a)$ together with these inclusion maps define a filtered simplicial complex $\mathbb{R}\text{ips}(X)$ on X , the **Vietoris–Rips complex**.

Lemma 4.3 (Vietoris–Rips interleaving) *Let $(X, d_X), (Y, d_Y)$ be metric spaces. For any $\varepsilon > 2d_{\text{GH}}(X, Y)$ the persistence modules $\mathbb{H}(\mathbb{R}\text{ips}(X))$ and $\mathbb{H}(\mathbb{R}\text{ips}(Y))$ are ε -interleaved.*

Proof Let $C : X \rightrightarrows Y$ be a correspondence with distortion at most ε .

If $\sigma \in \text{Rips}(X, a)$ then $d_X(x, x') \leq a$ for all $x, x' \in \sigma$. Let $\tau \subseteq C(\sigma)$ be any finite subset. For any $y, y' \in \tau$ there exist $x, x' \in \sigma$ such that $y \in C(x), y' \in C(x')$, and therefore:

$$d_Y(y, y') \leq d_X(x, x') \leq a + \varepsilon$$

It follows that $\tau \in \text{Rips}(Y, a + \varepsilon)$.

We have shown that C is ε -simplicial from $\mathbb{R}\text{ips}(X)$ to $\mathbb{R}\text{ips}(Y)$. Symmetrically, C^T is ε -simplicial from $\mathbb{R}\text{ips}(Y)$ to $\mathbb{R}\text{ips}(X)$. The result now follows from Proposition 4.2. \square

4.2.2 The intrinsic Čech complex

Let (X, d_X) be a metric space. For $a \in \mathbf{R}$ we define a simplicial complex $\check{\text{Cech}}(X, a)$ on the vertex set X by the following condition:

$$\{x_0, x_1, \dots, x_k\} \in \check{\text{Cech}}(X, a) \Leftrightarrow \bigcap_{i=0}^k B(x_i, a) \neq \emptyset$$

Here $B(x, a) = \{x' \in X : d_X(x, x') \leq a\}$ denotes the closed ball with centre $x \in X$ and radius a . Any point \bar{x} in the intersection $\bigcap_i B(x_i, a)$ is called an a -**centre** for the simplex $[x_0, \dots, x_k]$.

For $a \leq 0$, note that $\check{C}ech(X, a)$ consists of the vertex set X alone. There is a natural inclusion $\check{C}ech(X, a) \subseteq \check{C}ech(X, b)$ whenever $a \leq b$. Thus, the simplicial complexes $\check{C}ech(X, a)$ together with these inclusion maps define a filtered simplicial complex $\check{C}ech(X)$ on X , the **(intrinsic) Čech complex**.³

Lemma 4.4 (Čech interleaving) *Let $(X, d_X), (Y, d_Y)$ be metric spaces. For any $\varepsilon > 2d_{GH}(X, Y)$ the persistence modules $H(\check{C}ech(X))$ and $H(\check{C}ech(Y))$ are ε -interleaved.*

Proof Let $C : X \rightrightarrows Y$ be a correspondence with distortion at most ε .

Consider $\sigma \in \check{C}ech(X, a)$. Let \bar{x} be an a -centre for σ , so $d_X(\bar{x}, x) \leq a$ for all $x \in \sigma$. Pick $\bar{y} \in C(\bar{x})$. Now for any $y \in C(\sigma)$ we have $y \in C(x)$ for some $x \in \sigma$, and therefore:

$$d_Y(\bar{y}, y) \leq d_X(\bar{x}, x) + \varepsilon \leq a + \varepsilon$$

Let $\tau \subseteq C(\sigma)$ be any finite subset; then \bar{y} is an $(a + \varepsilon)$ -centre for τ and hence $\tau \in \check{C}ech(Y, a + \varepsilon)$.

We have shown that C is ε -simplicial from $\check{C}ech(X)$ to $\check{C}ech(Y)$. Symmetrically, C^T is ε -simplicial from $\check{C}ech(Y)$ to $\check{C}ech(X)$. The result now follows from Proposition 4.2. \square

4.2.3 Ambient Čech complexes and Dowker complexes

The reader should be aware that there are two distinct uses of the phrase ‘Čech complex’. The first is the intrinsic Čech complex described in the preceding subsection, that is constructed from a single metric space. The second, very commonly used in topological data analysis, is built from a pair or triple of spaces. These ‘ambient’ Čech complexes belong to a much more general family, the Dowker complexes. We consider these now.

Let $X \subseteq \mathbf{R}^n$. If X is assumed to be sampled from some unknown object, we can attempt to recover the structure of the object by thickening each point to a closed ball of radius a , say. By the Nerve Lemma [14, Section 4.G], the homotopy type of the thickened set can be retrieved by constructing the nerve of the collection of balls. This has vertex set X and a simplex for every finite subset of X for which the corresponding balls have nonempty intersection in \mathbf{R}^n .

We write $\check{C}ech(X, \mathbf{R}^n; a)$ for this nerve, and $\check{C}ech(X, \mathbf{R}^n)$ for the filtered complex obtained by varying a . The second argument \mathbf{R}^n is often omitted in certain literatures, being regarded as implicit. For us, however, $\check{C}ech(X)$ refers to the intrinsic Čech complex, where a simplex is included only if the corresponding balls meet at a point in X itself.

In general, an **ambient Čech complex** is defined as follows. Let L, W be subsets (‘landmarks’ and ‘witnesses’) of an unnamed metric space. For $a \in \mathbf{R}$, consider the complex with vertices L and simplices determined by:

$$\sigma \in \check{C}ech(L, W; a) \Leftrightarrow \exists w \in W \text{ such that } d(w, l) \leq a \text{ for all } l \in \sigma$$

The resulting filtered complex is denoted $\check{C}ech(L, W)$.

Example 4.5 The intrinsic Čech complex $\check{C}ech(X)$ for a metric space X is equal to $\check{C}ech(X, X)$, where the ambient space is X itself.

³ We will usually drop the word ‘intrinsic’ unless we are contrasting it with ‘ambient’.

More generally, a **Dowker complex** is defined as follows. Let L, W be two sets and let $\Lambda : L \times W \rightarrow \mathbf{R}$ be any function at all. For $a \in \mathbf{R}$, consider the complex with vertices L and simplices determined by:

$$\sigma \in \text{Dow}(\Lambda, a) \Leftrightarrow \exists w \in W \text{ such that } \Lambda(l, w) \leq a \text{ for all } l \in \sigma$$

The resulting filtered complex is denoted $\mathbb{D}\text{ow}(\Lambda)$.

Example 4.6 The intrinsic Čech complex $\check{\text{Cech}}(X)$ for a metric space X is equal to $\mathbb{D}\text{ow}(d_X)$, where $d_X : X \times X \rightarrow \mathbf{R}$ is the metric.

Example 4.7 The ambient Čech complex $\check{\text{Cech}}(L, W)$ for a pair of subsets of a metric space is equal to $\mathbb{D}\text{ow}(d|_{L \times W})$, where d is the ambient metric.

Remark 4.8 We name this complex in honour of C. H. Dowker [8], who compared two simplicial complexes constructed from a binary relation. Dowker’s theorem implies that $\text{Dow}(\Lambda, a)$ and $\text{Dow}(\Lambda^T, a)$ have the same homotopy type, where

$$\Lambda^T : W \times L \rightarrow \mathbf{R}; (w, l) \mapsto \Lambda(l, w)$$

is the ‘transpose’ of Λ (thus changing the vertex set to W). This is essentially an instance of the Nerve Lemma, since each complex can be interpreted as the nerve of a suitable covering of the other. The function Λ can be thought of as a filtered binary relation. Since the Nerve Lemma is functorial (i.e. respects maps) [6, Lemma 3.4] we obtain the stronger conclusion that $\mathbb{D}\text{ow}(\Lambda)$ and $\mathbb{D}\text{ow}(\Lambda^T)$ have the same filtered homotopy type. It follows that they have the same persistent homology, and therefore, where defined, the same persistence diagrams. We call this phenomenon **Dowker duality**.

Two sets of data (L, W, Λ) and (L', W', Λ') may be compared using a pair of correspondences $C : L \rightrightarrows L'$ and $D : W \rightrightarrows W'$. We define the distortion for such a pair to be:

$$\text{dis}(C, D) = \sup_{(l, l') \in C} \sup_{(w, w') \in D} |\Lambda(l, w) - \Lambda'(l', w')|$$

Lemma 4.9 (Dowker interleaving) *Let L, L', W, W' be sets with functions $\Lambda : L \times W \rightarrow \mathbf{R}$ and $\Lambda' : L' \times W' \rightarrow \mathbf{R}$. If $C : L \rightrightarrows L'$ and $D : W \rightrightarrows W'$ are correspondences and $\varepsilon \geq \text{dis}(C, D)$ then the persistence modules $\text{H}(\mathbb{D}\text{ow}(\Lambda))$ and $\text{H}(\mathbb{D}\text{ow}(\Lambda'))$ are ε -interleaved.*

Proof Consider $\sigma \in \text{Dow}(\Lambda, a)$. Let w be an a -centre for σ , so $\Lambda(l, w) \leq a$ for all $l \in \sigma$. Pick $w' \in C(w)$. For any $l' \in C(\sigma)$ we have $l' \in C(l)$ for some $l \in \sigma$, so:

$$\Lambda'(l', w') \leq \Lambda(l, w) + \varepsilon \leq a + \varepsilon$$

It follows that each finite $\sigma' \subseteq C(\sigma)$ belongs to $\text{Dow}(\Lambda', a + \varepsilon)$.

We have shown that C is ε -simplicial from $\mathbb{D}\text{ow}(\Lambda)$ to $\mathbb{D}\text{ow}(\Lambda')$. Symmetrically, C^T is ε -simplicial from $\mathbb{D}\text{ow}(\Lambda')$ to $\mathbb{D}\text{ow}(\Lambda)$. The result now follows from Proposition 4.2. \square

Let d_H denote the Hausdorff distance between subsets of a metric space.

Corollary 4.10 (ambient Čech interleaving) *Let L, L' and W be subsets of a metric space. For any $\varepsilon > d_H(L, L')$ the ambient Čech persistence modules $\text{H}(\check{\text{Cech}}(L, W))$ and $\text{H}(\check{\text{Cech}}(L', W))$ are ε -interleaved.*

Proof We regard the complexes as $\mathbb{D}\text{ow}(\Lambda)$, $\mathbb{D}\text{ow}(\Lambda')$, where $\Lambda = d|_{L \times W}$ and $\Lambda' = d|_{L' \times W}$. Since $\varepsilon > d_{\mathbb{H}}(L, L')$ the sets

$$C = \{(l, l') : l \in L, l' \in L' \text{ and } d(l, l') < \varepsilon\},$$

$$D = \{(w, w) : w \in W\}$$

are correspondences, and $\text{dis}(C, D) \leq \varepsilon$. It follows from Lemma 4.9 that the two persistence modules are ε -interleaved. \square

It is worth pointing out that Corollary 4.10 is well known in the special case $L, L' \subseteq W = \mathbf{R}^n$. The usual argument is based on the Nerve Lemma, so it relies on the local topological properties of Euclidean space and does not work in general. The elementary proof here shows that the dependence on the Nerve Lemma is unnecessary.

4.2.4 The witness complex

Let L, W be two sets ('landmarks' and 'witnesses') and let $\Lambda : L \times W \rightarrow \mathbf{R}$ be any function. For any finite subset $\sigma \subseteq L$, and any $w \in W$ and $a \in \mathbf{R}$, we say that w is an a -witness for the simplex σ iff

$$\Lambda(l, w) \leq \Lambda(l', w) + a \quad \text{for all } l \in \sigma \text{ and } l' \in L \setminus \sigma.$$

Given L, W and Λ , we can then define for any $a \in \mathbf{R}$ a simplicial complex $\text{Wit}(L, W; a)$ by

$$\sigma \in \text{Wit}(L, W; a) \Leftrightarrow \forall \tau \subseteq \sigma, \exists w \in W \text{ such that } w \text{ is an } a\text{-witness for } \tau.$$

An a -witness is obviously a b -witness for any $b \geq a$, so there is a natural inclusion $\text{Wit}(L, W; a) \subseteq \text{Wit}(L, W; b)$. The simplicial complexes $\text{Wit}(L, W; a)$ together with these inclusion maps define a filtered simplicial complex $\mathbb{W}\text{it}(L, W)$ with vertex set L , called the **witness complex filtration**.

Remark 4.11 We alert the reader that the witness complex defined here has nontrivial behaviour for $a < 0$, unlike the Vietoris–Rips and Čech complexes. This can be suppressed if necessary.

We now show that the witness complex filtration is stable with respect to varying the witness set while keeping the landmark set fixed. Let L be a set and let W, W' be witness sets for L with respect to maps $\Lambda : L \times W \rightarrow \mathbf{R}$ and $\Lambda' : L \times W' \rightarrow \mathbf{R}$. The distortion of a correspondence $C : W \rightrightarrows W'$ is defined:

$$\text{dis}(C) = \sup_{l \in L} \sup_{(w, w') \in C} |\Lambda(l, w) - \Lambda'(l, w')|$$

Lemma 4.12 (witness complex interleaving) *Let L be a set, and let W, W' be two witness sets of L with respect to maps $\Lambda : L \times W \rightarrow \mathbf{R}$ and $\Lambda' : L \times W' \rightarrow \mathbf{R}$. If $C : W \rightrightarrows W'$ is a correspondence and $\varepsilon \geq 2 \text{dis}(C)$ then the persistence modules $\mathbb{H}(\mathbb{W}\text{it}(L, W))$ and $\mathbb{H}(\mathbb{W}\text{it}(L, W'))$ are ε -interleaved.*

Proof Let $\sigma \in \text{Wit}(L, W; a)$. For every $\tau \subseteq \sigma$ we argue as follows. Let $w \in W$ be an a -witness for τ , and select $w' \in C(w)$. For all $l \in \tau$ and $l' \in L \setminus \tau$ we have

$$\Lambda'(l, w') \leq \Lambda(l, w) + \frac{1}{2}\varepsilon \leq \Lambda(l', w) + a + \frac{1}{2}\varepsilon \leq \Lambda'(l', w') + a + \varepsilon$$

so $w' \in W'$ is an $(a + \varepsilon)$ -witness for τ . It follows that $\sigma \in \text{Wit}(L, W'; a + \varepsilon)$.

Thus, the identity 1_L is an ε -simplicial map $\mathbb{W}it(L, W) \rightarrow \mathbb{W}it(L, W')$, and $1_L^T = 1_L$ is an ε -simplicial map $\mathbb{W}it(L, W') \rightarrow \mathbb{W}it(L, W)$ by symmetry. We conclude that $H(\mathbb{W}it(L, W))$ and $H(\mathbb{W}it(L, W'))$ are ε -interleaved. \square

The most common form of witness complex takes L, W to be subsets of a metric space with $\Lambda = d|_{L \times W}$ restricted from the ambient metric. Different witness sets may be compared using the Hausdorff distance d_H .

Corollary 4.13 *Let L, W, W' be subsets of a metric space, where W, W' are witness sets for L with respect to $\Lambda = d|_{L \times W}$ and $\Lambda' = d|_{L \times W'}$. For any $\varepsilon > 2d_H(W, W')$ the persistence modules $H(\mathbb{W}it(L, W))$ and $H(\mathbb{W}it(L, W'))$ are ε -interleaved.*

Proof Since $d_H(W, W') < \frac{1}{2}\varepsilon$, the set

$$C = \{(w, w') \in W \times W' : d_X(w, w') < \frac{1}{2}\varepsilon\}$$

is a correspondence with $\text{dis}(C) \leq \frac{1}{2}\varepsilon$. Now apply Lemma 4.12. \square

Unfortunately, in full generality there is no equivalent of Lemma 4.12 in the case where the set L is perturbed, even if the set of witnesses is constrained to stay fixed ($W = W'$). In contrast to ambient Čech complexes, the vertices in a witness complex interfere with one another. Here is an explicit counterexample:

Example 4.14 On the real line, consider the sets $W = L = \{0, 1\}$ and $L' = \{-\delta, 0, 1, 1 + \delta\}$, where $\delta \in (0, 1/2)$ is arbitrary. Then

$$\mathbb{W}it(L, W; a) = \{[0], [1], [0, 1]\}$$

for all $a \geq 0$, whereas

$$\mathbb{W}it(L', W; a) = \{[-\delta], [0], [1], [1 + \delta], [-\delta, 0], [1, 1 + \delta]\}$$

for all $a \in [\delta, 1 - \delta)$. Thus, $H(\mathbb{W}it(L, W))$ and $H(\mathbb{W}it(L', W))$ are not ε -interleaved for any $\varepsilon < 1 - 2\delta$, whereas $d_H(L, L') = \delta$ can be made arbitrarily small compared to $1 - 2\delta$.

Note that the set of witnesses in this example is fairly sparse compared to the set of landmarks. This raises several interesting questions, such as whether densifying W (e.g. taking the full real line) would allow to regain some stability. These questions lie beyond the scope of the paper.

4.2.5 Generalisation to dissimilarity spaces

In data analysis one often considers data sets X equipped with a dissimilarity measure, i.e. a map $\tilde{d}_X : X \times X \rightarrow \mathbf{R}$ that satisfies $\tilde{d}_X(x, x) \leq \tilde{d}_X(x, y) = \tilde{d}_X(y, x)$ for all $x, y \in X$ but is not required to satisfy any of the other metric space axioms. It is easily seen that the definitions for Vietoris–Rips, Čech and witness complexes continue to make sense for such spaces, and that the distortion of a correspondence $C : X \rightrightarrows Y$ is well-defined. Moreover, since the proofs of our interleaving results do not make use of any other distance axiom (triangle inequality, non-negativity, zero property), they remain valid in this more general context.

5 Regularity of Rips and Čech filtrations

5.1 Stability of Rips and Čech persistence for totally bounded spaces

The stability theorem for persistent homology is often expressed in terms of persistence diagrams. In this section we show that the Vietoris–Rips and Čech complexes of a totally bounded metric space have sufficiently tame persistent homology that their persistence diagrams are well defined. There is a similar result for Dowker complexes. The interleaving results of the previous section immediately imply a stability theorem for the persistence diagrams.

We recall the definitions. Given a positive real number $\varepsilon > 0$, a subset $F \subseteq X$ of a metric space (X, d_X) is an ε -sample of X if for any $x \in X$ there exists $f \in F$ such that $d_X(x, f) < \varepsilon$. A metric space (X, d_X) is **totally bounded** if it has a finite ε -sample for every $\varepsilon > 0$. Bounded subsets of Euclidean space are totally bounded. In general a metric space is totally bounded if and only if its completion is compact.

Proposition 5.1 *If (X, d_X) is a totally bounded metric space then the persistence modules $H(\mathbb{R}\text{rips}(X))$ and $H(\check{\text{Cech}}(X))$ are q -tame.*

Proof Let us first consider the case of the Vietoris–Rips persistence module. We must show that the map $I_a^b : H(\text{Rips}(X, a)) \rightarrow H(\text{Rips}(X, b))$ induced by the inclusion has finite rank whenever $a < b$. Let $\varepsilon = (b - a)/2$. Since X is totally bounded there exists a finite $\frac{1}{2}\varepsilon$ -sample F of X . The set $C = \{(x, f) \in X \times F : d_X(x, f) < \frac{1}{2}\varepsilon\}$ is an ε -correspondence, so the Gromov–Hausdorff distance between F and X is upper-bounded by $\frac{1}{2}\varepsilon$. It follows from Lemma 4.3 that there exists an ε -interleaving between $H(\mathbb{R}\text{rips}(X))$ and $H(\mathbb{R}\text{rips}(F))$. Using the interleaving maps, I_a^b factorises as

$$H(\text{Rips}(X, a)) \rightarrow H(\text{Rips}(F, a + \varepsilon)) \rightarrow H(\text{Rips}(X, a + 2\varepsilon)) = H(\text{Rips}(X, b)).$$

The second term is finite dimensional since $\text{Rips}(F; a + \varepsilon)$ is a finite simplicial complex, so I_a^b has finite rank.

The proof for the Čech persistence module is the same. □

The above proposition implies that the persistence diagrams of $H(\mathbb{R}\text{rips}(X))$ and $H(\check{\text{Cech}}(X))$ are well-defined for totally bounded metric spaces. We may now apply the persistence stability theorem to get the following result, which relates the Gromov–Hausdorff distance between two spaces to the bottleneck distance between the persistence diagrams of their Vietoris–Rips and Čech filtrations.

Theorem 5.2 *Let X, Y be totally bounded metric spaces. Then*

$$\begin{aligned} d_b(\text{dgm}(H(\mathbb{R}\text{rips}(X))), \text{dgm}(H(\mathbb{R}\text{rips}(Y)))) &\leq 2d_{\text{GH}}(X, Y), \\ d_b(\text{dgm}(H(\check{\text{Cech}}(X))), \text{dgm}(H(\check{\text{Cech}}(Y)))) &\leq 2d_{\text{GH}}(X, Y), \end{aligned}$$

Proof This is a consequence of Theorem 2.3 and Lemmas 4.3 and 4.4. □

Remark 5.3 The first inequality of Theorem 5.2 was earlier proved in [5] in the special case of finite metric spaces, using a different approach based on embedding the spaces into l^∞ and invoking the functorial Nerve Lemma.

Here are the corresponding results for ambient Čech complexes.

Proposition 5.4 *Let L, W be subsets of a metric space. If at least one of L, W is totally bounded, then the ambient Čech persistence $H(\check{\text{Cech}}(L, W))$ is q -tame.*

Proof We assume that L is totally bounded. The case where W is totally bounded follows by Dowker duality (Remark 4.8). It is enough to show for every $\varepsilon > 0$ that $H(\check{Cech}(L, W))$ is ε -interleaved with the persistent homology of a finite complex. To do this, let F be a finite ε -sample of L , so that $d_H(L, F) \leq \varepsilon$. Then $H(\check{Cech}(L, W))$ and $H(\check{Cech}(F, W))$ are ε -interleaved by Corollary 4.10, and $\check{Cech}(F, W)$ is finite as required. \square

Remark 5.5 The hypothesis in the proposition is most easily checked when the ambient metric space is a ‘proper’ space, meaning that its closed balls are compact. (For instance, \mathbf{R}^n is proper.) Then a subset L is totally bounded if and only if it is bounded.

Proposition 5.4 implies that the persistence diagram $\text{dgm}(H(\check{Cech}(L, W)))$ is well defined when at least one of L, W is totally bounded.

Theorem 5.6 *Let L, L' and W be subsets of a metric space. Suppose L, L' are totally bounded, or that W is totally bounded. Then*

$$d_b(\text{dgm}(H(\check{Cech}(L, W))), \text{dgm}(H(\check{Cech}(L', W)))) \leq d_H(L, L')$$

Proof This is a consequence of Theorem 2.3 and Corollary 4.10. \square

We finish with a tameness result for Dowker complexes.

Proposition 5.7 *Let L, W be sets and $\Lambda : L \times W \rightarrow \mathbf{R}$ be a function. Suppose the collection $(\lambda_l)_{l \in L}$ of functions $\lambda_l(w) = \Lambda(l, w)$ is bounded and totally bounded with respect to the supremum norm on functions $W \rightarrow \mathbf{R}$. Then $H(\mathbb{Dow}(\Lambda))$ is q -tame.*

Remark 5.8 Dowker duality implies that the same conclusion holds if the roles of L, W are interchanged in the hypothesis.

Proof As before, it is enough to show that for any $\varepsilon > 0$ the persistence module $H(\mathbb{Dow}(\Lambda))$ is ε -interleaved with the persistent homology of a finite complex. To do this, let F be a finite subset of L such that $(\lambda_l)_{l \in F}$ is an ε -sample of $(\lambda_l)_{l \in L}$, and let Λ_F be the restriction of Λ to $W \times F$. Then

$$C = \{(l, l') \mid l \in L, l' \in F \text{ and } \|\lambda_l - \lambda_{l'}\|_\infty < \varepsilon\}$$

$$D = \{(w, w) \mid w \in W\}$$

are correspondences $L \rightrightarrows F$ and $W \rightrightarrows W$, respectively, with $\text{dis}(C, D) \leq \varepsilon$. It follows from Lemma 4.9 that $H(\mathbb{Dow}(\Lambda))$ and $H(\mathbb{Dow}(\Lambda_F))$ are ε -interleaved, with $\mathbb{Dow}(\Lambda_F)$ being finite as required. \square

One can most straightforwardly use Proposition 5.7 in situations where the Arzelà–Ascoli theorem guarantees total boundedness. For instance, if L, W are totally bounded metric spaces and Λ is Lipschitz then $H(\mathbb{Dow}(\Lambda))$ is q -tame.

5.2 Non-persistent homology of Rips and Čech complexes

The good behaviour of Rips and Čech filtrations on compact metric spaces in their persistent homology stands in marked contrast to bad behaviour that can be found in the homology groups at particular parameter values. Whereas the persistence modules $H(\check{Cech}(X))$ and $H(\text{Rips}(X))$ are q -tame when X is a totally bounded metric space, the individual homology groups $H(\check{Cech}(X, a))$ and $H(\text{Rips}(X, a))$ may well be infinite dimensional for some or many values of a .

In the next few sections we present both positive and negative results in this direction.

We briskly remark that homology in dimension zero is easily handled when X is totally bounded: a generating set for both $H_0(\check{C}ech(X, a))$ and $H_0(Rips(X, a))$ is provided by any a -sample of X , so these vector spaces are finite-dimensional when $a > 0$.

5.2.1 The homology groups of a Rips filtration

It is easy to construct an example of a compact metric space X such that the homology group $H_1(Rips(X, 1))$ has an uncountable infinite dimension. For example consider the union X of two parallel segments in \mathbf{R}^2 defined by

$$X = \{(t, 0) \in \mathbf{R}^2 : t \in [0, 1]\} \cup \{(t, 1) \in \mathbf{R}^2 : t \in [0, 1]\}$$

with metric restricted from the Euclidean metric in \mathbf{R}^2 . Then for any $t \in [0, 1]$, the edge $e_t = [(t, 0), (t, 1)]$ belongs to $Rips(X, 1)$ but there is no triangle in $Rips(X, 1)$ that contains e_t in its boundary. As a consequence, for $t \in (0, 1]$ the cycles $\gamma_t = [(0, 0), (t, 0)] + e_t + [(t, 1), (0, 1)] - e_0$ are linearly independent in $H_1(Rips(X, 1))$.

Here, $a = 1$ is the only value of the Rips parameter for which the homology group $H_1(Rips(X, a))$ fails to be finite-dimensional. In fact, it is possible to construct examples where the set of ‘bad’ values is arbitrarily large.

Proposition 5.9 *For any $\alpha, \beta \in \mathbf{R}$ such that $0 < \alpha \leq \beta$ and any integer k there exists a compact metric space X such that for any $a \in [\alpha, \beta]$, $H_k(Rips(X, a))$ has an uncountable infinite dimension.*

Proof The following example was obtained with the help of J.-M. Droz who also proved that a similar example can be realised as a subset of \mathbf{R}^4 endowed with the Euclidean metric [9].

Without loss of generality, we can assume that $\alpha = 1$ and $\beta = 2$. Let us first consider the case $k = 1$. Consider the union X of two non-parallel rectangles in \mathbf{R}^3 , defined as

$$X = R_1 \cup R_2 = \{(t, 0, z) \in \mathbf{R}^3 : t \in [0, 2], z \in [0, 1]\} \\ \cup \{(t, 1 + \frac{1}{2}t, z) \in \mathbf{R}^3 : t \in [0, 2], z \in [0, 1]\}$$

and endowed with the restriction of the ℓ^1 -norm in \mathbf{R}^3 (see Fig. 1).

Since we are using the ℓ^1 -norm, for $a \in [1, 2]$ and $z \in [0, 1]$, the point $(2(a-1), 0, z) \in R_1$ is at distance a from R_2 and its unique closest point on R_2 is $(2(a-1), a, z)$. As a consequence $e_z = [(2(a-1), 0, z), (2(a-1), a, z)]$ is an edge of $Rips(X, a)$ for all $z \in [0, 1]$ but there is no (non degenerate) triangle in $Rips(X, a)$ that contains e_z in its boundary. Therefore, for $z \in (0, 1]$, the cycles $\gamma_z = [(2(a-1), 0, 0), (2(a-1), 0, z)] + e_z + [(2(a-1), a, z), (2(a-1), a, 0)] - e_0$ are not homologous to 0 and are linearly independent in $H_1(Rips(X, a))$.

To prove the result for $k > 1$, just consider the product of X with a $(k-1)$ -dimensional sphere of sufficiently large radius (to prevent the Rips construction from killing the $(k-1)$ -homology), and apply the Künneth formula [14, Theorem 3.16, p. 219]. □

5.2.2 The open Vietoris–Rips filtration

The examples given above, of Vietoris–Rips complexes with infinite-dimensional homology, rely strongly on the fact that $Rips(X, a)$ is defined using a non-strict inequality; that is to say, a closed condition: $[x_0, x_1, \dots, x_k] \in Rips(X, a)$ if and only if $d_X(x_i, x_j) \leq a$ for all i, j .




Fig. 1 The union X of the above 2 rectangles endowed with the restriction of the L^1 metric in \mathbf{R}^3 provides an example of a compact metric space such that $\dim H_1(\text{Rips}(X, a)) = \infty$ for any $a \in [1, 2]$

It is natural to ask what happens if strict inequality—an open condition—is used. Given a metric space (X, d_X) and a real number $a \in \mathbf{R}$, the open Vietoris–Rips complex is the simplicial complex $\text{Rips}(X, a^-)$ with vertex set X defined by the following condition:

$$[x_0, x_1, \dots, x_k] \in \text{Rips}(X, a^-) \Leftrightarrow d_X(x_i, x_j) < a, \text{ for all } i, j.$$

The reader may easily confirm that the examples in Sect. 5.2.1 dissolve when the open condition is used. The existence of other constructions is constrained by the following mild regularity result.

Proposition 5.10 *For any totally bounded metric space X and real number $a > 0$, the total homology $H(\text{Rips}(X, a^-))$ has a countable basis.*

Proof Any homology class in $H_k(\text{Rips}(X, a^-))$ is represented by a cycle, which by definition is a finite linear combination of simplices of diameter strictly less than a and therefore less than some $a - \frac{1}{n}$. It follows that the class lies in the image of $H_k(\text{Rips}(X, a - \frac{1}{n})) \rightarrow H_k(\text{Rips}(X, a^-))$. Since $H_k(\mathbb{R}\text{ips}(X))$ is q -tame, by Proposition 5.1, this image is finite dimensional. Since $H_k(\text{Rips}(X, a^-))$ is the union of these finite dimensional images for $n \rightarrow +\infty$, it has a countable basis. The full result follows by summing over k . \square

We cannot guarantee finite dimensionality. However, Proposition 5.10 suggests that we must proceed discretely if we are to find a counterexample.

Proposition 5.11 *For any given $a > 0$ there exists a totally bounded metric space X such that $H_1(\text{Rips}(X, a^-))$ has infinite dimension.*

Proof We may assume that $a = 1$. We will construct a bounded subset $X \subset \mathbf{R}^2$ whose open Vietoris–Rips complex $\text{Rips}(X, 1^-)$ has infinite-dimensional 1-dimensional homology. We will construct it as the union of two infinite sets L and R such that $\text{Rips}(X, 1^-)$ contains the complete graphs on L and R , and otherwise each vertex in L shares an edge with precisely one vertex in R , and vice versa. Following the same argument as the one used for the example before Proposition 5.9 we will deduce that $H_1(\text{Rips}(X, 1^-))$ is infinite dimensional.

We define X in terms of an auxiliary function $f(x) : [0, \infty) \rightarrow [0, \infty)$, which will be identified later. We suppose initially that f is continuous, increasing, and positive except at $f(0) = 0$. Specifically, let

$$L_x = (f(x), x), \quad R_x = (1 - f(x), x)$$

for $0 \leq x < f^{-1}(1/2)$. We define

$$X = \{L_x, R_x \mid x = \varepsilon_n, n \gg 0\}$$

where (ε_n) is a decreasing positive sequence with limit 0.

Clearly $|L_x - R_x| < 1$ for all $x > 0$. We must arrange that $|L_x - R_y| \geq 1$, for x, y distinct elements of the sequence (ε_n) . Suppose $x > y$. Then:

$$\begin{aligned} |L_x - R_y|^2 &= (1 - f(x) - f(y))^2 + (x - y)^2 \\ &\geq (1 - 2f(x))^2 + (x - y)^2 \\ &\geq 1 - 4f(x) + (x - y)^2 \end{aligned}$$

Suppose we have chosen our sequence so that $x > y$ implies $x \geq 2y$; for instance, by setting $(\varepsilon_n) = (2^{-n})$. Then $(x - y)^2 \geq \frac{1}{4}x^2$. Now we choose $f(x) = \frac{1}{16}x^2$. For $\sqrt{32} > x > y$ in the sequence (2^{-n}) we have

$$|L_x - R_y|^2 \geq 1 - 4f(x) + (x - y)^2 \geq 1 - \frac{1}{4}x^2 + \frac{1}{4}x^2 = 1$$

as required.

It follows that if we define

$$X = L \cup R = \{(2^{-2n-4}, 2^{-n}) \mid n \geq 1\} \cup \{(1 - 2^{-2n-4}, 2^{-n}) \mid n \geq 1\}$$

then $\text{Rips}(X, 1^-)$ contains the complete graphs on L and R , and otherwise each vertex in L shares an edge with precisely one vertex in R , and vice versa. As a consequence there is no non-degenerate triangle in $\text{Rips}(X, 1^-)$ that contains one of these edges connecting L to R . For each point $x_n = (2^{-2n-4}, 2^{-n}) \in L$, let $y_n \in R$ be the corresponding point in R , so the edge $e_n = [x_n, y_n]$ is in $\text{Rips}(X, 1^-)$. It follows that the cycles $\gamma_n = e_1 + [y_1, y_n] - e_n + [x_n, x_1]$, $n \geq 2$, are not homologous to 0 and are linearly independent. So $\dim H_1(\text{Rips}(X, 1^-)) = \infty$. □

5.2.3 The first homology group of a Čech filtration

Individual Čech complexes are almost as badly behaved as individual Vietoris–Rips complexes. It was shown in [2, Appendix B] that the homology groups $H_k(\check{\text{Cech}}(X; a))$ of a compact metric space X can be infinite dimensional for any $k \geq 2$.

However, the first homology is better behaved. The following result was originally obtained by Smale et al. [2, Theorem 8] using a different argument.

Proposition 5.12 *Let (X, d) be a totally bounded metric space, and let $a \geq 0$. Then, over any coefficient ring \mathbf{A} and any $a \in \mathbf{R}$, the 1-dimensional homology $H_1(\check{\text{Cech}}(X, a); \mathbf{A})$ is finitely generated over \mathbf{A} . In particular, over a field \mathbf{k} ,*

$$\dim_{\mathbf{k}}(H_1(\check{\text{Cech}}(X, a); \mathbf{k})) < \infty.$$

Proof The proposition follows from a sequence of elementary remarks.

1. Every 1-cycle in $\check{\text{Cech}}(X, a)$ is homologous to a 1-cycle whose edges have length at most a .

Proof Any edge $[x, x']$ belonging to $\check{C}ech(X, a)$ has an a -centre; that is, a point y which satisfies $d(x, y) \leq a$ and $d(x', y) \leq a$. Since $d(y, y) = 0 \leq a$, the point y is also an a -centre for the triangle $[x, y, x']$ and the edges $[x, y]$, $[y, x']$. It follows that any 1-cycle

$$\gamma = \sum_i a_i [x_i, x'_i]$$

can be replaced by a homologous 1-cycle

$$\hat{\gamma} = \gamma + \partial \sum_i a_i [x_i, y_i, x'_i] = \sum_i a_i ([x_i, y_i] + [y_i, x'_i])$$

all of whose edges $[x_i, y_i]$, $[y_i, x'_i]$ have length at most a . □

2. There exists a finite set E_a of edges of length at most a , with the following property: for any edge $[x, y]$ of length at most a , there exists an edge $[x', y']$ in E_a such that $d(x, x') \leq a$ and $d(y, y') \leq a$.

Proof Since X is totally bounded, so is $X \times X$ with the l^∞ product metric

$$d((x, y), (x', y')) = \max(d(x, x'), d(y, y')).$$

Since $X \times X$ is totally bounded, so is its subspace

$$[X \times X]_a = \{(x, y) \in X \times X \mid d(x, y) \leq a\}.$$

Let $(x'_1, y'_1), \dots, (x'_N, y'_N)$ be an a -sample for $[X \times X]_a$. Then

$$E_a = \{[x'_1, y'_1], \dots, [x'_N, y'_N]\}$$

satisfies the required condition. □

3. Any 1-cycle can be written as a finite linear combination of cycles of the form

$$[x_1, x_2] + [x_2, x_3] + \dots + [x_{k-1}, x_k] + [x_k, x_1] \tag{*}$$

(k may vary).

Proof This is standard, but we give the proof explicitly. Certainly any 1-cycle γ can be written as a finite linear combination of cycles (as above) and paths of the form

$$[x_1, x_2] + [x_2, x_3] + \dots + [x_{k-1}, x_k] + [x_k, x_{k+1}], \quad x_1 \neq x_{k+1},$$

(the trivial solution is to use paths of length 1 and no cycles). Consider the ‘free’ vertices in such a decomposition for γ : that is, vertices that occur as endpoints of the paths in the decomposition. We can eliminate the free vertices one by one as follows. Pick a free vertex and enumerate the paths which terminate there: P_1, P_2, \dots, P_m . Since $\partial\gamma = 0$, we must have $m \geq 2$. We can decrease m strictly by concatenating P_m with the appropriate multiple of P_{m-1} or its reverse. This creates a new, longer path (or cycle, if the other endpoints coincide) in place of P_m , and rescales or annihilates P_{m-1} . Eventually $m = 0$ and the free vertex is eliminated. Finally, when there are no free vertices the decomposition involves only cycles, and we are done. □

4. Consider a cycle of the form

$$\gamma = [x_1, x_2] + [x_2, x_3] + \dots + [x_{k-1}, x_k] + [x_k, x_1]$$

whose edges have length at most a . Then γ is homologous in $\check{C}ech(X, a)$ to a cycle whose edges belong to E_a .




Fig. 2 Construction of the 1-cycle $\hat{\gamma}$ (red and dashed blue edges) from the 1-cycle γ (black edges) in the proof of Proposition 5.12. (Color figure online)

Proof Approximate each edge $[x_i, x_{i+1}]$ by an edge $[x'_i, y'_i] \in E_a$ according to remark 2 (interpreting x_{k+1} as x_1 , cyclically). We claim that

$$\hat{\gamma} = [x'_1, y'_1] + [y'_1, x'_2] + [x'_2, y'_2] + [y'_2, x'_3] + \dots + [x'_k, y'_k] + [y'_k, x'_1]$$

is homologous to γ in $\check{C}ech(X, a)$. Indeed

$$\gamma - \hat{\gamma} = \partial \sum_{i=1}^k \left([x'_i, x_i, y'_i] + [y'_i, x_i, x_{i+1}] + [y'_i, x_{i+1}, x'_{i+1}] \right)$$

(see Fig. 2). To verify that the right-hand side of the equation belongs to $\check{C}ech(X, a)$, note that the triangles $[x'_i, x_i, y'_i]$, $[y'_i, x_i, x_{i+1}]$ and $[y'_i, x_{i+1}, x'_{i+1}]$ have a -centres x'_i , x_{i+1} and x_{i+1} respectively. \square

Combining remarks 1, 3 and 4, we see that in $\check{C}ech(X, a)$ any 1-cycle γ is homologous to a 1-cycle involving only edges in the finite set E_a . It follows that the homology $H_1(\check{C}ech(X, a); \mathbf{A})$ is finitely generated. \square

6 Special classes of metric spaces

Up this point we have been considering metric spaces in full generality, occasionally assuming total boundedness. In the last part of this paper, we show that within certain classes of metric space there are constraints on the persistence diagrams of their Rips complexes. The theme is ‘no new cycles’ beyond a certain Rips diameter. For path metric spaces, there are no new 1-cycles once the diameter is positive (Theorem 6.3). For δ -hyperbolic spaces, there are no new 2-cycles after the diameter exceeds 2δ (Theorem 6.7). It seems to us that these two results should be the beginning of a much richer story.

6.1 The persistence diagram of $H_1(\mathbb{R}ips)$ for a path metric space

Recall that a metric space (X, d_X) is a path metric space if the distance between each pair of points is equal to the infimum of the lengths of the curves joining these two points.⁴ Every path metric space (X, d_X) satisfies the following property (see [13, Theorem 1.8]): for any $x, x' \in X$ and any $\varepsilon > 0$ there exists a point $z \in X$ such that

$$\sup(d_X(x, z), d_X(x'z)) \leq \frac{1}{2}d_X(x, x') + \varepsilon.$$

Lemma 6.1 *Let (X, d_X) be a path metric space and let $q > 0$. Then the map*

$$H_1(\mathbb{R}ips(X, \frac{2}{3}q)) \rightarrow H_1(\mathbb{R}ips(X, q))$$

is surjective.

Proof Any 1-cycle γ in $\mathbb{R}ips(X, q)$ can be written as a finite linear combination $\lambda_1 e_1 + \dots + \lambda_n e_n$ of edges $e_i = [x_i, y_i]$ of length at most q . For any i , there exists $m_i \in X$ such that:

$$\begin{aligned} d_X(x_i, m_i) &\leq \frac{1}{2}d_X(x_i, y_i) + \varepsilon \leq \frac{2}{3}q \\ d_X(y_i, m_i) &\leq \frac{1}{2}d_X(x_i, y_i) + \varepsilon \leq \frac{2}{3}q \end{aligned}$$

As a consequence the triangle $[x_i, m_i, y_i]$ is contained in $\mathbb{R}ips(X, q)$ so γ is homologous in $\mathbb{R}ips(X, q)$ to

$$\gamma + \partial \left[\sum_{i=1}^n \lambda_i [x_i, m_i, y_i] \right] = \sum_{i=1}^n \lambda_i [x_i, m_i] + \lambda_i [m_i, y_i]$$

which is a 1-cycle whose edges belong to $\mathbb{R}ips(X, \frac{2}{3}q)$. □

Corollary 6.2 *Let (X, d_X) be a path metric space and $0 < p < q$. Then the map $H_1(\mathbb{R}ips(X, p)) \rightarrow H_1(\mathbb{R}ips(X, q))$ is surjective.*

Proof Iterate the surjectivity of $H_1(\mathbb{R}ips(X, \frac{2}{3}q)) \rightarrow H_1(\mathbb{R}ips(X, q))$. Eventually $(\frac{2}{3})^n q \leq p$. □

Theorem 6.3 *Let (X, d_X) be a path metric space. Then the persistence diagram of $H_1(\mathbb{R}ips(X))$ is contained in the vertical line $\{0\} \times [0, +\infty)$.*

Essentially, this is equivalent to Corollary 6.2 in asserting that there are no new 1-cycles once the Rips diameter is strictly positive. The formal deduction depends on how the persistence diagram is constructed. Since we are using the measure-theoretic construction of [7], we must argue in terms of the notation and concepts of that paper. Readers unfamiliar with [7] are encouraged to skip the proof and arrive at their own understanding of the relationship between Corollary 6.2 and Theorem 6.3.

Proof It is enough to show that $\mu(R) = 0$ for any rectangle which does not meet the line. Write $R = [p, q] \times [r, s]$, where $-\infty \leq p < q < r < s \leq +\infty$. Either $q < 0$, in which case

$$\mu(R) = \langle \circ_p - \bullet_q - \bullet_r - \circ_s \rangle \leq \langle \bullet_q \rangle = \dim H_1(\mathbb{R}ips(X, q)) = 0$$

because there are no edges when $q < 0$. Or else $0 < p$, in which case

$$\mu(R) = \langle \circ_p - \bullet_q - \bullet_r - \circ_s \rangle \leq \langle \circ_p - \bullet_q \rangle = 0$$

because $H_1(\mathbb{R}ips(X, p)) \rightarrow H_1(\mathbb{R}ips(X, q))$ is surjective. □

⁴ See [13] chap.1, def 1.2 for a definition of the length of a curve in a metric space.

Remark 6.4 We have not required X to be totally bounded in this theorem. This would seem necessary for invoking the persistence diagram. In fact, it is shown in [7] that the persistence diagram is defined wherever the persistence measure takes finite values. Above we see that the measure is zero away from the vertical line, so the diagram is defined, and empty, away from that line. A similar remark applies to Theorem 6.7.

Theorem 6.3 relies on the fact that any 1-dimensional simplex in a path metric space can be ‘subdivided’ into a sum of smaller simplices; in other words is homologous to a sum of simplices of strictly smaller diameter. Without further assumptions, this property does not hold for higher-dimensional simplices. For example if (X, d_X) is a circle of length 1, then any triple of points x, y, z such that $d_X(x, y) = d_X(y, z) = d_X(z, x) = \frac{1}{3}$ spans a triangle in $\text{Rips}(X, \frac{1}{3})$ that is not homologous to any finite sum of triangles of diameter strictly less than $\frac{1}{3}$.

In the next section, we show that there is an analogous result in H_2 for metric spaces which are δ -hyperbolic.

6.2 The persistence diagram of $H_2(\text{Rips})$ for a δ -hyperbolic space

Let X be a geodesic space, that is a metric space where any $x, y \in X$ are connected by geodesic of minimal length $d = d_X(x, y)$. If we choose a minimising geodesic $\llbracket x, y \rrbracket \subset X$, then $\llbracket x, y \rrbracket_t$ denotes the point on it which lies at distance t from x . Taking $\llbracket y, x \rrbracket$ to be the reverse geodesic, we clearly have $\llbracket x, y \rrbracket_t = \llbracket y, x \rrbracket_{d-t}$.

A geodesic space X is said to be δ -**hyperbolic** (see [12, chap. 2]) if the sides of every triangle run very close to each other in the following sense. Given $x, y, z \in X$, let $\llbracket y, z \rrbracket, \llbracket z, x \rrbracket, \llbracket x, y \rrbracket$ be minimising geodesics with lengths a, b, c respectively. The triangle inequality implies that there are non-negative numbers α, β, γ such that

$$a = \beta + \gamma, \quad b = \gamma + \alpha, \quad c = \alpha + \beta,$$

namely

$$\alpha = \frac{1}{2}(b + c - a), \quad \beta = \frac{1}{2}(c + a - b), \quad \gamma = \frac{1}{2}(a + b - c).$$

The δ -hyperbolicity condition for the triangle is:

$$\begin{aligned} d_X(\llbracket x, y \rrbracket_t, \llbracket x, z \rrbracket_t) &\leq \delta && \text{for } 0 \leq t \leq \alpha \\ d_X(\llbracket y, z \rrbracket_t, \llbracket y, x \rrbracket_t) &\leq \delta && \text{for } 0 \leq t \leq \beta \\ d_X(\llbracket z, x \rrbracket_t, \llbracket z, y \rrbracket_t) &\leq \delta && \text{for } 0 \leq t \leq \gamma \end{aligned}$$

If this holds for all triangles $[x, y, z]$, then X is δ -hyperbolic. See Fig. 3 (left). The points

$$u = \llbracket y, z \rrbracket_\beta = \llbracket z, y \rrbracket_\gamma \quad v = \llbracket z, x \rrbracket_\gamma = \llbracket x, z \rrbracket_\alpha \quad w = \llbracket x, y \rrbracket_\alpha = \llbracket y, x \rrbracket_\beta$$

have special importance. If X is a tree, then u, v, w coincide. In Euclidean or hyperbolic space u, v, w are the points of tangency of the incircle of xyz .

Lemma 6.5 *Let (X, d_X) be a δ -hyperbolic geodesic space and let $q > 2\delta$. Then the map $H_2(\text{Rips}(X, \frac{1}{2}q + \delta)) \rightarrow H_2(\text{Rips}(Z, q))$ is surjective.*

Proof Any class $\sigma \in H_2(\text{Rips}(X, q))$ is represented by a 2-cycle which is a linear combination of triangles in $\text{Rips}(X, q)$. We must break these triangles into smaller triangles.




Fig. 3 *Left* The triangle xyz is δ -hyperbolic iff the indicated secants have length at most δ . *Right* The triangle is split into four, with new vertices at the midpoints ℓ, m, n of the sides

We begin by selecting a geodesic $\llbracket x, y \rrbracket$ for each pair x, y which occurs as an edge of a triangle in σ .

Let $T = [x, y, z]$ be a triangle of σ , with side-lengths a, b, c and α, β, γ defined as before. We will break up the triangle using the midpoints

$$\ell = \llbracket y, z \rrbracket_{a/2} \quad m = \llbracket z, x \rrbracket_{b/2} \quad n = \llbracket x, y \rrbracket_{c/2}$$

of the sides. To help estimate the new edges we consider the ‘incircle’ points u, v, w defined above. See Fig. 3 (right).

Along the three geodesic sides of the triangle, one easily calculates:

$$d_X(u, \ell) = \frac{1}{2}|b - c| \quad d_X(v, m) = \frac{1}{2}|c - a| \quad d_X(w, n) = \frac{1}{2}|a - b|$$

We estimate the distances between ℓ, m, n . Since $[x, y, z] \in \text{Rips}(X, q)$ we have $\max(a, b, c) \leq q$. Moreover, we suppose that $a \leq b \leq c$. Then:

$$\begin{aligned} d_X(\ell, m) &\leq d_X(\ell, u) + d_X(u, v) + d_X(v, m) \\ &\leq \frac{1}{2}(c - b) + \delta + \frac{1}{2}(c - a) \\ &\leq \frac{1}{2}c - \gamma + \delta \\ &\leq \frac{1}{2}q + \delta \end{aligned}$$

and:

$$\begin{aligned} d_X(\ell, n) &\leq d_X(\ell, u) + d_X(u, w) + d_X(w, n) \\ &\leq \frac{1}{2}(c - b) + \delta + \frac{1}{2}(b - a) \\ &\leq \frac{1}{2}c - \frac{1}{2}a + \delta \\ &\leq \frac{1}{2}q + \delta \end{aligned}$$

For the third edge we introduce $m' = \llbracket x, y \rrbracket_{b/2}$. Since $\frac{1}{2}b \leq \alpha$ we can invoke the δ -hyperbolicity condition for m, m' to get:

$$\begin{aligned} d_X(m, n) &\leq d_X(m, m') + d_X(m', n) \\ &\leq \delta + \frac{1}{2}(c - b) \\ &\leq \frac{1}{2}q + \delta \end{aligned}$$

From this, we see that if we replace each triangle $[x, y, z]$ of σ with the corresponding sum

$$-[x, m, n] + [y, \ell, n] - [z, \ell, m] + [\ell, m, n]$$

we get a 2-cycle $\hat{\sigma}$ whose triangles belong to $\text{Rips}(X, \frac{1}{2}q + \delta)$.

We need to do one more thing, which is to show that $\sigma, \hat{\sigma}$ are homologous through a 3-cycle whose tetrahedra belong to $\text{Rips}(X, q)$. We require an additional edge to do this, the median zn :

$$\begin{aligned} d_X(z, n) &\leq d_X(z, u) + d_X(u, w) + d_X(w, n) \\ &\leq \frac{1}{2}(b + a - c) + \delta + \frac{1}{2}(b - a) \\ &\leq \frac{1}{2}c + \delta \\ &\leq q \end{aligned}$$

Now define

$$H[x, y, z] = [x, y, z, n] + [x, z, m, n] - [y, z, \ell, n] + [z, \ell, m, n]$$

for each triangle $[x, y, z]$, and extend linearly over all triangles in σ . One can check that $\hat{\sigma} = \sigma + \partial H\sigma$ when σ is a cycle. All edges of the tetrahedra in $H\sigma$ have length at most q , so $[\sigma] = [\hat{\sigma}]$ in $H_2(\text{Rips}(X, q))$. \square

Iterating the lemma, we get:

Corollary 6.6 *Let (X, d_X) be a δ -hyperbolic geodesic space and let $q > p > 2\delta$. Then the map $H_2(\text{Rips}(X, p)) \rightarrow H_2(\text{Rips}(X, q))$ is surjective.*

Proof If $f(t) = \frac{1}{2}t + \delta$ and $q > 2\delta$ then the iterates $f^n(q)$ are a decreasing sequence converging to 2δ , so eventually $f^n(q) < p$. \square

Theorem 6.7 *Let (X, d_X) be a δ -hyperbolic geodesic space. Then the persistence diagram of $H_2(\mathbb{R}\text{ips}(X))$ is confined to the vertical strip $[0, 2\delta] \times [0, +\infty)$.*

Proof It is enough to show that $\mu(R) = 0$ for any rectangle which does not meet the strip. Write $R = [p, q] \times [r, s]$, where $-\infty \leq p < q < r < s \leq +\infty$. Either $q < 0$, in which case

$$\mu(R) = \langle \circ_p \text{---} \bullet_q \text{---} \bullet_r \text{---} \circ_s \rangle \leq \langle \bullet_q \rangle = \dim H_2(\text{Rips}(X, q)) = 0$$

because there are no triangles when $q < 0$. Or else $2\delta < p$, in which case

$$\mu(R) = \langle \circ_p \text{---} \bullet_q \text{---} \bullet_r \text{---} \circ_s \rangle \leq \langle \circ_p \text{---} \bullet_q \rangle = 0$$

because $H_2(\text{Rips}(X, p)) \rightarrow H_2(\text{Rips}(X, q))$ is surjective. \square

As with Theorem 6.3, the result is essentially equivalent to the corollary that precedes it: there are no new H_2 classes once the Rips diameter exceeds 2δ so the persistence diagram is empty outside the vertical strip. Our formal proof is written in the language of [7], and again the reader unfamiliar with the concepts may safely skip it, and simply regard Theorem 6.7 as another way to state Corollary 6.6.

Acknowledgments The authors are grateful to Steve Smale for fruitful discussions that motivated the results Sect. 5.2, and to J.-M. Droz for suggesting the idea of the proof of Proposition 5.9.

References

1. Attali, D., Lieutier, A., Salinas, D.: Vietoris–Rips complexes also provide topologically correct reconstructions of sampled shapes. In: Proceedings of the 27th Annual ACM Symposium on Computational geometry, SoCG '11, pp. 491–500. ACM, New York, NY, USA (2011). doi:[10.1145/1998196.1998276](https://doi.org/10.1145/1998196.1998276)
2. Bartholdi, L., Schick, T., Smale, N., Smale, S., Baker, A.W.: Hodge theory on metric spaces. *Found. Comput. Math.* **12**(1), 1–48 (2012)
3. Burago, D., Burago, Y., Ivanov, S.: *A Course in Metric Geometry*, Graduate Studies in Mathematics, vol. 33. American Mathematical Society, Providence, RI (2001)
4. Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L., Oudot, S.: Proximity of persistence modules and their diagrams. In: SCG, pp. 237–246 (2009). doi:[10.1145/1542362.1542407](https://doi.org/10.1145/1542362.1542407)
5. Chazal, F., Cohen-Steiner, D., Guibas, L.J., Mémoli, F., Oudot, S.Y.: Gromov–Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum (Proceedings of the SGP 2009)* pp. 1393–1403 (2009)
6. Chazal, F., Oudot, S.Y.: Towards persistence-based reconstruction in euclidean spaces. In: Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry, SCG '08, pp. 232–241. ACM, New York, NY, USA (2008). doi:[10.1145/1377676.1377719](https://doi.org/10.1145/1377676.1377719)
7. Chazal, F., de Silva, V., Glisse, M., Oudot, S.: The structure and stability of persistence modules (2012). ArXiv:1207.3674 [math.AT]
8. Dowker, C.H.: Homology groups of relations. *Ann. Math.* **56**(1), 84–95 (1952)
9. Droz, J.M.: A subset of Euclidean space with large Vietoris–Rips homology (2012). ArXiv:1210.4097 [math.GT]
10. Edelsbrunner, H., Harer, J.: *Computational Topology: An Introduction*. American Mathematical Society, Providence, RI (2010)
11. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. *Discret. Comput. Geom.* **28**, 511–533 (2002)
12. Ghys, E., de la Harpe, P.: *Sur les groupes hyperboliques d’après Mikhael Gromov*, vol. 83. Birkhäuser, Basel (1990)
13. Gromov, M.: *Metric Structures for Riemannian and Non-Riemannian Spaces*, 2nd edn. Birkhäuser, Basel (2007)
14. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge, MA (2001). <http://www.math.cornell.edu/~hatcher/>
15. Hausmann, J.C.: On the Vietoris–Rips complexes and a cohomology theory for metric spaces. *Ann. Math. Stud.* **138**, 175–188 (1995)
16. Latschev, J.: Vietoris–Rips complexes of metric spaces near a closed Riemannian manifold. *Archiv der Mathematik* **77**(6), 522–528 (2001)
17. Munkres, J.R.: *Elements of Algebraic Topology*. Westview Press, Boulder, CO (1984)
18. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discret. Comput. Geom.* **33**(2), 249–274 (2005)