

# Robustness of Persistent Topological Features and Minimum Homological Cuts

Pepijn Roos Hoefgeest ✉

KTH Royal Institute of Technology Stockholm, Sweden

Lucas Slot ✉ 

University of Amsterdam, The Netherlands

---

## Abstract

Persistent homology is a popular method for computing topological features of (metric) data. Standard approaches based on the Čech or Rips filtration are stable under small perturbations of the data, but highly sensitive to outliers. This lack of robustness has been frequently addressed in the literature. In this paper, we take a novel perspective by asking the following question: When can we guarantee that an observed persistent feature (a bar) is inherent to the underlying data in the presence of a limited number of unknown, arbitrary outliers. We formalize this question by introducing the notion of *adversarial robustness*, and study the problem of deciding whether a given bar in the barcode of a filtered simplicial complex is adversarially robust. We show that this problem is essentially equivalent to a homological variant of the minimum cut problem in simplicial complexes, which we believe to be of independent interest. As our main technical contribution, we provide the first computational complexity results for this problem, consisting of an efficient algorithm in 0-dimensional homology, NP-hardness for the general problem, and an efficient algorithm for codimension-1 in  $n$ -dimensional complexes embedded in  $\mathbb{R}^n$ . We also analyze its natural linear programming relaxation, whose dual defines a homological analog of the max-flow problem in graphs. We show that a max-flow/min-cut theorem does not hold in our setting, implying that the LP relaxation is not tight in general. Finally, in the special case of the Rips filtration, we provide a global heuristic based on the Hausdorff distance that guarantees adversarial robustness of sufficiently long bars. This connects adversarial robustness to standard stability theorems in persistent homology.

**2012 ACM Subject Classification** Mathematics of computing → Algebraic topology

**Keywords and phrases** Topological Data Analysis, Persistent Homology, Min-cut Max-flow, Robustness, Vietoris-Rips Filtration

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2026.87

**Related Version** *Full Version*: <https://arxiv.org/abs/2602.23154> [28]

**Funding** *Pepijn Roos Hoefgeest*: This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

*Lucas Slot*: This work was completed while the author was at ETH Zurich, and supported by the Swiss National Science Foundation (SNSF), grant no. 10004947.

## 1 Introduction

Persistent homology is a central method in topological data analysis. It is used to extract topological features from (metric) data across a range of spatial scales. At a high level, it works as follows. First, we represent our data by a nested sequence of simplicial complexes, called a filtration. An important example is the (Vietoris-)Rips filtration, which arises from pairwise distances between points in a metric space. The evolution of the homology groups of the complexes in this filtration (i.e., its *persistent* homology) can be captured by a *barcode*, which consists of a (multi)set of intervals (or *bars*) whose endpoints represent the



© Pepijn Roos Hoefgeest and Lucas Slot;

licensed under Creative Commons License CC-BY 4.0


42nd International Symposium on Computational Geometry (SoCG 2026).

Editors: Hee-Kap Ahn, Michael Hoffmann, and Amir Nayyeri; Article No. 87; pp. 87:1–87:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



appearance (*birth*) and disappearance (*death*) of a homology class at a particular step in the filtration. Barcodes serve as a topological signature of the underlying data set. We refer to [19, 21, 26, 32] for surveys on persistent homology and its many applications.

**Stability and robustness.** An important property of barcodes is their *stability* under small perturbations: For example, the *Bottleneck distance* between the barcodes of Rips filtrations may be bounded in terms of the (*Gromov-*)*Hausdorff distance* between their underlying metric spaces [16]. On the other hand, barcodes are (in)famously sensitive to *outliers*. Indeed, even a single outlier may cause arbitrarily large changes in the barcode, which makes persistent homology unreliable in the presence of noise. This issue has been frequently addressed in the literature; we highlight three approaches. First, one may consider alternative metric interpretations of the data which take density into account, and whose resulting barcodes are (hopefully) more robust. For instance, a filtration based on the *distance-to-measure* function [14] achieves stability with respect to the *Wasserstein distance*, which tolerates small amounts of outliers [8]. See [24, 27] for similar approaches using kernel functions. A downside is that these filtrations are difficult to interpret geometrically, often depend on some choice of secondary parameters, and are computationally expensive, making them less suited for topological inference than the Rips filtration. A second approach is to (cleverly) subsample the data to determine so-called *landmarks*, and then construct a filtration based on these landmarks. An example is the (*lazy*) *witness filtration* [29]. Empirically, it appears that the landmarks may be chosen in a way that reduces the sensitivity to outliers (in the original data); see [30]. However, this effect is hard to quantify theoretically. Third, one may consider filtrations indexed by multiple parameters, modeling for example both scale and density. The persistent homology of such bifiltrations can be provably robust to certain types of noise [4]. However, multiparameter persistence modules are significantly more difficult to represent than their one-parameter counterparts. In particular, they generally do not have a barcode, which is a serious theoretical and practical drawback; see [7].

**A new notion of robustness.** In this paper, we take a different perspective. Rather than modify existing filtrations to increase their tolerance to noise, we ask the following question.

► **Question 1.** *When can we guarantee that an observed persistent feature (a bar) arising from a filtration is inherent to the underlying data in the presence of (adversarial) noise?*

To address this question, we introduce a new notion of outlier-robustness of persistent topological features (bars) arising from a filtration of a simplicial complex, which we call *adversarial robustness*. Before giving a formal definition, it is illustrative to first consider the Rips filtration on a (finite) set of metric data. Intuitively speaking, we say that a persistent feature of the Rips filtration is  $k$ -adversarially robust if it “continues to exist” after removing any  $k$  points from the data set. In this way,  $k$ -adversarial robustness certifies that a feature is *inherent* to an underlying data set even after adding (at most  $k$ ) unknown, arbitrary outliers. Alternatively, the largest  $k$  for which a feature is  $k$ -adversarially robust can be thought of as a *measure of robustness* of that feature (possibly after dividing by the total number of data points). With respect to the earlier work outlined above, a key advantage of our approach is that we make no structural assumptions on the outliers.

In what follows, we give a formal definition of adversarial robustness. The key ingredient is the induced matching between barcodes [2], which allows us to relate the persistent features of a filtered simplicial complex to those of its filtered subcomplexes. Then, we study the computational problem of deciding whether a given bar is adversarially robust. To this end,

we first show that adversarial robustness of a bar can be determined by solving a homological variant of the minimum cut problem in a single (unfiltered) simplicial complex. Thus, it suffices to study the complexity of that problem, which we believe to be of independent interest. We give efficient algorithms for zero-dimensional homology and complexes embedded in  $\mathbb{R}^n$ . On the other hand, we show that the general problem is NP-hard. We also analyze its natural linear programming relaxation, whose dual defines a homological analog of the max-flow problem in graphs. We conclude by giving an efficiently computable heuristic for adversarial robustness in Rips filtrations, which is connected to Hausdorff stability.

## 1.1 Adversarially robust persistent features

Let  $K$  be a (finite) simplicial complex, and let  $\mathcal{K} = (K_i)_{0 \leq i \leq m}$  be a filtration of  $K$ , i.e., a sequence of simplicial complexes  $\emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_m = K$ . Throughout, we assume that  $\mathcal{K}$  is a *simplex-wise* filtration, i.e., for any  $i < m$ , the complex  $K_{i+1}$  is obtained by adding (at most) a single simplex to  $K_i$ . Any filtration can be made simplex-wise by breaking ties (arbitrarily, but so that face relations are preserved) whenever multiple simplices are added at once, which is also what is done in practice. For  $p \geq 0$ , we denote the  $p$ -dimensional persistent homology of  $\mathcal{K}$  by  $\text{PH}_p(\mathcal{K})$ , and we write  $\mathcal{B}(\text{PH}_p(\mathcal{K}))$  for its barcode. For  $A \subseteq K$ , we write  $K - A$  for the largest subcomplex of  $K$  contained in  $K \setminus A$ . That is,  $K - A$  is obtained from  $K$  by removing all simplices that have a face in  $A$ . Similarly, we write  $\mathcal{K} - A$  for the filtration of  $K - A$  given by  $(K_i - A)_{0 \leq i \leq m}$ . There is a natural way to relate the barcodes associated with  $\mathcal{K}$  and  $\mathcal{K} - A$ . Namely, the inclusion  $K - A \hookrightarrow K$  induces a map  $\text{PH}_p(\mathcal{K} - A) \rightarrow \text{PH}_p(\mathcal{K})$ . In turn, this map induces a (partial) matching between the respective barcodes, which we denote  $\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}}$ . This is a special case of the so-called *induced matching*, which plays a crucial role in a proof of the algebraic stability theorem [2]. Generally, the induced matching is *not* functorial, but in the context of simplex-wise filtrations it allows us to unambiguously relate bars in the barcodes of  $\text{PH}_p(\mathcal{K})$  and  $\text{PH}_p(\mathcal{K} - A)$ ; see Section 3 for more details. This allows us to give the central definition of this paper. For  $s \in \mathbb{N}$ , let  $K^{(s)} \subseteq K$  denote the set of  $s$ -simplices in  $K$ .


► **Definition 2** (Adversarial robustness). *Let  $\mathcal{K}$  be a simplex-wise filtration of a simplicial complex  $K$ . Let  $p \geq 0$ , and let  $s \leq p$ . A bar  $B \in \mathcal{B}(\text{PH}_p(\mathcal{K}))$  is  $k$ -adversarially robust (in degree  $s$ ) if, for each subset  $A \subseteq K^{(s)}$  of size at most  $k$ , we have  $B \in \text{Im}(\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}})$ .*

Definition 2 guarantees that a bar is “present” in any (filtered) subcomplex of  $K$  missing at most  $k$   $s$ -simplices. If  $\mathcal{K}$  is a Rips filtration, the 0-simplices (vertices) of  $K$  correspond to metric data points. Thus,  $k$ -adversarial robustness (in degree 0) of a bar means that it is present in the Rips filtration of any subset of the original data obtained by removing at most  $k$  points, matching our earlier intuitive description.

## 1.2 (Minimum) homological cuts

Adversarial robustness of bars is closely related to a novel homological variant of the min-cut problem in simplicial complexes, which we introduce in this work. For  $C \subseteq K^{(s)}$ , we write  $\iota_{K-C \hookrightarrow K}^* : H_p(K - C) \rightarrow H_p(K)$  for the map induced by the inclusion  $K - C \hookrightarrow K$ .

► **Definition 3** (homological cuts). *Let  $K$  be a simplicial complex. Let  $p \geq 0$ , and  $s \leq p$ . We say that  $C \subseteq K^{(s)}$  is a homological  $s$ -cut for  $\gamma \in H_p(K)$  if  $\gamma \notin \text{Im}(\iota_{K-C \hookrightarrow K}^*)$ . We refer to the special cases  $s = 0, s = 1$  as homological vertex and edge cuts, respectively; see Figure 1.*



■ **Figure 1** A simplicial complex  $K$  with cycles  $c_{\text{left}}, c_{\text{right}} \in C_1(K; \mathbb{R})$  drawn in blue single arrows and red double arrows, respectively (all coefficients equal to 1). The classes  $[c_{\text{left}}]$  and  $[c_{\text{right}}]$  generate  $H_1(K; \mathbb{R}) \cong \mathbb{R}^2$ . On the right: two subcomplexes obtained by removing subsets  $C_1, C_2 \subseteq K^{(1)}$  (dashed) from  $K$ , respectively. Note that  $C_1$  is an edge cut for  $[c_{\text{left}}]$  and  $[c_{\text{left}} + c_{\text{right}}]$ , but not for  $[c_{\text{right}}]$ . On the other hand,  $C_2$  is a 1-cut for both  $[c_{\text{left}}]$  and  $[c_{\text{right}}]$ , but not for  $[c_{\text{left}} + c_{\text{right}}]$ .

To connect homological cuts to adversarial robustness, we use the fact that, for a simplex-wise filtration  $\mathcal{K}$ , each bar  $B \in \mathcal{B}(\text{PH}_p(\mathcal{K}))$  corresponds to a pair  $(\sigma_B, \tau_B)$  of simplices, whose insertions at steps  $i = b$  and  $i = d$ , respectively, represent the birth and death of any cycle representing  $B$ . These are called *persistence pairs*. We call the complex  $K_B := K_{d-1}$  the *predeath complex* of  $B$ . Now, the following proposition shows that adversarial robustness of  $B$  is characterized by homological cuts of  $[\tau_B]$  in  $K_B$ . We give its proof in Section 3.2.

► **Proposition 4.** *Let  $p \geq 0$ , and let  $B \in \mathcal{B}(\text{PH}_p(\mathcal{K}))$  be a bar in the barcode of a simplex-wise filtration  $\mathcal{K}$  of a simplicial complex  $K$ . Let  $K_B, \tau_B$  be its pre-death complex and death simplex. Then,  $B$  is  $k$ -adversarially robust (in degree  $s$ ) if, and only if,*

$$[\partial\tau_B] \in \text{Im}(\iota_{K_B - A \hookrightarrow K_B}^*) \text{ for all } A \subseteq K^{(s)} \text{ with } |A| \leq k.$$

Thus,  $B$  is  $k$ -adversarially robust iff all homological  $s$ -cuts of  $[\tau_B]$  in  $K_B$  have size at least  $k$ .

The above proposition motivates our study of the following problem.

► **Problem 5.** *The minimum homological  $s$ -cut problem asks to compute*

$$\text{mhc}(K, \gamma, s) := \min_{C \subseteq K^{(s)}} \{|C| : C \text{ is a homological } s\text{-cut for } \gamma\}. \quad (\text{MHC})$$

We call a set  $C$  attaining the minimum above a minimum homological  $s$ -cut for  $\gamma$  (in  $K$ ).

Apart from its connection to adversarial robustness, we believe that the minimum homological cut problem is of intrinsic interest. It is related to, *but distinct from*, two types of well-studied problems in computational topology. On the one hand, there are problems related to finding a smallest (or otherwise “optimal”) representing cycle of a homology class in a simplicial complex [3, 5, 12]; this is often referred to as *homology localization*. By contrast, the problem we consider can be thought of as *cohomology localization*. The computational complexity of homology localization varies based on the choice of coefficient field, the dimension of the homology group, and additional assumptions on the underlying complex. For example, while the general problem is hard [15], a linear programming relaxation yields an efficient algorithm for finding minimum weight homologous cycles over integer coefficients for complexes whose boundary matrix is totally unimodular [17]. On the other hand, there are ordinary min-cut (and max-flow) problems in graphs with topological structure, e.g., graphs that can be embedded in a surface of low genus. There, (co)homological properties of cuts (or flows) can be used to achieve algorithmic speedups with respect to the general case [10, 11, 22].

**Linear programming relaxations.** As we explain in the full version of this paper [28], homological edge cuts of a class  $\gamma \in H_1(K; \mathbb{R})$  are naturally related to (ordinary) cuts in graphs. Namely, we show that a minimum homological edge cut can be found by optimizing the number  $\|\varphi\|_0$  of non-zero coefficients of a vector  $\varphi$  (indexed by the edges of  $K$ ) under a set of linear constraints involving the boundary matrix of  $K$ . The relaxation of this problem obtained by optimizing the 1-norm  $\|\varphi\|_1$  instead is a linear program. Its dual may be interpreted as a homological analog of the max-flow problem in graphs. Contrary to the graph setting, we show that there is no max-flow min-cut theorem in our case: homological max-flows are not necessarily integral, meaning the LP relaxation is not tight.

### 1.3 Main contributions

**Complexity of the minimum homological cut problem.** We prove positive and negative results on the computational complexity of the minimum homological cut problem (MHC) introduced above, in terms of the number of simplices  $|K|$  in the complex  $K$ . By Proposition 4, these results have immediate implications on the complexity of determining adversarial robustness of bars. Our first contribution is an efficient algorithm for the case where  $\gamma \in H_0(K)$  is a 0-dimensional homology class, i.e.,  $p = 0$ .

► **Theorem 6.** *Let  $K$  be a simplicial complex, and let  $\gamma \in H_0(K)$ . We can compute a minimum homological vertex cut for  $\gamma$  in time  $O(|K^{(0)}| + |K^{(1)}|)$ .*

Theorem 6 follows from the fact that a minimum homological cut of a class in  $H_0$  is always equal to the vertex set of a connected component of  $K$ , as we show in Section 4.1.

Next, we show that finding a minimum homological cut is NP-hard, already when  $p = 1$ .

► **Theorem 7.** *For  $p = 1$  and  $s = 1$ , the minimum homological  $s$ -cut problem is NP-hard. This is true in particular when homology is taken with coefficients in  $\mathbb{R}$  or  $\mathbb{Z}/2\mathbb{Z}$ .*

We give the proof of Theorem 7 in Section 4.2. It relies on a reduction from EXACT COVER BY 3-SETS (X3C). In [23, pp. 246], X3C is used to show hardness of the problem of finding minimum weight solutions to linear equations. It was also used in [25] to show hardness of finding the *sparsest* approximate solution to a set of linear equations. As mentioned, we show in the full version of this paper [28] that (MHC) can be solved by finding the sparsest solution to a particular set of linear equations. In light of these observations, X3C is a natural candidate for showing hardness of (MHC). In our proof, we construct for any X3C-instance an equivalent instance of (MHC). Our construction relies on “gluing together” punctured discs, each representing a 3-set of the X3C-instance, to obtain a topological space which is then triangulated carefully to ensure minimum homological cuts correspond to exact covers. This approach resembles earlier work on NP-hardness in homology localization [1, 15, 20]. A distinction is that these mostly rely on reductions from SAT-problems; to the best of our knowledge, a reduction from X3C was not considered before in this context. We remark that our construction can be extended to cover the case  $s = 0$ . Via Proposition 4, our result also implies that testing  $k$ -adversarial robustness of bars (in degree 0 and 1) is hard in general.

Finally, we give an efficient algorithm for the following special case, which covers, e.g., subcomplexes of the Delauney triangulation of a finite metric data set in  $\mathbb{R}^n$ .


► **Theorem 8.** *Suppose that  $K$  is an  $n$ -dimensional simplicial complex embedded in  $\mathbb{R}^n$ , and let  $\gamma \in H_{n-1}(K)$ . We can compute a minimum homological  $(n - 1)$ -cut for  $\gamma$  in polynomial time in the number of simplices of  $K$ .*

Our proof of Theorem 8 in Section 4.3 relies on Alexander duality, which can be used to construct the so-called *extended dual graph* of  $K$ . In [18], this graph was used to find optimal generators for certain homology classes: these correspond to minimum cuts in the graph. For us, the situation is precisely the opposite: as we will show, minimum homological cuts in  $K$  correspond to shortest paths in the extended dual graph (allowing for efficient computation). The use of the extended dual graph in homology localization dates back to [31]; see also [9].

**The Rips filtration.** Complementing our hardness result, we give an efficiently computable *heuristic* based on the Hausdorff distance to test adversarial robustness of bars arising from the Rips filtration of a metric data set  $X$ . Namely, we show that any bar in the barcode of the Rips filtration of length at least

$$\mathcal{H}_{X,k} := \max_{A \subseteq X, |A| \leq k} d_H(X \setminus A, X)$$

is automatically  $k$ -adversarily robust (in degree 0)<sup>1</sup>; see Section 5 for a precise statement. Moreover, we show that the parameter  $\mathcal{H}_{X,k}$  can be computed in time  $O(|X|^2 \log |X|)$ . This result links adversarial robustness to classical stability results in persistent homology. It shows that our definition is at least as expressive as a naive definition based only on the length of bars and the Hausdorff stability of the Rips filtration. That is, sufficiently long bars are outlier-robust according to our definition (as one would expect). But importantly, our definition is also capable of recognizing relatively short bars as robust; see Figure 2.



**Figure 2** Four data sets in  $\mathbb{R}^2$ , each of size 100, whose Rips filtrations each induce a 1-dimensional persistent feature. For  $X_1$  and  $X_2$ , these features are not  $k$ -adversarily robust for  $k = 10$ , evidenced by the subsets  $A_1 \subseteq X_1$  and  $A_2 \subseteq X_2$  marked in red. The points in  $A_2$  are quite dense, and so standard subsampling techniques will likely not remove them. On the other hand, the feature is 10-robust in  $X_3$ , as its length exceeds  $\mathcal{H}_{X_3,10}$ , which equals  $d_H(X_3 \setminus A_3, X_3)$  for the subset  $A_3 \subseteq X_3$  marked in red (cf. Theorem 20). In  $X_4$ , the feature induced by the densely sampled circle is 10-robust, even though its bar has length strictly less than  $\mathcal{H}_{X_4,10}$  (evidenced by the set  $A_4$  marked in red).

## 2 Preliminaries

### 2.1 Simplicial complexes

An (abstract) simplicial complex (on a set  $Y$ ) is a collection  $K$  of non-empty subsets  $\sigma \subseteq Y$ , called *simplices*, such that if  $\tau \in K$  and  $\sigma \subseteq \tau$ , then  $\sigma \in K$ . In that case,  $\sigma$  is called a *face* of  $\tau$ , and  $\tau$  is called a *coface* of  $\sigma$ . For  $A \subseteq K$ , we denote by  $K - A$  the largest subcomplex

<sup>1</sup> The Rips filtration is not simplex-wise, therefore, as mentioned, we must break ties to obtain a simplex wise filtration to apply Definition 2. This result does not depend on the choice of tiebreaker.

of  $K$  which does not contain any simplices in  $A$  (which is obtained by removing all simplices from  $K$  which have a face in  $A$ ). For  $s \in \mathbb{N}$ , we write  $K^{(s)}$  for the set of  $s$ -simplices of  $K$ , being the simplices  $\sigma \in K$  with  $|\sigma| = s + 1$ . We write  $V(K) := K^{(0)}$  for the vertices of  $K$ , and  $E(K) := K^{(1)}$  for its edges. If there is a  $p$  such that  $K^{(p)} \neq \emptyset$  and  $K^{(q)} = \emptyset$  for all  $q > p$ , then we say that  $K$  is  $p$ -dimensional, and we call  $K$  a  $p$ -complex. If  $K$  is a  $p$ -complex and each simplex of  $K$  is contained in a  $p$ -simplex, then  $K$  is called a *pure  $p$ -complex*.

**Embedded complexes.** A *geometric  $p$ -simplex*  $\sigma$  in  $\mathbb{R}^n$  is the convex hull of  $p + 1$  affinely independent vectors  $v_0, \dots, v_p \in \mathbb{R}^n$ . The elements of  $V(\sigma) := \{v_0, \dots, v_p\}$  are called the vertices of  $\sigma$ . If  $\tau$  is the convex hull of a subset of the vertices of  $\sigma$ , then  $\tau$  is called a *face* of  $\sigma$ . A *geometric simplicial complex*  $K$  in  $\mathbb{R}^n$  is a collection of geometric simplices in  $\mathbb{R}^n$  with the properties that 1) the face of every simplex is in  $K$ ; 2) for every pair of simplices  $\sigma, \tau \in K$  with non-empty intersection,  $\sigma \cap \tau$  is in  $K$ . We denote the *geometric realization* of  $K$  by  $\|K\| := \bigcup_{\sigma \in K} \sigma \subseteq \mathbb{R}^n$ , i.e. the subset of  $\mathbb{R}^n$  consisting of the union of simplices that make up  $K$ . If  $K$  is a geometric simplicial complex, then  $A(K) := \{V(\sigma) \mid \sigma \in K\}$  is an abstract simplicial complex. We say that an abstract simplicial complex  $L$  is *embedded in  $\mathbb{R}^n$*  if  $L \cong A(K)$  for some geometric simplicial complex  $K$  in  $\mathbb{R}^n$ , and we define the geometric realization of the embedded simplicial complex  $L$  to be  $\|L\| := \|K\|$ .

**Simplicial homology.** For a field  $\mathbb{F}$ , and  $p \in \mathbb{N}$ , we denote by  $C_p(K; \mathbb{F})$  the simplicial  $p$ -chains on  $K$ . We write  $\partial_p : C_p(K; \mathbb{F}) \rightarrow C_{p-1}(K; \mathbb{F})$  for the boundary operator. We have subspaces of cycles  $Z_p(K; \mathbb{F}) = \text{Ker}(\partial_p)$  and boundaries  $B_p(K; \mathbb{F}) = \text{Im}(\partial_{p+1})$ , and the simplicial homology of  $K$  is defined as  $H_p(K; \mathbb{F}) = Z_p(K; \mathbb{F}) / B_p(K; \mathbb{F})$ . Throughout, if we omit  $\mathbb{F}$  from the notation, we mean homology over any field. For two complexes  $K \subseteq K'$ , we denote the inclusion  $K \hookrightarrow K'$  by  $\iota_{K \hookrightarrow K'}$ . Then, we write  $\iota_{K \hookrightarrow K'}^* : H_p(K) \rightarrow H_p(K')$ ,  $[c] \mapsto [\iota_{K \hookrightarrow K'}(c)]$  for the induced maps in homology (suppressing  $p$  in the notation).

## 2.2 Persistent homology

A persistence module is a functor  $M : (\mathbb{R}, \leq) \rightarrow \mathbf{vect}_{\mathbb{F}}$ , where  $\mathbf{vect}_{\mathbb{F}}$  denotes the category of finitely generated vector spaces. A morphism between persistence modules  $M$  and  $N$  is a natural transformation of functors  $f : M \rightarrow N$ . There exists a multiset of intervals  $\mathcal{B}(M)$ , called the barcode of  $M$ , that completely captures  $M$  up to isomorphism. We write  $\mathbb{I}_I$  for the interval module over  $I$ , i.e. the persistence module which is equal to  $\mathbb{F}$  on  $I$ , connected by identity morphisms, and 0 elsewhere. Then there is an isomorphism  $M \cong \bigoplus_{I \in \mathcal{B}(M)} \mathbb{I}_I$  [6].

A filtration  $\mathcal{K} = (K_i)_{0 \leq i \leq m}$  of simplicial complexes naturally defines a persistence module  $\text{PH}_p(\mathcal{K})$ , called the persistent homology of  $\mathcal{K}$ , which is given at  $t \in \mathbb{R}$  by  $\text{PH}_p(\mathcal{K})_t = H_p(K_{\lfloor t \rfloor})$ , and whose structure maps are maps induced by inclusions. If  $\mathcal{K}$  is a simplex-wise filtration, there is at most one birth or death of a bar at each step of the filtration. Hence, any bar  $B = [i, j) \in \mathcal{B}(\text{PH}_p(\mathcal{K}))$  then uniquely corresponds to a pair of simplices  $(\sigma_B, \tau_B)$ , inserted at steps  $i$  and  $j$  respectively. This is called a *persistence pairing*, and we sometimes refer to the bar as  $B = [\sigma_B, \tau_B)$ . We denote by  $K_B := K_{d-1} = K_d \setminus \{\tau_B\}$  the *predeath complex* of  $B$ . For a filtration  $\mathcal{K}$  and a non-decreasing function  $f : \mathbb{N} \rightarrow \mathbb{R}$ , we get a persistence module  $\text{PH}_p(\mathcal{K}, f)$  given at time  $t \in \mathbb{R}$  by  $H_p(K_{\sup\{i: f(i) \leq t\}})$ . We think of  $f$  as a reparametrization of the filtration. Note that if  $f$  is the identity,  $\text{PH}_p(\mathcal{K}, f) = \text{PH}_p(\mathcal{K})$ . For each bar  $B$  in the barcode of  $\text{PH}_p(\mathcal{K}, f)$  there is a bar  $[i, j)$  in the barcode of  $\text{PH}_p(\mathcal{K})$  such that  $B = [f(i), f(j))$ .

### 3 Induced matchings, adversarial robustness and homological cuts

#### 3.1 Induced matchings

While the barcode of a persistence module  $M$  is unique, the isomorphism  $M \cong \bigoplus_{I \in \mathcal{B}(M)} \mathbb{I}_I$  is not. For a morphism  $f : M \rightarrow N$ , this makes it difficult to relate the barcodes  $\mathcal{B}(M)$  and  $\mathcal{B}(N)$  to one another in a way that reflects the algebraic structure of  $f$ . In [2], a procedure is described that associates a partial matching  $\mathcal{X}_f : \mathcal{B}(M) \leftrightarrow \mathcal{B}(N)$  with  $f$ , in a way that only depends on the barcodes of  $M$ ,  $N$ , and  $\text{Im}(f)$ , the image persistence module of  $f$ . Denote by  $\langle \cdot, d \rangle_M \subseteq \mathcal{B}(M)$  the bars in the barcode of the form  $[b, d)$  (for simplicity, we assume all intervals in the barcodes of  $M$  and  $N$  are half-open). If  $f$  is injective, it can be shown that  $|\langle \cdot, d \rangle_M| \leq |\langle \cdot, d \rangle_N|$ . Moreover, if we order the bars in both  $\langle \cdot, d \rangle_M$  and  $\langle \cdot, d \rangle_N$  by their length, there is an order preserving injection  $\mathcal{X}_f^d : \langle \cdot, d \rangle_M \hookrightarrow \langle \cdot, d \rangle_N$ , with the property that if  $\mathcal{X}_f^d([b, d)) = [b', d)$ , then  $b' \leq b$ . Note that if  $M$  and  $N$  come from simplex-wise filtrations, the sets  $\langle \cdot, d \rangle_M$  and  $\langle \cdot, d \rangle_N$  consist of at most one element. The induced matching  $\mathcal{X}_f$  is then defined as the union of the matchings  $\mathcal{X}_f^d$ . If  $f$  is surjective, there is a similar way to match bars with the same birth in an order preserving way. For arbitrary  $f : M \rightarrow N$ , there is a factorization  $f : M \twoheadrightarrow \text{Im}(f) \hookrightarrow N$  of  $f$  into a composition of a surjective and an injective map. The induced matching  $\mathcal{X}_f$  is then defined as the composition of the induced matchings of  $M \twoheadrightarrow \text{Im}(f)$  and  $\text{Im}(f) \hookrightarrow N$ . The induced matching is *not* functorial in  $f$ .

A persistence module  $M$  is called  $\varepsilon$ -trivial if for each  $t \in \mathbb{R}$ , the internal morphism  $\phi_{t \rightarrow t+\varepsilon}^M : M_t \rightarrow M_{t+\varepsilon}$  is the zero morphism. The following theorem is the main result of [2].

► **Theorem 9** (Induced matching theorem [2]). *Let  $f : M \rightarrow N$  be a morphism with  $\varepsilon$ -trivial kernel. Then each bar  $[b, d) \in \mathcal{B}(M)$  with  $d - b > \varepsilon$  is matched to a bar in  $\mathcal{B}(N)$  by  $\mathcal{X}_f$ . If  $\text{coker}(f)$  is  $\varepsilon$ -trivial,  $\mathcal{X}_f$  matches each bar  $[b, d) \in \mathcal{B}(N)$  with  $d - b > \varepsilon$  to a bar in  $\mathcal{B}(M)$ .*

We will need this result for our proof of Theorem 20 in Section 5.

#### 3.2 Adversarial robustness and minimum homological cuts

In this section, we show that  $k$ -adversarial robustness of a bar can be determined by solving a minimum homological vertex cut problem. That is, we prove Proposition 4.

**Proof of Proposition 4.** Let  $B = [b, d)$  be a bar in the barcode of  $\text{PH}_p(\mathcal{K})$ , thus with birth simplex  $\sigma_B$  inserted at step  $b$ , and death simplex  $\tau_B$  inserted at time  $d$ . Let  $K_B = K_{d-1}$  be its predeath complex. Let  $A \subseteq K^{(s)}$ , and write  $\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}}$  for the matching induced by the inclusion  $\mathcal{K} - A \hookrightarrow \mathcal{K}$ . We show that  $B \in \text{Im}(\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}})$  if and only  $[\partial\tau_B] \in \text{Im}(\iota_{K_B-A \hookrightarrow K_B}^*)$ .

Note that any representative cycle for  $B$  is homologous to a scalar multiple of  $\partial\tau_B$  in  $C_p(K_B)$ , as  $[\partial\tau_B]$  generates the kernel of  $H_p(K_B) \rightarrow H_p(K_B \cup \{\tau_B\})$ . If  $B \in \text{Im}(\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}})$ , there is thus a bar  $B' = [b', d)$  in the barcode of the image module  $\text{Im}(\text{PH}_p(\mathcal{K} - A) \rightarrow \text{PH}_p(\mathcal{K}))$ . Then, any representative cycle  $\xi \in C_p(K_{b'} - A) \subseteq C_p(K_{b'})$  for  $[b', d)$ , when considered as a homology class  $[\xi] \in H_p(K_B)$  through the inclusions  $K_{b'} - A \hookrightarrow K_{b'} \hookrightarrow K_B$ , is in the kernel of  $H_p(K_B) \rightarrow H_p(K_B \cup \{\tau_B\})$ , and is hence homologous to a scalar multiple of  $[\partial\tau_B]$ .

Conversely, suppose that  $[\partial\tau_B]$  is in the image of  $H_p(K_B - A) \rightarrow H_p(K_B)$ , and let  $\xi \in C_p(K_B - A)$  be a cycle such that  $[\xi] = [\partial\tau_B] \in H_p(K_B)$ . Let  $\xi_1, \dots, \xi_k$  denote representative cycles for summands of  $\text{Im}(\text{PH}_p(\mathcal{K} - A) \rightarrow \text{PH}_p(\mathcal{K}))$  so that the  $[\xi_i]$  are linearly independent in  $H_p(K_B)$ , and  $[\xi] = \lambda_1[\xi_1] + \dots + \lambda_k[\xi_k]$ . Since  $[\xi]$  is a non-zero element in the kernel of the map

$$\text{Im}\left(H_p(K_B - A) \rightarrow H_p(K_B)\right) \rightarrow \text{Im}\left(H_p((K_B \cup \{\tau_B\}) - A) \rightarrow H_p(K_B \cup \{\tau_B\})\right),$$

we see that one of the generators  $\xi_j$  must be in this kernel, and thus generate a summand isomorphic to  $\mathbb{I}_{[b', d)}$  for some  $b \leq b' < d$ . This shows that  $B$  is in the image of  $\mathcal{X}_{\mathcal{K}-A \hookrightarrow \mathcal{K}}$ . ◀

## 4 Complexity of the minimum homological cut problem

### 4.1 An efficient algorithm for zero-dimensional cuts

Let  $\gamma \in H_0(K)$ . We show that minimum homological vertex cuts for  $\gamma$  have a very particular structure, allowing us to compute them efficiently. Let  $\Gamma_1, \dots, \Gamma_q \subseteq V(K)$  be the vertex sets of the connected components of  $K$ , and let  $c_1 \in V(\Gamma_1), \dots, c_q \in V(\Gamma_q)$ . Then,  $H_0(K)$  is generated by the classes  $[c_1], \dots, [c_q]$ , meaning we can write  $\gamma = \sum_{i=1}^q \lambda_i [c_i]$ , with  $\lambda_i \in \mathbb{F}$ .

► **Proposition 10.** *Let  $\gamma = \sum_{i=1}^q \lambda_i [c_i] \in H_0(K)$  as in the above. A subset  $C \subseteq V(K)$  is a homological vertex cut for  $\gamma$  iff  $C \supseteq V(\Gamma_\ell)$  for some  $1 \leq \ell \leq q$  with  $\lambda_\ell \neq 0$ . In particular, a homological mincut for  $\gamma$  is of the form  $C = V(\Gamma_\ell)$ , where  $\ell = \arg \min_i \{ |V(\Gamma_i)| : \lambda_i \neq 0 \}$ .*

**Proof.** Write  $V = V(K)$ . For any  $A \subseteq V$ , and any  $v \in V \setminus A$  we have  $\iota_{K-A \hookrightarrow K}^*([v]) = [v]$ . Furthermore, for any  $v, w \in V$ , we have  $[v] = [w]$  (in  $H_0(K)$ ) if and only if  $v$  and  $w$  lie in the same connected component of  $K$ . Since the classes  $[v]$ ,  $v \in V \setminus A$ , generate  $H_0(K - A)$ , we thus have  $\gamma \in \text{Im}(\iota_{K-A \hookrightarrow K}^*)$  if and only if  $V \setminus A$  contains a vertex in each connected component  $\Gamma_i$  of  $K$  with  $\lambda_i \neq 0$ . Thus,  $\gamma \notin \text{Im}(\iota_{K-A \hookrightarrow K}^*)$  (i.e.,  $A$  is a homological cut) iff  $A$  contains all the vertices of at least one connected component  $\Gamma_\ell$  of  $K$  with  $\lambda_\ell \neq 0$ . ◀

**Proof of Theorem 6.** To determine a minimum homological vertex cut of a 0-dimensional class in a simplicial complex  $K$ , it suffices by Proposition 10 to find its connected components. We can do so in time linear in  $|V(K)| + |E(K)|$  using a depth-first search. ◀

### 4.2 NP-hardness of the general problem


In this section, we prove Theorem 7. That is, we show that the minimum homological cut problem is hard when  $p = s = 1$ . Our arguments apply for homology taken over any field, and even over  $\mathbb{Z}$ . We use a reduction from the following NP-hard problem.

► **Problem 11** (Exact cover by 3-sets (X3C)). *Given a set  $U$ , and a collection  $\mathcal{S} \subseteq \mathcal{P}(U)$  of subsets of  $U$ , each of size exactly 3, determine whether  $\mathcal{S}$  contains an exact cover of  $U$ , i.e., whether there is an  $\mathcal{S}' \subseteq \mathcal{S}$  so that each element of  $U$  occurs in exactly one element of  $\mathcal{S}'$ .*

We proceed as follows: To a given instance  $(U, \mathcal{S})$  of the X3C problem, we associate a topological space  $X = X(U, \mathcal{S})$ , together with a specific homology class  $\gamma \in H_1(X)$  (Construction 12). Then, we construct a triangulation  $K = K(U, \mathcal{S})$  of this space, which satisfies useful technical properties (Lemma 13). Finally, we show that if there exists a minimum cut for  $\gamma$  in  $K$  of a particular size, there exists a solution to the X3C instance, and vice versa.

► **Construction 12.** *Let  $(U, \mathcal{S})$  be an instance of exact cover by 3-sets. For each  $u \in U$ , denote by  $\mathcal{S}_u = \{S_{u,1}, \dots, S_{u,m_u}\} \subseteq \mathcal{S}$  the elements of  $\mathcal{S}$  containing  $u$ . Let  $F_u$  denote the surface obtained from a 2-dimensional disk by removing  $|\mathcal{S}_u|$  open disks from its interior. Label the outer boundary of  $F_u$  by  $u$ , and the interior boundary components by the respective elements of  $\mathcal{S}$  they correspond to. Choose an orientation for  $F_u$ , inducing an orientation on its boundary components. Choose generators for the homology groups of the boundary components such that  $\partial[F_u] = [u] + [S_{u,1}] + \dots + [S_{u,m_u}]$ ; see Figure 3.*

*We obtain  $X = X(U, \mathcal{S})$  from  $\bigsqcup_{u \in U} F_u$  by gluing together the surfaces in the following way: We identify all outer boundary components with each other along orientation preserving homeomorphisms, so that  $[u] = [v]$  in  $H_1(X)$  for each  $u, v \in U$ . We also identify all interior boundary components with each other that have the same label, again in an orientation preserving way. This means that if  $S = \{u, v, w\}$ , then after gluing  $S = S_{u,i_u} = S_{v,i_v} = S_{w,i_w}$ , and  $[S_{u,i_u}] = [S_{v,i_v}] = [S_{w,i_w}]$  in  $H_1(X)$ . Finally, we set  $\gamma = [u] \in H_1(X)$  for some  $u \in U$ .*



■ **Figure 3** Surfaces  $F_{u_1}$  and  $F_{u_2}$  involved in the construction of  $X(U, \mathcal{S})$  for  $(U, \mathcal{S}) = (\{u_1, u_2, u_3, u_4\}, \{S_1, S_2, S_3\})$  with  $S_1 = \{u_1, u_2, u_3\}$ ,  $S_2 = \{u_1, u_2, u_4\}$ ,  $S_3 = \{u_1, u_3, u_4\}$ . The surfaces are glued along boundary components with the same label, in an orientation preserving way.

► **Lemma 13.** *Let  $(U, \mathcal{S})$  be an instance of X3C, and let  $k = \max_{u \in U} |\{S \in \mathcal{S} \mid u \in S\}|$ . Set  $c = c(U, \mathcal{S}) := 5 \cdot \lceil \log_3(k) \rceil + 3$ . There exists a triangulation  $K = K(U, \mathcal{S})$  for the space  $X(U, \mathcal{S})$  of Construction 12, obtained by constructing a triangulation  $L_u$  for each surface  $F_u$  and gluing them together, that has the following properties:*

1.  $K$  is a pure 2-complex consisting of at most  $O(|U| \cdot |\mathcal{S}|)$  2-simplices;
2. A minimum homology edge cut for  $[u] \in H_1(L_u)$  consists of precisely  $c + 2$  edges;
3. A minimum homology edge cut for  $[u]$  contains precisely one edge from  $u$ , and precisely one edge from any of the interior boundary components  $S_{u,i}$ . For any edge in  $u$  and any edge in any  $S_{u,i}$ , there is a minimum cut for  $[u]$  containing both those edges.

**Proof.** We give an explicit construction in the full version of this paper [28]. ◀

Using Lemma 13, we can sketch our reduction; a complete proof is given in the full version [28].


**Proof sketch of Theorem 7.** We show that there exists a solution to the X3C-instance  $(U, \mathcal{S})$  if and only if there exists a cut for  $\gamma \in H_1(K(U, \mathcal{S}))$  consisting of at most  $c \cdot |U| + \frac{|U|}{3} + 1$  edges. If  $\mathcal{S}^* \subseteq \mathcal{S}$  is a solution to  $(U, \mathcal{S})$ , for each  $u \in U$ , we pick a cut  $C_u$  for  $[u] \in H_1(L_u)$  with the property that it contains one edge in the boundary labelled by  $u$ , and one in the boundary labeled by the set  $S_u \in \mathcal{S}^*$  for which  $u \in S_u$ . If we consistently choose the same edges in the intersections between surfaces, we can verify that  $C = \bigcup_{u \in U} C_u$  is a cut for  $\gamma$  consisting of precisely  $c \cdot |U| + \frac{|U|}{3} + 1$  edges. Conversely, suppose  $C$  is any cut for  $\gamma$ . Then it can be shown that for each  $u \in U$ ,  $C \cap L_u$  is a cut for  $[u] \in H_1(L_u)$ . By point 3 of Lemma 13, we then obtain a cover  $\mathcal{C}_C$  for  $U$ . It can be checked that  $|C| \geq c \cdot |U| + |\mathcal{C}_C| + 1$ . It follows that if  $|C| \leq c \cdot |U| + \frac{|U|}{3} + 1$ , then  $|\mathcal{C}_C| = \frac{|U|}{3}$ , and so  $\mathcal{C}_C$  is an exact cover. ◀

### 4.3 An efficient algorithm for embedded complexes

In this section, we prove Theorem 8. To construct an efficient algorithm, we rely on *Alexander duality*. Recall that, for an embedded complex  $K \subseteq \mathbb{R}^n$ , Alexander duality gives an isomorphism  $H_{n-1}(K) \cong \tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$  between  $(n-1)$ -dimensional homology and 0-dimensional (reduced) singular cohomology. For a subcomplex  $L \subseteq K$ , a class  $\gamma \in H_{n-1}(K)$  lies in the image of  $H_{n-1}(L) \rightarrow H_{n-1}(K)$  if and only if its Alexander dual lies in the image of the corresponding map  $\tilde{H}^0(\mathbb{R}^n \setminus \|L\|) \rightarrow \tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$ . The latter map can be understood via an extension of the *dual graph* of  $K$ .

► **Definition 14** (Dual graph). Let  $K$  be a simplicial complex. We denote by  $\mathcal{G}_K$  the dual graph of  $K$ , which has a vertex for each  $n$ -simplex of  $K$ , and which has an edge between each pair of vertices whose corresponding simplices share an  $(n-1)$ -dimensional face.

► **Definition 15** (Extended dual graph). For an embedded simplicial complex  $K$ , the extended dual graph  $\mathcal{G}_K^*$  of  $K$  has one vertex for each  $n$ -simplex of  $K$  and one vertex for each component of  $\mathbb{R}^n \setminus \|K\|$ . For each vertex  $v$ , we denote by  $R_v \subseteq \mathbb{R}^n$  the interior of either the simplex or the connected component of  $\mathbb{R}^n \setminus \|K\|$  to which  $v$  corresponds. For each  $(n-1)$ -simplex  $\tau$  of  $K$ ,  $\mathcal{G}_K^*$  has an edge  $(v, w)$  between the vertices for which  $\tau \subseteq \overline{R_v} \cap \overline{R_w}$  (possibly allowing loops and multi edges). We denote the set of vertices corresponding to connected components of  $\mathbb{R}^n \setminus \|K\|$  by  $V_{\mathbb{R}^n \setminus K}$ , and the vertex corresponding to the unbounded component by  $v_\infty$ .



■ **Figure 4** Left: The embedded simplicial complex  $K$  of Figure 1 with a class  $\gamma \in H_1(K)$  (blue, single arrows). Center: The (extended) dual graph of  $K$ . The vertices  $V_{\mathbb{R}^n \setminus K} = \{v_1, v_2, v_\infty\}$  and edges added to obtain  $\mathcal{G}_K^*$  from  $\mathcal{G}_K$  are in red. Only some of the edges incident to  $v_\infty$  are drawn. Right: A shortest path  $v_1 \rightarrow v_\infty$  (green, single arrows) and a shortest path  $v_1 \rightarrow v_2$  (purple, double arrows) in  $\mathcal{G}_K^*$ . These correspond to the (minimum) 1-cuts for  $\gamma$  depicted in Figure 1, see Proposition 18. The paths are directed for visual clarity only.

The outline for the rest of this section is as follows. First, we show that one can associate a subgraph  $\mathcal{G}_C$  of  $\mathcal{G}_K^*$  to any  $C \subseteq K^{(n-1)}$  so that  $H_{n-1}(K - C) \cong \tilde{H}^0(\mathcal{G}_C)$  (and this association is functorial). In case  $C = \emptyset$ ,  $\mathcal{G}_C = \mathcal{G}_K$  and this comes down to Alexander duality, and we see that any  $\gamma \in H_{n-1}(K)$  can be written as a sum of homology classes corresponding to vertices in  $V_{\mathbb{R}^n \setminus K}$ . We then show that whether  $C$  is a homology cut for  $\gamma$  depends on the existence of paths in  $\mathcal{G}_C$  between specific vertices in  $V_{\mathbb{R}^n \setminus K}$ . This allows us to reduce (MHC) to a (number of) shortest path problems in  $\mathcal{G}_K^*$ . See Figure 4 for an illustration.

► **Lemma 16.** Let  $K \subseteq \mathbb{R}^n$  be an embedded simplicial complex, and let  $C \subseteq K^{(n-1)}$ . Denote by  $\mathcal{G}_C \subseteq \mathcal{G}_K^*$  the smallest subgraph of  $\mathcal{G}_K^*$  containing both  $V_{\mathbb{R}^n \setminus K}$  and the dual edges corresponding to the simplices in  $C$ . Then there is an isomorphism  $H_{n-1}(K - C) \cong \tilde{H}^0(\mathcal{G}_C)$ . Moreover, if  $C \subseteq C'$ , then the following diagram commutes:

$$\begin{array}{ccc} H_{n-1}(K - C') & \xrightarrow{\cong} & \tilde{H}^0(\mathcal{G}_{C'}) \\ \downarrow & & \downarrow \\ H_{n-1}(K - C) & \xrightarrow{\cong} & \tilde{H}^0(\mathcal{G}_C) \end{array}$$

**Proof.** See the full version of this paper [28]. ◀

Recall that  $H^0(\mathbb{R}^n \setminus \|K\|)$  consists of functions  $f : \mathbb{R}^n \setminus \|K\| \rightarrow \mathbb{R}$  that are constant on each connected component of  $\mathbb{R}^n \setminus \|K\|$ . We represent  $\tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$  by locally constant functions that are 0 on the unbounded connected component. Similarly, we think of  $\tilde{H}^0(\mathcal{G}_C)$

as functions that are constant on the components of  $\mathcal{G}_C$ , and 0 on the component containing  $v_\infty$ . Moreover, if  $C \subseteq C'$ , the map  $\tilde{H}^0(\mathcal{G}_{C'}) \rightarrow \tilde{H}^0(\mathcal{G}_C)$  is given by a restriction of functions. The following proposition allows us to relate cuts for specific elements of  $H_{n-1}(K)$  to paths in the extended dual graph  $\mathcal{G}_K^*$  of  $K$ .

► **Proposition 17.** *Let  $v \in V_{\mathbb{R}^n \setminus K} \setminus \{v_\infty\}$ . Consider the map  $\mathbf{1}_v \in \tilde{H}^0(\mathcal{G}_\emptyset) \cong \tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$ . A set  $C \subseteq K^{(n-1)}$  is a homological cut for the Alexander dual  $\text{AD}(\mathbf{1}_v) \in H_{n-1}(K)$  iff the set of edges in  $\mathcal{G}_K^*$  dual to the simplices of  $C$  contains a path from  $v$  to another vertex in  $V_{\mathbb{R}^n \setminus K}$ . (Here,  $\text{AD}(\mathbf{1}_v)$  is meant to denote the Alexander dual of  $\mathbf{1}_{R_v} \in \tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$ .)*

**Proof.** By Lemma 16,  $\text{AD}(\mathbf{1}_v)$  is in the image of  $H_{n-1}(K - C) \rightarrow H_{n-1}(K)$  if and only if  $\mathbf{1}_v$  is in the image of  $\tilde{H}^0(\mathcal{G}_C) \rightarrow \tilde{H}^0(\mathcal{G}_\emptyset)$ . This is the case if and only if the function  $\mathbf{1}_v$  extends to a function  $\varphi \in \tilde{H}^0(\mathcal{G}_C)$ , which is possible if and only if  $v$  is in a different component from any other vertex of  $V_{\mathbb{R}^n \setminus K}$  inside  $\mathcal{G}_C$ . That means that  $C$  is a homology cut for  $\text{AD}(\mathbf{1}_v)$  if and only if there is a path in  $\mathcal{G}_C$  from  $v$  to some other vertex of  $V_{\mathbb{R}^n \setminus K}$ , proving our claim. ◀

We need the following generalization of the above proposition, that allows us to relate cuts for arbitrary elements of  $H_1(K)$  to paths in  $\mathcal{G}_K^*$ ; see Figure 4 for an example.

► **Proposition 18.** *Let  $\gamma = \sum_{v \in V_{\mathbb{R}^n \setminus K}} \text{AD}(\alpha_v \mathbf{1}_v)$  with  $\alpha_{v_\infty} = 0$  be any element of  $H_{n-1}(K)$ . Then  $C$  is a homology cut for  $\gamma$  if and only if the set of edges dual to the simplices of  $C$  contains a path between a pair of vertices  $v, w \in V_{\mathbb{R}^n \setminus K}$  for which  $\alpha_v \neq \alpha_w$ .*

**Proof.** The proof is similar to that of Proposition 17, with the following adjustment: The function  $\psi = \sum_{v \in V_{\mathbb{R}^n \setminus K}} \alpha_v \mathbf{1}_v \in \tilde{H}^0(\mathcal{G}_\emptyset)$  extends to a function  $\varphi \in \tilde{H}^0(\mathcal{G}_C)$  iff no pair of vertices  $v, w \in V_{\mathbb{R}^n \setminus K}$  for which  $\psi(v) = \alpha_v \neq \alpha_w = \psi(w)$  is in the same component of  $\mathcal{G}_C$ . ◀

**Proof of Theorem 8.** To compute a minimum homological cut for a class  $\gamma \in H_{n-1}(K)$  of an  $n$ -complex  $K$  embedded in  $\mathbb{R}^n$ , we proceed in three steps:

1. **Constructing the extended dual graph.** First, we construct  $\mathcal{G}_K^*$ . For this we use the Void Boundary Reconstruction algorithm [18, Section 4], which outputs sets of oriented simplices  $(\vec{\xi}_1, \dots, \vec{\xi}_k)$  that constitute the boundaries of the components of  $\mathbb{R}^n \setminus \|K\|$  (cf. Figure 4) in time bounded by  $O(|K|^2)$ . We assume wlog that  $\vec{\xi}_k$  constitutes the boundary of the unbounded component of  $\mathbb{R}^n \setminus \|K\|$ . The sum of oriented simplices in any of the  $\vec{\xi}_i$  with  $i < k$  bounding a component  $R_i \subseteq \mathbb{R}^n \setminus \|K\|$ , which we denote by  $[\vec{\xi}_i]$ , represents the Alexander dual to the function  $f_i \in \tilde{H}^0(\mathbb{R}^n \setminus \|K\|)$  that is constant 1 on  $R_i$ , and 0 everywhere else. From the  $(\vec{\xi}_1, \dots, \vec{\xi}_k)$  we can construct  $\mathcal{G}_K^*$  in time linear in  $k + |K|$ . Each  $\vec{\xi}_i$  corresponds to a vertex  $v_i = v_{R_i} \in V_{\mathbb{R}^n \setminus K} \subseteq V(\mathcal{G}_K^*)$ .
2. **Expressing  $\gamma$  in an appropriate basis.** We write  $\gamma$  as  $\gamma = \sum_{i=1}^k \alpha_i [\vec{\xi}_i]$  with  $\alpha_i \in \mathbb{F}$ ,  $\alpha_k = 0$ . This comes down to a change of basis, which can be done in matrix-multiplication time. The representation of  $\gamma$  is precisely the representation as assumed in Proposition 18, as  $[\vec{\xi}_i]$  is the Alexander dual of  $\mathbf{1}_{v_i}$ .
3. **Determining shortest paths.** Using breadth-first search, we can find the minimum distance between any two vertices  $v_i, v_j \in V_{\mathbb{R}^n \setminus K}$  inside  $\mathcal{G}_K^*$  for which  $\alpha_i \neq \alpha_j$  in time bounded by  $O(k \cdot |\mathcal{G}_K^*|)$ . Let  $P$  denote a path in  $\mathcal{G}_K^*$  achieving this minimum. The set of  $(n-1)$ -simplices  $C_P$  dual to the edges constituting  $P$  form a cut for  $\gamma$  by Proposition 18. Since the edges dual to any cut  $C$  for  $\gamma$  must contain a path between vertices of  $V_{\mathbb{R}^n \setminus K}$  with differing coefficients, the cut  $C_P$  is a minimum cut by minimality of  $P$ . ◀

## 5 Adversarial robustness for the Rips filtration

Let  $X = (X, d)$  be a (finite) metric space. For a subset  $\sigma \subseteq X$ , write  $\mathbf{r}(\sigma)$  for its diameter, i.e., the largest pairwise distance between points in  $\sigma$ . The Rips filtration of  $X$  is typically defined as the  $\mathbb{R}$ -indexed family of simplicial complexes  $\mathcal{R}(X) := (\{\sigma \subseteq X : \mathbf{r}(\sigma) \leq r\})_{r \in \mathbb{R}}$ . If we choose a total order  $\preceq$  on the points in  $X$ , we can transform  $\mathcal{R}(X)$  into a discrete, simplex-wise filtration. Namely, we insert simplices one-by-one according to their diameter, breaking ties first by simplex dimension, and then by the lexicographic order on subsets of  $X$  induced by  $\preceq$ . We denote the resulting *simplex-wise Rips filtration* by  $\mathcal{R}^\preceq(X)$ . The *length* of a bar  $B = [\sigma_B, \tau_B)$  in the barcode of  $\mathrm{PH}_p(\mathcal{R}^\preceq(X))$  is  $\ell(B) := \mathbf{r}(\tau_B) - \mathbf{r}(\sigma_B)$ .

We consider the following heuristic to determine adversarial robustness of bars in the barcode of  $\mathrm{PH}_p(\mathcal{R}^\preceq(X))$  based on their length. We denote the Hausdorff distance between  $A, B \subseteq X$  by  $d_H(A, B) := \max\{\sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A)\}$ .

► **Definition 19** (Hausdorff heuristic). *Let  $X$  be a finite metric space. For  $k \in \mathbb{N}$ , we define*

$$\mathcal{H}_{X,k} := \max_{A \subseteq X, |A| \leq k} d_H(X \setminus A, X). \quad (1)$$

The Hausdorff heuristic has the following two key properties.

► **Theorem 20.** *Let  $p \geq 0$ , and let  $B$  be a bar in the barcode of  $\mathrm{PH}_p(\mathcal{R}^\preceq(X))$ . If  $B$  has length  $\ell(B) \geq \mathcal{H}_{X,k}$ , then  $B$  is  $k$ -adversarially robust in degree 0.*

► **Proposition 21.** *The parameter  $\mathcal{H}_{X,k}$  can be computed in time  $O(|X|^2 \log |X|)$ .*

**Proof sketch of Theorem 20.** Consider the  $\mathbb{R}$ -indexed Rips filtration  $\mathcal{R}(X)$ . Let  $A \subseteq X$ , and set  $\delta = d_H(X \setminus A, X)$ . For any  $r \in \mathbb{R}$ , there is a map  $H_p(\mathcal{R}_r(X)) \rightarrow H_p(\mathcal{R}_{r+\delta}(X - A))$  making the following diagram commute:

$$\begin{array}{ccc} H_p(\mathcal{R}_r(X \setminus A)) & \longrightarrow & H_p(\mathcal{R}_{r+\delta}(X \setminus A)) \\ \downarrow & \nearrow & \downarrow \\ H_p(\mathcal{R}_r(X)) & \longrightarrow & H_p(\mathcal{R}_{r+\delta}(X)) \end{array}$$

This can be shown similarly to how stability of Rips persistence is proved, e.g., as in [13]. It follows that the cokernel of  $H_p(\mathcal{R}(X \setminus A)) \rightarrow H_p(\mathcal{R}(X))$  is  $\delta$ -trivial, which would allow us to apply the induced matching theorem (Theorem 9).

To transport this idea to the simplex-wise Rips filtration  $\mathcal{R}^\preceq(X)$ , we make use of an “ $\varepsilon$ -smoothing”. That is, an  $\mathbb{R}$ -indexed simplex-wise filtration  $\mathcal{R}^\varepsilon(X)$ , which inserts simplices in the same order as  $\mathcal{R}^\preceq(X)$ , and which is  $\varepsilon$ -interleaved with  $\mathcal{R}(X)$ , i.e.,

$$\mathcal{R}_r^\varepsilon(X) \subseteq \mathcal{R}_r(X) \subseteq \mathcal{R}_{r+\varepsilon}^\varepsilon(X) \quad \forall r \in \mathbb{R}.$$

Such a filtration exists for any  $\varepsilon > 0$  small enough. Importantly, the persistence pairings and induced matchings of  $\mathcal{R}^\preceq(X)$  and  $\mathcal{R}^\varepsilon(X)$  agree. By the previous, the cokernel of  $H_p(\mathcal{R}^\varepsilon(X \setminus A)) \rightarrow H_p(\mathcal{R}^\varepsilon(X))$  is  $\delta + 2\varepsilon$ -trivial. By Theorem 9, every bar of length greater than  $\delta + 2\varepsilon$  in the barcode of  $\mathcal{R}^\varepsilon(X)$  is in the image of the matching induced by  $\mathcal{R}^\varepsilon(X \setminus A) \hookrightarrow \mathcal{R}^\varepsilon(X)$ . Making  $\varepsilon$  arbitrarily small, the theorem follows. For details, see the full version [28]. ◀

**Proof of Proposition 21.** To compute  $\mathcal{H}_{X,k}$  efficiently, we use the following structural lemma on the sets  $A$  that attain the maximum in (1); for its proof, see the full version [28].

► **Lemma 22.** For  $i \in \mathbb{N}$  and  $x \in X$ , let  $\nu_i(x)$  denote the  $i$ th nearest neighbor of  $x$ . The set  $A \subseteq X$  attaining the maximum in (1) is of the form  $\{x\} \cup \{\nu_i(x) : i = 1, \dots, k\}$ , with  $x \in X$ . In particular, we have  $\mathcal{H}_{X,k} = \max_{x \in X} d(x, \nu_k(x))$ .

For each  $x \in X$ , we can compute  $\nu_k(x)$  by sorting the points in  $X$  by their distance to  $x$ , which can be done in  $O(|X| \log |X|)$ . This immediately proves Proposition 21. ◀

---

## References

- 1 Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(9):3821–3850, 2006. doi:10.1090/S0002-9947-05-03919-X.
- 2 Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *J. Comput. Geom.*, 6:162–191, 2013.
- 3 Nello Blaser, Morten Brun, Lars M. Salbu, and Erlend Raa Vågset. The parameterized complexity of finding minimum bounded chains. *Computational Geometry*, 122:102102, 2024. doi:10.1016/j.comgeo.2024.102102.
- 4 Andrew J. Blumberg and Michael Lesnick. Stability of 2-parameter persistent homology. *Foundations of Computational Mathematics*, 24(2):385–427, 2022. doi:10.1007/s10208-022-09576-6.
- 5 Glencora Borradaile, William Maxwell, and Amir Nayyeri. Minimum Bounded Chains and Minimum Homologous Chains in Embedded Simplicial Complexes. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2020.21.
- 6 Magnus Bakke Botnan and William Crawley-Boevey. Decomposition of persistence modules. *Proc. Am. Math. Soc.*, 148(11):4581–4596, 2020.
- 7 Magnus Bakke Botnan and Michael Lesnick. An introduction to multiparameter persistence. *Preprint, arXiv:2203.14289*, 2023. doi:10.48550/arXiv.2203.14289.
- 8 Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. *Comput. Geom.*, 58:70–96, 2016. doi:10.1016/J.COMGEO.2016.07.001.
- 9 Chris Buehler, Steven J. Gortler, Michael F. Cohen, and Leonard McMillan. *Minimal Surfaces for Stereo*, pages 885–899. Springer Berlin Heidelberg, 2002. doi:10.1007/3-540-47977-5\_58.
- 10 Erin W. Chambers, Jeff Erickson, Kyle Fox, and Amir Nayyeri. Minimum cuts in surface graphs. *SIAM Journal on Computing*, 52(1):156–195, 2023. doi:10.1137/19M1291820.
- 11 Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Homology flows, cohomology cuts. *SIAM Journal on Computing*, 41(6):1605–1634, 2012. doi:10.1137/090766863.
- 12 Erin W. Chambers and Mikael Vejdemo-Johansson. Computing minimum area homologies. *Computer Graphics Forum*, 34(6):13–21, 2015. doi:10.1111/cgf.12514.
- 13 Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Memoli, and Steve Y. Oudot. Gromov-Hausdorff Stable Signatures for Shapes using Persistence. *Computer Graphics Forum*, 2009.
- 14 Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Found. Comput. Math.*, 11:733–751, 2011. doi:10.1007/S10208-011-9098-0.
- 15 Chao Chen and Daniel Freedman. Hardness results for homology localization. *Discrete Comput. Geom.*, 45(3):425–448, 2011. doi:10.1007/S00454-010-9322-8.
- 16 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37:103–120, 2007. doi:10.1007/S00454-006-1276-5.
- 17 Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM J. Comput.*, 40(4):1026–1044, 2011. doi:10.1137/100800245.

- 18 Tamal K. Dey, Tao Hou, and Sayan Mandal. Computing minimal persistent cycles: polynomial and hard cases. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, pages 2587–2606, USA, 2020. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611975994.158.
- 19 Tamal K. Dey and Yusu Wang. *Computational Topology for Data Analysis*. Cambridge University Press, 2022.
- 20 Nathan M. Dunfield and Anil N. Hirani. The least spanning area of a knot and the optimal bounding chain problem. In *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry*, SoCG '11, pages 135–144, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1998196.1998218.
- 21 Herbert Edelsbrunner and John Harer. Persistent homology—a survey. *Discrete Comput. Geom.*, 453, 2008.
- 22 Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. In *Proceedings of the 2011 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1166–1176, 2011. doi:10.1137/1.9781611973082.88.
- 23 Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990.
- 24 Pepijn Roos Hoefgeest and Lucas Slot. The Christoffel-Darboux kernel for topological data analysis. In *39th International Symposium on Computational Geometry*, volume 258 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 38, 20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/lipics.socg.2023.38.
- 25 B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, 1995. doi:10.1137/S0097539792240406.
- 26 Steve Y. Oudot. *Persistence Theory - From Quiver Representations to Data Analysis*, volume 209 of *Mathematical surveys and monographs*. American Mathematical Society, 2015. URL: <http://bookstore.ams.org/surv-209/>.
- 27 Jeff M. Phillips, Bei Wang, and Yan Zheng. Geometric inference on kernel density estimates. In *31st International Symposium on Computational Geometry*, volume 34 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 857–871. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.SOCG.2015.857.
- 28 Pepijn Roos Hoefgeest and Lucas Slot. Robustness of persistent topological features and minimum homological cuts (full version), 2026. URL: <https://arxiv.org/abs/2602.23154>.
- 29 Vin Silva and Gunnar Carlsson. Topological estimation using witness complexes. *Proc. Sympos. Point-Based Graphics*, 2004.
- 30 Bernadette J. Stolz. Outlier-robust subsampling techniques for persistent homology. *J. Mach. Learn. Res.*, 24:Paper No. [90], 35, 2023.
- 31 J. Sullivan. *A Crystalline Approximation Theorem for Hypersurfaces*. PhD thesis, Princeton University, 1990.
- 32 Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33:249–274, 2005. doi:10.1007/S00454-004-1146-Y.