

---

# On the Expressivity of Persistent Homology in Graph Learning

---

**Rubén Ballester**

Departament de Matemàtiques i Informàtica  
Universitat de Barcelona  
ruben.ballester@ub.edu

**Bastian Rieck**

Department of Computer Science  
University of Fribourg  
bastian.grossenbacher@unifr.ch

## Abstract

Persistent homology, a technique from computational topology, has recently shown strong empirical performance in the context of graph classification. Being able to capture long range graph properties via higher-order topological features, such as cycles of arbitrary length, in combination with multi-scale topological descriptors, has improved predictive performance for data sets with prominent topological structures, such as molecules. At the same time, the *theoretical properties* of persistent homology have not been formally assessed in this context. This paper intends to bridge the gap between computational topology and graph machine learning by providing a brief introduction to persistent homology in the context of graphs, as well as a theoretical discussion and empirical analysis of its expressivity for graph learning tasks.

## 1 Introduction

Graph learning is a highly-active research domain in machine learning, fuelled in large parts by the *geometric deep learning* [17, 18] paradigm as well as the resurgence of new neural network architectures for handling graph data. Methods from computational topology, by contrast, have not yet been applied in this domain at large scales. Even though a large amount of prior work employs topological features to solve graph learning tasks [22, 25, 44, 45, 46, 48, 78, 97, 99, 102, 103], a formal investigation relating expressivity in graph learning and topological machine learning is still lacking.<sup>1</sup> We believe that this is largely a deficit in communication between the communities. This paper provides an introduction to topological methods for graph learning, while also showing new theoretical and empirical results about the *expressivity* of topological graph learning methods. Here, we understand expressivity as a general concept to signify which graph properties can be captured by a method. This includes being aware of certain substructures in graphs [27], for instance, but also being able to distinguish large classes of non-isomorphic graphs [16, 51, 83]. While graph neural networks have demonstrated substantial gains in this area, our paper focuses on topology-based algorithms, aiming to provide a better understanding of their theoretical and empirical properties.

**Contributions.** Our main theoretical contribution is a full characterisation of the expressivity of *persistent homology* in terms of the Weisfeiler–Leman hierarchy [66, 96]. We prove that persistent homology is *at least as expressive* as a corresponding Weisfeiler–Leman test for graph isomorphism. Moreover, we show that there exist graphs that cannot be distinguished using  $k$ -FWL, the *folklore Weisfeiler–Leman algorithm* [65], for a specific dimension  $k$  but that can be distinguished by persistent homology (with or without access to  $k$ -cliques in the graph). Along the way, we also prove new properties of filtrations, i.e. descriptor functions that are commonly employed to obtain topological representations of graphs, hinting at their ability to capture information about graph substructures. We complement our theoretical expressivity discussions by an experimental suite that highlights the capabilities of different filtrations for (i) distinguishing certain types of graphs, (ii) predicting characteristic graph properties, and (iii) serving as a baseline for classification tasks.<sup>2</sup>

<sup>1</sup>A notable exception is recent work by Immonen et al. [50], which analyses the expressivity of topological methods for graph classification tasks.

<sup>2</sup>Our code is available at [https://github.com/aidos-lab/PH\\_expressivity](https://github.com/aidos-lab/PH_expressivity).




FIGURE 1: An example graph and three different steps of a degree-based filtration. The respective caption indicates the pre-image of the corresponding filtration function.

**Guide for readers.** Section 2 briefly summarises the main concepts in graph learning. Section 2 and Section 2, may be skipped by readers that are already well versed in computational topology. Section 3 outlines advantageous properties of filtrations in the context of graph learning, while Section 4 discusses the expressivity of persistent homology with respect to the Weisfeiler–Leman hierarchy of graph isomorphism tests. Section 5 concludes the paper with an experimental suite that highlights the performance of topological methods in a variety of graph-learning tasks.

## 2 Background & Notation




FIGURE 2: Persistence diagrams of the filtration depicted in Fig. 1. Points can have different multiplicities; we show *essential features*, i.e. topological features that persist over the full filtration using an  $\infty$  symbol.

We deal with undirected graphs in this paper. An undirected graph  $G$  is a pair  $G = (V, E)$  of finite sets of  $n$  vertices and  $m$  edges, with  $E \subseteq \{\{u, v\} \mid u, v \in V, u \neq v\}$ . We will also refer to an edge using tuple notation, with the understanding that  $(u, v)$  and  $(v, u)$  refer to the same edge. Furthermore, we denote the set of all such graphs by  $\mathcal{G}$ . Two graphs  $G = (V, E)$  and  $G' = (V', E')$  are *isomorphic*, denoted by  $G \simeq G'$ , if there is a bijective function  $\varphi: V \rightarrow V'$  that preserves adjacency, i.e.  $(u, v) \in E$  if and only if  $(\varphi(u), \varphi(v)) \in E'$ . Since  $\varphi$  is bijective, it has an inverse function, which we will denote by  $\varphi^{-1}$ . The problem of figuring out whether two graphs are isomorphic or not is referred to as the *graph isomorphism problem*. Presently, there is no known algorithm that solves this problem in polynomial time—efficient algorithms exist only for special families of graphs [30, 52]. Hence, all subsequently-discussed graph isomorphism tests are perforce limited with respect to their expressivity, i.e. there exist classes of non-isomorphic graphs that they cannot distinguish. In the context of graph isomorphism tests, we will often require the definition of a *multiset*, which is a set whose elements are included with multiplicities. We will denote such a multiset by  $\{\!\!\{\}$ .

In this paper, we use the term *graph expressivity* to denote two different concepts: first, the ability of a method to distinguish non-isomorphic graphs, and second, the ability of a method to capture certain graph properties.

**Definition 1** (Expressivity as distinguishing non-isomorphic graphs). *Let  $f_1: \mathcal{G} \rightarrow \mathcal{Y}_1$  and  $f_2: \mathcal{G} \rightarrow \mathcal{Y}_2$  be two functions such that for two isomorphic graphs  $G, G'$  we have that  $f_1(G) = f_1(G')$  and  $f_2(G) = f_2(G')$ . We say that  $f_1$  is at least as expressive as  $f_2$  if for any pair of non-isomorphic graphs  $G, G'$  we have that  $f_2(G) \neq f_2(G')$  implies  $f_1(G) \neq f_1(G')$ .*

**Definition 2** (Expressivity as detecting graph properties). *Let  $f_1, f_2, t: \mathcal{G} \rightarrow \mathbb{R}^e$  be functions for some  $e$ . We say that  $f_1$  is at least as expressive as  $f_2$  according to property  $t$  if for any graph  $G$  we have that  $f_2(G) = t$  implies  $f_1(G) = t$ .*

In both cases, if  $f_1$  is at least as expressive as  $f_2$  and  $f_2$  is at least as expressive as  $f_1$ , we say that  $f_1$  and  $f_2$  are equally expressive. A more complete characterization of expressivity can be found in Li and Leskovec [58, Definition 5.5].

**Topological Features of Graphs.** The simplest kind of topological features to prescribe to graphs are *connected components* and *cycles*. Formally, we refer to the number of connected components and the number of (independent) cycles in a graph as the first two *Betti numbers* of a graph, denoted

as  $\beta_0$  and  $\beta_1$ , respectively. While the expressive power of these two numbers is limited, it can be improved by evaluating them alongside a *filtration*, i.e. a sequence of nested subgraphs of the form  $\emptyset \subseteq G_0 \subseteq G_1 \dots \subseteq G_{k-1} \subseteq G_k = G$ . Filtrations typically arise from scalar-valued functions of the form  $f: G \rightarrow \mathbb{R}$ , which assign vertices and edges a value. Concretely, a filtration sequence of indices requires sorting the values of the image of the filtration function  $f$ , with subgraphs  $G_a$  being assigned the *pre-images* of intervals  $(-\infty, a]$  under  $f$ , i.e.  $G_a = f^{-1}((-\infty, a])$ . Filtrations are usually obtained from more general maps  $F: \mathcal{G} \rightarrow \mathbb{R}^{\mathcal{G}}, G \mapsto f_G$ , which assign to each graph a filtration function. In this case, we say that  $F$  is *equivariant* if, given any two graphs  $G, G'$  and any graph isomorphism  $\varphi: G \rightarrow G'$ , we have that  $f_G = f_{G'} \circ \varphi$ . Changes in the Betti numbers can be tracked over the course of the filtration. This leads to *persistent Betti numbers*, which are typically summarised in a *persistence diagram*, i.e. a topological descriptor, consisting of tuples  $(a_i, a_j) \in \mathbb{R}^2$ , with  $a_i$  referring to the value at which a topological feature was *created*, and  $a_j$  referring to the value at which a topological feature was destroyed (for instance, because two connected components are being merged). The absolute difference in function values  $|a_j - a_i|$  is called the *persistence* of a topological feature; it indicates the prominence or relevance of said feature. Fig. 1 depicts an example filtration for a simple graph, where we use the degree of each vertex to filter the graph. The more complex structure of persistence diagrams (in comparison to the simple Betti numbers) already hints at their capabilities in providing expressive graph descriptors.

**Topological Features of Simplicial Complexes.** Filtrations and persistence diagrams generalise to higher dimensions as well, with the understanding that the connectivity of higher-order structures of a graph—cliques—is being modelled. The resulting framework is referred to as *persistent homology*. There are different constructions for obtaining simplicial complexes from graphs; we will subsequently use *clique complexes* (in which a  $k$ -clique is represented by a  $(k - 1)$ -simplex) because it is (i) straightforward to implement, and (ii) cliques are known to be characteristic graph structures [16]. Persistent homology is typically calculated using matrix reduction algorithms, whose optimisation is still an ongoing topic of research [10]. We refer readers to Otter et al. [73] for a comprehensive introduction of computational strategies. Readers are invited to read Appendix C for a more detailed exposition of concepts in computational topology.

### 3 Properties of Filtrations

As in the graph case, it is common to obtain simplicial complex filtrations from a general map  $F: \mathcal{K} \mapsto f_{\mathcal{K}} \in \mathbb{R}^{\mathcal{K}}$  where  $\mathcal{K}$  is the set of all finite simplicial complexes. In the scenarios where we construct simplicial complexes from graphs, we can promote  $F$  to a function  $\mathcal{F}: \mathcal{G} \mapsto (K_G, f_{K_G}) \in \mathcal{K} \times \mathbb{R}^{\mathcal{K}}$  that assigns to each graph a simplicial complex with vertex set equal to the set of vertices of the graph and a filtration function for this simplicial complex. We say that  $\mathcal{F}$  is *equivariant* if for each isomorphism  $\varphi: G \rightarrow G'$  we have that  $\varphi$  induces a simplicial complex isomorphism between  $K_G$  and  $K_{G'}$  and  $f_{K_G} = f_{K_{G'}} \circ \varphi$ . If  $K_G = G$  for all  $G \in \mathcal{G}$ , then we recover the previous definition of equivariance for the graph case. A crucial result of the previous maps  $\mathcal{F}$  is that, if they are equivariant, then the persistence diagrams of the filtrations they induce are invariant under graph isomorphism.

**Proposition 1.** *For  $\mathcal{F}$  equivariant and  $G \simeq G'$ , the persistence diagrams of  $f_{K_G}$  and  $f_{K_{G'}}$  coincide.*

Proposition 1 shows that it is impossible to ‘adversarially’ pick an equivariant filtration generator function  $\mathcal{F}$  that leads to a dissimilarity between two non-isomorphic graphs. Dropping this condition incurs a substantial loss of expressive power. The proposition thus guarantees that persistent homology is compatible with equivariant function learning, pointing towards the utility of hybrid models that leverage different types of structural properties of graphs. To finish this discussion of general properties, we remark that the calculation of persistent homology carries information about the *diameter* (the length of the longest shortest path) and *girth* (the length of the shortest cycle) of a graph.

**Proposition 2.** *Given any filtration  $f$  of a graph  $G$  with a single connected component such that the values of  $f$  of the endpoints of edges are strictly lower than the values of their corresponding edges, Algorithm 1, used to compute  $\beta_0$ , yields an upper bound of  $\text{diam}(G)$ .*

**Proposition 3.** *Given any filtration  $f$  of a graph  $G$ , Algorithm 1 yields an upper bound of the girth of  $G$ .*

While the theoretical upper bounds are not tight, our empirical analysis in Section 5.3 shows that persistent homology and its algorithms captures more than ‘just’ topological information about a

graph. Our work thus corroborates recent results in the setting of point clouds, where persistence diagrams permit inferring additional properties about the input data [19, 60, 88].

## 4 The Weisfeiler–Leman Hierarchy

Having discussed the properties of specific filtrations, we now analyse the expressivity of persistent homology in the context of the Weisfeiler–Leman hierarchy of graph isomorphism tests. To this end, we categorise our filtrations (and their corresponding generators) into two distinct classes:

1. Filtrations operating on clique complexes derived from graphs
2. Filtrations that consider all possible cliques of the vertex set up to a specified maximum dimension.

In the first class, we incorporate only those cliques present in the original graph, constrained to a predetermined maximum dimension. Conversely, in the second class, we include *all* feasible cliques up to a specified dimension, irrespective of their presence in the original graph structure.

1-WL, also known as *colour refinement*, constitutes a simple method for addressing the graph isomorphism problem. It is the backbone of graph expressivity research; readers are referred to Morris et al. [65] for a comprehensive survey of 1-WL and its higher-order variants. It is already a known result that *any* 1-WL colouring can be reproduced by creating a special filtration [48, Theorem 4]; we restate this result as [Theorem 3](#) in the appendix. For a longer discussion on the 1-WL test, we refer readers to [Appendix D](#). Since 1-WL is also unable to distinguish between graphs with different triangle counts or graphs with cycle information, it was generalised to include information about labelling *tuples* of  $k$  nodes (as opposed to only labelling a single node), leading to a hierarchy of algorithms. The variant we shall subsequently describe is also known as the *folklore Weisfeiler–Leman algorithm* [65]. It can be shown that there are non-isomorphic graphs that cannot be distinguished by  $k$ -FWL, but that can be distinguished by  $(k + 1)$ -FWL.<sup>3</sup> A well-known family of such non-isomorphic, non-distinguishable graphs are commonly known as CFI graphs [20]. Following Morris et al. [65],  $k$ -FWL is based on the idea of assigning colours to *subgraphs* as opposed to assigning colours to *vertices*. To achieve this,  $k$ -FWL operates on  $k$ -tuples of vertices; for iteration  $i = 0$ , two tuples  $\mathbf{v} = (v_1, \dots, v_k)$  and  $\mathbf{w} = (w_1, \dots, w_k)$  are assigned the same colour if the map  $v_j \mapsto w_j$  induces a graph isomorphism between the subgraphs induced by  $\mathbf{v}$  and  $\mathbf{w}$ , respectively. For subsequent iterations with  $i > 0$ , we relabel the tuples similar to 1-WL, i.e.

$$C_i^{(k)}(\mathbf{v}) := \text{RELABEL} \left( \left( C_{i-1}^{(k)}(\mathbf{v}), \left\{ \left\{ C_i^{(k)}(\phi_1(\mathbf{v}, u)), \dots, C_i^{(k)}(\phi_k(\mathbf{v}, u)) \mid u \in \mathcal{N}(v) \right\} \right) \right), \quad (1)$$

where  $\phi_j(\mathbf{v}, u) := (v_1, \dots, v_{j-1}, u, v_{j+1}, \dots, v_k)$  refers to the function that replaces the  $j$ th element of the  $k$ -tuple  $\mathbf{v}$  with  $u$ . This induces a neighbourhood relation between tuples and just as in the case of 1-WL, we run the algorithm until the assigned colours of tuples stabilise for one graph. Similarly, if the colour sequences of two graphs differ, the graphs are non-isomorphic. As a generalisation of previous work [48], we can show that *any*  $k$ -FWL colouring can be reproduced with one equivariant filtration of type 2, thus proving that persistent homology is *at least* as expressive as  $k$ -FWL. We achieve this by showing that zero-dimensional persistence diagrams can *recover* any isomorphism-invariant real-valued function  $f: \mathcal{G} \rightarrow \mathbb{R}$  on the space of finite graphs by setting, for a given graph  $G$ , a constant filtration with value equal to the function value in  $G$ , i.e.  $f(G)$ .

**Theorem 1.** *For  $k \geq 2$ , there exists an equivariant filtration generator  $\mathcal{F}$  of type 2 such that its zero-dimensional persistence diagrams are at least as expressive as  $k$ -FWL.*

Both [Theorem 1](#) and [Theorem 3](#) may not be completely satisfactory because they only show the *existence* of such a filtration, but make no claims about the theoretical expressivity of filtrations of type 1 filtrations. Hence, it is particularly interesting to understand the expressivity of the Vietoris–Rips filtration [1], which is a popular choice in the context of persistent homology. Given a graph  $G = (V, E)$  equipped with the shortest-path distance  $d_G$ , the Vietoris–Rips filtration generator yields the pair  $(K, f_{\mathbf{v}})$ , where  $K$  is the set of all non-empty subsets  $\sigma$  of  $V$  such that  $\text{diam}(\sigma) \neq \infty$  and  $f_{\mathbf{v}}$  is the Vietoris–Rips filtration function given by  $f_{\mathbf{v}}(\sigma) = \max_{u, v \in \sigma} d_G(u, v)$ . Notice that the Vietoris–Rips filtration generator is equivariant. This is because graph isomorphisms preserve

<sup>3</sup>There are also other variants, for instance the *oblivious Weisfeiler–Leman algorithm*. It slightly differs in the way tuples are being relabelled, but a paper by Grohe [41] shows that the variant is essentially as powerful as  $k$ -FWL (with a minor shift in indices). The reader is referred to Morris et al. [65] and the references therein for an extended discussion of these aspects.

distances between vertices, making the induced simplicial morphisms preserve diameters and filtration values, and inducing an isomorphic simplicial complex morphism.

Vietoris–Rips persistent homology cannot be directly compared to the WL hierarchy, as no method is at least as expressive as the other. For example, Vietoris–Rips persistence diagrams of dimension zero can capture the number of connected components of graphs, while the 1-WL hierarchy cannot always capture them. See [Theorem 4](#) for a proof of this fact. By contrast, the 1-WL test can capture the difference between a path graph and a cycle graph both of length 3, while Vietoris–Rips persistent homology cannot. The fact that both methods are not comparable in terms of expressivity suggests that persistent homology methods based on Vietoris–Rips filtrations and WL/message passing methods are complimentary, since many standard graph neural networks are exactly as expressive as 1-WL [[66](#), [Theorem 2](#)]. Regarding filtrations of type 1, we do not expect cliques and high-dimensional persistent homology to be crucial in distinguishing pairs of non-isomorphic graphs. On the one hand, equivariant graph filtration generators of type 1 are already *at least as expressive* as 1-WL, thus they can distinguish almost all non-isomorphic graph pairs [[53](#), [Theorem 3.3](#)]. On the other hand, to prove that persistent homology for filtrations of type 1 is at least as expressive as  $k$ -FWL for  $k \geq 2$ , we would need to distinguish non-isomorphic CFI pairs for all  $k \geq 2$ . However, by definition of CFI graphs, they do not contain cliques of size three or greater, see [Lemma 4](#), meaning that most of the expressive power of these filtrations can only arise from the values in the original graph structure and low-dimensional persistent homology. We would ideally want to extend [Theorem 1](#) to state that persistent homology is *strictly* more expressive than  $k$ -FWL, although this is not as straightforward as for the case  $k = 1$ . Currently, we can provide one such counterexample, described in [Table 1](#), consisting of the  $4 \times 4$  rook’s graph and the Shrikhande graph. With an appropriate filtration, persistent homology can distinguish these two graphs *without* requiring more than vertices and edges, whereas 2-FWL is unable to distinguish them. We leave a general result for future work.

## 5 Experiments

The previous sections discussed the theoretical properties of filtrations. Here, we analyse their *empirical* performance as a complement to the theoretical discussion to demonstrate the high expressivity of persistent homology in graph tasks. We first analyse the *expressivity* of five different well-known filtrations by letting them distinguish non-isomorphic graphs. This is followed by experiments on graph property prediction. Please refer to [Appendices G](#) and [H](#) for additional expressivity and graph-classification experiments.

**Experimental setup.** We use different data sets containing strongly-regular graphs [[64](#)], minimal Cayley graphs, as well as benchmark data sets for graph-learning tasks [[95](#), BREC]. In the following, we will use five different filtrations for each graph:

1. A *degree* filtration (denoted by **D**), i.e.  $v \mapsto \deg(v)$ . The degree filtration is the most basic non-trivial filtration of a graph, showing nevertheless surprising empirical performance in graph classification tasks [[45](#), [71](#), [80](#)].
2. A filtration based on the eigenvalues of the *undirected graph Laplacian* (denoted by **L**), i.e.  $v \mapsto \lambda_v$ , where  $\lambda_v$  indicates the eigenvalue of the undirected graph Laplacian corresponding to vertex  $v$ . The graph Laplacian is known to capture characteristic properties of a graph; in the context of persistent homology it is often used in the form of a *heat kernel signature* [[22](#)].
3. A filtration based on the *Ollivier–Ricci curvature* [[72](#)] (denoted by **O**) in the graph, setting  $v \mapsto -1$  and  $(u, v) \mapsto \kappa(u, v)$ , with  $\kappa$  denoting the Ollivier–Ricci curvature

$$\kappa(u, v) := 1 - W_1(\mu_u^\alpha, \mu_v^\alpha), \quad (2)$$

where  $W_1$  denotes the first Wasserstein distance,<sup>4</sup> and  $\mu_u^\alpha, \mu_v^\alpha$  denote *probability measures* based on a lazy random walk in the graph, i.e.  $\mu_u^\alpha(u) := \alpha$ , indicating the probability of staying at the same vertex,  $\mu_u^\alpha(v) := (1 - \alpha)^{1/\deg(u)}$  for a neighbour  $v$  of  $u$ , and  $\mu_u^\alpha(\cdot) := 0$  otherwise. The probability measures in [Eq. \(2\)](#), i.e.  $\mu_u^\alpha$ , may be adjusted; recent work investigates the utility of this perspective [[32](#), [87](#)]. We set  $\alpha = 0$  for our subsequent experiments, thus obtaining a non-lazy random walk, and leave the investigation of the impact of other values for future work.

<sup>4</sup>This metric is also known as the *Earth Mover’s Distance* [[55](#)]. The Wasserstein distance is a fundamental concept in optimal transport; the monograph by Villani [[92](#)] contains a comprehensive introduction to this topic.

TABLE 1: Success rate ( $\uparrow$ ) for distinguishing pairs of *strongly-regular graphs* when using different filtrations at varying expansion levels of the graph (denoted by  $k$ ). 2-FWL cannot distinguish between any of these pairs.

Data	$k = 1$					$k = 2$					$k = 3$				
	Filtration														
	D	O	F	L	V	D	O	F	L	V	D	O	F	L	V
16622	0.00	<b>1.00</b>	0.00	0.00	0.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
251256	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	<b>0.90</b>	<b>0.90</b>	<b>0.90</b>	<b>0.90</b>	<b>0.90</b>
261034	0.00	0.00	0.00	0.00	0.00	0.20	0.20	0.20	0.76	0.20	0.93	0.93	0.93	<b>0.98</b>	0.93
281264	0.00	0.83	0.00	0.00	0.00	0.00	<b>1.00</b>	0.00	0.00	0.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
291467	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>	<b>0.77</b>
351668	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.00	0.95	0.95	0.95	<b>0.99</b>	0.95
351899	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	<b>0.81</b>	<b>0.81</b>	<b>0.81</b>	<b>0.81</b>	<b>0.81</b>
361446	0.00	0.00	0.00	0.00	0.00	0.02	0.02	0.02	0.83	0.02	0.92	0.92	0.92	<b>0.99</b>	0.92
401224	0.00	0.00	0.00	0.00	0.00	0.93	0.93	0.93	0.99	0.93	0.94	0.94	0.94	<b>0.99</b>	0.94

4. A filtration based on the *augmented Forman–Ricci curvature* [82] (denoted by **F**), where we again set  $v \mapsto -1$  and  $(u, v) \mapsto \mathcal{F}(u, v)$ , with  $\mathcal{F}(u, v) := 4 - \deg(u) - \deg(v) + 3|\mathcal{N}(u) \cap \mathcal{N}(v)|$ .
5. A *Vietoris–Rips* filtration (denoted by **V**) over the metric space defined by the shortest-path distance between its nodes [1].

We exclude colouring filtrations, as used in [Theorem 1](#), from our analysis due to their practical infeasibility. For each input graph  $G$ , this approach would require generating  $k!$  filtrations and constructing simplicial complexes containing all cliques up to dimension  $k$ . After picking a filtration (except for the Vietoris–Rips one), we expand the graph by filling in all  $(k + 1)$ -cliques, with filtration value for a clique  $\sigma$  given recursively by the maximum filtration value of its proper subcliques, i.e.  $\sigma \mapsto \max_{\tau \subsetneq \sigma} f(\tau)$ , and calculating persistent homology up to dimension  $k$ . Hence, for  $k = 1$ , we leave the graph ‘as-is,’ making use of connected components and cycles only.<sup>5</sup> Our persistent homology calculations result in a set of persistence diagrams for each graph, which we compare in a pairwise manner using the bottleneck distance described by [Eq. \(11\)](#). We consider two graphs to be different whenever the distance between their persistence diagrams is  $> 1 \times 10^{-8}$ , i.e. above machine precision. This setup has the advantage that no additional classifier is required; when dealing with data sets of pairs of non-isomorphic graphs, we may thus simply *count* the number of non-zero distance pairs, showing the utility of persistent homology as a powerful baseline for graph expressivity analysis. We omit comparisons between persistent homology and isolated filtration values, as persistent homology is at least as expressive as filtration values alone. This is a consequence of the *pairing lemma*, which ensures all filtration values, with multiplicities, are distributed as births or deaths in persistence diagrams up to the dimension of the simplicial complex. Hence, the filtration values can always be reconstructed from the persistence diagrams.

## 5.1 Strongly-Regular Graphs and Minimal Cayley Graphs

We start our investigation by analysing *strongly-regular graphs*, which are known to be extremely challenging to distinguish. 2-FWL, for instance, cannot distinguish *any* of these graphs [65, Section 3.3]. [Table 1](#) summarises the performance of our selected filtrations. We first observe that for  $k = 1$ , i.e. for the original graph without any cliques, few pairs of graphs can be distinguished by the five filtrations. Notably, a curvature-based filtration *is* sufficient to distinguish the two graphs in the 16622 data set, colloquially known as the  $4 \times 4$  rook’s graph and the Shrikhande graph. Distinguishing between these two graphs is usually said to require knowledge about cliques [13], but it turns out that a suitable filtration is sufficient. However, the empirical expressivity of curvature-based filtrations appears limited for  $k = 1$ , improving only for higher-order clique complexes. The Laplacian filtration, by contrast, exhibits strong empirical performance for  $k = 2$  on almost half of the data sets, increasing to near-perfect performance for  $k = 3$  in almost all data sets. Vietoris–Rips and degree filtrations obtain the exact same success rates for every  $k$  and data set, exhibiting the lowest

<sup>5</sup>Notice the shift in dimension: a  $k$ -simplex has  $k + 1$  vertices, meaning that persistent homology in dimension  $k$  contains information about  $(k + 1)$ -cliques.

TABLE 2: Success rate ( $\uparrow$ ) for distinguishing pairs of *minimal Cayley graphs* when using five different filtrations at varying expansion levels of the graph (denoted by  $k$ ). Values for 1-WL are shown as a baseline.

Data	$k = 1$					$k = 2$					
	1-WL	Filtration					D	O	F	L	V
		D	O	F	L	V					
cay12	0.67	0.67	0.71	0.95	0.86	0.00	0.95	0.95	0.95	<b>1.00</b>	0.95
cay16	0.83	0.83	0.42	0.83	0.58	0.00	0.83	0.92	0.83	<b>1.00</b>	0.94
cay20	0.61	0.61	0.46	0.61	0.79	0.00	0.61	0.79	0.61	<b>1.00</b>	0.89
cay24	0.65	0.65	0.82	0.86	0.98	0.00	0.83	0.93	0.86	<b>1.00</b>	0.93
cay32	0.76	0.76	0.81	0.76	0.90	0.00	0.76	0.94	0.76	<b>1.00</b>	0.90
cay36	0.69	0.69	0.87	0.84	0.99	0.00	0.84	0.95	0.84	<b>1.00</b>	0.94
cay60	0.69	0.69	0.90	0.78	1.00	0.00	0.77	0.95	0.78	<b>1.00</b>	0.97
cay63	0.49	0.49	0.89	0.73	0.88	0.00	0.73	0.93	0.73	<b>1.00</b>	0.96

 TABLE 3: Success rate ( $\uparrow$ ) for distinguishing pairs of instances of the *BREC data set* when using different filtrations at varying expansion levels of the graph (denoted  $k$ ). Due to combinatorial constraints, we did not calculate the Vietoris–Rips filtration for  $k = 4$ . Legend and number of graphs per category: B (Basic, 60), R (Regular, 100), E (Extension, 100), C (CFI, 100), 4, 20 (4-Vertex Condition), D (Distance-Regular, 20) graphs, respectively and A (average over full data set, 400 graphs).

Data	$k = 1$					$k = 2$					$k = 3$					$k = 4$			
	Filtration																		
	D	O	F	L	V	D	O	F	L	V	D	O	F	L	V	D	O	F	L
B	0.03	0.93	0.87	<b>1.00</b>	0.00	0.78	<b>1.00</b>	0.98	<b>1.00</b>	0.52	0.83	<b>1.00</b>	0.98	<b>1.00</b>	0.58	0.83	<b>1.00</b>	0.98	<b>1.00</b>
R	0.00	0.42	0.32	0.00	0.00	0.39	0.54	0.50	0.48	0.39	0.85	0.93	0.91	0.93	0.85	0.89	<b>0.97</b>	0.95	<b>0.97</b>
E	0.07	0.76	0.44	0.94	0.00	0.26	0.92	0.59	<b>1.00</b>	0.11	0.29	0.92	0.59	<b>1.00</b>	0.16	0.29	0.92	0.59	<b>1.00</b>
C	0.03	0.03	0.03	<b>0.06</b>	0.03	0.03	0.03	0.03	<b>0.06</b>	0.03	0.03	0.03	0.03	<b>0.06</b>	0.03	0.03	0.03	0.03	0.06
4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
D	0.00	0.00	0.00	<b>0.05</b>	0.00	0.00	0.00	0.00	<b>0.05</b>	0.00	0.00	0.00	0.00	<b>0.05</b>	<b>0.05</b>	0.00	0.00	0.00	0.05
A	0.03	0.44	0.33	0.40	0.01	0.29	0.52	0.43	0.54	0.21	0.47	0.67	0.58	0.70	0.40	0.48	0.68	0.59	<b>0.71</b>

success rates for  $k = 1$  and  $k = 2$ , and the same ones to the curvature filtrations for  $k = 3$ . It is clear that knowledge about higher-order cliques helps in driving performance here. Notice that in contrast to other algorithms [13], no additional embedding of the graphs is required; we are comparing ‘raw’ persistence diagrams directly.

As an additional class of complex graphs, we analyse *minimal Cayley graphs*, i.e. Cayley graphs that encode a group with a minimal generating set. Minimal Cayley graphs are still a topic of active research in graph theory, with several conjectures yet to be proven [5, 6]. Specifically, isomorphisms of Cayley graphs have been extensively studied [3, 56, 57, 67] and can be associated to isomorphisms of subsets of groups and interesting questions about their structure; see [56, Section 3]. We follow the same experimental setup as described above but also show the performance of 1-WL, calculated via a subtree Weisfeiler–Leman graph kernel [85]. Table 2 shows the results. We observe that the Laplacian filtration is trivially able to distinguish between all these graphs for  $k = 2$  and has a strong performance for  $k = 1$ ; this is not surprising since spectra of the graphs are known to be characteristic [61]. The performance of the curvature filtrations also points towards the utility of this formulation in practice. We also observe that Vietoris–Rips filtrations are the ones that most benefit from the availability of higher-order information, with a significant improvement in performance, going from a 0% success rate for  $k = 1$  to  $> 90\%$  for  $k = 2$ .

## 5.2 BREC Data Set

BREC [95] is a novel graph expressivity data set focused on providing a robust, challenging benchmark for graph isomorphism detection, containing particularly difficult graph classes. The data set consists of 400 graphs, divided into 6 different categories, Table 3 provides an overview of them

TABLE 4: Accuracy ( $\uparrow$ ) when predicting the properties of graphs in the ogbg-molhiv molecular graph data set [49] using different filtrations at varying expansion levels of the graph (denoted by  $k$ ). Column R contains the average probability of successfully predicting a property at random over all possible values of the property in the data, with the probability of choosing a label being proportional to the number of graphs with that label.

Data	$k = 1$						$k = 2$					
	Filtration											
	R	D	O	F	L	V	D	O	F	L	V	
Diameter	0.02	0.09	<b>0.11</b>	0.05	0.08	0.07	0.10	0.08	0.06	0.09	-	
Girth	0.04	0.00	0.11	0.34	0.46	<b>0.48</b>	0.14	0.21	0.33	0.45	-	
Radius	0.03	0.16	<b>0.21</b>	0.06	0.15	0.14	0.20	0.18	0.10	0.17	-	

and shows our experimentally-observed success rates. We observe that the Vietoris–Rips filtration is the worst performing filtration for all tested  $k$  values, followed by the degree filtration. The most informative curvature-based filtration is the Ollivier–Ricci one, which obtained higher or equal success rates than the Forman–Ricci curvature filtration in all the subsets of data, with strictly higher average success rates for all  $k$  values. The most effective filtration overall is the Laplacian filtration for  $k = 4$ , surpassing almost all the algorithms, including both graph neural networks and classical methods, described in Wang and Zhang [95, Table 2]. The only exception was the  $N_2$  algorithm [75] that obtained a success rate of 74.5%, compared to the 71% obtained by the Laplacian filtration for  $k = 4$ . Given the fact that  $N_2$  requires knowledge of the isomorphism class of *all* 2-hop-induced subgraphs of a graph, the computational complexity makes it infeasible to apply for many graph sizes in practices [75, Section 4.3], whereas the Laplacian filtration remains computable. Overall, these results experiments underscore the high expressivity of persistent homology, making it a strong baseline for graph-learning tasks. Please refer to [Appendix G.2](#) for additional results.

### 5.3 Predicting Graph Properties

As our final expressivity experiments, we assess the capability of persistent homology to predict graph properties. Here, we focus on the *diameter*, the *radius*, and the *girth*.

**Predicting the diameter of random graphs.** We predict the diameter of Erdős–Rényi and Watts–Strogatz graphs. For the Erdős–Rényi graphs, we generate  $N = 100$  graphs with  $n = 100$  vertices and  $p = 0.1$ . This edge probability corresponds to the critical connectivity regime, for which closed-form solutions of the diameter distribution are not readily available [42]. We assess the utility of persistent homology by specifying a regression task;<sup>6</sup> to this end, we vectorise the persistence diagrams for each filtration using *Betti curves* [71, 80], a simple curve-based topological representation. While this representation is technically a function, we represent it as a histogram of 10 bins and train a *ridge regression classifier* via leave-one-out cross-validation to predict the diameter of each graph. We deliberately focus only on zero-dimensional and one-dimensional persistent homology, i.e. we leave  $k = 1$ . Using the *mean absolute error* (MAE) for evaluation, we find that Ollivier–Ricci curvature performs best (0.057), followed by the Laplacian spectrum (0.061), and the degree filtration (0.065). We observe similar patterns when calculating  $N = 100$  Watts–Strogatz graphs of type  $(100, 5, 0.1)$ , i.e. we keep the same number of vertices and the same edge rewiring probability  $p$ , but connect each node with its 5 nearest neighbours in a ring neighbourhood. Using the same classifier, we again find that Ollivier–Ricci curvature achieves the lowest MAE (0.851), followed by the degree filtration (0.865), and the Laplacian spectrum (1.072).

**Predicting graph properties of the ogbg-molhiv [49] data set.** We predict the maximum radius and diameter among the radii and diameters of the connected components of each graph as well as their girth. See [Fig. S.3](#) for a visualisation of the distribution of these properties across all graphs in the data set. We use a random forest regression model and *persistence images* [2] as its input, computed from the persistence diagrams of the graphs calculated as in the previous sections. The model is trained on the usual train and test splits of the data set, with the final prediction obtained by rounding

<sup>6</sup>Following Hartmann and Mézard [42], we take the diameter of an Erdős–Rényi graph, which might consist of different connected components, to be the largest diameter of all connected components of the graph.



to the nearest integer.<sup>7</sup> Overall, this is a challenging task, with 24, 15, and 5 labels having more than 100 examples for the diameter, radius and girth, respectively. Thus, the average probabilities of guessing the correct label by random choice is 0.02, 0.03, and 0.04. However, we find that persistent homology, with a suitable filtration, can predict the maximum radius, the maximum diameter, and the girth of the graphs on unseen examples (i.e. on the test data set) in approximately 11%, 20%, and 48% of the cases, respectively, exhibiting substantially-improved results over a random baseline. For predicting radii and diameters, we find that the Ollivier–Ricci curvature outperforms the other filtrations for  $k = 1$  but gets worse for  $k = 2$ , where the degree filtration performs the best. It is surprising that the Vietoris–Rips filtration is *not* able to predict the radii of the graphs in this data set accurately ( $\approx 14\%$  accuracy), being the only filtration that works with explicit distances. However, this filtration proves most effective for predicting the girth, having an accuracy of almost 50% as compared to the 4% of the baseline. These results complement prior work on predicting properties of point clouds [19, 88], showing that topology-based graph-learning approaches carry a large degree of additional information about graphs.

## 6 Discussion

We discussed various aspects of the computation and provided evidence of the advantageous properties of persistent homology in the context of graph learning. Our primary theoretical insight is that persistent homology is *at least as expressive* as a corresponding 1-WL or  $k$ -FWL test, in some cases surpassing their discriminative power. Experiments underscore these theoretical expressivity properties, while also demonstrating that persistent homology is able to capture additional properties of a graph. In light of the performance differences among the different filtration functions in our experiments, we suggest that future work should focus on elucidating properties of classes of such functions, aiming to strike a balance between expressivity and efficiency. Another limitation involves the calculation *per se*, which requires costly clique-finding operations and is thus not scalable to large, dense graphs. If node features are present, alternative geometrical-topological approaches could potentially be used [62, 63, 68, 81, 90], necessitating additional research.

In the future, we would like to formally prove which classes of filtrations make  $k$ -dimensional persistent homology strictly more expressive than  $k$ -FWL (if any). Follow-up research could also focus on identifying other properties (next to the diameter and girth) that can be captured by persistent homology in graph learning. Such research has both theoretical and empirical components; a first step would be a formalisation of which substructures are captured by models [16, 27, 87]. Moreover, the success of the Laplacian filtration at the experimental tasks may hint at new filtrations based on spectral graph theory that provide a trade-off between utility and computational efficiency. We find that this research direction is overlooked by the computational topology research community, with most of the expressivity/stability results focusing on describing the stability of distance-based filtrations under perturbations [23, 24], and few works focusing on graphs [9]. Based on our experiments, we thus envision that persistent homology will constitute a strong baseline for graph-learning applications. As previous work shows, even topology-inspired approaches, making use of concepts such as filtrations, can approximate the performance of highly-parametrised models at a fraction of the computational cost [71]. All insights obtained using such topological methods hint at the overall utility of graph-structural information for graph learning tasks, but it is not clear whether current graph benchmark data actually exhibit such structures [74]. We thus hope that persistent homology and related techniques will also find more applications in *hybrid models*, which are able to incorporate geometrical-topological information about graphs. This is an emerging research topic of crucial relevance since there are now numerous graph data sets that combine geometrical information (node coordinates) with topological information [51].

Our theoretical analysis of the properties of persistent homology for graph learning tasks show the potential and benefits of a topology-based perspective. We are confident that additional computational topology concepts will enrich and augment machine learning models, leading to new insights about their theoretical and empirical capabilities. This paper is but a first attempt at elucidating the theoretical utility of computational topology in a graph learning context; advancing the field will require many more insights.

<sup>7</sup>We predict a girth of  $\infty$  in case the output value is larger than the number of nodes in the graph.

## Acknowledgements

This paper was motivated in large parts by discussions with participants of the BIRS 2022 Workshop on ‘Deep Exploration of Non-Euclidean Data with Geometric and Topological Representation Learning.’ The authors are also indebted to Leslie O’Bray for comments and discussions that helped substantially improve the main arguments of the paper. Finally, the authors are grateful for the stimulating discussions with the anonymous reviewers and the area chair, who believed in the merits of this work. Rubén Ballester was supported by the Ministry of Science, Innovation and Universities through projects PID2019-105093GB-I00, PID2020-117971GB-C22, and PID2022-136436NB-I00 and through the FPU contract FPU21/00968, and by the Departament de Recerca i Universitats de la Generalitat de Catalunya (2021 SGR 00697). Bastian Rieck was partially supported by the Bavarian state government with funds from the *Hightech Agenda Bavaria*. This work has received funding from the Swiss State Secretariat for Education, Research, and Innovation (SERI). The funders had no role in the preparation of the manuscript or the decision to publish.

## References

- [1] H. Adams and B. Coskunuzer. ‘Geometric Approaches to Persistent Homology’. In: *SIAM Journal on Applied Algebra and Geometry* 6.4 (2022), pp. 685–710. DOI: [10.1137/21M1422914](https://doi.org/10.1137/21M1422914)
- [2] H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta, and L. Ziegelmeier. ‘Persistence Images: A Stable Vector Representation of Persistent Homology’. In: *Journal of Machine Learning Research* 18.8 (2017), pp. 1–35
- [3] B. Alspach. ‘Isomorphism and Cayley graphs on abelian groups’. In: *Graph Symmetry: Algebraic Methods and Applications*. Ed. by G. Hahn and G. Sabidussi. Dordrecht, Netherlands: Springer, 1997, pp. 1–22. DOI: [10.1007/978-94-015-8937-6\\_1](https://doi.org/10.1007/978-94-015-8937-6_1)
- [4] V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. ‘On the Power of Color Refinement’. In: *Fundamentals of Computation Theory*. Ed. by A. Kosowski and I. Walukiewicz. Cham, Switzerland: Springer, 2015, pp. 339–350
- [5] L. Babai. ‘Automorphism Groups, Isomorphism, Reconstruction’. In: *Handbook of Combinatorics*. MIT Press, 1996, pp. 1447–1540
- [6] L. Babai. ‘Chromatic number and subgraphs of Cayley graphs’. In: *Theory and Applications of Graphs*. Ed. by Y. Alavi and D. R. Lick. 1978, pp. 10–22
- [7] R. Ballester, C. Casacuberta, and S. Escalera. *Topological Data Analysis for Neural Network Analysis: A Comprehensive Survey*. 2024. arXiv: [2312.05840](https://arxiv.org/abs/2312.05840) [cs.LG]
- [8] U. Bauer. ‘Ripser: Efficient computation of Vietoris–Rips persistence barcodes’. In: *Journal of Applied and Computational Topology* 5.3 (2021), pp. 391–423. DOI: [10.1007/s41468-021-00071-5](https://doi.org/10.1007/s41468-021-00071-5)
- [9] U. Bauer. ‘Ripser: efficient computation of Vietoris–Rips persistence barcodes’. In: *Journal of Applied and Computational Topology* 5.3 (2021), pp. 391–423. DOI: [10.1007/s41468-021-00071-5](https://doi.org/10.1007/s41468-021-00071-5)
- [10] U. Bauer, T. B. Masood, B. Giunti, G. Houry, M. Kerber, and A. Rathod. ‘Keeping it sparse: Computing Persistent Homology revisited’. 2022. arXiv: [2211.09075](https://arxiv.org/abs/2211.09075) [cs.CG]
- [11] M. Belkin and P. Niyogi. ‘Laplacian Eigenmaps for Dimensionality Reduction and Data Representation’. In: *Neural Computation* 15.6 (2003), pp. 1373–1396. DOI: [10.1162/089976603321780317](https://doi.org/10.1162/089976603321780317)
- [12] P. Bendich, J. S. Marron, E. Miller, A. Pieloch, and S. Skwerer. ‘Persistent homology analysis of brain artery trees’. In: *The Annals of Applied Statistics* 10.1 (2016), pp. 198–218. DOI: [10.1214/15-AOAS886](https://doi.org/10.1214/15-AOAS886)
- [13] C. Bodnar, F. Frasca, Y. Wang, N. Otter, G. F. Montufar, P. Lió, and M. Bronstein. ‘Weisfeiler and Lehman Go Topological: Message Passing Simplicial Networks’. In: *International Conference on Machine Learning (ICML)*. Ed. by M. Meila and T. Zhang. Proceedings of Machine Learning Research 139. PMLR, 2021, pp. 1026–1037
- [14] B. Bollobás. ‘Distinguishing Vertices of Random Graphs’. In: *Graph Theory*. Ed. by B. Bollobás. North–Holland Mathematics Studies 62. Amsterdam, Netherlands: North–Holland, 1982, pp. 33–49. DOI: [10.1016/S0304-0208\(08\)73545-X](https://doi.org/10.1016/S0304-0208(08)73545-X)

- [15] K. M. Borgwardt, C. S. Ong, S. Schönauer, S. V. N. Vishwanathan, A. J. Smola, and H.-P. Kriegel. ‘Protein function prediction via graph kernels’. In: *Bioinformatics* 21.suppl 1 (June 2005), pp. i47–i56. DOI: [10.1093/bioinformatics/bti1007](https://doi.org/10.1093/bioinformatics/bti1007)
- [16] G. Bouritsas, F. Frasca, S. Zafeiriou, and M. M. Bronstein. ‘Improving Graph Neural Network Expressivity via Subgraph Isomorphism Counting’. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.1 (2023), pp. 657–668. DOI: [10.1109/TPAMI.2022.3154319](https://doi.org/10.1109/TPAMI.2022.3154319)
- [17] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. ‘Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges’. 2021. arXiv: [2104.13478](https://arxiv.org/abs/2104.13478) [cs.LG]
- [18] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. ‘Geometric Deep Learning: Going beyond Euclidean data’. In: *IEEE Signal Processing Magazine* 34.4 (2017), pp. 18–42. DOI: [10.1109/MSP.2017.2693418](https://doi.org/10.1109/MSP.2017.2693418)
- [19] P. Bubenik, M. Hull, D. Patel, and B. Whittle. ‘Persistent homology detects curvature’. In: *Inverse Problems* 36.2 (2020), p. 025008. DOI: [10.1088/1361-6420/ab4ac0](https://doi.org/10.1088/1361-6420/ab4ac0)
- [20] J.-y. Cai, M. Furer, and N. Immerman. ‘An optimal lower bound on the number of variables for graph identification’. In: *30th Annual Symposium on Foundations of Computer Science*. 1989, pp. 612–617. DOI: [10.1109/SFCS.1989.63543](https://doi.org/10.1109/SFCS.1989.63543)
- [21] J.-Y. Cai, M. Fürer, and N. Immerman. ‘An optimal lower bound on the number of variables for graph identification’. In: *Combinatorica* 12.4 (1992), pp. 389–410. DOI: [10.1007/BF01305232](https://doi.org/10.1007/BF01305232)
- [22] M. Carrière, F. Chazal, Y. Ike, T. Lacombe, M. Royer, and Y. Umeda. ‘PersLay: A Neural Network Layer for Persistence Diagrams and New Graph Topological Signatures’. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. Ed. by S. Chiappa and R. Calandra. Proceedings of Machine Learning Research 108. PMLR, 2020, pp. 2786–2796
- [23] F. Chazal, V. De Silva, M. Glisse, and S. Oudot. *The structure and stability of persistence modules*. Vol. 10. SpringerBriefs in Mathematics. Cham, Switzerland: Springer, 2016. DOI: [10.1007/978-3-319-42545-0](https://doi.org/10.1007/978-3-319-42545-0)
- [24] F. Chazal, V. de Silva, and S. Oudot. ‘Persistence stability for geometric complexes’. In: *Geometriae Dedicata* 173.1 (2014), pp. 193–214. DOI: [10.1007/s10711-013-9937-z](https://doi.org/10.1007/s10711-013-9937-z)
- [25] Y. Chen, B. Coskunuzer, and Y. Gel. ‘Topological Relational Learning on Graphs’. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 27029–27042
- [26] Y. Chen and Y. R. Gel. ‘Topological Pooling on Graphs’. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 37.6 (2023), pp. 7096–7103. DOI: [10.1609/aaai.v37i6.25866](https://doi.org/10.1609/aaai.v37i6.25866)
- [27] Z. Chen, L. Chen, S. Villar, and J. Bruna. ‘Can Graph Neural Networks Count Substructures?’ In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., 2020, pp. 10383–10395
- [28] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. ‘Stability of persistence diagrams’. In: *Discrete & Computational Geometry* 37.1 (2007), pp. 103–120. DOI: [10.1007/s00454-006-1276-5](https://doi.org/10.1007/s00454-006-1276-5)
- [29] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. ‘Lipschitz functions have  $L_p$ -stable persistence’. In: *Foundations of Computational Mathematics* 10.2 (2010), pp. 127–139. DOI: [10.1007/s10208-010-9060-6](https://doi.org/10.1007/s10208-010-9060-6)
- [30] C. J. Colbourn. ‘On testing isomorphism of permutation graphs’. In: *Networks* 11.1 (1981), pp. 13–21. DOI: [10.1002/net.3230110103](https://doi.org/10.1002/net.3230110103)
- [31] K. Coolsaet, S. D’hondt, and J. Goedgebeur. ‘House of Graphs 2.0: A database of interesting graphs and more’. In: *Discrete Applied Mathematics* 325 (2023), pp. 97–107. DOI: [10.1016/j.dam.2022.10.013](https://doi.org/10.1016/j.dam.2022.10.013). URL: <https://houseofgraphs.org>
- [32] C. Coupette, S. Dalleiger, and B. Rieck. ‘Ollivier–Ricci Curvature for Hypergraphs: A Unified Framework’. In: *International Conference on Learning Representations*. 2023. URL: <https://openreview.net/forum?id=sPCKN15qDps>
- [33] M. Cuturi. ‘Sinkhorn Distances: Lightspeed Computation of Optimal Transport’. In: *Advances in Neural Information Processing Systems*. Ed. by C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger. Vol. 26. 2013

- [34] A. K. Debnath, R. L. Lopez de Compadre, G. Debnath, A. J. Shusterman, and C. Hansch. ‘Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. Correlation with molecular orbital energies and hydrophobicity’. In: *Journal of Medicinal Chemistry* 34.2 (1991), pp. 786–797. DOI: [10.1021/jm00106a046](https://doi.org/10.1021/jm00106a046)
- [35] T. K. Dey and Y. Wang. *Computational Topology for Data Analysis*. Cambridge, United Kingdom: Cambridge University Press, 2022. DOI: [10.1017/9781009099950](https://doi.org/10.1017/9781009099950)
- [36] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. Applied Mathematics. American Mathematical Society, 2010. ISBN: 9780821849255
- [37] H. Edelsbrunner, D. Letscher, and A. J. Zomorodian. ‘Topological persistence and simplification’. In: *Discrete & Computational Geometry* 28.4 (2002), pp. 511–533. DOI: [10.1007/s00454-002-2885-2](https://doi.org/10.1007/s00454-002-2885-2)
- [38] Y. Elkin and V. Kurlin. ‘The Mergegram of a Dendrogram and Its Stability’. In: *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*. Ed. by J. Esparza and D. Král’. Vol. 170. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, 32:1–32:13. ISBN: 978-3-95977-159-7. DOI: [10.4230/LIPIcs.MFCS.2020.32](https://doi.org/10.4230/LIPIcs.MFCS.2020.32)
- [39] F. Errica, M. Podda, D. Bacciu, and A. Micheli. ‘A Fair Comparison of Graph Neural Networks for Graph Classification’. In: *International Conference on Learning Representations*. 2020. URL: <https://openreview.net/forum?id=HygDF6NFPB>
- [40] R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser, and E. I. Moser. ‘Toroidal topology of population activity in grid cells’. In: *Nature* 602.7895 (2022), pp. 123–128. DOI: [10.1038/s41586-021-04268-7](https://doi.org/10.1038/s41586-021-04268-7)
- [41] M. Grohe. ‘The Logic of Graph Neural Networks’. In: *Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 2021, pp. 1–17. DOI: [10.1109/LICS52264.2021.9470677](https://doi.org/10.1109/LICS52264.2021.9470677)
- [42] A. K. Hartmann and M. Mézard. ‘Distribution of diameters for Erdős–Rényi random graphs’. In: *Physical Review E* 97 (3 2018), p. 032128. DOI: [10.1103/PhysRevE.97.032128](https://doi.org/10.1103/PhysRevE.97.032128)
- [43] F. Hensel, M. Moor, and B. Rieck. ‘A Survey of Topological Machine Learning Methods’. In: *Frontiers in Artificial Intelligence* 4 (2021). ISSN: 2624-8212. DOI: [10.3389/frai.2021.681108](https://doi.org/10.3389/frai.2021.681108)
- [44] C. Hofer, R. Kwitt, M. Niethammer, and A. Uhl. ‘Deep Learning with Topological Signatures’. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. 2017, pp. 1634–1644
- [45] C. D. Hofer, F. Graf, B. Rieck, M. Niethammer, and R. Kwitt. ‘Graph Filtration Learning’. In: *International Conference on Machine Learning (ICML)*. Ed. by H. Daumé III and A. Singh. Proceedings of Machine Learning Research 119. PMLR, 2020, pp. 4314–4323. arXiv: [1905.10996](https://arxiv.org/abs/1905.10996) [cs.LG]
- [46] C. D. Hofer, R. Kwitt, and M. Niethammer. ‘Learning Representations of Persistence Barcodes’. In: *Journal of Machine Learning Research* 20.126 (2019), pp. 1–45
- [47] D. Horak, S. Maletić, and M. Rajković. ‘Persistent homology of complex networks’. In: *Journal of Statistical Mechanics: Theory and Experiment* 2009.3 (2009), P03034. DOI: [10.1088/1742-5468/2009/03/P03034](https://doi.org/10.1088/1742-5468/2009/03/P03034)
- [48] M. Horn, E. De Brouwer, M. Moor, Y. Moreau, B. Rieck, and K. Borgwardt. ‘Topological Graph Neural Networks’. In: *International Conference on Learning Representations*. 2022. URL: <https://openreview.net/forum?id=oxxUMeFwEHd>
- [49] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. ‘Open Graph Benchmark: Datasets for Machine Learning on Graphs’. In: *arXiv preprint arXiv:2005.00687* (2020)
- [50] J. Immonen, A. Souza, and V. Garg. ‘Going beyond persistent homology using persistent homology’. In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 63150–63173
- [51] C. K. Joshi, C. Bodnar, S. V. Mathis, T. Cohen, and P. Lio. ‘On the Expressive Power of Geometric Graph Neural Networks’. In: *International Conference on Machine Learning (ICML)*. Ed. by A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett. Proceedings of Machine Learning Research 202. PMLR, 2023, pp. 15330–15355

- [52] P. J. Kelly. ‘A congruence theorem for trees’. In: *Pacific Journal of Mathematics* 7.1 (1957), pp. 961–968. DOI: [pjm/1103043674](https://doi.org/10.2307/2372442)
- [53] S. Kiefer. ‘Power and limits of the Weisfeiler-Leman algorithm’. PhD thesis. Dissertation, RWTH Aachen University, 2020, 2020
- [54] R. Kwitt, S. Huber, M. Niethammer, W. Lin, and U. Bauer. ‘Statistical Topological Data Analysis - A Kernel Perspective’. In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett. Vol. 28. 2015
- [55] E. Levina and P. Bickel. ‘The Earth Mover’s distance is the Mallows distance: some insights from statistics’. In: *IEEE International Conference on Computer Vision (ICCV)*. Vol. 2. 2001, pp. 251–256. DOI: [10.1109/ICCV.2001.937632](https://doi.org/10.1109/ICCV.2001.937632)
- [56] C. Li. ‘On isomorphisms of connected Cayley graphs’. In: *Discrete Mathematics* 178.1 (1998), pp. 109–122. DOI: [10.1016/S0012-365X\(97\)81821-3](https://doi.org/10.1016/S0012-365X(97)81821-3)
- [57] C. Li. ‘On isomorphisms of finite Cayley graphs—a survey’. In: *Discrete Mathematics* 256.1 (2002), pp. 301–334. DOI: [10.1016/S0012-365X\(01\)00438-1](https://doi.org/10.1016/S0012-365X(01)00438-1)
- [58] P. Li and J. Leskovec. ‘The Expressive Power of Graph Neural Networks’. In: *Graph Neural Networks: Foundations, Frontiers, and Applications*. Ed. by L. Wu, P. Cui, J. Pei, and L. Zhao. Singapore: Springer Singapore, 2022, pp. 63–98
- [59] M. Lichter, S. Raßmann, and P. Schweitzer. *Computational complexity of the Weisfeiler-Leman dimension*. 2024. arXiv: [2402.11531](https://arxiv.org/abs/2402.11531) [cs.CC]. URL: <https://arxiv.org/abs/2402.11531>
- [60] S. Lim, F. Memoli, and O. B. Okutan. *Vietoris–Rips Persistent Homology, Injective Metric Spaces, and The Filling Radius*. 2020. arXiv: [2001.07588](https://arxiv.org/abs/2001.07588) [math.AT]
- [61] L. Lovász. ‘Spectra of graphs with transitive groups’. In: *Periodica Mathematica Hungarica* 6.2 (1975), pp. 191–195
- [62] K. Maggs, C. Hacker, and B. Rieck. ‘Simplicial Representation Learning with Neural  $k$ -forms’. In: *International Conference on Learning Representations*. 2024. URL: <https://openreview.net/forum?id=DjwOXhjHZb>
- [63] L. Marsh and D. Beers. ‘Stability and Inference of the Euler Characteristic Transform’. Preprint. 2023. arXiv: [2303.13200](https://arxiv.org/abs/2303.13200) [math.ST]
- [64] B. McKay. *Collection of strongly-regular graphs*. URL: <https://users.cecs.anu.edu.au/~bdkm/data/graphs.html>
- [65] C. Morris, Y. Lipman, H. Maron, B. Rieck, N. M. Kriege, M. Grohe, M. Fey, and K. Borgwardt. ‘Weisfeiler and Leman go Machine Learning: The Story so far’. 2021. arXiv: [2112.09992](https://arxiv.org/abs/2112.09992) [cs.LG]
- [66] C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. E. Lenssen, G. Rattan, and M. Grohe. ‘Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks’. In: *AAAI Conference on Artificial Intelligence*. 2019. DOI: [10.1609/aaai.v33i01.33014602](https://doi.org/10.1609/aaai.v33i01.33014602)
- [67] D. W. Morris, J. Morris, and G. Verret. *Isomorphisms of Cayley graphs on nilpotent groups*. 2016. arXiv: [1603.01883](https://arxiv.org/abs/1603.01883) [math.CO]
- [68] E. Munch. ‘An Invitation to the Euler Characteristic Transform’. Preprint. 2023. arXiv: [2310.10395](https://arxiv.org/abs/2310.10395) [cs.CG]
- [69] J. R. Munkres. *Elements of Algebraic Topology*. Menlo Park, California, USA: Addison–Wesley Publishing Company, 1984
- [70] V. Nanda. *Computational Algebraic Topology Lecture Notes*. <https://people.maths.ox.ac.uk/nanda/cat/TDANotes.pdf>. 2022
- [71] L. O’Bray, B. Rieck, and K. Borgwardt. ‘Filtration Curves for Graph Representation’. In: *International Conference on Knowledge Discovery & Data Mining (KDD)*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 1267–1275. DOI: [10.1145/3447548.3467442](https://doi.org/10.1145/3447548.3467442)
- [72] Y. Ollivier. ‘Ricci curvature of metric spaces’. In: *Comptes Rendus Mathématique* 345.11 (2007), pp. 643–646. DOI: [10.1016/j.crma.2007.10.041](https://doi.org/10.1016/j.crma.2007.10.041)
- [73] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington. ‘A roadmap for the computation of persistent homology’. In: *EPJ Data Science* 6.1 (2017), 17:1–17:38. DOI: [10.1140/epjds/s13688-017-0109-5](https://doi.org/10.1140/epjds/s13688-017-0109-5)

- [74] J. Palowitch, A. Tsitsulin, B. Mayer, and B. Perozzi. ‘GraphWorld: Fake Graphs Bring Real Insights for GNNs’. In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2022, pp. 3691–3701. DOI: [10.1145/3534678.3539203](https://doi.org/10.1145/3534678.3539203)
- [75] P. A. Papp and R. Wattenhofer. ‘A Theoretical Comparison of Graph Neural Network Extensions’. In: *International Conference on Machine Learning*. Ed. by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato. Proceedings of Machine Learning Research 162. PMLR, 2022, pp. 17323–17345
- [76] J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt. ‘A stable multi-scale kernel for topological machine learning’. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 4741–4748. DOI: [10.1109/CVPR.2015.7299106](https://doi.org/10.1109/CVPR.2015.7299106)
- [77] B. Rieck. ‘Persistent Homology in Multivariate Data Visualization’. PhD thesis. Heidelberg University, 2017. DOI: [10.11588/heidok.00022914](https://doi.org/10.11588/heidok.00022914)
- [78] B. Rieck, C. Bock, and K. Borgwardt. ‘A Persistent Weisfeiler–Lehman Procedure for Graph Classification’. In: *International Conference on Machine Learning (ICML)*. Ed. by K. Chaudhuri and R. Salakhutdinov. Proceedings of Machine Learning Research 97. PMLR, 2019, pp. 5448–5458
- [79] B. Rieck, U. Fugacci, J. Lukasczyk, and H. Leitte. ‘Clique Community Persistence: A Topological Visual Analysis Approach for Complex Networks’. In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2018), pp. 822–831. DOI: [10.1109/TVCG.2017.2744321](https://doi.org/10.1109/TVCG.2017.2744321)
- [80] B. Rieck, F. Sadlo, and H. Leitte. ‘Topological Machine Learning with Persistence Indicator Functions’. In: *Topological Methods in Data Analysis and Visualization V*. Ed. by H. Carr, I. Fujishiro, F. Sadlo, and S. Takahashi. Cham, Switzerland: Springer, 2020, pp. 87–101. DOI: [10.1007/978-3-030-43036-8\\_6](https://doi.org/10.1007/978-3-030-43036-8_6)
- [81] E. Röell and B. Rieck. ‘Differentiable Euler Characteristic Transforms for Shape Classification’. In: *International Conference on Learning Representations*. 2024. arXiv: [2310.07630](https://arxiv.org/abs/2310.07630) [cs.LG]. URL: <https://openreview.net/forum?id=M0632iPq3I>
- [82] A. Samal, R. P. Sreejith, J. Gu, S. Liu, E. Saucan, and J. Jost. ‘Comparative analysis of two discretizations of Ricci curvature for complex networks’. In: *Scientific Reports* 8.1 (2018), p. 8650. DOI: [10.1038/s41598-018-27001-3](https://doi.org/10.1038/s41598-018-27001-3)
- [83] R. Sato. ‘A Survey on The Expressive Power of Graph Neural Networks’. 2020. arXiv: [2003.04078](https://arxiv.org/abs/2003.04078) [cs.LG]
- [84] D. R. Sheehy. ‘Linear-Size Approximations to the Vietoris–Rips Filtration’. In: *Discrete & Computational Geometry* 49.4 (2013), pp. 778–796. DOI: [10.1007/s00454-013-9513-1](https://doi.org/10.1007/s00454-013-9513-1)
- [85] N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. ‘Weisfeiler–Lehman Graph Kernels’. In: *Journal of Machine Learning Research* 12.77 (2011), pp. 2539–2561
- [86] P. Škraba and K. Turner. ‘Wasserstein Stability for Persistence Diagrams’. 2020. arXiv: [2006.16824](https://arxiv.org/abs/2006.16824) [math.AT]
- [87] J. Southern, J. Wayland, M. Bronstein, and B. Rieck. ‘Curvature Filtrations for Graph Generative Model Evaluation’. In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 63036–63061
- [88] R. Turkeš, G. Montúfar, and N. Otter. ‘On the Effectiveness of Persistent Homology’. In: *Advances in Neural Information Processing Systems*. Ed. by A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho. 2022. URL: <https://openreview.net/forum?id=DRjUkfExCix>
- [89] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer. ‘Fréchet Means for Distributions of Persistence Diagrams’. In: *Discrete & Computational Geometry* 52.1 (2014), pp. 44–70. DOI: [10.1007/s00454-014-9604-7](https://doi.org/10.1007/s00454-014-9604-7)
- [90] K. Turner, S. Mukherjee, and D. M. Boyer. ‘Persistent homology transform for modeling shapes and surfaces’. In: *Information and Inference: A Journal of the IMA* 3.4 (2014), pp. 310–344. DOI: [10.1093/imaiai/iau011](https://doi.org/10.1093/imaiai/iau011)
- [91] Y. Verma, A. H. Souza, and V. Garg. ‘Topological Neural Networks go Persistent, Equivariant, and Continuous’. In: *Proceedings of the 41st International Conference on Machine Learning*. Ed. by R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp. Vol. 235. Proceedings of Machine Learning Research. PMLR, 21–27 Jul 2024, pp. 49388–49407

- [92] C. Villani. *Optimal Transport. Old and New*. Grundlehren der mathematischen Wissenschaften 338. Springer, 2009
- [93] C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2021
- [94] N. Wale and G. Karypis. ‘Comparison of Descriptor Spaces for Chemical Compound Retrieval and Classification’. In: *Sixth International Conference on Data Mining (ICDM’06)*. 2006, pp. 678–689. DOI: [10.1109/ICDM.2006.39](https://doi.org/10.1109/ICDM.2006.39)
- [95] Y. Wang and M. Zhang. *Towards Better Evaluation of GNN Expressiveness with BREC Dataset*. 2023. arXiv: [2304.07702](https://arxiv.org/abs/2304.07702) [cs.LG]
- [96] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. ‘How Powerful are Graph Neural Networks?’ In: *International Conference on Learning Representations*. 2019. URL: <https://openreview.net/forum?id=ryGs6iA5Km>
- [97] Z. Yan, T. Ma, L. Gao, Z. Tang, Y. Wang, and C. Chen. ‘Neural Approximation of Graph Topological Features’. In: *Advances in Neural Information Processing Systems*. Ed. by A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho. 2022. URL: <https://openreview.net/forum?id=qwjr07Rewqy>
- [98] P. Yanardag and S. Vishwanathan. ‘Deep graph kernels’. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2015, pp. 1365–1374
- [99] X. Ye, F. Sun, and S. Xiang. ‘TREP: A Plug-In Topological Layer for Graph Neural Networks’. In: *Entropy* 25.2 (2023). DOI: [10.3390/e25020331](https://doi.org/10.3390/e25020331)
- [100] C. Ying, X. Zhao, and T. Yu. ‘Boosting Graph Pooling with Persistent Homology’. In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024. URL: <https://openreview.net/forum?id=WcmqdY2AKu>
- [101] S. Zhang, S. Mukherjee, and T. K. Dey. ‘GEFL: Extended Filtration Learning for Graph Classification’. In: *Proceedings of the First Learning on Graphs Conference*. Ed. by B. Rieck and R. Pascanu. Vol. 198. Proceedings of Machine Learning Research. PMLR, 2022, 16:1–16:26
- [102] Q. Zhao and Y. Wang. ‘Learning metrics for persistence-based summaries and applications for graph classification’. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Vol. 32. 2019, pp. 9855–9866
- [103] Q. Zhao, Z. Ye, C. Chen, and Y. Wang. ‘Persistence Enhanced Graph Neural Network’. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. Ed. by S. Chiappa and R. Calandra. Proceedings of Machine Learning Research 108. PMLR, 2020, pp. 2896–2906

---

**Appendix (Supplementary Materials)**

---

<b>A</b>	<b>Extended literature review</b>	<b>17</b>
<b>B</b>	<b>Counting Connected Components</b>	<b>17</b>
<b>C</b>	<b>A Primer in Computational Topology</b>	<b>17</b>
	C.1 Simplicial Homology . . . . .	17
	C.2 Persistent Homology . . . . .	19
	C.3 Computational complexity of persistent homology . . . . .	21
<b>D</b>	<b>1-WL and Persistent Homology</b>	<b>22</b>
<b>E</b>	<b>Theorems &amp; Proofs</b>	<b>22</b>
<b>F</b>	<b>Topology and the Weisfeiler–Leman Hierarchy</b>	<b>25</b>
<b>G</b>	<b>Additional Expressivity Experiments</b>	<b>26</b>
	G.1 Additional Results for Connected Cubic Graphs . . . . .	26
	G.2 Additional Results for the BREC Data Set . . . . .	26
	G.3 Additional Figures for Graph-Property Prediction Tasks . . . . .	26
<b>H</b>	<b>Additional Graph Classification Experiments</b>	<b>26</b>
<b>I</b>	<b>Additional results on alpha complexes</b>	<b>28</b>
<b>J</b>	<b>Implementation and Hardware details</b>	<b>29</b>
<b>K</b>	<b>Licenses</b>	<b>31</b>



## A Extended literature review

Topological Data Analysis has been used in a variety of applications in machine learning, such as in the development of topological input features, the analysis of learning algorithms, or the development of new, topology-aware models; see [43] for a general introduction of TDA in machine learning and [7] for a more specific review of TDA in neural networks.

In the context of graph learning, the use of topological data analysis, and particularly persistent homology has contributed many new insights and tools. In the context of architectures, persistent homology has been used as pooling layers [26, 100], readout layers [101], and regular layers [48, 91, 99]. In the context of expressivity and graph neural network property analysis, [48] also proved that persistent homology is at least as expressive as 1-WL and [50] extended this result and our results for persistent homology for vertex- and edge-based filtrations.

## B Counting Connected Components

Since the main text deals with higher-order topological features, and such features afford a substantially less intuitive grasp, we want to briefly comment on how to obtain  $\beta_0$ , the number of connected components. As with many problems in computer science, this procedure turns out to be simple if we pick our data structures correctly. Here, we need a *union-find* data structure, also known as a disjoint set forest. This data structure is built on the vertices of a graph and affords two operations, viz. `union` (or `merge`) and `find`. The `merge` operation assigns two vertices to the same connected component, while the `find` operation returns the current connected component of a vertex. Building such a data structure is reasonably easy in programming languages like Python, which offer *associative arrays*. Algorithm 1 shows one particular pseudo-code implementation of a simple union-find data structure. The pseudo-code assumes that all operations are changing objects ‘in place.’ Notice that the `find` operation is implemented implicitly via a lookup in the `merge` function. A proper object-oriented implementation of a union-find data structure should have these two operations in its public interface.

---

ALGORITHM 1 Using associative arrays to find connected components

---

```

1: function GET_CONNECTED_COMPONENTS( $V, E$ )
2:   UF  $\leftarrow$  {}
3:   for  $v \in V$  do
4:     UF[ $v$ ]  $\leftarrow$   $v$ 
5:   end for
6:   for  $e = (v, w) \in E$  do
7:     MERGE(UF,  $v, w$ )
8:   end for
9:   return { $v \mid$  UF[ $v$ ] =  $v$ }
10: end function

11: function MERGE(UF,  $v, w$ )
12:   if UF[ $v$ ]  $\neq$  UF[ $w$ ] then
13:     UF[ $v$ ]  $\leftarrow$   $w$ 
14:   end if
15: end function

```


---

## C A Primer in Computational Topology

With the understanding that our readers have different backgrounds, the following section provides a primer of the most relevant concepts in computational topology.

### C.1 Simplicial Homology

The Betti numbers of a graph are actually a special instance of a more generic concept, known as *simplicial homology*. We will see that under this framework, the Betti numbers are the ranks of the zeroth and first homology group, respectively. Simplicial homology is not required in order to



The triangle is a simple simplicial complex, consisting of one 2-simplex, three 1-simplices and three 0-simplices, respectively. The boundary of the 2-simplex is non-zero: we have  $\partial_2\{a, b, c\} = \{b, c\} + \{a, c\} + \{a, b\}$ . The set of edges, on the other hand, does not have a boundary, i.e.  $\partial_1(\{b, c\} + \{a, c\} + \{a, b\}) = \{c\} + \{b\} + \{c\} + \{a\} + \{b\} + \{a\} = 0$ , because the simplices cancel each other out.

FIGURE S.1: Calculating the boundaries of a 2-simplex and the boundary of a simplicial chain consisting of 1-simplices. Notice that the boundary of a boundary is always zero. This is a fundamental property of persistent homology. The figure is slightly adapted from Rieck [77].

understand most of the results of this paper, but an appreciation for some of the concepts will be helpful in understanding connections to other concepts. We try to provide a self-contained introduction to the most relevant concepts and refer to the textbook by Munkres [69] for a more in-depth exposition of these concepts. We start by introducing the central object of algebraic topology—the *simplicial complex*.<sup>8</sup>

**Definition 3** (Simplicial complex). *A simplicial complex  $K$  is a system of sets that is closed under the subset operation. Thus, for any  $\sigma \in K$  and  $\tau \subseteq \sigma$ , we have  $\tau \in K$ . An element  $\sigma \in K$  with  $|\sigma| = k + 1$  is also referred to as a  $k$ -simplex. We also express this by writing  $\dim \sigma = k$ . Moreover, if  $k$  is maximal among all the simplices of  $K$ , we say that the  $K$  is a  $k$ -dimensional simplicial complex.*

Note that there is an unfortunate shift in dimensions: a  $k$ -simplex has indeed  $k + 1$  elements. This convention makes sense when we relate it to the concept of *dimension*. A 0-simplex, i.e. a point or a vertex, should be assigned a dimension of 0. The reader should thus mentally equate the dimension of a simplex with its dimension. The text will aim to quell any confusion about such shifts. The quintessential example of a simplicial complex is a graph  $G = (V, E)$ . Setting  $K := V \cup E$ , we obtain a 1-dimensional simplicial complex. We may calculate additional types of simplicial complexes from a graph, for instance by *expanding* each  $(k + 1)$ -clique into a  $k$ -simplex [47, 79].

The simplicial complex on its own is only a set system; to perform calculations with this type of data structure, we need to imbue it with additional operations. One of the most common operations involves defining homomorphisms between the subsets of a simplicial complex  $K$ .

**Definition 4** (Chain group of a simplicial complex). *Given a simplicial complex  $K$ , the vector space generated over  $\mathbb{Z}_2$  coefficients whose elements are the  $k$ -simplices of  $K$  is called the  $k$ th chain group, denoted by  $C_k(K)$ . The elements of a chain group are also referred to as simplicial chains.*

Elements of the chain group are thus sums of simplices of a compatible dimension. For instance, we may write the sum of all edges of a graph to obtain a valid simplicial chain. Operating over  $\mathbb{Z}_2$  coefficients means that  $\sigma + \sigma = 0$ , the empty chain, for all  $\sigma \in K$ .<sup>9</sup> Simplicial chains permit us to define homomorphisms between chain groups, which will ultimately permit us to treat topological questions with tools of linear algebra.

**Definition 5** (Boundary homomorphism). *Given  $\sigma = (v_0, \dots, v_k) \in K$ , we define the  $k$ th boundary homomorphism  $\partial_k : C_k(K) \rightarrow C_{k-1}(K)$  as*

$$\partial_k(\sigma) := \sum_{i=0}^k (v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k), \quad (3)$$

*i.e. a sum of simplices with the  $i$ th entry—vertex—of the simplex missing, respectively.*

It is sufficient to define  $\partial_k$  on individual simplices; since it is a homomorphism, it extends to arbitrary simplicial chains. Fig. S.1 shows an example of this calculation. The boundary operator already

<sup>8</sup>Technically, we will be working with *abstract simplicial complexes*. A definition of a simplicial complex in terms of convex subsets is also possible, but this necessitates understanding certain nuances that are irrelevant to this paper.

<sup>9</sup>Readers familiar with algebraic topology will recognise  $\mathbb{Z}_2$  as a deliberate choice of *coefficient field* for the subsequent calculations. Other choices are possible, but the computational topology community predominantly uses  $\mathbb{Z}_2$  coefficients in practice, with very few exceptions [40]. However, all the proofs and concepts introduced in this paper apply, *mutatis mutandis*, for other coefficient sets as well.

assigns some algebraic structure to  $K$ , but it turns out that we can use it to assign a set of groups to the simplicial complex.

**Definition 6** (Homology group). *We define the  $k$ th homology group of a simplicial complex  $K$  as*

$$H_k(K) := \ker \partial_k / \text{im } \partial_{k+1}, \quad (4)$$

*i.e. a quotient group that we obtain from the two subgroups  $\ker \partial_k$  and  $\text{im } \partial_{k+1}$  of  $C_k$ .*

The  $k$ th homology group of a simplicial complex contains its  $k$ -dimensional topological features in the form of *equivalence classes* of simplicial chains, also known as *homology classes*. This rather abstract definition is best understood by an additional simplification step that involves calculating the rank of a homology group.

**Definition 7** (Betti number). *The rank of the  $k$ th homology group  $H_k(K)$  is known as the  $k$ th Betti number, denoted by  $\beta_k$ .*

Despite the rank being a rather coarse summary, Betti numbers turn out to be of immense utility in comparing different simplicial complexes. We may even reproduce the equation for the cyclomatic number,  $\beta_1 = m + \beta_0 - n$ , by noting that the *Euler characteristic*  $\chi(K) := \sum_i (-1)^i |\{\sigma \mid \dim \sigma = i\}|$  can also be expressed as a sum of alternating Betti numbers, i.e.  $\chi(K) := \sum_i (-1)^i \beta_i$ . For a proof of this surprising fact, see e.g. Munkres [69, p. 124]. Using this equivalence, we see that we can calculate  $\beta_1$  by reshuffling some of the terms, thus also explaining why the equation exhibits alternating signs.

At this point, we have introduced a large amount of algebraic machinery. Changing our focus back to graphs, we may reap some advantageous properties by noting that homology groups are somewhat preserved under graph isomorphism.<sup>10</sup>

**Lemma 1.** *Let  $G, G'$  be two isomorphic graphs. Then the homology groups of  $G$  and  $G'$  are isomorphic, i.e.  $H_k(G) \simeq H_k(G')$  for all  $k \geq 0$ .*

Lemma 1 is a direct consequence of the *functoriality* of homology. Functoriality of homology implies that, given any map  $f: G \rightarrow G'$ , there is an induced map  $H_k(f): H_p(G) \rightarrow H_p(G')$  between homology groups for all  $p \geq 0$ , and that, given two maps  $f: G \rightarrow G'$  and  $g: G' \rightarrow G''$ , we have  $H_k(g \circ f) = H_k(g) \circ H_k(f)$ . In particular, the identity map  $\text{id}: G \rightarrow G$  induces the identity map on the homology groups of  $G$ , making that an isomorphism between  $G$  and  $G'$  induces an isomorphism between their homology groups. Thus, the proof of the lemma is straightforward.

*Proof.* Let  $k \geq 0$  and let  $\varphi: G \rightarrow G'$  be an isomorphism between  $G$  and  $G'$ . By *functoriality* of homology,  $H_k(\varphi)$  is an isomorphism between  $H_k(G)$  and  $H_k(G')$ . ■

As a direct corollary, the Betti numbers of  $G$  and  $G'$  do not change, and in fact, a similar property holds for isomorphic simplicial complexes.

**Corollary 1.** *The Betti numbers of isomorphic graphs are equal, i.e.  $\beta_p(G) = \beta_p(G')$  for all  $p$ .*

This may be seen as a hint about the popularity of simplicial homology in algebraic topology: the framework leads directly to characteristic descriptions that remain invariant under (graph) isomorphism.

## C.2 Persistent Homology

Because of their conceptual simplicity—their calculation in low dimensions only involves knowledge about the connected components of a graph—Betti numbers are somewhat limited in their expressivity. Taking any graph  $G = (V, E)$ , even the addition of a single edge to  $G$  will change its Betti numbers, either by merging two connected components (thus decreasing  $\beta_0$ ) or by creating an additional cycle (thus increasing  $\beta_1$ ). This is a direct consequence of the definition of Euler’s formula, which effectively states that the insertion of a new edge  $e = (u, v)$  with  $u, v \in V$  either causes  $\beta_1$  to increase by 1 because  $m$  changes, or remain the same in case the number of connected components  $\beta_0$  changes. However, a single edge may only merge two connected components into one, so  $\beta_0$  may also at most decrease by 1. This indicates that Betti numbers are too coarse to be practically useful in large-scale graph analysis. It is possible to turn Betti numbers into a *multi-scale descriptor* of a graph.

<sup>10</sup>The reader well-versed in algebraic topology may be aware of this property directly, but we find it useful to mention this fact briefly.

This requires certain modifications to the previously-introduced concepts. Similar to [Appendix C.1](#), we will formulate everything in terms of simplicial complexes, again pointing out that this results in a more general description.

**Definition 8** (Filtration). *Given a simplicial complex  $K$ , we call a sequence of simplicial complexes filtration if it affords a nesting property of the form*

$$\emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_{m-1} \subseteq K_m = K. \quad (5)$$

*Since each element of this sequence is a valid simplicial complex, we can also think of this construction as ‘growing’  $K$  by adding simplices one after the other.*

Filtrations arise naturally when building simplicial complexes from point cloud data, but even in the context of graphs, we can imagine filtrations as *filtering* a graph based on some type of data, or function, assigned to its vertices. For instance, we may build a filtration of a graph based on the degree of its vertices, defining  $K_i$  to be the subgraph consisting of all vertices satisfying the degree condition, plus all edges whose endpoints satisfy it, i.e.

$$K_i := \{v \in V \mid \deg(v) \leq i\} \cup \{\{u, v\} \in E \mid \deg(u) \leq i \wedge \deg(v) \leq i\}. \quad (6)$$

Notice that we could also express the second condition more compactly by assigning to each 1-simplex (each edge) the *maximum* of the weight of its vertices. This construction is sometimes also referred to as a *lower-star filtration* since it extends a node-level function to higher-order simplices [35]. Not all filtrations have to be defined on the vertex level; as long as each edge in the filtration is preceded by its vertices, we can also build valid filtrations from functions that are primarily defined on edges.<sup>11</sup>

Setting aside further discussions about how to obtain filtrations for now, filtrations are compatible with the simplicial homology framework introduced above. The boundary operators  $\partial(\cdot)$ , together with the inclusion homomorphism between consecutive simplicial complexes, induce a homomorphism between corresponding homology groups of any filtration of  $m$  simplicial complexes. Given  $i \leq j$ , we write  $\iota^{i,j} : H_k(K_i) \rightarrow H_k(K_j)$  to denote this homomorphism. This construction yields a sequence of homology groups

$$0 = H_k(K_0) \xrightarrow{\iota_k^{0,1}} H_k(K_1) \xrightarrow{\iota_k^{1,2}} \dots \xrightarrow{\iota_k^{m-2,m-1}} H_k(K_{m-1}) \xrightarrow{\iota_k^{m-1,m}} H_k(K_m) = H_k(K) \quad (7)$$

for every dimension  $k$ . We then define the *kth persistent homology group* as

$$H_d^{i,j} := \ker \partial_k(K_i) / (\text{im } \partial_{k+1}(K_j) \cap \ker \partial_k(K_i)), \quad (8)$$

containing all topological features—homology classes—created in  $K_i$  that still exist in  $K_j$ . Following the definition of the ordinary Betti number, we then define the *kth persistent Betti number* to be the rank of this group, leading to  $\beta_k^{i,j} := \text{rank } H_d^{i,j}$ . It should be noted that this type of construction makes use of numerous deep mathematical concepts; for the sake of an expository article, we heavily summarise and compress everything to the most pertinent results.

The appeal of persistent homology can be seen when we start to make use of the features it captures. If we assume that our filtration is associated with a set of values  $a_0 \leq a_1 \leq \dots \leq a_{m-1} \leq a_m$ , such as the function values on the vertices, we can calculate *persistence diagrams*, i.e. simple topological feature descriptors.

**Definition 9** (Persistence diagram). *The  $k$ -dimensional persistence diagram of a filtration is the multiset of points in  $\mathbb{R}^2$  that, for each pair  $i, j$  with  $i \leq j$ , stores the tuple  $(a_i, a_j)$  with multiplicity*

$$\mu_{i,j}^{(k)} := \left( \beta_k^{i,j-1} - \beta_k^{i,j} \right) - \left( \beta_k^{i-1,j-1} - \beta_k^{i-1,j} \right). \quad (9)$$

*We will also assign a multiplicity to essential topological features of the simplicial complex, setting*

$$\mu_{i,\infty}^{(k)} := \beta_k^{i,m} - \beta_k^{i-1,m}, \quad (10)$$

*which denotes all features that are still present in the last simplicial complex of the filtration, i.e. in  $K_m = K$ . The persistence diagram thus contain all the information carried by Betti numbers.*

<sup>11</sup>In [Section 5](#), we will make use of a filtration defined on edge-based curvature values.

Persistence diagrams summarise the topological activity of a filtration. Given a persistence diagram  $\mathcal{D}$ , for any tuple  $(a_i, a_j)$ , the quantity  $|a_j - a_i|$  is called the *persistence* of the respective topological feature. Persistence indicates whether a feature, created in some simplicial complex during the filtration, is prominent or not. This notion was originally introduced by Edelsbrunner et al. [37] to analyse the relevance of topological features of a distance function; the terminology is supposed to indicate the prominence of a topological feature. Features with a high persistence are commonly taken to be relevant, whereas features with a low persistence used to be considered as noise; this assumption is changing as, depending on the filtration, low persistence may also just imply ‘low reliability.’ [12]

Persistence diagrams can be endowed with different metrics and kernels [54, 76], and it is known that the space of persistence diagrams is an Alexandrov space with curvature bounded from below [89]. The most common metric to compare two persistence diagrams is the *bottleneck distance*, defined as

$$d_B(\mathcal{D}, \mathcal{D}') := \inf_{\eta: \mathcal{D} \rightarrow \mathcal{D}'} \sup_{x \in \mathcal{D}} \|x - \eta(x)\|_\infty, \quad (11)$$

where  $\eta$  ranges over all bijections between the two persistence diagrams. Eq. (11) is solved using optimal transport; different cardinalities are handled by permitting points in one diagram to be transported to their corresponding projection on the diagonal. Another metric is the *Wasserstein distance* [93, Theorem 7.3], in which the sup calculation is replaced by a weighted sum over all distances between points in a diagram.

Stability properties of filtrations are a crucial aspect of research in computational topology [86]. Given two filtrations  $f, g$  of the same simplicial complex, a seminal result by Cohen-Steiner et al. [28] proves the following bound:

**Theorem 2** (Bottleneck stability). *Let  $f, g$  refer to filtrations of a simplicial complex  $K$ , and let  $\mathcal{D}_f$  and  $\mathcal{D}_g$  denote their respective persistence diagrams. The bottleneck distance is upper-bounded by  $d_B(\mathcal{D}_f, \mathcal{D}_g) \leq \|f - g\|_\infty$ , where  $\|\cdot\|_\infty$  refers to the supremum norm.*

An extension of this theorem, with  $d_B$  being replaced by the Wasserstein distance, shows that persistence diagrams are also stable in the Lipschitz sense [29]. These stability properties are remarkable because they link a topological quantity with a geometrical one, thus underscoring how persistent homology itself incorporates both geometrical and topological aspects of input data.

### C.3 Computational complexity of persistent homology

In this section, we briefly discuss the computational complexity of computing persistence diagrams. For a more detailed discussion, we refer the reader to [73].

Recall that persistence diagrams are obtained by first constructing a filtration of a simplicial complex. Then, persistence diagrams are computed from the homology groups of the simplicial complexes and the induced maps from inclusion between them. Therefore, the computational complexity of computing persistence diagrams is the sum of the complexity of building the filtration, and the complexity of the algorithm constructing the persistence diagrams.

Depending on the filtration type, the former complexity can greatly vary. For instance, the complexity of building a filtration based on the degree of vertices is quadratic in the number of vertices using the adjacency matrix if we do not add simplices of higher dimension than edges. On the other hand, the complexity of building a Vietoris–Rips filtration is high in general, having  $\mathcal{O}(n^{k+1})$  simplices up to dimension  $k$  for a set of  $n$  vertices, leading to research on more efficient filtrations similar to the original Vietoris–Rips one [84]. For the filtration of Theorem 1, computing the filtration is equivalent to compute the  $k$ -WL colouring of the graph, which can be done in  $\mathcal{O}(n^{k+1} \log n)$  where  $n$  is the number of vertices [59, Section 1].

For the latter, several algorithms exist depending on the filtration type and the homology coefficients used. A notable example is Ripser [8], which computes persistence diagrams for Vietoris–Rips filtrations in an efficient way. Using the standard algorithm, the worst case complexity is cubic in the number of simplices  $m$ , being this a sharp bound. For dimension zero, alternatives based on the connection between the single linkage clustering algorithm and zero-dimensional persistence diagrams [38, Claim 5.1] allows computation in  $\mathcal{O}(n^2)$ .

## D 1-WL and Persistent Homology

Formally, 1-WL proceeds by iteratively calculating a node colouring function  $C_i^{(1)} : V \rightarrow \mathbb{N}$ . The output of this function depends on the neighbours of a given node. For a vertex  $v$  at iteration  $i > 0$ , we have

$$C_i^{(1)}(v) := \text{RELABEL} \left( \left( C_{i-1}^{(1)}(v), \left\{ \left\{ C_i^{(1)}(u) \mid u \in \mathcal{N}(v) \right\} \right\} \right) \right), \quad (12)$$

where RELABEL refers to an injective function that maps the tuple of colours to a unique colour, i.e. a unique number and  $\mathcal{N}(v)$  refers to the neighbors of  $v$ , i.e. its adjacent vertices. The algorithm is initialised by either using existing labels or the degree of vertices.<sup>12</sup> After a finite number of steps, the colour assignments generated using Eq. (12) stabilise. If two graphs give rise to different colour sequences, the graphs are guaranteed to be non-isomorphic. The 1-WL test is computationally easy and constitutes an upper bound for the expressivity of many graph neural network (GNN) architectures [66, 96]. In other words, if 1-WL cannot distinguish two non-isomorphic graphs, GNNs will also not be able to distinguish them. As stated in Section 4, it is already a known result that *any* 1-WL colouring can be reproduced by creating a special filtration. The implication is that persistent homology is *at least as expressive* as 1-WL because there is a filtration that distinguishes all the graphs 1-WL can distinguish. In fact, Theorem 4 demonstrates that the 1-WL algorithm fails to detect even the most fundamental topological property of graphs: the number of connected components. This limitation stands in contrast to the capabilities of (persistent) homology, which can readily capture such basic topological features. Appendix F shows examples of graphs that can be distinguished based on their topological features but not by the 1-WL algorithm. These examples demonstrate that a topological perspective is strictly more expressive than 1-WL.

## E Theorems & Proofs

For the convenience of the reader, we restate all results again before providing their proof.

When working with filtrations in the subsequent proofs, it would be ideal to have filtrations that satisfy *injectivity* on the level of vertices, i.e.  $f(v) \neq f(v')$  if  $v \neq v'$ . Such injective filtrations have the advantage of permitting gradient-based optimisation schemes [45]. The following lemma, first proved in Horn et al. [48], demonstrates that injectivity is not a strict requirement, though, as it is always possible to find an injective filtration function that is arbitrarily close (in the Hausdorff sense) to a non-injective filtration function.

**Lemma 2.** *For all  $\epsilon > 0$  and a filtration function  $f$  defined on the vertices, i.e.  $f : V \rightarrow \mathbb{R}^d$ , there is an injective function  $\tilde{f} : V \rightarrow \mathbb{R}^d$  such that  $\|f - \tilde{f}\|_\infty < \epsilon$ .*

*Proof.* Let  $V = \{v_1, \dots, v_n\}$  be the vertices of a graph and  $\text{im } f = \{u_1, \dots, u_m\}$  be their images under  $f$ . Since  $f$  is not injective, we have  $m < n$ . We resolve non-injective vertex pairs iteratively. For  $u \in \text{im } f$ , let  $V' := \{v \in V \mid f(v) = u\}$ . If  $V'$  only contains a single element, we do not have to do anything. Otherwise, for each  $v' \in V'$ , pick a new value from  $B_\epsilon(u) \setminus \text{im } f$ , where  $B_r(x) \subset \mathbb{R}^d$  refers to the open ball of radius  $r$  around a point  $x$  (for  $d = 1$ , this becomes an open interval in  $\mathbb{R}$ , but the same reasoning applies in higher dimensions). Since we only ever remove a finite number of points, such a new value always exists, and we can modify  $\text{im } f$  accordingly. The number of vertex pairs for which  $f$  is non-injective decreases by at least one in every iteration, hence after a finite number of iterations, we have modified  $f$  to obtain  $\tilde{f}$ , an *injective* approximation to  $f$ . By always picking new values from balls of radius  $\epsilon$ , we ensure that  $\|f - \tilde{f}\|_\infty < \epsilon$ , as required. ■

**Proposition 1.** *For  $\mathcal{F}$  equivariant and  $G \simeq G'$ , the persistence diagrams of  $f_{K_G}$  and  $f_{K_{G'}}$  coincide.*

The proof of Proposition 1 is a direct consequence of the following lemma.

**Lemma 3.** *Let  $\mathbb{K}$  be any field. If  $\mathcal{F}$  is equivariant and  $G \simeq G'$ , then the persistence modules over  $\mathbb{K}$  obtained by applying the homology functor  $H_k$  to the chain complexes generated by the sublevel set filtrations  $f_{K_G}$  and  $f_{K_{G'}}$  are isomorphic for any  $k \geq 0$ .*

*Proof.* Let  $\varphi : G \rightarrow G'$  be an isomorphism between  $G$  and  $G'$ , and let  $a \in \mathbb{R}$ . We claim that  $\varphi$  induces a simplicial isomorphism between  $K_G(a) = f_{K_G}^{-1}((-\infty, a])$  and  $K_{G'}(a) = f_{K_{G'}}^{-1}((-\infty, a])$ . First,

<sup>12</sup>Note that Eq. (12) does not recognise an ordering of labels. Initialising 1-WL with a constant value thus leads to the *same* colouring—up to renaming—after the first iteration.

note that  $\sigma \in K_G(a)$  implies  $\varphi(\sigma) \in K_{G'}(a)$ . Then, the restriction  $\varphi_a$  of the map  $\varphi: K_G \rightarrow K_{G'}$  to  $K_G(a)$  is a well-defined simplicial complex morphism between  $K_G(a)$  and  $K_{G'}(a)$ . In particular,  $\varphi_a$  is an isomorphism. The injectivity of  $\varphi_a$  stems from the fact that  $\varphi$  is an isomorphism between simplicial complexes. The surjectivity of  $\varphi_a$  comes from the fact that  $f_{K_G}(\varphi^{-1}(\tau)) = f_{K_{G'}}(\tau)$  for all  $\tau \in K_{G'}(a)$ , and thus if  $\tau \in K_{G'}(a)$ , then  $\varphi^{-1}(\tau) \in K_G(a)$  and  $\varphi_a(\varphi^{-1}(\tau)) = \tau$ . Having prove that  $\varphi_a$  is an isomorphism, let  $C_\bullet(K_G(a))$  and  $C_\bullet(K_{G'}(a))$  be the chain complexes induced by the simplicial complexes  $K_G(a)$  and  $K_{G'}(a)$ , respectively. The previous function  $\varphi_a$  induces a chain map  $C_\bullet(\varphi_a)$  between  $C_\bullet(K_G(a))$  and  $C_\bullet(K_{G'}(a))$  defined as the linear map  $C_k(\varphi_a)$  sending an element  $\sigma \in C_k(K_G(a))$  to  $\varphi_a(\sigma)$  for  $k \geq 0$ . The proof that  $C_\bullet(\varphi_a)$  is a chain map can be found in Nanda [70, Proposition 4.5]. Concretely,  $C_k(\varphi_a)$  is an isomorphism because  $\varphi_a$  is a simplicial complex isomorphism that generates a one-to-one correspondence between the  $k$ -simplices of  $K_G(a)$  and  $K_{G'}(a)$ . As  $C_\bullet(\varphi_a)$  is a chain isomorphism,  $C_\bullet(\varphi_a)$  induces isomorphisms between the homology groups  $H_k(C_\bullet(K_G(a)))$  and  $H_k(C_\bullet(K_{G'}(a)))$  for all  $k \geq 0$ . Moreover, these isomorphisms  $C_\bullet(\varphi_a)$  constitute isomorphisms between the persistence modules given by  $(H_k(C_\bullet(K_G(a))))_{a \in \mathbb{R}}$  and  $(H_k(C_\bullet(K_{G'}(a))))_{a \in \mathbb{R}}$  for all  $k \geq 0$ . To prove it, we only need to show that the diagram

$$\begin{array}{ccc} H_k(C_\bullet(K_G(a_i))) & \xleftarrow{i} & H_k(C_\bullet(K_G(a_j))) \\ \downarrow C_k(\varphi_{a_i}) & & \downarrow C_k(\varphi_{a_j}) \\ H_k(C_\bullet(K_{G'}(a_i))) & \xleftarrow{i} & H_k(C_\bullet(K_{G'}(a_j))) \end{array} \quad (13)$$

commutes for  $a_i \leq a_j$ . This is a consequence of the definition of the isomorphisms  $C_\bullet(\varphi_a)$ . Note that  $C_k(\varphi_{a_i})$  and  $C_k(\varphi_{a_j})$  are the same map for the elements of  $C_k(K_G(a_i))$  for all  $k \geq 0$  and  $a_i \leq a_j$  due to the fact that  $\varphi_{a_i} = \varphi_{a_j}$  for  $K_G(a_i)$  for  $a_i \leq a_j$  by definition. Thus, given  $[c] \in H_k(C_\bullet(K_G(a_i)))$ , we have  $(C_k(\varphi_{a_j}) \circ i)([c]) = [\varphi_{a_j}(c)] = (i \circ C_k(\varphi_{a_i}))([c])$  for all  $a_i \leq a_j$  and  $k \geq 0$ , as we wanted to prove. ■

**Theorem 1.** *For  $k \geq 2$ , there exists an equivariant filtration generator  $\mathcal{F}$  of type 2 such that its zero-dimensional persistence diagrams are at least as expressive as  $k$ -FWL.*

*Proof.* The main idea involves harnessing the colours of  $k$ -tuples. Denote by  $\mathcal{F}(G) = (K_G, f_{K_G})$  the output of our filtration generation. For a given graph  $G$ , we set  $K_G$  to be the simplicial complex of dimension zero with 0-simplices given by the vertices of  $G$ . Let  $C_G$  be the multiset of colours assigned to the graph  $G$  by the  $k$ -FWL algorithm. Without loss of generality, we assume that the  $k$ -FWL algorithm assigns colours that are always greater or equal than one.

Now, given a finite multiset  $S$  of colours, this is, a finite multiset of elements in  $\mathbb{N} \cap [2, +\infty)$ , we define the representation of  $S$  as the concatenation

$$R(S) = c_1 \parallel \left\| \begin{array}{c} |S| \\ 0 \\ \parallel \\ i=2 \end{array} \right\| r_9(c_i),$$

where  $S = \{\{c_1, \dots, c_{|S|}\}\}$  is the multiset of colours ordered in non-decreasing order, and  $r_9(c_i)$  is the representation of the number  $c_i$  in bijective base-9 numeration. The map from the set of multisets of colours to the set of natural numbers given by  $R$  is injective because the representation of individual natural numbers in bijective base-9 is injective and because the digit zero does not appear in the base-9 representation of the numbers, allowing for the separation of the different colours, and for the recovery of the original multiset.

Set now  $f_{K_G}$  to be the constant filtration function that assigns to each simplex the value  $R(C_G)$ . Thus, the filtration generator  $\mathcal{F}$  is equivariant because the  $k$ -FWL algorithm outputs the same multiset of colours for isomorphic graphs  $G \simeq G'$ , and then the representation of the multiset of colours is the same for both graphs and yield the same constant filtration function.

Finally, take two non-isomorphic graphs  $G$  and  $G'$ . By definition of  $\mathcal{F}$ , the simplicial complexes  $K_G$  and  $K_{G'}$  contain only vertices, yielding non-empty persistence diagrams only in dimension zero. Particularly, the zero-dimensional persistence diagrams of  $G$  and  $G'$  are multisets containing as many non-diagonal points as the number of vertices in the graphs of the form  $(R(C_G), +\infty)$  and  $(R(C_{G'}), +\infty)$ , respectively. Since the graphs are non-isomorphic, the multisets and their representations are different and thus the persistence diagrams are different. ■

**Theorem 3.** *Given 1-WL colourings of two graphs  $G$  and  $G'$  that are different, there exists a filtration of  $G$  and  $G'$  such that their persistence diagrams in dimension 0 are also different.*

*Proof.* Since the colourings are different, there is an iteration  $h$  of 1-WL such that the label sequences of  $G$  and  $G'$  are different. We thus have at least one colour—equivalently, one label—whose count is different. Let  $\mathcal{L}^{(h)} := \{l_1, l_2, \dots\}$  be an enumeration of the finitely many hashed labels at iteration  $h$ . We can build a filtration function  $f$  by assigning a vertex  $v$  with label  $l_i$  to its index, i.e.  $f(v) := i$ , and setting  $f(v, w) := \max\{f(v), f(w)\}$  for an edge  $(v, w)$ . The resulting 0-dimensional persistence diagrams for  $G$  and  $G'$ , denoted by  $\mathcal{D}_0$  and  $\mathcal{D}'_0$ , respectively, now contain tuples of the form  $(i, j)$ . Moreover, each vertex is guaranteed to give rise to *exactly* one such pair since each vertex creates a connected component in 0-dimensional persistent homology. Letting  $\mu^{(i,j)}(\mathcal{D}_0)$  refer to the multiplicity of a tuple in  $\mathcal{D}_0$ , we know that, since the label count is different, there is *at least* one tuple  $(k, l)$  with  $\mu^{(k,l)}(\mathcal{D}_0) \neq \mu^{(k,l)}(\mathcal{D}'_0)$ . Hence,  $\mathcal{D}_0 \neq \mathcal{D}'_0$ . ■

Prior to stating [Corollary 2](#) concerning CFI graphs, we must first provide a precise definition of these structures. These definitions can also be found in Cai et al. [21, Section 6]. CFI graphs are constructed from a family of graphs  $\{X_k\}_{k \geq 2}$ . For each  $k \geq 2$ , we construct a pair of non-isomorphic CFI graphs  $(G_k, H_k)$  such that  $G_k$  and  $H_k$  cannot be distinguished by the  $(k-1)$ -FWL algorithm but can be distinguished by the  $k$ -FWL algorithm. Concretely,  $X_k$  is defined as the graph  $(V_k, E_k)$  where (i)  $V_k = A_k \cup B_k \cup M_k$  such that  $A_k = \{a_i\}_{i=1}^k$ ,  $B_k = \{b_i\}_{i=1}^k$ , and  $M_k = \{m_S : S \subseteq \{1, \dots, k\}, |S| \text{ is even}\}$ , and (ii)  $E_k = \{(m_S, a_i) : i \in S\} \cup \{(m_S, b_i) : i \notin S\}$ . Given any finite and connected graph  $G$  with at least degree two for all vertices, we can use the graphs  $X_k$  to build new graphs  $X(G)$  and  $\tilde{X}(G)$  that allows us to construct the CFI graphs.  $X(G)$  is built as follows. For each vertex  $v$  of  $G$ , we replace  $v$  by a copy of  $X_{\deg v}$ , called  $X(v)$ . Then, for each edge  $\{u, v\}$  of  $G$ , we associate to each of its endpoints  $u$  and  $v$  one of the pairs  $\{a_i, b_i\}$  from  $X(u)$  and  $X(v)$ , denoting the pairs as  $\{a(\{u, v\}, u), b(\{u, v\}, u)\}$  and  $\{a(\{u, v\}, v), b(\{u, v\}, v)\}$ , respectively. Finally, for each of edges  $\{u, v\}$  of  $G$ , we add the edges  $\{a(\{u, v\}, u), a(\{u, v\}, v)\}$  and  $\{b(\{u, v\}, u), b(\{u, v\}, v)\}$  to  $G$ , discarding the original edge. The graph  $\tilde{X}(G)$  is constructed in the same way as  $X(G)$ , but we arbitrarily choose one edge  $\{u, v\}$  of  $G$  and, instead of adding the previous two edges, we add the edges  $\{a(\{u, v\}, u), b(\{u, v\}, v)\}$  and  $\{a(\{u, v\}, v), b(\{u, v\}, u)\}$ . Finally,  $G_k$  and  $H_k$  are given by  $G_k = X(T_k)$  and  $H_k = \tilde{X}(T_k)$  where  $T_k$  is a degree-three graph with separator size  $k$ . In the original construction, these graphs are also coloured, but we will not need this information for the proof of [Corollary 2](#).

**Lemma 4.** *For any finite and connected graph  $G$ ,  $X(G)$  and  $\tilde{X}(G)$  have no cycles of length 3, and thus, no cliques of size 3 or more.*

*Proof.* Let  $\gamma$  be a cycle of  $X(G)$  of length three. First, note that there are no cycles of length 3 in the graphs  $X_k$  for any  $k \geq 1$ . This is because all edges go from  $M_k$  to a subset  $A_k$  or  $B_k$ . This implies that, in order to have a cycle, one needs at least 4 edges, passing twice by the set  $M_k$ . Therefore, the cycle  $\gamma$  must contain at least one edge  $e$  connecting two subgraphs  $X(u)$  and  $X(v)$  for  $u$  and  $v$  original vertices of  $G$ . The endpoints of this edge must be connected by design to two disjoint sets of type  $M_k$  of the subgraphs  $X(u)$  and  $X(v)$ , respectively. However, if the cycle  $\gamma$  is of length three, this means that the two endpoints of  $e$  are connected to the same element, implying that both sets  $M_k$  are not disjoint and arriving at a contradiction. The proof for  $\tilde{X}(G)$  is analogous. ■

**Corollary 2.** *For any  $k \geq 2$ , the CFI graphs  $G_k$  and  $H_k$  have no cycles of length 3, and thus, no cliques of size 3 or more.*

*Proof.* This is a direct consequence of [Lemma 4](#). ■

**Theorem 4.** *Let  $n = ab$  with integers  $a, b \geq 3$ . Then, the 1-WL test cannot distinguish between a cycles of length  $b$ ,  $b$  cycles of length  $a$ , and a single cycle of length  $n$ .*

*Proof.* In all three cases, each graph consists of  $n$  vertices, with each vertex having a degree of two. Initially, all vertices across the three graphs are assigned the same colour  $c$ . During the first iteration, each node receives as input the tuple  $(c, \{\{c, c\}\})$ . Consequently, after this iteration, all vertices in the three graphs obtain the same colour assignment. Since all vertices have the same colour at the end of the first iteration, the algorithm terminates. This occurs because the partition of nodes generated






FIGURE S.2: The image shows three graphs composed of three disjoint cycles of length four, four disjoint cycles of length three, and a single cycle of length twelve. For 1-WL, all vertices are initially assigned the same colour, in this case blue. During the first iteration, the RELABEL function for each vertex receives the tuple  $(\bullet, \{\{\bullet, \bullet\}\})$ . Consequently, all vertices in the three graphs receive the same new colour, keeping intact the initial partitions, finishing the algorithm. Thus, the final multiset of colours is the same for all three cases (twelve blue colours), making 1-WL unable to distinguish between these three graphs.

by the colours after the first iteration is identical to the partition from the initial assignment. As a result, the 1-WL test fails to distinguish between the three cases. ■

An example for  $a = 3, b = 4$  is provided in Fig. S.2. We end this section with two results concerning other graph properties that are being captured by persistent homology.

**Proposition 2.** *Given any filtration  $f$  of a graph  $G$  with a single connected component such that the values of  $f$  of the endpoints of edges are strictly lower than the values of their corresponding edges, Algorithm 1, used to compute  $\beta_0$ , yields an upper bound of  $\text{diam}(G)$ .*

*Proof.* We can provide a procedure to obtain an upper bound  $d$  of  $\text{diam}(G)$  alongside the calculation of  $\mathcal{D}_0$ . To this end, we set  $d = 0$ . While calculating  $\mathcal{D}_0$  with Algorithm 1, we check for each edge whether it is a destroyer, or a regular edge. If the edge is a destroyer, we increase  $d$  by one; otherwise, we do nothing. This procedure works because  $\text{diam}(G)$  is upper bounded by the diameter of the minimum spanning tree of  $G$ . Our estimate  $d$  counts the number of edges in such a tree. We thus have  $\text{diam}(G) \leq d$ . The bound is tight for some graphs, such as line graphs. ■

**Proposition 3.** *Given any filtration  $f$  of a graph  $G$ , Algorithm 1 yields an upper bound of the girth of  $G$ .*

*Proof.* Similar to Proposition 2, we can use the calculation of  $\mathcal{D}_0$ , the persistence diagram in dimension 0, of the filtration function to obtain an upper bound of the girth of the graph. We calculate  $\mathcal{D}_0$  with Algorithm 1, checking once again for each edge whether it is a creator or destroyer. If the edge is a creator, we stop and set our upper bound of the girth  $g$  to be the number of vertices in the connected component that contains the endpoints of the edge. Since the addition of the respective edge is guaranteed to create a cycle—the edge is a creator—the cycle has to make use of at most as many vertices as the number of vertices in the connected component to which it belongs. Our estimate  $g$  thus constitutes an upper bound of the girth of the graph. ■

Propositions 2 and 3 indicate that persistent homology captures more than ‘just’ topological information about a graph. We substantiate these theorems with empirical results concerning different filtrations and other graph properties in Section 5.3, thus showing how a topological perspective complements and enriches graph-learning tasks.

## F Topology and the Weisfeiler–Leman Hierarchy

In the following, we want to briefly extend the argumentation of the expressivity of persistent homology with respect to the Weisfeiler–Leman hierarchy. Since 1-WL is oblivious to certain topological structures such as cycles [4], the existence of graphs with different Betti number counts proves that persistent homology is *strictly more expressive* than 1-WL. For example, consider a graph consisting of the union of two triangles, i.e.  $\triangle \cup \triangle$ . This graph has  $\beta_0 = \beta_1 = 2$  since it consists of two connected components and two cycles. If we change the connectivity slightly to obtain a hexagon, i.e.  $\hexagon$ , we obtain a graph with  $\beta_0 = \beta_1 = 1$ . 1-WL is not able to distinguish between these graphs, but persistent homology can, since the Betti numbers of the graph are still encoded in the persistence diagram as essential features. Note that Theorem 3 does *not* apply to arbitrary filtrations since the theorem requires knowing the correct labels assigned by 1-WL. Finding filtration functions that are able to split graphs in a manner that is provably equivalent to 1-WL remains an open research question.

TABLE S.1: Success rate ( $\uparrow$ ) of distinguishing pairs of *connected cubic graphs* when using five different filtrations at varying expansion levels of the graph (denoted by  $k$ ). Due to the regularity of each graph,  $k = 3$  is omitted since the clique complex of the graph is exactly the same as for  $k = 2$ . These graphs can be distinguished by 2-FWL but not by 1-WL.

Data	$k = 1$					$k = 2$				
	Filtration									
	D	O	F	L	V	D	O	F	L	V
cub06	0.00	<b>1.00</b>	<b>1.00</b>	0.00	0.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
cub08	0.00	0.90	0.70	0.00	0.00	0.90	0.90	0.90	<b>1.00</b>	0.90
cub10	0.00	0.98	0.66	0.00	0.00	0.81	0.99	0.88	<b>1.00</b>	0.85
cub12	0.00	0.99	0.64	0.00	0.00	0.80	1.00	0.87	<b>1.00</b>	0.87
cub14	0.00	0.99	0.62	0.00	0.00	0.79	1.00	0.86	<b>1.00</b>	0.89

## G Additional Expressivity Experiments

This section contains additional expressivity experiments that had to be excluded from the main paper for reasons of space.

### G.1 Additional Results for Connected Cubic Graphs

We start our supplementary experiments by distinguishing connected cubic graphs, i.e. 3-regular graphs. These graphs cannot be distinguished by 1-WL, but they can be distinguished by 2-FWL [14, 65]. As such, they provide a good example of how different filtrations harness different types of graph information. Table S.1 shows the results. We first observe that the degree filtration is incapable of distinguishing graphs for  $k = 1$ . This is a direct consequence of the regularity—since the function is constant on the graph, persistent homology cannot capture any variability. This changes, however, when higher-order structures—triangles—are included for  $k = 2$ . We also observe that the Laplacian-based filtration for  $k = 2$  exhibits strong empirical performance; in the absence of additional information, the spectral properties captured by the Laplacian help in distinguishing graphs. The Ollivier–Ricci curvature filtration is performing similarly well also for the same dimension, while it outperforms the Laplacian filtration for  $k = 1$ . In contrast to the Laplacian-based filtration, it does not require the calculation of eigenvalues, which may be prohibitive for larger graphs.<sup>13</sup> Finally, we see that the Forman–Ricci curvature filtration, a purely combinatorial quantity depending only on local counts, is consistently outperformed by the Ollivier–Ricci curvature for  $k = 1$ . However, for  $k = 1$ , the Forman–Ricci filtration outperforms the Laplacian filtration. Finally, the Vietoris–Rips filtration is incapable of distinguishing the graphs for  $k = 1$ , and performs similarly to the Forman–Ricci curvature for  $k = 2$ , suggesting that filtrations based on distances are not capable of capturing the graph structure in an optimal way.

### G.2 Additional Results for the BREC Data Set

Table S.2 shows BREC data set results, itemised by the value of  $k$ , while Table S.3 provides a summary and performance comparison of the persistent homology results with other baselines.

### G.3 Additional Figures for Graph-Property Prediction Tasks

Fig. S.3 shows the distributions of graph properties that we predict in the main paper in Section 5.3.

## H Additional Graph Classification Experiments

To assess the capacity of persistent homology in graph learning tasks, we have designed a set of graph classification experiments that use the filtrations introduced in Section 5 to extract persistence diagrams from graph classification data set to perform inference on them. Two fundamental differences between our approach and the one used in [45] are that, we do not learn an optimal filtration

<sup>13</sup>Ollivier–Ricci curvature requires solving optimal transport problems, for which highly efficient approximative solvers are available [33].

TABLE S.2: Success rate ( $\uparrow$ ) for distinguishing pairs of instances of the *BREC data set* when using five different filtrations at varying expansion levels of the graph (denoted by  $k$ ). Due to combinatorial constraints, we did not calculate the Vietoris–Rips filtration for  $k = 4$ . Legend and number of graphs per category: B (Basic, 60), R (Regular, 100), E (Extension, 100), C (CFI, 100), 4, 20 (4-Vertex Condition), D (Distance-Regular, 20) graphs, respectively and A (average over full data set, 400).

$k = 1$					
<i>Data</i>	<i>Filtration</i>				
	D	O	F	L	V
Basic (60)	0.03	0.93	0.87	<b>1.00</b>	0.00
Regular (100)	0.00	0.42	0.32	0.00	0.00
Extension (100)	0.07	0.76	0.44	0.94	0.00
CFI (100)	0.03	0.03	0.03	<b>0.06</b>	0.03
4-VC (20)	0.00	0.00	0.00	0.00	0.00
DR (20)	0.00	0.00	0.00	<b>0.05</b>	0.00
Average (400)	0.03	0.44	0.33	0.40	0.01
$k = 2$					
<i>Data</i>	<i>Filtration</i>				
	D	O	F	L	V
Basic (60)	0.78	<b>1.00</b>	0.98	<b>1.00</b>	0.52
Regular (100)	0.39	0.54	0.50	0.48	0.39
Extension (100)	0.26	0.92	0.59	<b>1.00</b>	0.11
CFI (100)	0.03	0.03	0.03	<b>0.06</b>	0.03
4-VC (20)	0.00	0.00	0.00	0.00	0.00
DR (20)	0.00	0.00	0.00	<b>0.05</b>	0.00
Average (400)	0.29	0.52	0.43	0.54	0.21
$k = 3$					
<i>Data</i>	<i>Filtration</i>				
	D	O	F	L	V
Basic (60)	0.83	<b>1.00</b>	0.98	<b>1.00</b>	0.58
Regular (100)	0.85	0.93	0.91	0.93	0.85
Extension (100)	0.29	0.92	0.59	<b>1.00</b>	0.16
CFI (100)	0.03	0.03	0.03	<b>0.06</b>	0.03
4-VC (20)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
DR (20)	0.00	0.00	0.00	<b>0.05</b>	<b>0.05</b>
Average (400)	0.47	0.67	0.58	0.70	0.40
$k = 4$					
<i>Data</i>	<i>Filtration</i>				
	D	O	F	L	V
Basic (60)	0.83	<b>1.00</b>	0.98	<b>1.00</b>	—
Regular (100)	0.89	0.97	0.95	0.97	—
Extension (100)	0.29	0.92	0.59	<b>1.00</b>	—
CFI (100)	0.03	0.03	0.03	0.06	—
4-VC (20)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	—
DR (20)	0.00	0.00	0.00	0.05	—
Average (400)	0.48	0.68	0.59	0.71	—

TABLE S.3: Success rate ( $\uparrow$ ) for distinguishing pairs of instances of the *BREC data set* when using four different filtrations for  $k = 4$ . Vietoris–Rips filtration not computed for  $k = 4$  due to computational constraints. The data sets *Regular*, *4-Vertex Condition*, and *Distance Regular*, from Table 3 are merged into the *All regular* data set. Green indicates the best performing algorithm, while orange indicates the second best.

Data	SOTA		Filtration ( $k = 4$ )			
	3-WL	$N_2$	D	O	F	L
Basic (60)	1.000	1.000	0.83	1.000	0.983	1.000
All regular (140)	0.36	0.986	0.78	0.84	0.82	0.843
Extension (100)	1.000	1.000	0.29	0.920	0.59	1.000
CFI (100)	0.600	0.00	0.03	0.03	0.03	0.060
Average (400)	0.68	0.745	0.48	0.68	0.59	0.710




FIGURE S.3: Distribution of maximum radii, maximum diameters, and girths in the *ogbg-molhiv* molecular graph data set [49]. The values are concentrated on the lower end of the spectrum, with medians of 6, 11, and 5 and maximums of 47, 93, and 36, respectively. The maximum value for the girth is computed without taking into account infinite values.

and also we do not use deep learning models to perform the classification, as we are interested into the capacity of persistent homology to perform classification, and not in its combination with neural networks. For this endeavour, we train a Random Forest classifier on the persistent images [2] computed from the persistence diagrams extracted from the input graphs using the previous filtrations up to dimension  $k = 2$ . We test our approach in the MUTAG [34], IMDB-Binary [98], PROTEINS [15], NC1 [94], and *ogbg-molhiv* [49] data set. Except for *ogbg-molhiv*, we perform a Stratified 3-Fold experiment and provide the average and standard deviation of the multiple runs. For *ogbg-molhiv*, we perform only one experiment instance with the official train and test splits. The results are shown in Table S.4. We observe that, although we discard the data set features when using persistent homology, we are able to achieve suitably good results on some of the data sets, such as MUTAG and IMDB-Binary [39], even surpassing state-of-the-art results reported in Hofer et al. [45] and O’Bray et al. [71] for IMDB-Binary, MUTAG, and NCI1. For the *ogbg-molhiv*, however, we obtain consistent ROC-AUC values of 0.5, which is the same as random guessing. Our hypothesis is that molecular graphs rely on the annotated data set features to perform well, and the isolated topological information is not enough to perform the classification. The success of the experiments suggests that persistent homology is capable of capturing essential structural information about the graphs to classify and thus, suggests that persistent homology can be expressive in practice. We leave a more structured and comprehensive benchmark for persistent homology-like methods for graph-learning tasks to future work.

## I Additional results on alpha complexes

As the last set of experiments, we repeat the experiments from Section 5 except for the the property prediction tasks on random graphs, but using an alpha complex filtration. *Alpha complex filtrations* [36,

TABLE S.4: Graph classification average accuracy ( $\uparrow$ ) and standard deviation in test for the experiments with data set IMDB-Binary (I), PROTEINS (P), MUTAG (M), and NCI1 (N), and ROC-AUC ( $\uparrow$ ) for the ogbg-molhiv (O) test data set. The results are averaged over three runs using a stratified 3-means for all data set except for ogbg-molhiv, where only one run is performed. The best results are highlighted in bold. Experiments with the Vietoris-Rips filtration for the PROTEINS and molhiv are not reported for  $k = 2$  due to out-of-resources errors during execution.  $S_1$  and  $S_2$  refer to accuracy of the most effective methods reported in [71] and [45], respectively. The abbreviation DS stands for data set.

$k = 1$							
DS	SOTA		Filtration				
	$S_1$	$S_2$	D	O	F	L	V
P	<b>0.75 ± 0.05</b>	0.74 ± 0.03	0.70 ± 0.01	0.74 ± 0.01	0.72 ± 0.02	0.72 ± 0.02	0.70 ± 0.00
I	0.74 ± 0.04	<b>0.75 ± 0.05</b>	<b>0.75 ± 0.02</b>	0.71 ± 0.01	0.69 ± 0.03	0.65 ± 0.01	0.68 ± 0.03
M	0.87 ± 0.08	-	0.86 ± 0.01	0.87 ± 0.04	<b>0.90 ± 0.04</b>	0.79 ± 0.06	0.87 ± 0.01
N	-	0.71 ± 0.02	0.68 ± 0.00	0.73 ± 0.00	0.69 ± 0.01	0.67 ± 0.00	0.66 ± 0.01
O	-	-	<b>0.50</b>	0.50	<b>0.50</b>	0.50	0.50
$k = 2$							
DS	SOTA		Filtration				
	$S_1$	$S_2$	D	O	F	L	V
P	0.75 ± 0.05	0.74 ± 0.03	0.70 ± 0.02	0.73 ± 0.02	0.70 ± 0.01	0.68 ± 0.01	-
I	0.74 ± 0.04	0.75 ± 0.05	0.73 ± 0.03	0.71 ± 0.02	0.72 ± 0.03	0.68 ± 0.03	0.66 ± 0.02
M	0.87 ± 0.08	-	0.86 ± 0.01	0.89 ± 0.04	0.89 ± 0.03	0.81 ± 0.04	0.87 ± 0.03
N	-	0.71 ± 0.02	0.68 ± 0.00	<b>0.74 ± 0.01</b>	0.70 ± 0.00	0.68 ± 0.00	0.69 ± 0.01
O	-	-	<b>0.50</b>	<b>0.50</b>	0.50	0.50	-

Section III.4] are *geometric* filtrations, meaning that they are computed from point clouds in the Euclidean space, in our case, coming from embeddings of graph vertices. Concretely, we embed the vertices of our graphs in  $\mathbb{R}^3$  using *Laplacian eigenmaps* [11]. Note that the embedding map is an arbitrary choice, and we leave the exploration of other embeddings for future work, making the use of alpha complexes in this context a proof of concept. Tables S.5 to S.10 contain the BREC, classification, connected cubic, minimal Cayley, strongly-regular, and property prediction experiments, respectively, using the alpha complex filtration.

For BREC, we observe that the alpha complex filtration obtains almost perfect success rate in distinguishing the non-isomorphic graphs, failing only on the CFI category, where the alpha complex filtration still shows a strong success rate, outperforming the other filtrations and non-topological methods by a large margin. For the other isomorphism datasets (connected cubic, minimal Cayley, strongly-regular), alpha complexes obtain perfect accuracy in all cases.

For the classification tasks, though, we observe that the alpha complex filtration performs worse than the best filtrations at each task worse than the other state-of-the-art methods. Due to the good results in the graph isomorphism tasks, we hypothesize that the alpha complex is highly expressive (and sensitive) and that the Random Forest with the persistence images overfits the data without further regularization.

Finally, for the property prediction experiments, we observe that the alpha complex filtration, as in the classification tasks, performs worse than the other filtrations except for the girth prediction, where it obtains better results than curvature-based and degree filtrations.

Overall, our preliminary results suggest that the alpha complex filtration is highly expressive and capable of capturing the graph structure, but it is also sensitive to noise and overfitting, at least for Laplacian eigenmaps embeddings.

## J Implementation and Hardware details

Our implementation is based on Python 3. We plan on releasing the code under a BSD-3-Clause license. For review purposes, the code has been attached to the supplementary materials. Please

TABLE S.5: Success rate ( $\uparrow$ ) for distinguishing pairs of instances of the *BREC data set* when using the alpha complex filtration described at [Appendix I](#) and persistence diagrams up to varying dimensions (denoted by  $k$ ). Legend and number of graphs per category: B (Basic, 60), R (Regular, 100), E (Extension, 100), C (CFI, 100), 4, 20 (4-Vertex Condition), D (Distance-Regular, 20) graphs, respectively and A (average over full data set, 400). For the results using the other filtrations, see [Table S.2](#).

Data	$k = 1$	$k = 2$	$k = 3$
Basic (60)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
Regular (100)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
Extension (100)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
CFI (100)	0.63	0.80	<b>0.87</b>
4-Vertex_Condition (20)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
Distance_Regular (20)	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
Average (400)	0.91	0.95	<b>0.97</b>

TABLE S.6: Graph classification average accuracy ( $\uparrow$ ) and standard deviation in test for the experiments with data set IMDB-Binary (I), PROTEINS (P), MUTAG (M), and NCI1 (N), and ROC-AUC ( $\uparrow$ ) for the ogbg-molhiv (O) test data set when using the alpha complex filtration described at [Appendix I](#) and persistence diagrams up to varying dimensions (denoted by  $k$ ). The results are averaged over three runs using a stratified 3-means for all data set except for ogbg-molhiv, where only one run is performed. The best results are highlighted in bold. For the results using the other filtrations, see [Table S.4](#).

Data set	$k = 1$	$k = 2$
P	<b>0.66 <math>\pm</math> 0.02</b>	0.66 $\pm$ 0.01
I	<b>0.59 <math>\pm</math> 0.04</b>	0.57 $\pm$ 0.04
M	0.79 $\pm$ 0.03	<b>0.79 <math>\pm</math> 0.03</b>
N	<b>0.62 <math>\pm</math> 0.01</b>	0.60 $\pm$ 0.02
O	0.48 $\pm$ 0.00	<b>0.50 <math>\pm</math> 0.00</b>

TABLE S.7: Success rate ( $\uparrow$ ) of distinguishing pairs of *connected cubic graphs* when using the alpha complex filtration described at [Appendix I](#) and persistence diagrams up to varying dimensions (denoted by  $k$ ). These graphs can be distinguished by 2-FWL but not by 1-WL. For the results using the other filtrations, see [Table S.1](#).

Data	$k = 1$	$k = 2$
cub06	<b>1.00</b>	<b>1.00</b>
cub08	<b>1.00</b>	<b>1.00</b>
cub10	<b>1.00</b>	<b>1.00</b>
cub12	<b>1.00</b>	<b>1.00</b>
cub14	<b>1.00</b>	<b>1.00</b>

TABLE S.8: Success rate ( $\uparrow$ ) for distinguishing pairs of *minimal Cayley graphs* when using the alpha complex filtration described at [Appendix I](#) and persistence diagrams up to varying dimensions (denoted by  $k$ ). For the results using the other filtrations, see [Table 2](#).

Data	$k = 1$	$k = 2$
cay12	<b>1.00</b>	<b>1.00</b>
cay16	<b>1.00</b>	<b>1.00</b>
cay20	<b>1.00</b>	<b>1.00</b>
cay24	<b>1.00</b>	<b>1.00</b>
cay32	<b>1.00</b>	<b>1.00</b>
cay36	<b>1.00</b>	<b>1.00</b>
cay60	<b>1.00</b>	<b>1.00</b>
cay63	<b>1.00</b>	<b>1.00</b>

TABLE S.9: Success rate ( $\uparrow$ ) for distinguishing pairs of *strongly-regular graphs* when using the alpha complex filtration described at Appendix I and persistence diagrams up to varying dimensions (denoted by  $k$ ). 2-FWL cannot distinguish between any of these pairs. For the results using the other filtrations, see Table 1.

Data	$k = 1$	$k = 2$	$k = 3$
16622	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
251256	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
261034	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
281264	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
291467	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
351668	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
351899	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
361446	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
401224	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>

TABLE S.10: Accuracy ( $\uparrow$ ) when predicting the properties of graphs in the ogbg-molhiv molecular graph data set [49] using the alpha complex filtration described at Appendix I and persistence diagrams up to varying dimensions (denoted by  $k$ ). For the results using the other filtrations, see Table 4.

Data	$k = 1$	$k = 2$
Diameter	0.00	<b>0.01</b>
Girth	<b>0.41</b>	0.40
Radius	0.00	<b>0.01</b>

use `pip install -r requirements.txt` on the root folder of the code to set up the environment. The experiments were executed on a server with an AMD EPYC 7452 (128) @ 2.350GHz CPU, 503GiB of RAM memory, no GPU acceleration, and Ubuntu 22.04.4 LTS with the 6.5.0-28-generic Linux kernel. Each experiment was executed on a single process. Not completed experiments in the main text were due to process termination by the operating system with the previous constraints.

- To perform the BREC experiments, use the script `BREC_experiments.py`.
- To perform the classification experiments, use the script `classification_tasks.py`.
- To perform the three non-BREC expressivity experiments (cubic, regular, and Cayley graphs), use the script `simple_isomorphism_experiments.py`.
- To perform the graph properties prediction experiments, use the code `property_prediction_experiments.py`.
- To perform the graph properties prediction experiments for the Watts–Strogatz and Erdős–Rényi random graphs, use the code `predict_diameter_random_graphs.py`.

The easiest way to perform the experiments in all the cases is using the flag `-a`, that executes all the experiments without the need of specifying specific parameters about the datasets and filtrations. Once the experiments are performed, the tables from the paper can be generated using the flag `-t` on the same scripts.

## K Licenses

We do not redistribute any existing data sets, but briefly mention their licenses here. Expressivity experiments on known graphs make use of an existing database [31], which does not directly specify any licensing requirements but asks for a citation. The BREC data set and the ogbg-molhiv data set are distributed under a MIT license. All other graph-classification data sets, i.e. IMDB-Binary, PROTEINS, MUTAG, and NCI1, do not specify a license. They are distributed as part of the ‘TUDatasets’ repository, and can be accessed via <https://chrsmrrs.github.io/datasets/>. We make our code available under a 3-Clause BSD License.