
WL meet VC

Christopher Morris^{* 1} Floris Geerts^{* 2} Jan Tönshoff¹ Martin Grohe¹

Abstract

Recently, many works studied the expressive power of graph neural networks (GNNs) by linking it to the 1-dimensional Weisfeiler–Leman algorithm (1-WL). Here, the 1-WL is a well-studied heuristic for the graph isomorphism problem, which iteratively colors or partitions a graph’s vertex set. While this connection has led to significant advances in understanding and enhancing GNNs’ expressive power, it does not provide insights into their generalization performance, i.e., their ability to make meaningful predictions beyond the training set. In this paper, we study GNNs’ generalization ability through the lens of Vapnik–Chervonenkis (VC) dimension theory in two settings, focusing on graph-level predictions. First, when no upper bound on the graphs’ order is known, we show that the bitlength of GNNs’ weights tightly bounds their VC dimension. Further, we derive an upper bound for GNNs’ VC dimension using the number of colors produced by the 1-WL. Secondly, when an upper bound on the graphs’ order is known, we show a tight connection between the number of graphs distinguishable by the 1-WL and GNNs’ VC dimension. Our empirical study confirms the validity of our theoretical findings.

1. Introduction

Graph-structured data are prevalent across application domains ranging from chemo- and bioinformatics (Barabasi & Oltvai, 2004; Jumper et al., 2021; Stokes et al., 2020) to image (Simonovsky & Komodakis, 2017) and social-network analysis (Easley & Kleinberg, 2010), indicating the importance of machine learning methods for such data. Nowadays, there are numerous approaches for machine

^{*}Equal contribution ¹Department of Computer Science, RWTH Aachen University, Aachen, Germany ²Department of Computer Science, University of Antwerp, Antwerp, Belgium. Correspondence to: Christopher Morris <morris@cs.rwth-aachen.de>.

learning for graph-structured, most notably those based on *graph kernels* (Borgwardt et al., 2020; Kriege et al., 2020) or *graph neural networks* (GNNs) (Chami et al., 2020; Gilmer et al., 2017; Morris et al., 2021). Here, graph kernels (Sher-vashidze et al., 2011) based on the *1-dimensional Weisfeiler–Leman algorithm* (1-WL) (Weisfeiler & Leman, 1968), a well-studied heuristic for the graph isomorphism problem, and corresponding GNNs (Morris et al., 2019; Xu et al., 2019), have recently advanced the state-of-the-art in supervised vertex- and graph-level learning (Morris et al., 2021). Further, based on the *k-dimensional Weisfeiler–Leman algorithm* (*k*-WL), 1-WL’s more powerful generalization, several works generalized GNNs to *higher-order GNNs* (*k*-GNNs), resulting in provably more expressive architectures, e.g., Azizian & Lelarge (2021); Geerts & Reutter (2022); Maron et al. (2019); Morris et al. (2019; 2020b; 2021; 2022).

While devising provably expressive GNN-like architectures is a meaningful endeavor, it only partially addresses the challenges of machine learning with graphs. That is, expressiveness results reveal little about an architecture’s ability to generalize to graphs outside the training set. Surprisingly, only a few notable contributions study GNNs’ generalization behaviors, e.g., Garg et al. (2020); Kriege et al. (2018); Liao et al. (2021); Maskey et al. (2022); Scarselli et al. (2018). However, these approaches express GNN’s generalization ability using only classical graph parameters, e.g., maximum degree, number of vertices, or edges, which cannot fully capture the complex structure of real-world graphs. Further, most approaches study generalization in the *non-uniform regime*, i.e., assuming that the GNNs operate on graphs of a pre-specified order. Further, they only investigate the case $k = 1$, i.e., standard GNNs, ignoring more expressive generalizations; see the previous paragraph.

1.1. Present work

This paper investigates the influence of graph structure and the parameters’ encoding lengths on GNNs’ generalization by tightly connecting 1-WL’s expressivity and GNNs’ Vapnik–Chervonenkis (VC) dimension. Specifically, our contributions are:

1. In the non-uniform regime, we prove *tight* bounds on GNNs’ VC dimension. We show that GNNs’ VC dimension depends tightly on the number of equivalence

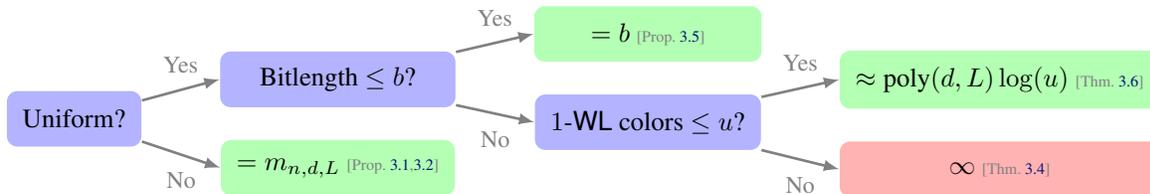


Figure 1: Overview of our results for bounded-width GNNs. Green and red boxes denote VC dimension bounds. Here, $m_{n,d,L}$ denotes the number of graphs of order at most n with boolean d -dimensional features distinguishable by 1-WL after L iterations.

classes computed by the 1-WL over a set of graphs; see Propositions 3.1 and 3.2. Moreover, our results easily extend to the k -WL and many recent expressive GNN extensions.

2. In the uniform regime, i.e., when graphs can have arbitrary order, we show that GNNs’ VC dimension is *lower and upper bounded* by the largest bitlength of its weights; see Proposition 3.5.
3. In both the uniform and non-uniform regimes, GNNs’ VC dimension depends *logarithmically on the number of colors* computed by the 1-WL and polynomially on the number of parameters; see Theorem 3.6.
4. Empirically, we show that our theoretical findings hold in practice.

Overall, our results provide new insights into GNNs’ generalization behavior and how graph structure and parameters influence it. Specifically, our results imply that a complex graph structure, captured by 1-WL, results in worse generalization performance. The same holds for increasing the encoding length of the GNN’s parameters. *Importantly, our theory provides the first link between expressivity results and generalization ability.* Moreover, our results establish the *first* lower bounds for GNNs’ VC dimension. See Figure 1 for a high-level overview of our results.

1.2. Related work

In the following, we discuss relevant related work.

GNNs Recently, GNNs (Gilmer et al., 2017; Scarselli et al., 2009) emerged as the most prominent graph representation learning architecture. Notable instances of this architecture include, e.g., Duvenaud et al. (2015); Hamilton et al. (2017), and Veličković et al. (2018), which can be subsumed under the message-passing framework introduced in Gilmer et al. (2017). In parallel, approaches based on spectral information were introduced in, e.g., Bruna et al. (2014); Defferrard et al. (2016); Gama et al. (2019); Kipf & Welling (2017); Levie et al. (2019), and Monti et al. (2017)—all of which descend from early work in Baskin

et al. (1997); Goller & Küchler (1996); Kireev (1995); Merkwirth & Lengauer (2005); Micheli & Sestito (2005); Micheli (2009); Scarselli et al. (2009), and Sperduti & Starita (1997).

Limits of GNNs and more expressive architectures Recently, connections between GNNs and Weisfeiler–Leman type algorithms have been shown (Barceló et al., 2020; Geerts et al., 2021; Morris et al., 2019; Xu et al., 2019). Specifically, Morris et al. (2019) and Xu et al. (2019) showed that the 1-WL limits the expressive power of any possible GNN architecture in terms of distinguishing non-isomorphic graphs. In turn, these results have been generalized to the k -WL, see, e.g., Azizian & Lelarge (2021); Geerts (2020); Maron et al. (2019); Morris et al. (2019; 2020b; 2022), and connected to permutation-equivariant functions approximation over graphs, see, e.g., Chen et al. (2019); Maehara & NT (2019); Azizian & Lelarge (2021); Geerts & Reutter (2022). Further, Aamand et al. (2022) devised an improved analysis using randomization. Recent works have extended the expressive power of GNNs, e.g., by encoding vertex identifiers (Murphy et al., 2019; Vignac et al., 2020), using random features (Abboud et al., 2021; Dasoulas et al., 2020; Sato et al., 2021), equivariant graph polynomials (Puny et al., 2023), homomorphism and subgraph counts (Barceló et al., 2021; Bouritsas et al., 2020; Nguyen & Maehara, 2020), spectral information (Balcilar et al., 2021), simplicial (Bodnar et al., 2021b) and cellular complexes (Bodnar et al., 2021a), persistent homology (Horn et al., 2022), random walks (Tönshoff et al., 2021; Martinkus et al., 2022), graph decompositions (Talak et al., 2021), relational (Barceló et al., 2022), distance (Li et al., 2020) and directional information (Beaini et al., 2021), subgraph information (Bevilacqua et al., 2022; Cotta et al., 2021; Feng et al., 2022; Frasca et al., 2022; Huang et al., 2022; Morris et al., 2021; Papp et al., 2021; Papp & Wattenhofer, 2022; Qian et al., 2022; Thiede et al., 2021; Wijesinghe & Wang, 2022; You et al., 2021; Zhang & Li, 2021; Zhao et al., 2022; Zhang et al., 2023a), and biconnectivity (Zhang et al., 2023b). See Morris et al. (2021) for an in-depth survey on this topic. Geerts & Reutter (2022) devised a general approach for bounding the expressive power of a large variety of GNNs utilizing the 1-WL or k -WL.

Recently, Kim et al. (2022) showed that transformer architectures (Müller et al., 2023) can simulate the 2-WL. Grohe (2023) showed tight connections between GNNs’ expressivity and circuit complexity. Moreover, Rosenbluth et al. (2023) investigated the expressive power of different aggregation functions beyond sum aggregation.

GNN’s generalization capabilities Scarselli et al. (2018) used classical techniques from learning theory (Karpinski & Macintyre, 1997) to show that GNNs’ VC dimension (Vapnik, 1995) with piece-wise polynomial activation functions on a *fixed* graph, under various assumptions, is in $\mathcal{O}(P^2 n \log n)$, where P is the number of parameters and n is the order of the input graph. We note here that Scarselli et al. (2018) analyzed a different type of GNN not aligned with modern GNN architectures (Gilmer et al., 2017). Garg et al. (2020) showed that the empirical Rademacher complexity, e.g., (Mohri et al., 2018), of a specific, simple GNN architecture, using sum aggregation, is bounded in the maximum degree, the number of layers, Lipschitz constants of activation functions, and parameter matrices’ norms. We note here that their analysis assumes weight sharing across layers. Liao et al. (2021) refined these results via a PAC-Bayesian approach, further refined in Ju et al. (2023). Maskey et al. (2022) used random graphs models to show that GNNs’ generalization ability depends on the (average) number of vertices in the resulting graphs. Verma & Zhang (2019) studied the generalization abilities of 1-layer GNNs in a transductive setting based on algorithmic stability. Similarly, Esser et al. (2021) used stochastic block models to study the transductive Rademacher complexity (El-Yaniv & Pechyony, 2007; Tolstikhin & Lopez-Paz, 2016) of standard GNNs. Moreover, (Kriege et al., 2018) leveraged results from graph property testing (Goldreich, 2010) to study the sample complexity of learning to distinguish various graph properties, e.g., planarity or triangle freeness, using graph kernels (Borgwardt et al., 2020; Kriege et al., 2020). We stress that all of the above approaches only consider classical graph parameters to bound the generalization abilities of GNNs. Finally, (Yehudai et al., 2021) showed negative results for GNNs’ ability to generalize to larger graphs. However, the generalization properties of GNNs and their connection to expressivity is understood to a lesser extent.

See Appendix A for an overview of the Weisfeiler–Leman algorithm’s theoretical properties.

2. Preliminaries

Let $\mathbb{N} := \{1, 2, 3, \dots\}$. For $n \geq 1$, let $[n] := \{1, \dots, n\} \subset \mathbb{N}$. We use $\{\{\dots\}\}$ to denote multisets, i.e., the generalization of sets allowing for multiple instances for each of its elements.

Graphs A graph G is a pair $(V(G), E(G))$ with *finite* sets of *vertices* or *nodes* $V(G)$ and *edges* $E(G) \subseteq \{\{u, v\} \subseteq V(G) \mid u \neq v\}$. If not otherwise stated, we set $n := |V(G)|$, and the graph is of *order* n . We also call the graph G an n -order graph. For ease of notation, we denote the edge $\{u, v\}$ in $E(G)$ by (u, v) or (v, u) . In the case of *directed graphs*, the set $E(G) \subseteq \{(u, v) \in V(G) \times V(G) \mid u \neq v\}$ and a *directed acyclic graph* (DAG) is a directed graph with no directed cycles. A *(vertex-)labeled graph* G is a triple $(V(G), E(G), \ell)$ with a (vertex-)label function $\ell: V(G) \rightarrow \mathbb{N}$. Then $\ell(v)$ is a *label* of v , for v in $V(G)$. An *attributed graph* G is a triple $(V(G), E(G), a)$ with a graph $(V(G), E(G))$ and (vertex-)attribute function $a: V(G) \rightarrow \mathbb{R}^{1 \times d}$, for some $d > 0$. That is, contrary to labeled graphs, we allow for vertex annotations from an uncountable set. Then $a(v)$ is an *attribute* or *feature* of v , for v in $V(G)$. Equivalently, we define an n -order attributed graph $G := (V(G), E(G), a)$ as a pair $\mathbf{G} = (G, \mathbf{L})$, where $G = (V(G), E(G))$ and \mathbf{L} in $\mathbb{R}^{n \times d}$ is a *vertex feature matrix*. Here, we identify $V(G)$ with $[n]$. For a matrix \mathbf{L} in $\mathbb{R}^{n \times d}$ and v in $[n]$, we denote by \mathbf{L}_v in $\mathbb{R}^{1 \times d}$ the v th row of \mathbf{L} such that $\mathbf{L}_v := a(v)$. We also write \mathbb{R}^d for $\mathbb{R}^{1 \times d}$.

The *neighborhood* of v in $V(G)$ is denoted by $N(v) := \{u \in V(G) \mid (v, u) \in E(G)\}$ and the *degree* of a vertex v is $|N(v)|$. In case of directed graphs, $N^+(u) := \{v \in V(G) \mid (v, u) \in E(G)\}$ and $N^-(u) := \{v \in V(G) \mid (u, v) \in E(G)\}$. The *in-degree* and *out-degree* of a vertex v are $|N^+(v)|$ and $|N^-(v)|$, respectively. Two graphs G and H are *isomorphic* and we write $G \simeq H$ if there exists a bijection $\varphi: V(G) \rightarrow V(H)$ preserving the adjacency relation, i.e., (u, v) is in $E(G)$ if and only if $(\varphi(u), \varphi(v))$ is in $E(H)$. Then φ is an *isomorphism* between G and H . In the case of labeled graphs, we additionally require that $\ell(v) = \ell(\varphi(v))$ for v in $V(G)$, and similarly for attributed graphs.

2.1. The Weisfeiler–Leman algorithm

We here describe the 1-WL and refer to Appendix B for the k -WL. The 1-WL or color refinement is a well-studied heuristic for the graph isomorphism problem, originally proposed by Weisfeiler & Leman (1968).¹ Intuitively, the algorithm determines if two graphs are non-isomorphic by iteratively coloring or labeling vertices. Given an initial coloring or labeling of the vertices of both graphs, e.g., their degree or application-specific information, in each iteration, two vertices with the same label get different labels if the number of identically labeled neighbors is unequal. These labels induce a vertex partition, and the algorithm

¹Strictly speaking, the 1-WL and color refinement are two different algorithms. That is, the 1-WL considers neighbors and non-neighbors to update the coloring, resulting in a slightly higher expressive power when distinguishing vertices in a given graph; see (Grohe, 2021) for details. For brevity, we consider both algorithms to be equivalent.

terminates when after some iteration, the algorithm does not refine the current partition, i.e., when a *stable coloring* or *stable partition* is obtained. Then, if the number of vertices annotated with a specific label is different in both graphs, we can conclude that the two graphs are not isomorphic. It is easy to see that the algorithm cannot distinguish all non-isomorphic graphs (Cai et al., 1992). Nonetheless, it is a powerful heuristic that can successfully test isomorphism for a broad class of graphs (Babai & Kucera, 1979).

Formally, let $G = (V(G), E(G), \ell)$ be a labeled graph. In each iteration, $t > 0$, the 1-WL computes a vertex coloring $C_t^1: V(G) \rightarrow \mathbb{N}$, depending on the coloring of the neighbors. That is, in iteration $t > 0$, we set

$$C_t^1(v) := \text{RELABEL} \left((C_{t-1}^1(v), \{\{C_{t-1}^1(u) \mid u \in N(v)\}\}) \right),$$

for all vertices v in $V(G)$, where RELABEL injectively maps the above pair to a unique natural number, which has not been used in previous iterations. In iteration 0, the coloring $C_0^1 := \ell$. To test if two graphs G and H are non-isomorphic, we run the above algorithm in “parallel” on both graphs. If the two graphs have a different number of vertices colored c in \mathbb{N} at some iteration, the 1-WL *distinguishes* the graphs as non-isomorphic. Moreover, if the number of colors between two iterations, t and $(t+1)$, does not change, i.e., the cardinalities of the images of C_t^1 and C_{t+1}^1 are equal, or, equivalently,

$$C_t^1(v) = C_t^1(w) \iff C_{t+1}^1(v) = C_{t+1}^1(w),$$

for all vertices v and w in $V(G)$, the algorithm terminates. For such t , we define the *stable coloring* $C_\infty^1(v) = C_t^1(v)$, for v in $V(G)$. The stable coloring is reached after at most $\max\{|V(G)|, |V(H)|\}$ iterations (Grohe, 2017). We define the *color complexity* of a graph G as the number of colors computed by the 1-WL after $|V(G)|$ iterations on G .

Due to the shortcomings of the 1-WL or color refinement in distinguishing non-isomorphic graphs, several researchers, e.g., Babai (1979); Cai et al. (1992), devised a more powerful generalization of the former, today known as the k -dimensional Weisfeiler–Leman algorithm, operating on k -tuples of vertices rather than single vertices; see Appendix B for details.

2.2. Graph Neural Networks

Intuitively, GNNs learn a vectorial representation, i.e., a d -dimensional real-valued vector, representing each vertex in a graph by aggregating information from neighboring vertices. Formally, let $G = (V(G), E(G), \ell)$ be a labeled graph with initial vertex features $\mathbf{h}_v^{(0)}$ in \mathbb{R}^d that are *consistent* with ℓ . That is, each vertex v is annotated with a feature $\mathbf{h}_v^{(0)}$ in \mathbb{R}^d such that $\mathbf{h}_v^{(0)} = \mathbf{h}_u^{(0)}$ if and only if $\ell(v) = \ell(u)$, e.g., a one-hot encoding of the labels $\ell(u)$ and $\ell(v)$. Alternatively, $\mathbf{h}_v^{(0)}$

can be an attribute or a feature of the vertex v , e.g., physical measurements in the case of chemical molecules. A GNN architecture consists of a stack of neural network layers, i.e., a composition of permutation-equivariant parameterized functions. Similarly to the 1-WL, each layer aggregates local neighborhood information, i.e., the neighbors’ features around each vertex, and then passes this aggregated information on to the next layer.

Following, Gilmer et al. (2017) and Scarselli et al. (2009), in each layer, $t > 0$, we compute vertex features $\mathbf{h}_v^{(t)} :=$

$$\text{UPD}^{(t)} \left(\mathbf{h}_v^{(t-1)}, \text{AGG}^{(t)} \left(\{\{\mathbf{h}_u^{(t-1)} \mid u \in N(v)\}\} \right) \right) \quad (1)$$

in \mathbb{R}^d , where $\text{UPD}^{(t)}$ and $\text{AGG}^{(t)}$ may be differentiable parameterized functions, e.g., neural networks.² In the case of graph-level tasks, e.g., graph classification, one uses

$$\mathbf{h}_G := \text{READOUT} \left(\{\{\mathbf{h}_v^{(L)} \mid v \in V(G)\}\} \right) \in \mathbb{R}, \quad (2)$$

to compute a single vectorial representation based on learned vertex features after iteration L .³ Again, READOUT may be a differentiable parameterized function. To adapt the parameters of the above three functions, they are optimized end-to-end, usually through a variant of stochastic gradient descent, e.g., Kingma & Ba (2015), together with the parameters of a neural network used for classification or regression. See Appendix C for a definition of (higher-order) k -GNNs.

Notation In the subsequent sections, we use the following notation. We denote the class of all (labeled) graphs by \mathcal{G} , the class of all graphs with d -dimensional, real-valued vertex features by \mathcal{G}_d , the class of all graphs with d -dimensional boolean vertex features by $\mathcal{G}_d^{\mathbb{B}}$, the class of all graphs with an order of at most n and d -dimensional vertex features by $\mathcal{G}_{d,n}$, and the class of all graphs with d -dimensional vertex features and of color complexity at most u by $\mathcal{G}_{d,\leq u}$.

Further, we consider the following classes of GNNs. We denote the class of all GNNs consisting of L layers with $(L+1)$ th layer readout layer by $\text{GNN}(L)$, the subset of $\text{GNN}(L)$ but whose aggregate, update and readout functions have a width at most d by $\text{GNN}(d, L)$, and the subset of $\text{GNN}(L)$ but whose aggregation function is a summation and update and readout functions are single layer perceptrons of width at most d by $\text{GNN}_{\text{slp}}(d, L)$. More generally,

²Strictly speaking, Gilmer et al. (2017) consider a slightly more general setting in which vertex features are computed by $\mathbf{h}_v^{(t)} := \text{UPD}^{(t)} \left(\mathbf{h}_v^{(t-1)}, \text{AGG}^{(t)} \left(\{\{\mathbf{h}_u^{(t-1)}, \mathbf{h}_u^{(t-1)}, \ell(v, u) \mid u \in N(v)\}\} \right) \right)$, where $\ell(v, u)$ denotes the edge label of the edge (v, u) .

³For simplicity, we assume GNNs to return scalars on graphs. This makes the definition of VC dimension more concise.

we consider the class $\text{GNN}_{\text{mlp}}(d, L)$ of GNNs using summation for aggregation and such that update and readout functions are *multilayer perceptrons* (MLPs), all of width of at most d . We refer to elements in $\text{GNN}_{\text{mlp}}(d, L)$ as *simple GNNs*. See Appendix E for details. We stress that simple GNNs are already expressive enough to be equivalent to the 1-WL in distinguishing non-isomorphic graphs.

VC dimension of GNNs For a class \mathcal{C} of GNNs and \mathcal{X} of graphs, $\text{VC-dim}_{\mathcal{X}}(\mathcal{C})$ is the maximal number m of graphs $\mathbf{G}_1, \dots, \mathbf{G}_m$ in \mathcal{X} that can be shattered by \mathcal{C} . Here, $\mathbf{G}_1, \dots, \mathbf{G}_m$ are *shattered* if for any τ in $\{0, 1\}^m$ there exists a GNN gnn in \mathcal{C} such that for all i in $[m]$:

$$\text{gnn}(\mathbf{G}_i) = \begin{cases} \geq 2/3 & \text{if } \tau_i = 1, \text{ and} \\ \leq 1/3 & \text{if } \tau_i = 0. \end{cases} \quad (3)$$

The above definition can straightforwardly be generalized to k -GNNs. Bounding the VC dimension directly implies an upper bound on the generalization error; see Appendix D and (Vapnik, 1995; Mohri et al., 2018) for details.

Bitlength of GNNs Below we study the dependence of GNNs’ VC dimension on the *bitlength* of its weights. Assume an L -layered GNN with a set of parameters Θ , then the GNN’s bitlength is the maximum number of bits needed to encode each weight in Θ and the parameters specifying the activation functions. We define the bitlength of a class of GNNs as the maximum bitlength across all GNNs in the class.

3. WL meet VC

We first consider the non-uniform regime, i.e., we assume an upper bound on the graphs’ order. Given the connection between GNNs and the 1-WL (Morris et al., 2019; Xu et al., 2018), GNNs’ ability to shatter a set of graphs can easily be related to distinguishability by the 1-WL. Indeed, let \mathcal{S} be a collection of graphs that GNNs can shatter. Hence, for each pair of graphs in \mathcal{S} , we have a GNN that distinguishes them. By the results of (Morris et al., 2019; Xu et al., 2019), this implies that the graphs in \mathcal{S} are pairwise 1-WL distinguishable. In other words, when considering the VC dimension of GNNs on a class \mathcal{S} of graphs with a bounded number m of 1-WL distinguishable graphs, then m is also an upper bound on the VC dimension of GNNs on graphs in \mathcal{S} . For example, let us first consider the VC dimension of GNNs on the class $\mathcal{G}_{d,n}^{\mathbb{B}}$ consisting of graphs of an order of at most n with d -dimensional boolean features. Let $m_{n,d,L}$ be the maximal number of graphs in $\mathcal{G}_{d,n}^{\mathbb{B}}$ distinguishable by 1-WL after L iterations. Then, the same argument as above implies that $m_{n,d,L}$ is also the maximal number of graphs in $\mathcal{G}_{d,n}^{\mathbb{B}}$ that can be shattered by L -layer GNNs, as is stated next.

Proposition 3.1. *For all n, d and L , it holds that*

$$\text{VC-dim}_{\mathcal{G}_{d,n}^{\mathbb{B}}}(\text{GNN}(L)) \leq m_{n,d,L}.$$

This upper bound holds regardless of the choice of aggregation, update, and readout functions used in the GNNs. We next show a matching lower bound for the VC dimension of GNNs on graphs in $\mathcal{G}_{d,n}^{\mathbb{B}}$. In fact, the lower bound already holds for simple GNNs of width $\mathcal{O}(nm_{n,d,L})$.

Proposition 3.2. *For all n, d , and L , all $m_{n,d,L}$ 1-WL-distinguishable graphs of order at most n with d -dimensional boolean features can be shattered by sufficiently wide L -layer GNNs. Hence,*

$$\text{VC-dim}_{\mathcal{G}_{d,n}^{\mathbb{B}}}(\text{GNN}(L)) = m_{n,d,L}.$$

The lower bound follows from the fact that GNNs are as powerful as the 1-WL (Morris et al., 2019; Xu et al., 2018). Indeed, Morris et al. (2019) showed that L iterations of the 1-WL on graphs in $\mathcal{G}_{d,n}^{\mathbb{B}}$ can be simulated by a simple L -layered GNN of width $\mathcal{O}(n)$. To shatter all $m_{n,d,L}$ graphs, we first simulate 1-WL on all $m_{n,d,L}$ graphs combined using a simple L -layer GNN of width in $\mathcal{O}(nm_{n,d,L})$. We then define a readout layer whose weights can be used to shatter the input graphs based on the computed 1-WL vertex colors in the graphs. We note that the above two results can be straightforwardly generalized to k -GNNs, i.e., their VC dimension is tightly connected to the expressive power of the k -WL, and other recent extensions of GNNs; see Appendix C.1.

We now consider the uniform regime, i.e., we assume no upper bound on the graphs’ order. Since the number $m_{n,d,L}$ of 1-WL distinguishable graphs increases for growing n , Proposition 3.2 implies that the VC dimension of L -layered GNNs on the class $\mathcal{G}_d^{\mathbb{B}}$ of all graphs with d -dimensional boolean features but of arbitrary order is unbounded.

Corollary 3.3. *For all d and $L \geq 1$, it holds that $\text{VC-dim}_{\mathcal{G}_d^{\mathbb{B}}}(\text{GNN}(L)) = \infty$.*

The proof of this result requires update and readout functions in GNNs of *unbounded width*. Using a different “bit extraction” proof technique, we can strengthen the previous result such that *fixed-width* GNNs can be considered.

Theorem 3.4. *For all d, L at least two, it holds that $\text{VC-dim}_{\mathcal{G}_d^{\mathbb{B}}}(\text{GNN}(d, L)) = \infty$.*

Again, this result holds even for the class of simple GNNs. The theorem relies on the following result, which is of independent interest.

Proposition 3.5. *There exists a family \mathcal{F}_b of simple 2-layer GNNs of width two and bitlength $\mathcal{O}(b)$ using piece-wise linear activation functions such that its VC dimension is exactly b .*

Theorem 3.4 directly follows from Proposition 3.5. That is, to show that the VC dimension is infinite for GNNs in $\text{GNN}(d, L)$ with d and $L \geq 2$, we leverage Proposition 3.5, implying that we can shatter an arbitrary number of graphs by such GNNs, provided that they have bit precision $\mathcal{O}(b)$. The GNNs in Theorem 3.4 have arbitrary precision reals, so they can also shatter these graphs. Since this works for any b , Theorem 3.4 follows.

Proposition 3.5 in turn is proved as follows. For the upper bound, we observe that there are only exponentially many (in b) GNNs in \mathcal{F}_b , from which the upper bound immediately follows. Indeed, classical VC theory implies a bound on the VC dimension for finite classes of GNNs, logarithmic in the number of GNNs in the class. The lower bound proof is more challenging and requires constructing the collection \mathcal{F}_b of simple GNNs that can shatter b graphs belonging to $\mathcal{G}_d^{\mathbb{B}}$. We remark that the b graphs used are of order $\mathcal{O}(2^b)$ and have $\mathcal{O}(2^b)$ 1-WL vertex colors.

Finally, we show bounded VC dimension when both the width and the number of layers of GNNs, and input graphs are restricted. We obtain the bound by leveraging state-of-the-art VC dimension bounds for feedforward neural networks (FNNs) by Bartlett et al. (2019a). To control the number of parameters, we consider the class $\text{GNN}_{\text{slp}}(d, L)$ of simple GNNs in which update and readout functions are single-layer perceptrons of bounded width d . We specify this class of GNNs using $P = d(2dL + L + 1) + 1$ parameters and the choice of activation functions in the perceptrons. Regarding the class of input graphs, we consider the class $\mathcal{G}_{d, \leq u}$ consisting of graphs having d -dimensional features and color complexity at most u . Note that we only bound the number of colors appearing in a single graph and not the number of colors appearing in all graphs of the class. Intuitively, the parameter u comes into play because we can reduce input graphs by combining 1-WL equivalent vertices. The reduced graph has u vertices, and edges have weights. One can run extended GNNs, taking edge weights into account on the reduced graph without loss of information. Moreover, one can tie extended GNNs to FNNs. Hence, the parameter n can be replaced by u when analyzing the VC dimension of the GNN-related FNNs.

We now relate $\mathcal{G}_{d, n}$ and $\mathcal{G}_{d, \leq u}$. Since any graph of order at most n has at most n 1-WL colors, $\mathcal{G}_{d, n} \subseteq \mathcal{G}_{d, \leq n}$. The bound below thus also complements the upper bound on $\mathcal{G}_{d, n}^{\mathbb{B}}$ given earlier, but now for fixed-width GNNs. However, $\mathcal{G}_{d, \leq u}$ may contain graphs of arbitrary order. For example, all regular graphs (of the same degree) belong to $\mathcal{G}_{d, \leq 1}$. We also remark that there is no upper bound on the number of 1-WL distinguishable graphs for $\mathcal{G}_{d, \leq u}$, because we only bound the number of colors appearing in a single graph and not the number of colors appearing in all graphs of the class. For example, all regular graphs of arbitrary degrees

are in $\mathcal{G}_{0, \leq 1}$, but regular graphs of different degrees can be distinguished. As such, the bounds obtained earlier do not apply. Finally, in the bound below, input graphs can have real features.

Theorem 3.6. *Assume d and L in \mathbb{N} , and GNNs in $\text{GNN}_{\text{slp}}(d, L)$ using piece-wise polynomial activation functions with $p > 0$ pieces and degree $\delta \geq 0$. Let $P = d(2dL + L + 1) + 1$ be the number of parameters in the GNNs. For all u in \mathbb{N} , $\text{VC-dim}_{\mathcal{G}_{d, \leq u}}(\text{GNN}_{\text{slp}}(d, L)) \leq$*

$$\begin{cases} \mathcal{O}(P \log(puP)) & \text{if } \delta = 0, \\ \mathcal{O}(LP \log(puP)) & \text{if } \delta = 1, \\ \mathcal{O}(LP \log(puP) + L^2 P \log(\delta)) & \text{if } \delta > 1. \end{cases}$$

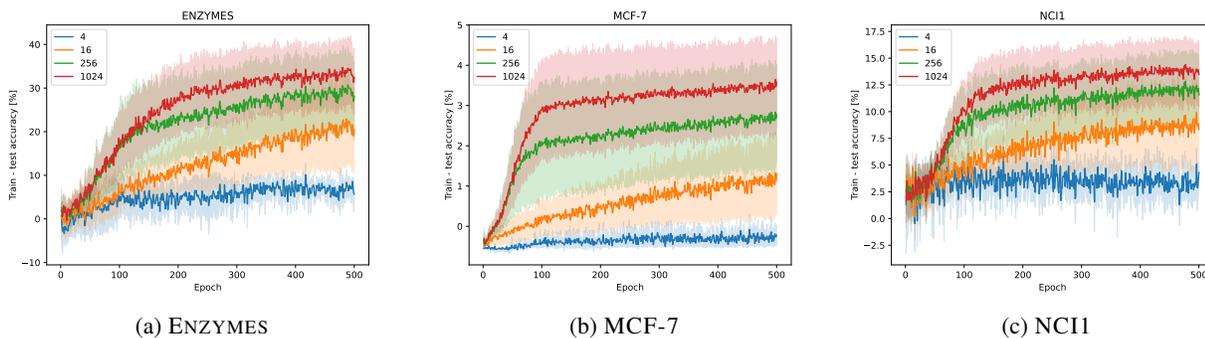
We note that the above result can be straightforwardly generalized to k -GNNs. These upper bounds, concerning the dependency on u , cannot be improved by more than a constant factor. Indeed, for the b graphs used in Proposition 3.5, $u = \mathcal{O}(2^b)$. Moreover, the simple GNNs in \mathcal{F}_b belong to $\text{GNN}_{\text{slp}}(2, 2)$ and use piecewise-linear activation functions with $p = 4$ and $\delta = 1$. Since they shatter $b = \mathcal{O}(\log(u))$ graphs, this matches with the upper bound up to constant factors.

3.1. Implications of your theoretical results

Our results include the first lower bounds of GNNs' VC dimension and the first inherent connection between GNNs' VC dimension and the expressive power of the 1-WL. In particular, we find that a larger bit precision leads to a higher VC dimension and sample complexity. Therefore, our findings provide an additional incentive for reducing the precision of GNNs besides the already-known benefits of reduced training and inference time. Furthermore, we show that the VC dimension bounds for GNNs can be tightened when considering graphs of low color complexity. This connection to the number of colors indicates that graphs with complex structures, captured by the 1-WL, may have a higher VC dimension. Moreover, if the 1-WL can distinguish a large set of graphs of a given dataset, our results imply that a sufficiently expressive GNN will require a large set of training samples to generalize well. Therefore, we can use the 1-WL to assess GNNs' generalization ability on a given dataset quickly. The same relation holds for k -GNNs and the k -WL. Moreover, our results extend easily to recent, more expressive GNN enhancements, e.g., subgraph-based (Bouritsas et al., 2020) or subgraph-enhanced GNNs (Bevilacqua et al., 2022; Qian et al., 2022). Hence, our results lead to a better understanding of how expressivity influences the learning performance of modern GNN architectures; see Appendix C.1.

Table 1: Train and test classification accuracies with different numbers of parameters, using five layers, studying how the number of parameters influences generalization.

Dimension	Split	Dataset					
		ENZYMES	MCF-7	MCF-7H	MUTAGENICITY	NCI1	NCI109
4	Train	31.0 \pm 3.1	92.1 \pm 0.4	92.1 \pm 0.4	79.7 \pm 0.9	76.7 \pm 7.3	66.4 \pm 9.5
	Test	25.3 \pm 5.2	92.4 \pm 0.2	92.3 \pm 0.3	75.8 \pm 0.9	72.4 \pm 7.1	63.6 \pm 9.1
16	Train	76.8 \pm 6.4	96.0 \pm 0.1	96.5 \pm 0.3	92.7 \pm 2.7	88.4 \pm 6.2	86.4 \pm 0.9
	Test	41.7 \pm 9.4	93.2 \pm 0.5	93.1 \pm 0.4	79.8 \pm 2.2	76.1 \pm 2.1	78.6 \pm 1.9
256	Train	98.2 \pm 3.6	99.7 < 0.1	99.9 < 0.1	100.0 \pm 0.0	99.8 < 0.1	97.5 \pm 2.1
	Test	54.7 \pm 2.4	94.0 \pm 0.2	93.6 \pm 0.2	80.7 \pm 1.0	81.8 \pm 1.5	82.1 \pm 1.0
1024	Train	99.8 \pm 0.2	99.8 < 0.1	99.8 \pm 0.1	99.9 \pm 0.2	99.8 \pm 0.1	98.6 \pm 1.0
	Test	54.3 \pm 2.3	93.8 \pm 0.3	93.6 \pm 0.2	81.7 \pm 0.8	80.5 \pm 1.0	82.9 \pm 0.9


 Figure 2: Difference between train and test accuracy for different feature dimensions in $\{4, 16, 256, 1024\}$. See Figure 5 in the appendix for additional results.

4. Limitations, possible road maps, and future work

Although the results are the first ones explicitly drawing a tight connection between expressivity and generalization, there are still many open questions. First, excluding Proposition 3.1, the results investigate specific GNN classes using sum aggregation. Hence, the results should be extended to specific GNN layers commonly used in practice, such as that in Xu et al. (2019), and the effect of different aggregation functions, such as max or mean, should be studied in detail. Moreover, the results only give meaningful results for discretely labeled graphs. Hence, the results should be extended to attributed graphs. Secondly, although the experimental results in Section 5 suggest that our VC dimension bounds hold in practice to some extent, it is well known that they do not explain the generalization behavior of deep neural networks in the over-parameterized regime Bartlett et al. (2019b), trained with variants of stochastic gradient descent. Therefore, it is a future challenge to understand how graph structure influences the generalization properties of over-parameterized GNNs, trained with variants of stochastic gradient descent, and what role the Weisfeiler–Leman algorithm plays in this context.

5. Experimental evaluation

In the following, we investigate how well the VC dimension bounds from the previous section hold in practice. Specifically, we answer the following questions.

- Q1** How does the number of parameters influence GNNs’ generalization performance?
- Q2** How does the number of 1-WL-distinguishable graphs influence GNNs’ generalization performance?
- Q3** How does the bitlength influence a GNN’s ability to fit random data?

The source code of all methods and evaluation procedures is available at https://www.github.com/chrsrrs/wl_vs_vc.

Datasets To investigate questions **Q1** and **Q2**, we used the datasets ENZYMES (Borgwardt et al., 2005; Schomburg et al., 2004), MCF-7 (Yan et al., 2008), MCF-7H (Yan et al., 2008), MUTAGENICITY (Kazius et al., 2005; Riesen & Bunke, 2008), and NCI1 and NCI109 (Wale et al., 2008; Shervashidze et al., 2011) provided by Morris et al. (2020a). See Table 3 for dataset statistics and properties. For ques-

Table 2: Train and test classification accuracies using different numbers of layers and a feature dimension of 64, studying how the number of different color histograms influences generalization.

Layers	Split	Dataset					
		ENZYMES	MCF-7	MCF-7H	MUTAGENICITY	NCI1	NCI109
0	Train	40.7 \pm 0.5	91.7 \pm 0.1	91.8 \pm 0.1	77.2 \pm 0.3	74.5 \pm 0.3	73.1 \pm 0.5
	Test	33.7 \pm 1.6	91.9 $<$ 0.1	91.2 \pm 0.1	75.7 \pm 1.2	67.9 \pm 1.3	71.5 \pm 0.7
	Difference	7.0 \pm 1.9	-0.2 \pm 0.1	1.0 \pm 0.1	1.5 \pm 1.3	6.5 \pm 1.2	1.6 \pm 0.6
	# Histograms	385	11 533	19 625	2 819	2 889	2 929
1	Train	66.7 \pm 3.6	91.8 \pm 0.1	92.1 $<$ 0.1	90.9 \pm 0.1	92.0 \pm 1.5	83.4 \pm 1.5
	Test	52.3 \pm 5.0	91.9 $<$ 0.1	91.4 \pm 0.1	82.0 \pm 1.0	78.6 \pm 1.3	76.1 \pm 1.0
	Difference	14.4 \pm 5.2	$<$ 0.1 \pm 0.1	0.1 \pm 0.1	8.9 \pm 1.3	13.4 \pm 0.9	7.3 \pm 0.7
	# Histograms	595	25 417	26 037	3 624	3 906	3 950
2	Train	93.5 \pm 2.1	92.0 \pm 0.2	91.9 \pm 0.3	96.9 \pm 1.9	98.3 \pm 0.5	91.1 \pm 0.5
	Test	62.7 \pm 7.2	92.1 \pm 0.1	91.0 \pm 0.6	82.5 \pm 1.0	80.5 \pm 1.3	78.1 \pm 1.5
	Difference	39.9 \pm 5.5	-0.1 \pm 0.2	1.0 \pm 0.3	14.4 \pm 1.0	17.8 \pm 1.0	13.0 \pm 1.5
	# Histograms	595	26 872	27 353	4 239	4 027	4 055
3	Train	98.0 \pm 2.5	92.1 \pm 0.3	92.1 \pm 0.2	99.4 \pm 0.9	99.8 \pm 0.1	93.6 \pm 1.2
	Test	58.7 \pm 5.3	92.1 \pm 0.2	91.5 \pm 0.2	82.8 \pm 1.0	83.5 \pm 0.7	77.8 \pm 1.8
	Difference	39.4 \pm 2.8	0.1 \pm 0.2	1.0 \pm 0.1	16.6 \pm 1.0	16.3 \pm 0.7	15.8 \pm 1.4
	# Histograms	595	27 048	27 524	4 317	4 039	4 067
4	Train	99.8 \pm 0.3	92.0 \pm 0.1	92.2 \pm 0.2	99.1 \pm 0.2	99.8 $<$ 0.1	96.9 \pm 1.0
	Test	62.7 \pm 2.5	92.1 \pm 0.1	91.5 \pm 0.2	82.7 \pm 0.8	83.2 \pm 0.4	79.8 \pm 1.2
	Difference	37.1 \pm 2.5	-0.1 \pm 0.1	1.0 \pm 0.1	16.4 \pm 0.7	16.6 \pm 0.4	17.2 \pm 0.8
	# Histograms	595	27 059	OOM	4 317	4 039	4 067
5	Train	98.9 \pm 1.9	92.1 \pm 0.2	92.4 \pm 0.2	99.9 \pm 0.2	99.8 \pm 0.0	97.7 \pm 0.9
	Test	57.0 \pm 3.9	92.3 \pm 0.2	91.6 \pm 0.2	83.0 \pm 0.8	84.1 \pm 1.1	79.6 \pm 0.5
	Difference	41.9 \pm 2.9	-0.2 \pm 0.2	1.0 \pm 0.2	16.9 \pm 0.7	15.7 \pm 1.1	18.1 \pm 0.5
	# Histograms	595	OOM	OOM	4 317	4 039	4 067
6	Train	99.4 \pm 0.8	92.0 \pm 0.2	92.2 \pm 0.2	99.1 \pm 1.9	99.6 \pm 0.6	95.2 \pm 1.9
	Test	54.0 \pm 2.3	92.2 \pm 0.2	91.4 \pm 0.4	83.5 \pm 1.0	83.4 \pm 1.3	79.2 \pm 1.3
	Difference	44.4 \pm 1.9	-0.2 \pm 0.1	1.0 \pm 0.2	15.6 \pm 1.2	16.2 \pm 0.9	16.0 \pm 2.1
	# Histograms	595	OOM	OOM	4 317	4 039	4 067

tion **Q3**, to investigate the influence of bitlength on GNN’s VC dimension, we probed how well GNNs can fit random data. Hence, the experiments on these datasets aim at empirically verifying the VC dimension bounds concerning bitlength. To that, we created a synthetic dataset; see Appendix G.3. Since it is challenging to simulate different bitlengths without specialized hardware, we resorted to simulating an increased bitlength via an increased feature dimension; see Appendix G.4.

All experiments are therefore conducted with standard 32-bit precision. We also experimented with 64-bit precision but observed no clear difference. Furthermore, 16-bit precision proved numerically unstable in this setting.

Neural architectures For the experiments regarding **Q1** and **Q2**, we used the simple GNN layer described in Appendix G.1 using a ReLU activation function, ignoring possible edge labels. To answer question **Q1**, we fixed the number of layers to five and chose the feature dimension d in $\{4, 16, 256, 1024\}$. To answer **Q2**, we set the feature dimension d to 64 and chose the number of layers from $\{0, \dots, 6\}$. We used sum pooling and a two-layer MLP for all experiments for the final classification. To investigate **Q3**, we used the architecture described in Appendix G.2. In essence, we used a 2-layer MLP for the message generation

function in each GNN layer and added batch normalization (Ioffe & Szegedy, 2015) before each non-linearity and fixed the number of layers to 3, and varied the feature dimension d in $\{4, 16, 64, 256\}$.

Experimental protocol and model configuration For the experiments regarding **Q1** and **Q2**, we uniformly and at random choose 90% of a dataset for training and the remaining 10% for testing. We repeated each experiment five times and report mean test accuracies and standard deviations. We optimized the standard cross entropy loss for 500 epochs using the ADAM optimizer (Kingma & Ba, 2015). Moreover, we used a learning rate of 0.001 across all experiments and no learning rate decay or dropout. For **Q3**, we set the learning rate to 10^{-4} and the number of epochs to 100 000, and repeated each experiment 50 times. All architectures were implemented using PYTORCH GEOMETRIC (Fey & Lenssen, 2019) and executed on a workstation with 128GB RAM and an NVIDIA Tesla V100 with 32GB memory.

5.1. Results and discussion

In the following, we answer questions **Q1** to **Q3**.

Q1 See Table 1 and Figure 2 (and Figure 5 in the appendix). Increasing the feature dimension d increases the

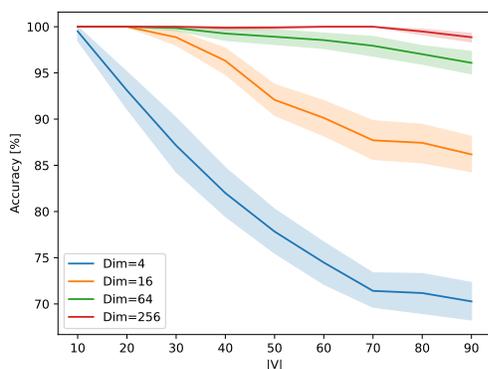


Figure 3: GNN’s ability to fit the synthetic dataset for different feature dimensions in $\{4, 16, 64, 256\}$.

average difference between train and test accuracies across all datasets. For example, on the ENZYMES dataset, the difference increases from around 5% for $d = 4$ to more than 45% for $d = 1024$. However, we also observe that the difference does not increase when reaching near-perfect training accuracies, i.e., going from $d = 256$ to $d = 1024$ does not increase the difference. Hence, the results show that the number of parameters plays a crucial role in GNNs’ generalization ability, in accordance with Theorem 3.6.

Q2 See Table 2. The results indicate that the number of 1-WL-distinguishable graphs ($m_{n,d,L}$) influence GNNs’ generalization properties. For example, on the MUTAGENICITY dataset, after two iterations, the number of unique histograms computed by 1-WL stabilizes, and similarly, the generalization error stabilizes as well. Similar effects can be observed for the ENZYMES, NCI1, and NCI109 datasets. Hence, our results largely confirm Propositions 3.1 and 3.2.

Q3 See Figure 3. Increasing the feature dimension boosts the model’s capacity to fit random class labels, indicating that increased bitlength implies an increased VC dimension. For example, for an order of 70, a GNN using a feature dimension of 4 cannot reach an accuracy of over 75%. In contrast, feature dimensions 64 and 256 can almost fit such data. Moreover, for larger graphs, up to order 90, a GNN with a feature dimension of 256 can almost perfectly fit random class labels, with a feature dimension of 64 only slightly worse, confirming Proposition 3.5.

6. Conclusion

We investigated GNNs’ generalization capabilities through the lens of VC dimension theory in different settings. Specifically, when not assuming a bound on the graphs’ order, we showed that the VC dimension tightly depends on the

bitlength of the GNNs’ weights. We further showed that the number of colors computed by the 1-WL, besides the number of parameters and layers, influences the VC dimension. When a bound on the graphs’ order is known, we upper and lower bounded GNNs’ VC dimension via the maximal number of graphs distinguishable by the 1-WL. Thus, our theory provides the first link between expressivity results and generalization. Further, our theory also applies to a large set of recently proposed GNN enhancements.

Acknowledgements

Christopher Morris is partially funded by a DFG Emmy Noether grant (468502433) and RWTH Junior Principal Investigator Fellowship under Germany’s Excellence Strategy. Martin Grohe is partially funded by the European Union (ERC, SymSim, 101054974). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- Aamand, A., Chen, J. Y., Indyk, P., Narayanan, S., Rubinfeld, R., Schiefer, N., Silwal, S., and Wagner, T. Exponentially improving the complexity of simulating the Weisfeiler-Lehman test with graph neural networks. *ArXiv preprint*, 2022.
- Abboud, R., Ceylan, İ. İ., Grohe, M., and Lukaszewicz, T. The surprising power of graph neural networks with random node initialization. In *Joint Conference on Artificial Intelligence*, pp. 2112–2118, 2021.
- Arvind, V., Köbler, J., Rattan, G., and Verbitsky, O. On the power of color refinement. In *International Symposium on Fundamentals of Computation Theory*, pp. 339–350, 2015.
- Arvind, V., Fuhlbrück, F., Köbler, J., and Verbitsky, O. On Weisfeiler-Leman invariance: Subgraph counts and related graph properties. In *International Symposium on Fundamentals of Computation Theory*, pp. 111–125, 2019.
- Atserias, A. and Maneva, E. N. Sherali-adams relaxations and indistinguishability in counting logics. *SIAM Journal on Computing*, 42(1):112–137, 2013.
- Atserias, A., Mancinska, L., Roberson, D. E., Sámal, R., Severini, S., and Varvitsiotis, A. Quantum and non-signalling graph isomorphisms. *Journal of Combinatorial Theory, Series B*, pp. 289–328, 2019.
- Azizian, W. and Lelarge, M. Characterizing the expressive power of invariant and equivariant graph neural networks.

- In *International Conference on Learning Representations*, 2021.
- Babai, L. Lectures on graph isomorphism. University of Toronto, Department of Computer Science. Mimeographed lecture notes, October 1979, 1979.
- Babai, L. Graph isomorphism in quasipolynomial time. In *Symposium on Theory of Computing*, pp. 684–697, 2016.
- Babai, L. and Kucera, L. Canonical labelling of graphs in linear average time. In *Symposium on Foundations of Computer Science*, pp. 39–46, 1979.
- Balcilar, M., Héroux, P., Gaüzère, B., Vasseur, P., Adam, S., and Honeine, P. Breaking the limits of message passing graph neural networks. In *International Conference on Machine Learning*, pp. 599–608, 2021.
- Barabasi, A.-L. and Oltvai, Z. N. Network biology: Understanding the cell’s functional organization. *Nature Reviews Genetics*, 5(2):101–113, 2004.
- Barceló, P., Kostylev, E. V., Monet, M., Pérez, J., Reutter, J. L., and Silva, J. P. The logical expressiveness of graph neural networks. In *International Conference on Learning Representations*, 2020.
- Barceló, P., Geerts, F., Reutter, J. L., and Ryschkov, M. Graph neural networks with local graph parameters. In *Advances in Neural Information Processing Systems*, pp. 25280–25293, 2021.
- Barceló, P., Galkin, M., Morris, C., and Orth, M. A. R. Weisfeiler and leman go relational. *ArXiv preprint*, 2022.
- Bartlett, P. L., Harvey, N., Liaw, C., and Mehrabian, A. Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20:63:1–63:17, 2019a.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. Benign overfitting in linear regression. *ArXiv preprint*, 2019b.
- Baskin, I. I., Palyulin, V. A., and Zefirov, N. S. A neural device for searching direct correlations between structures and properties of chemical compounds. *Journal of Chemical Information and Computer Sciences*, 37(4): 715–721, 1997.
- Beaini, D., Passaro, S., Létourneau, V., Hamilton, W. L., Corso, G., and Lió, P. Directional graph networks. In *International Conference on Machine Learning*, pp. 748–758, 2021.
- Berkholz, C., Bonsma, P. S., and Grohe, M. Tight lower and upper bounds for the complexity of canonical colour refinement. *Theory of Computing Systems*, 60(4):581–614, 2017.
- Bevilacqua, B., Frasca, F., Lim, D., Srinivasan, B., Cai, C., Balamurugan, G., Bronstein, M. M., and Maron, H. Equivariant subgraph aggregation networks. In *International Conference on Learning Representations*, 2022.
- Bodnar, C., Frasca, F., Otter, N., Wang, Y. G., Lió, P., Montúfar, G., and Bronstein, M. M. Weisfeiler and Lehman go cellular: CW networks. In *Advances in Neural Information Processing Systems*, pp. 2625–2640, 2021a.
- Bodnar, C., Frasca, F., Wang, Y., Otter, N., Montúfar, G. F., Lió, P., and Bronstein, M. M. Weisfeiler and Lehman go topological: Message passing simplicial networks. In *International Conference on Machine Learning*, pp. 1026–1037, 2021b.
- Borgwardt, K. M., Ong, C. S., Schönauer, S., Vishwanathan, S. V. N., Smola, A. J., and Kriegel, H.-P. Protein function prediction via graph kernels. *Bioinformatics*, 21 (Supplement 1):i47–i56, 2005.
- Borgwardt, K. M., Ghisu, M. E., Llinares-López, F., O’Bray, L., and Rieck, B. Graph kernels: State-of-the-art and future challenges. *Foundations and Trends in Machine Learning*, 13(5–6), 2020.
- Bouritsas, G., Frasca, F., Zafeiriou, S., and Bronstein, M. M. Improving graph neural network expressivity via subgraph isomorphism counting. *ArXiv preprint*, 2020.
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. Spectral networks and deep locally connected networks on graphs. In *International Conference on Learning Representation*, 2014.
- Cai, J., Fürer, M., and Immerman, N. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.
- Chami, I., Abu-El-Haija, S., Perozzi, B., Ré, C., and Murphy, K. Machine learning on graphs: A model and comprehensive taxonomy. *ArXiv preprint*, 2020.
- Chen, Z., Villar, S., Chen, L., and Bruna, J. On the equivalence between graph isomorphism testing and function approximation with gnns. In *Advances in Neural Information Processing Systems*, pp. 15868–15876, 2019.
- Chen, Z., Chen, L., Villar, S., and Bruna, J. Can graph neural networks count substructures? In *Advances in Neural Information Processing Systems*, 2020.
- Cotta, L., Morris, C., and Ribeiro, B. Reconstruction for powerful graph representations. In *Advances in Neural Information Processing Systems*, pp. 1713–1726, 2021.

- Dasoulas, G., Santos, L. D., Scaman, K., and Virmaux, A. Coloring graph neural networks for node disambiguation. In *International Joint Conference on Artificial Intelligence*, pp. 2126–2132, 2020.
- Defferrard, M., Bresson, X., and Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pp. 3837–3845, 2016.
- Dell, H., Grohe, M., and Rattan, G. Lovász meets Weisfeiler and Leman. In *International Colloquium on Automata, Languages, and Programming*, pp. 40:1–40:14, 2018.
- Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pp. 2224–2232, 2015.
- Easley, D. and Kleinberg, J. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- El-Yaniv, R. and Pechyony, D. Transductive rademacher complexity and its applications. In *Annual Conference on Learning Theory*, pp. 157–171, 2007.
- Esser, P. M., Vankadara, L. C., and Ghoshdastidar, D. Learning theory can (sometimes) explain generalisation in graph neural networks. In *Advances in Neural Information Processing Systems*, pp. 27043–27056, 2021.
- Feng, J., Chen, Y., Li, F., Sarkar, A., and Zhang, M. How powerful are k-hop message passing graph neural networks. In *Advances in Neural Information Processing Systems*, 2022.
- Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *International Conference on Learning Representations, Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Frasca, F., Bevilacqua, B., Bronstein, M. M., and Maron, H. Understanding and extending subgraph GNNs by rethinking their symmetries. *CoRR*, 2022.
- Fürer, M. On the combinatorial power of the Weisfeiler-Lehman algorithm. In *International Conference on Algorithms and Complexity*, pp. 260–271, 2017.
- Gama, F., Marques, A. G., Leus, G., and Ribeiro, A. Convolutional neural network architectures for signals supported on graphs. *IEEE Transactions on Signal Processing*, 67(4):1034–1049, 2019.
- Garg, V. K., Jegelka, S., and Jaakkola, T. S. Generalization and representational limits of graph neural networks. In *International Conference on Machine Learning*, pp. 3419–3430, 2020.
- Geerts, F. The expressive power of kth-order invariant graph networks. *ArXiv preprint*, 2020.
- Geerts, F. and Reutter, J. L. Expressiveness and approximation properties of graph neural networks. In *International Conference on Learning Representations*, 2022.
- Geerts, F., Mazowiecki, F., and Pérez, G. A. Let’s agree to degree: Comparing graph convolutional networks in the message-passing framework. In *International Conference on Machine Learning*, pp. 3640–3649, 2021.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pp. 1263–1272, 2017.
- Goldreich, O. Introduction to testing graph properties. In *Property Testing*. Springer, 2010.
- Goller, C. and Küchler, A. Learning task-dependent distributed representations by backpropagation through structure. In *International Conference on Neural Networks*, pp. 347–352, 1996.
- Grohe, M. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Cambridge University Press, 2017.
- Grohe, M. Word2vec, Node2vec, Graph2vec, X2vec: Towards a theory of vector embeddings of structured data. In *Symposium on Principles of Database Systems*, pp. 1–16, 2020.
- Grohe, M. The logic of graph neural networks. In *Symposium on Logic in Computer Science*, pp. 1–17, 2021.
- Grohe, M. The descriptive complexity of graph neural networks. *ArXiv preprint*, 2023.
- Grohe, M. and Otto, M. Pebble games and linear equations. *Journal of Symbolic Logic*, 80(3):797–844, 2015.
- Grohe, M., Schweitzer, P., and Wiebking, D. Deep Weisfeiler Leman. *ArXiv preprint*, 2020.
- Hamilton, W. L., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- Horn, M., Brouwer, E. D., Moor, M., Moreau, Y., Rieck, B., and Borgwardt, K. M. Topological graph neural networks. In *International Conference on Learning Representations*, 2022.

- Huang, Y., Peng, X., Ma, J., and Zhang, M. Boosting the cycle counting power of graph neural networks with I^2 -GNNs. *ArXiv preprint*, 2022.
- Immerman, N. and Lander, E. Describing graphs: A first-order approach to graph canonization. In *Complexity Theory Retrospective: In Honor of Juris Hartmanis on the Occasion of His Sixtieth Birthday, July 5, 1988*, pp. 59–81, 1990.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456, 2015.
- Ju, H., Li, D., Sharma, A., and Zhang, H. R. Generalization in graph neural networks: Improved pac-bayesian bounds on graph diffusion. *ArXiv preprint*, 2023.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohli, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., and Hassabis, D. Highly accurate protein structure prediction with AlphaFold. *Nature*, 2021.
- Karpinski, M. and Macintyre, A. Polynomial bounds for VC dimension of sigmoidal and general Pfaffian neural networks. *Journal of Computer and System Sciences*, 54(1):169–176, 1997.
- Kazius, J., McGuire, R., and Bursi, R. Derivation and validation of toxicophores for mutagenicity prediction. *Journal Medicinal Chemistry*, 48(13):312–320, 2005.
- Kiefer, S. *Power and Limits of the Weisfeiler-Leman Algorithm*. PhD thesis, Department of Computer Science, RWTH Aachen University, 2020.
- Kiefer, S. and McKay, B. D. The iteration number of Colour Refinement. In *International Colloquium on Automata, Languages, and Programming*, pp. 73:1–73:19, 2020.
- Kiefer, S. and Schweitzer, P. Upper bounds on the quantifier depth for graph differentiation in first-order logic. In *Symposium on Logic in Computer Science*, pp. 287–296, 2016.
- Kiefer, S., Schweitzer, P., and Selman, E. Graphs identified by logics with counting. In *International Symposium on Mathematical Foundations of Computer Science*, pp. 319–330, 2015.
- Kiefer, S., Ponomarenko, I., and Schweitzer, P. The Weisfeiler-Leman dimension of planar graphs is at most 3. *Journal of the ACM*, 66(6):44:1–44:31, 2019.
- Kim, J., Nguyen, T. D., Min, S., Cho, S., Lee, M., Lee, H., and Hong, S. Pure transformers are powerful graph learners. *ArXiv preprint*, 2022.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- Kireev, D. B. Chemnet: A novel neural network based method for graph/property mapping. *Journal of Chemical Information and Computer Sciences*, 35(2):175–180, 1995.
- Kriege, N. M., Morris, C., Rey, A., and Sohler, C. A property testing framework for the theoretical expressivity of graph kernels. In *International Joint Conference on Artificial Intelligence*, pp. 2348–2354, 2018.
- Kriege, N. M., Johansson, F. D., and Morris, C. A survey on graph kernels. *Applied Network Science*, 5(1):6, 2020.
- Levie, R., Monti, F., Bresson, X., and Bronstein, M. M. Cayleynets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Transactions on Signal Processing*, 67(1):97–109, 2019.
- Li, P., Wang, Y., Wang, H., and Leskovec, J. Distance encoding: Design provably more powerful neural networks for graph representation learning. In *Advances in Neural Information Processing Systems*, 2020.
- Liao, R., Urtasun, R., and Zemel, R. S. A PAC-Bayesian approach to generalization bounds for graph neural networks. In *International Conference on Learning Representations*, 2021.
- Lichter, M., Ponomarenko, I., and Schweitzer, P. Walk refinement, walk logic, and the iteration number of the Weisfeiler-Leman algorithm. In *Symposium on Logic in Computer Science*, pp. 1–13, 2019.
- Maehara, T. and NT, H. A simple proof of the universality of invariant/equivariant graph neural networks. *ArXiv preprint*, 2019.
- Malkin, P. N. Sherali–Adams relaxations of graph isomorphism polytopes. *Discrete Optimization*, pp. 73–97, 2014.
- Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y. Provably powerful graph networks. In *Advances in Neural Information Processing Systems*, pp. 2153–2164, 2019.

- Martinkus, K., Papp, P. A., Schesch, B., and Wattenhofer, R. Agent-based graph neural networks. *ArXiv preprint*, 2022.
- Maskey, S., Lee, Y., Levie, R., and Kutyniok, G. Generalization analysis of message passing neural networks on large random graphs. In *Advances in Neural Information Processing Systems*, 2022.
- Merkwirth, C. and Lengauer, T. Automatic generation of complementary descriptors with molecular graph networks. *Journal of Chemical Information and Modeling*, 45(5):1159–1168, 2005.
- Micheli, A. Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 20(3):498–511, 2009.
- Micheli, A. and Sestito, A. S. A new neural network model for contextual processing of graphs. In *Italian Workshop on Neural Nets Neural Nets and International Workshop on Natural and Artificial Immune Systems*, pp. 10–17, 2005.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. *Foundations of Machine Learning*. MIT Press, 2018.
- Monti, F., Boscaini, D., Masci, J., Rodolà, E., Svoboda, J., and Bronstein, M. M. Geometric deep learning on graphs and manifolds using mixture model cnns. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5425–5434, 2017.
- Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *AAAI Conference on Artificial Intelligence*, pp. 4602–4609, 2019.
- Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., and Neumann, M. TUDataset: A collection of benchmark datasets for learning with graphs. *ArXiv preprint*, 2020a.
- Morris, C., Rattan, G., and Mutzel, P. Weisfeiler and Leman go sparse: Towards higher-order graph embeddings. In *Advances in Neural Information Processing Systems*, 2020b.
- Morris, C., L., Y., Maron, H., Rieck, B., Kriege, N. M., Grohe, M., Fey, M., and Borgwardt, K. Weisfeiler and Leman go machine learning: The story so far. *ArXiv preprint*, 2021.
- Morris, C., Rattan, G., Kiefer, S., and Ravanbakhsh, S. SpqNets: Sparsity-aware permutation-equivariant graph networks. In *International Conference on Machine Learning*, pp. 16017–16042, 2022.
- Müller, L., Galkin, M., Morris, C., and Rampásek, L. Attending to graph transformers. *CoRR*, abs/2302.04181, 2023.
- Murphy, R. L., Srinivasan, B., Rao, V. A., and Ribeiro, B. Relational pooling for graph representations. In *International Conference on Machine Learning*, pp. 4663–4673, 2019.
- Nguyen, H. and Maehara, T. Graph homomorphism convolution. In *International Conference on Machine Learning*, pp. 7306–7316, 2020.
- Papp, P. A. and Wattenhofer, R. A theoretical comparison of graph neural network extensions. In *International Conference on Machine Learning*, pp. 17323–17345, 2022.
- Papp, P. A., K. Martinkus, L. F., and Wattenhofer, R. DropGNN: Random dropouts increase the expressiveness of graph neural networks. In *Advances in Neural Information Processing Systems*, 2021.
- Puny, O., Lim, D., Kiani, B. T., Maron, H., and Lipman, Y. Equivariant polynomials for graph neural networks. *ArXiv preprint*, 2023.
- Qian, C., Rattan, G., Geerts, F., Morris, C., and Niepert, M. Ordered subgraph aggregation networks. In *Advances in Neural Information Processing Systems*, 2022.
- Riesen, K. and Bunke, H. IAM graph database repository for graph based pattern recognition and machine learning. In *Structural, Syntactic, and Statistical Pattern Recognition, Joint IAPR International Workshop*, pp. 287–297, 2008.
- Rosenbluth, E., Tönshoff, J., and Grohe, M. Some might say all you need is sum. *ArXiv preprint*, 2023.
- Sato, R., Yamada, M., and Kashima, H. Random features strengthen graph neural networks. In *SIAM International Conference on Data Mining*, pp. 333–341, 2021.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Scarselli, F., Tsoi, A. C., and Hagenbuchner, M. The Vapnik-Chervonenkis dimension of graph and recursive neural networks. *Neural Networks*, pp. 248–259, 2018.
- Schomburg, I., Chang, A., Ebeling, C., Gremse, M., Heldt, C., Huhn, G., and Schomburg, D. Brenda, the enzyme database: updates and major new developments. *Nucleic acids research*, pp. D431–3, 2004.
- Shervashidze, N., Schweitzer, P., van Leeuwen, E. J., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-Lehman graph kernels. *Journal of Machine Learning Research*, pp. 2539–2561, 2011.

- Simonovsky, M. and Komodakis, N. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 29–38, 2017.
- Sperduti, A. and Starita, A. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–35, 1997.
- Stokes, J., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N., MacNair, C., French, S., Carfrae, L., Bloom-Ackerman, Z., Tran, V., Chiappino-Pepe, A., Badran, A., Andrews, I., Chory, E., Church, G., Brown, E., Jaakkola, T., Barzilay, R., and Collins, J. A deep learning approach to antibiotic discovery. *Cell*, pp. 688–702.e13, 2020.
- Talak, R., Hu, S., Peng, L., and Carlone, L. Neural trees for learning on graphs. *ArXiv preprint*, 2021.
- Thiede, E. H., Zhou, W., and Kondor, R. Autobahn: Automorphism-based graph neural nets. In *Advances in Neural Information Processing Systems*, pp. 29922–29934, 2021.
- Tolstikhin, I. O. and Lopez-Paz, D. Minimax lower bounds for realizable transductive classification. *ArXiv preprint*, 2016.
- Tönshoff, J., Ritzert, M., Wolf, H., and Grohe, M. Graph learning with 1D convolutions on random walks. *ArXiv preprint*, 2021.
- Vapnik, V. N. *The Nature of Statistical Learning Theory*. Springer, 1995.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- Verma, S. and Zhang, Z. Stability and generalization of graph convolutional neural networks. In *International Conference on Knowledge Discovery & Data Mining*, pp. 1539–1548, 2019.
- Vignac, C., Loukas, A., and Frossard, P. Building powerful and equivariant graph neural networks with structural message-passing. In *Advances in Neural Information Processing Systems*, 2020.
- Wale, N., Watson, I. A., and Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008.
- Weisfeiler, B. *On Construction and Identification of Graphs*. Springer, 1976.
- Weisfeiler, B. and Leman, A. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Tekhnicheskaya Informatsia*, 2(9):12–16, 1968. English translation by G. Ryabov is available at https://www.iti.zcu.cz/wl2018/pdf/wl_paper_translation.pdf.
- Wijesinghe, A. and Wang, Q. A new perspective on “how graph neural networks go beyond weisfeiler-lehman?”. In *International Conference on Learning Representations*, 2022.
- Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K., and Jegelka, S. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning*, pp. 5449–5458, 2018.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- Yan, X., Cheng, H., Han, J., and Yu, P. S. Mining significant graph patterns by leap search. In *International Conference on Management of Data*, pp. 433–444, 2008.
- Yehudai, G., Fetaya, E., Meir, E. A., Chechik, G., and Maron, H. From local structures to size generalization in graph neural networks. In *International Conference on Machine Learning*, pp. 11975–11986, 2021.
- You, J., Gomes-Selman, J., Ying, R., and Leskovec, J. Identity-aware graph neural networks. In *AAAI Conference on Artificial Intelligence*, pp. 10737–10745, 2021.
- Zhang, B., Feng, G., Du, Y., He, D., and Wang, L. A complete expressiveness hierarchy for subgraph gnn’s via subgraph weisfeiler-lehman tests. *CoRR*, abs/2302.07090, 2023a.
- Zhang, B., Luo, S., Wang, L., and He, D. Rethinking the expressive power of gnn’s via graph biconnectivity. *ArXiv preprint*, 2023b.
- Zhang, M. and Li, P. Nested graph neural networks. In *Advances in Neural Information Processing Systems*, pp. 15734–15747, 2021.
- Zhao, L., Jin, W., Akoglu, L., and Shah, N. From stars to subgraphs: Uplifting any GNN with local structure awareness. In *International Conference on Learning Representations*, 2022.

A. Extended related work

In the following, we discuss more related work.

Expressive power of k -WL The Weisfeiler–Leman algorithm constitutes one of the earliest and most natural approaches to isomorphism testing (Weisfeiler, 1976; Weisfeiler & Leman, 1968), and the theory community has heavily investigated it over the last few decades (Grohe, 2017). Moreover, the fundamental nature of the k -WL is evident from various connections to other fields such as logic, optimization, counting complexity, and quantum computing. The power and limitations of the k -WL can be neatly characterized in terms of logic and descriptive complexity (Babai, 1979; Immerman & Lander, 1990), Sherali-Adams relaxations of the natural integer linear optimization problem for the graph isomorphism problem (Atserias & Maneva, 2013; Grohe & Otto, 2015; Malkin, 2014), homomorphism counts (Dell et al., 2018), and quantum isomorphism games (Atserias et al., 2019). In their seminal paper, Cai et al. (1992) showed that, for each k , a pair of non-isomorphic graphs of size $\mathcal{O}(k)$ exists not distinguished by the k -WL. Kiefer (2020) gives a thorough survey of more background and related results concerning the expressive power of the k -WL. For $k = 1$, the power of the algorithm has been completely characterized (Arvind et al., 2015; Kiefer et al., 2015). Moreover, upper bounds on the running time (Berkholz et al., 2017) and the number of iterations for $k = 1$ (Kiefer & McKay, 2020) and the non-oblivious $k = 2$ (Kiefer & Schweitzer, 2016; Lichter et al., 2019) have been shown. For k in $\{1, 2\}$, Arvind et al. (2019) studied the abilities of the (non-oblivious) k -WL to detect and count fixed subgraphs, extending the work of Fürer (2017). The former was refined in (Chen et al., 2020). Kiefer et al. (2019) showed that the non-oblivious 3-WL completely captures the structure of planar graphs. The algorithm for logarithmic k plays a prominent role in the recent result of (Babai, 2016) improving the best-known running time for the graph isomorphism problem. Recently, Grohe et al. (2020) introduced the framework of Deep Weisfeiler–Leman algorithms, which allow the design of a more powerful graph isomorphism test than Weisfeiler–Leman type algorithms. Finally, the emerging connections between the Weisfeiler–Leman paradigm and graph learning are described in two recent surveys (Grohe, 2020; Morris et al., 2021).

B. Oblivious k -WL

Intuitively, to surpass the limitations of the 1-WL, the k -WL colors ordered subgraphs instead of a single vertex.⁴ More precisely, given a graph G , the k -WL colors the tuples from $V(G)^k$ for $k \geq 2$ instead of the vertices. By defining a neighborhood between these tuples, we can define a coloring similar to the 1-WL. Formally, let G be a graph, and let $k \geq 2$. In each iteration, $t \geq 0$, the algorithm, similarly to the 1-WL, computes a *coloring* $C_t^k : V(G)^k \rightarrow \mathbb{N}$. In the first iteration, $t = 0$, the tuples \mathbf{v} and \mathbf{w} in $V(G)^k$ get the same color if they have the same atomic type, i.e., $C_0^k(\mathbf{v}) := \text{atp}(\mathbf{v})$. Here, we define the atomic type $\text{atp} : V(G)^k \rightarrow \mathbb{N}$, for $k > 0$, such that $\text{atp}(\mathbf{v}) = \text{atp}(\mathbf{w})$ for \mathbf{v} and \mathbf{w} in $V(G)^k$ if and only if the mapping $\varphi : V(G)^k \rightarrow V(G)^k$ where $v_i \mapsto w_i$ induces a partial isomorphism, i.e., we have $v_i = v_j \iff w_i = w_j$ and $(v_i, v_j) \in E(G) \iff (\varphi(v_i), \varphi(v_j)) \in E(G)$. Then, for each layer, $t > 0$, C_t^k is defined by

$$C_t^k(\mathbf{v}) := \text{RELABEL}(C_{t-1}^k(\mathbf{v}), M_t(\mathbf{v})),$$

with $M_t(\mathbf{v})$ the multiset

$$M_t(\mathbf{v}) := (\{\{C_{t-1}^k(\phi_1(\mathbf{v}, w)) \mid w \in V(G)\}\}, \dots, \{\{C_{t-1}^k(\phi_k(\mathbf{v}, w)) \mid w \in V(G)\}\}),$$

and where

$$\phi_j(\mathbf{v}, w) := (v_1, \dots, v_{j-1}, w, v_{j+1}, \dots, v_k).$$

That is, $\phi_j(\mathbf{v}, w)$ replaces the j -th component of the tuple \mathbf{v} with the vertex w . Hence, two tuples are *adjacent* or *j -neighbors* if they are different in the j th component (or equal, in the case of self-loops). Hence, two tuples \mathbf{v} and \mathbf{w} with the same color in iteration $(t - 1)$ get different colors in iteration t if there exists a j in $[k]$ such that the number of j -neighbors of \mathbf{v} and \mathbf{w} , respectively, colored with a certain color is different.

We run the k -WL algorithm until convergence, i.e., until for t in \mathbb{N}

$$C_t^k(\mathbf{v}) = C_t^k(\mathbf{w}) \iff C_{t+1}^k(\mathbf{v}) = C_{t+1}^k(\mathbf{w}),$$

⁴There exists two definitions of the k -WL, the so-called oblivious k -WL and the folklore or non-oblivious k -WL; see Grohe (2021). There is a subtle difference in how they aggregate neighborhood information. Within the graph learning community, it is customary to abbreviate the oblivious k -WL as k -WL, a convention we follow in this paper.

for all \mathbf{v} and \mathbf{w} in $V(G)^k$, holds. For such t , we define $C_\infty^k(\mathbf{v}) = C_t^k(\mathbf{v})$ for \mathbf{v} in $V(G)^k$. At convergence, we call the partition of $V(G)^k$ induced by C_t^k the *stable partition*. We set $C_\infty^k(v) := C_\infty^k(v, \dots, v)$ and refer to this as the color of the vertex v .

Similarly to the 1-WL, to test whether two graphs G and H are non-isomorphic, we run the k -WL in “parallel” on both graphs. Then, if the two graphs have a different number of vertices colored c , for c in \mathbb{N} , the k -WL *distinguishes* the graphs as non-isomorphic. By increasing k , the algorithm gets more powerful in distinguishing non-isomorphic graphs, i.e., for each $k \geq 2$, there are non-isomorphic graphs distinguished by $(k + 1)$ -WL but not by k -WL (Cai et al., 1992). For a finite set of graphs $S \subset \mathcal{G}$, we run the algorithm in “parallel” over all graphs in the set S .

C. k -order GNNs

By generalizing Equation (1) in Section 2.2, following (Morris et al., 2019; 2020b; 2021), we can derive k -GNNs computing features for all k -tuples $V(G)^k$, for $k > 0$, defined over the set of vertices of an attributed graph $G = (V(G), E(G), a)$ with features from \mathbb{R}^d . Concretely, in each layer, $t > 0$, for each k -tuple $\mathbf{v} = (v_1, \dots, v_k)$ in $V(G)^k$, we compute a feature

$$\mathbf{h}_\mathbf{v}^{(t)} := \text{UPD}^{(t)}\left(\mathbf{h}_\mathbf{v}^{(t-1)}, \text{AGG}^{(t)}\left(\{\{\mathbf{h}_\mathbf{v}^{(t-1)}(\phi_1(\mathbf{v}, w)) \mid w \in V(G)\}\}, \dots, \{\{\mathbf{h}_\mathbf{v}^{(t-1)}(\phi_k(\mathbf{v}, w)) \mid w \in V(G)\}\}\right)\right).$$

Initially, for $t = 0$, we set

$$\mathbf{h}_\mathbf{v}^{(0)} := \text{UPD}([\text{atp}(\mathbf{v}), a(v_1), \dots, a(v_k)]) \in \mathbb{R}^d,$$

i.e., the atomic type and the attributes of a given k -tuple determine the initial feature of a k -tuple’s vertices. In the above, UPD, $\text{UPD}^{(t)}$, and $\text{AGG}^{(t)}$ may be differentiable parameterized functions, e.g., neural networks. In the case of graph-level tasks, e.g., graph classification, one additionally uses

$$\mathbf{h}_G := \text{READOUT}(\{\{\mathbf{h}_\mathbf{v}^{(L)} \mid \mathbf{v} = (v, \dots, v), v \in V(G)\}\}) \in \mathbb{R}^d,$$

to compute a single vectorial graph representation based on the learned k -tuple features after iteration L .

C.1. Transferring VC bounds from GNNs to k -order GNNs and other more expressive architectures

In the following, we briefly sketch how Propositions 3.1 and 3.2, Corollary 3.3, and Theorem 3.6 can be lifted to k -GNNs. First, observe that we can simulate the computation of a k -GNN via a GNN on a sufficiently defined auxiliary graph. That is, the auxiliary graph contains a vertex for each k -tuple, and an edge connects two k -tuples j if they are j -neighbors for j in $[k]$; see Morris et al. (2021) for details. Using a 1-WL equivalent GNN taking edge labels into account, we can extend Propositions 3.1 and 3.2 and Corollary 3.3 to k -GNNs. Similar reasoning applies to Theorem 3.6, i.e., we can apply the proof technique from Appendix F.3 to this auxiliary graph.

C.1.1. ARCHITECTURES BASED ON SUBGRAPH INFORMATION

Further, we note that Propositions 3.1 and 3.2 also easily extend to recent GNN enhancements, e.g., subgraph-based (Bouritsas et al., 2020) or subgraph-enhanced GNNs (Bevilacqua et al., 2022; Qian et al., 2022). Since suitably defined variations of the 1-WL, incorporating subgraph information at initialization, upper bound the architectures’ expressive power, we can easily apply the reasoning behind the proofs of Propositions 3.1 and 3.2 to these cases. Hence, the architectures’ VC dimensions are also tightly related to the number of graphs distinguishable by respective 1-WL variants.

D. Relationship between VC dimension and generalization error

If we can bound the VC dimension of a hypothesis class \mathcal{C} of GNNs, we directly get insights into its generalization ability, i.e., the difference of the empirical error $R_S(h)$ and the true error $R_{\mathcal{D}}(h)$ for $h \in \mathcal{C}$ and a data generating distribution \mathcal{D} .

Theorem D.1. *Let \mathcal{C} be a class of GNNs, with finite VC dimension $\text{VC-dim}(\mathcal{C}) = d$. Then for \mathcal{C} , for all $\varepsilon > 0$ and $\delta \in (0, 1)$, using*

$$m = \mathcal{O}\left(\frac{1}{\varepsilon^2} \left(d \ln\left(\frac{d}{\varepsilon}\right) + \ln\left(\frac{1}{\delta} + 1\right)\right)\right)$$

samples, for all data generating distributions \mathcal{D} , we have

$$\Pr_{S \sim \mathcal{D}^m} (\forall h \in \mathcal{C} : |R_S(h) - R_{\mathcal{D}}(h)| \leq \varepsilon) \geq 1 - \delta.$$

This result was first proven by Vladimir Vapnik and Alexey Chervonenkis in 1960's; see, e.g., [Mohri et al. \(2018\)](#) for a proof.

E. Simple GNNs

We here provide more detail on the simple GNNs mentioned in Section 2.2. That is, for given d and L in \mathbb{N} , we define the class $\text{GNN}_{\text{mlp}}(d, L)$ of simple GNNs as L -layer GNNs for which, according to Equation (1), for each t in $[L]$, the aggregation function $\text{AGG}^{(t)}$ is simply summation and the update function $\text{UPD}^{(t)}$ is a multilayer perceptron $\text{mlp}^{(t)} : \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ of width at most d . Similarly, the readout function in Equation (2) consists of a multilayer perceptron $\text{mlp} : \mathbb{R}^d \rightarrow \mathbb{R}$ applied on the sum of all vertex features computed in layer L .⁵ More specifically, GNNs in $\text{GNN}_{\text{mlp}}(d, L)$ compute on a graph (G, \mathbf{L}) in \mathcal{G}_d , for each $v \in V(G)$,

$$\mathbf{h}_v^{(t)} := \text{mlp}^{(t)}\left(\mathbf{h}_v^{(t-1)}, \sum_{u \in N(v)} \mathbf{h}_u^{(t-1)}\right) \in \mathbb{R}^d, \quad (4)$$

for t in $[L]$ and $\mathbf{h}_v^{(0)} := \mathbf{L}_v$, and

$$\mathbf{h}_G := \text{mlp}\left(\sum_{v \in V(G)} \mathbf{h}_v^{(L)}\right) \in \mathbb{R}. \quad (5)$$

We also consider an even simpler class $\text{GNN}_{\text{slp}}(d, L)$ of $\text{GNN}_{\text{mlp}}(d, L)$ in which the multilayer perceptrons are in fact single layer perceptrons. That is, Equation (4) is replaced by

$$\mathbf{h}_v^{(t)} := \sigma_t\left(\mathbf{h}_v^{(t-1)} \mathbf{W}_1^{(t)} + \sum_{u \in N(v)} \mathbf{h}_u^{(t-1)} \mathbf{W}_2^{(t)} + \mathbf{b}^{(t)}\right) \in \mathbb{R}^d, \quad (6)$$

where $\mathbf{W}_1^{(t)}$ in $\mathbb{R}^{d \times d}$ and $\mathbf{W}_2^{(t)}$ in $\mathbb{R}^{d \times d}$ are weight matrices, and $\mathbf{b}^{(t)}$ in $\mathbb{R}^{1 \times d}$ is a bias vector, and $\sigma_t : \mathbb{R} \rightarrow \mathbb{R}$ is an activation function, for t in $[L]$. Similarly, Equation (5) is replaced by

$$\mathbf{h}_G := \sigma_{L+1}\left(\sum_{v \in V(G)} \mathbf{h}_v^{(L)} \mathbf{w} + b\right) \in \mathbb{R}. \quad (7)$$

with \mathbf{w} in $\mathbb{R}^{d \times 1}$ a weight vector and b in \mathbb{R} a bias value of the final readout layer. Also, $\sigma_{L+1} : \mathbb{R} \rightarrow \mathbb{R}$ is an activation function. We can thus represent elements in $\text{GNN}_{\text{slp}}(d, L)$ more succinctly by the following tuple of parameters,

$$\Theta = \left(\mathbf{W}_1^{(1)}, \mathbf{W}_2^{(1)}, \mathbf{b}^{(1)}, \dots, \mathbf{W}_1^{(L)}, \mathbf{W}_2^{(L)}, \mathbf{b}^{(L)}, \mathbf{w}, b\right),$$

together with the tuple of activation functions $\sigma = (\sigma_1, \dots, \sigma_L, \sigma_{L+1})$. We can equivalently view Θ as an element in $\mathbb{R}^{d(2dL+L+1)+1}$. Each Θ in $\mathbb{R}^{d(2dL+L+1)+1}$ and $\sigma = (\sigma_1, \dots, \sigma_{L+1})$ induces a permutation-invariant *graph function*

$$\text{gnn}_{\Theta, \sigma} : \mathcal{G}_d \rightarrow \mathbb{R} : (G, \mathbf{L}) \mapsto \text{gnn}_{\Theta, \sigma}(G, \mathbf{L}) := \mathbf{h}_G,$$

with \mathbf{h}_G as defined in Equation (7).

F. Missing proofs

In the following, we outline missing proofs from the main paper.

⁵For simplicity we assume that all feature dimensions of the layers are fixed to d in \mathbb{N} and also assume that the readout layer returns a scalar.

F.1. Proofs of Proposition 3.1 and Proposition 3.2

We start with the general upper bound on the VC dimension in terms of the number of 1-WL-indistinguishable graphs.

Proposition F.1 (Proposition 3.1 in the main text). *For all n, d , and L , the maximal number of graphs of order at most n with d -dimensional boolean features that can be shattered by L -layer GNNs is bounded by the maximal number $(m_{n,d,L})$ of 1-WL-distinguishable graphs. That is,*

$$\text{VC-dim}_{\mathcal{G}_{d,n}^{\mathbb{B}}}(\text{GNN}(L)) \leq m_{n,d,L}.$$

Proof. Clearly, every set \mathcal{S} of $m_{n,d,L} + 1$ graphs from $\mathcal{G}_{d,n}^{\mathbb{B}}$ contains at least two graphs \mathbf{G} and \mathbf{G}' not distinguishable by the 1-WL. Since GNNs cannot distinguish 1-WL-indistinguishable graphs (Morris et al., 2019; Xu et al., 2019), they cannot tell \mathbf{G} and \mathbf{G}' apart and hence cannot shatter \mathcal{S} . Hence, the VC dimension can be at most $m_{n,d,L}$. \square

We next show a corresponding lower bound. In fact, the lower bound already holds for the class of simple GNNs of arbitrary width, that is for GNNs in $\text{GNN}_{\text{mlp}}(L) := \bigcup_{d \in \mathbb{N}} \text{GNN}_{\text{mlp}}(d, L)$.

Proposition F.2 (Proposition 3.2 in the main paper). *For all n, d , and L , all $m_{n,d,L}$ 1-WL-distinguishable graphs of order at most n with d -dimensional boolean features can be shattered by sufficiently wide L -layer GNNs. Hence,*

$$\text{VC-dim}_{\mathcal{G}_{d,n}^{\mathbb{B}}}(\text{GNN}(L)) = m_{n,d,L}.$$

Proof. For all i in $[m_{n,d,L}]$, choose \mathbf{G}_i in $\mathcal{G}_{d,n}^{\mathbb{B}}$ such that $\mathcal{S} = \{\mathbf{G}_1, \dots, \mathbf{G}_{m_{n,d,L}}\}$ consists of the maximum number of graphs in $\mathcal{G}_{d,n}^{\mathbb{B}}$ pairwise distinguishable by the 1-WL after L iterations.

We next show that the class of simple GNNs which are wide enough, that is, $\text{GNN}_{\text{mlp}}(d', L)$ for large enough d' , is sufficiently rich to shatter \mathcal{S} . That is, we show that for each $\mathcal{T} \subseteq \mathcal{S}$ there is a $\text{gnn}_{\mathcal{T}}$ in $\text{GNN}_{\text{mlp}}(d', L)$ such that for all i in $[m_{n,d,L}]$:

$$\text{gnn}_{\mathcal{T}}(\mathbf{G}_i) = \begin{cases} 1 & \text{if } \mathbf{G}_i \in \mathcal{T}, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

This shows that \mathcal{S} is shattered by $\text{GNN}_{\text{mlp}}(d', L)$ and hence its VC dimension is at least $|\mathcal{S}| = m_{n,d,L}$, as desired.

Overview of the construction Intuitively, we will show that $\text{GNN}_{\text{mlp}}(d', L)$, with d' large enough, is powerful enough to return a one-hot encoding of the color histograms of graphs in \mathcal{S} . That is, there is a simple GNN gnn in $\text{GNN}_{\text{mlp}}(d', L)$ which in the MLP in its readout layer embeds a graph \mathbf{G}_i in \mathcal{S} as a vector $\mathbf{h}_{\mathbf{G}}$ in $\{0, 1\}^{m_{n,d,L}}$ satisfying $(h_{\mathbf{G}})_i = 1$ if and only if \mathbf{G}_i in \mathcal{T} and $i \in [m_{n,d,L}]$. Then, we extend the readout multilayer perceptron of gnn by one more layer such that on input \mathbf{G} the revised GNN evaluates to the scalar

$$g_{\mathbf{G}} := \text{sign}(\mathbf{h}_{\mathbf{G}} \cdot \mathbf{w}^T - 1) \in \{0, 1\},$$

with \mathbf{w} in $\mathbb{R}^{d' \times 1}$. We observe that given $\mathcal{T} \subseteq \mathcal{S}$ it suffices to let the parameter vector \mathbf{w} be the indicator vector for \mathcal{T} . Indeed, this ensures that $g_{\mathbf{G}} = 1$ if and only if \mathbf{G} is in a color class included in \mathcal{T} . We can explore all such subsets \mathcal{T} of \mathcal{S} by varying \mathbf{w} ; hence, this GNN will shatter \mathcal{S} .

Encoding 1-WL colors via GNNs We proceed with the construction of the required GNN. For simplicity of exposition, in the description below we will construct GNN layers of non-uniform width. One can easily obtain uniform width by padding each layer. First, by Morris et al. (2019, Theorem 2), there exists a GNN architecture with feature dimension (at most) n and consisting of L layers such that for each \mathbf{G}_i in \mathcal{S} it computes 1-WL-equivalent vertex features \mathbf{f}_v in $\mathbb{R}^{1 \times d}$ for $v \in V(G_i)$. That is, for vertices v and w in $V(G_i)$ it holds that

$$\mathbf{f}_v = \mathbf{f}_w \iff C_L^1(v) = C_L^1(w).$$

We note here that we can construct a single GNN architecture for all graphs by applying (Morris et al., 2019, Theorem 2) over the disjoint union the graphs in \mathcal{S} . This increases the width from n to $nm_{n,d,L}$.

Encoding 1-WL histograms via GNNs Moreover, again by (Morris et al., 2020b, Theorem 2) there exists \mathbf{W} in $\mathbb{R}^{nm_{n,d,L} \times nm_{n,d,L}}$ and \mathbf{b} in $\mathbb{R}^{nm_{n,d,L}}$ such that

$$\sigma\left(\sum_{v \in V(G)} \mathbf{f}_v \mathbf{W} + \mathbf{b}\right) = \sigma\left(\sum_{v \in V(H)} \mathbf{f}_v \mathbf{W} + \mathbf{b}\right) \iff h_{\mathbf{G}} = h_{\mathbf{H}},$$

for graphs \mathbf{G} and \mathbf{H} in \mathcal{S} . We use ReLU as activation function σ here, just as in (Morris et al., 2019). Other activation functions could be used as well (Grohe, 2021). Hence, for each graph in \mathcal{S} , we have a vector in $\mathbb{R}^{1 \times nm_{n,d,L}}$ uniquely encoding it. Since the number of vertices n is fixed, there exists a number M in \mathbb{N} such that $M\sigma(\sum_{v \in V(G)} \mathbf{W}\mathbf{f}_v)$ is in $\mathbb{N}^{1 \times nm_{n,d,L}}$ for all \mathbf{G} in \mathcal{S} . Moreover, observe that there exists a matrix \mathbf{W}' in $\mathbb{N}^{nm_{n,d,L} \times 2m_{n,d,L}}$ such that

$$M\sigma\left(\sum_{v \in V(G)} \mathbf{f}_v \mathbf{W} + \mathbf{b}\right) \mathbf{W}' = M\sigma\left(\sum_{v \in V(H)} \mathbf{f}_v \mathbf{W} + \mathbf{b}\right) \mathbf{W}' \iff h_{\mathbf{G}} = h_{\mathbf{H}},$$

for graphs \mathbf{G} and \mathbf{H} in \mathcal{S} . For example, we can set

$$\mathbf{W}' = \begin{bmatrix} K^{nm_{n,d,L}-1} & \dots & K^{nm_{n,d,L}-1} \\ \vdots & \dots & \vdots \\ K^0 & \dots & K^0 \end{bmatrix} \in \mathbb{N}^{nm_{n,d,L} \times 2m_{n,d,L}}$$

for sufficiently large $K > 1$. Hence, the above GNN architecture computes a vector $\mathbf{k}_{\mathbf{G}}$ in $\mathbb{N}^{2m_{n,d,L}}$ containing $2m_{n,d,L}$ occurrences of a natural number uniquely encoding each color histogram for each graph \mathbf{G} in \mathcal{S} .

We next turn $\mathbf{k}_{\mathbf{G}}$ into our desired $\mathbf{h}_{\mathbf{G}}$ as follows. We first define an intermediate vector $\mathbf{h}'_{\mathbf{G}}$ whose entries will be used to check which color histogram is returned. More specifically, we define

$$\mathbf{h}'_{\mathbf{G}} = \text{lsig}(\mathbf{k}_{\mathbf{G}} \cdot (\mathbf{w}'')^T + \mathbf{b}),$$

with $\mathbf{w}'' = (1, -1, 1, -1, \dots, 1, -1) \in \mathbb{R}^{2m_{n,d,L}}$ and $\mathbf{b} = (-c_1 - 1, c_1 + 1, -c_2 - 1, c_2 + 1, \dots, -c_{m_{n,d,L}} - 1, c_{m_{n,d,L}} + 1) \in \mathbb{R}^{2m_{n,d,L}}$ with c_i the number encoding the i th color histogram. We note that for odd i ,

$$(h'_{\mathbf{G}})_i := \text{lsig}(\text{col}(G) - c_i - 1) = \begin{cases} 1 & \text{col}(G) \geq c_i \\ 0 & \text{otherwise.} \end{cases}$$

and for even i ,

$$(h'_{\mathbf{G}})_i := \text{lsig}(-\text{col}(G) + c_i + 1) = \begin{cases} 1 & \text{col}(G) \leq c_i \\ 0 & \text{otherwise.} \end{cases}$$

In other words, $((h'_{\mathbf{G}})_i, (h'_{\mathbf{G}})_{i+1})$ are both 1 if and only if $\text{col}(G) = c_i$. We thus obtain $\mathbf{h}_{\mathbf{G}}$ by combining $((h'_{\mathbf{G}})_i, (h'_{\mathbf{G}})_{i+1})$ using an ‘‘AND’’ encoding (e.g., $\text{lsig}(x + y - 1)$) applied to pairs of consecutive entries in $\mathbf{h}'_{\mathbf{G}}$. That is,

$$\mathbf{h}_{\mathbf{G}} := \text{lsig} \left(\mathbf{h}'_{\mathbf{G}} \cdot \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ 0 & 0 & \dots & 1 \end{pmatrix} - (1, 1, \dots, 1) \right) \in \mathbb{R}^{m_{d,n,L}}$$

We thus see that a 3-layer MLP suffices for the readout layer of the simple GNN, finishing the proof. We remark that the maximal width is $2nm_{n,d,L}$, so we can take $d' = 2nm_{n,d,L}$. \square

F.2. Proof of Proposition 3.5

We now prove Proposition 3.5.

Proposition F.3 (Proposition 3.5 in the main text). *There exists a family \mathcal{F}_b of simple 2-layer GNNs of width two and of bitlength $\mathcal{O}(b)$ using a piece-wise linear activation such that its VC dimension is exactly b .*

Proof. We first show the lower bound. We fix some $n \geq 1$. We shall construct a family of GNNs whose weights have bitlength $\mathcal{O}(n)$ and a family of n graphs shattered by these GNNs. Thereto, for all $\mathbf{x} = (x_1, \dots, x_n)$ in $\{0, 1\}^n$, we let

$$\rho(\mathbf{x}) := \sum_{i=1}^n (2^{-2i+1} + x_i 2^{-2i}).$$

Written in binary, we have

$$\rho(\mathbf{x}) = 0.1x_11x_21x_3\dots 1x_n.$$

Observe that

$$\frac{1}{2} \leq \rho(\mathbf{x}) \leq 1. \quad (8)$$

For $1 \leq k \leq n$, we let

$$\rho_k(\mathbf{x}) := \rho((x_{k+1}, \dots, x_{k+n})) = \sum_{i=1}^n (2^{-2i+1} + x_{k+i}2^{-2i}),$$

where $x_{k+i} := 0$ for $k+i > n$. Then it follows from (8) that

$$\frac{1}{2} \leq \rho_k(\mathbf{x}) \leq 1. \quad (9)$$

We claim that

$$\rho_k(\mathbf{x}) = 2^{2k} \rho(\mathbf{x}) - \underbrace{\left(\sum_{i=1}^k 2^{2(k-i)+1} - \sum_{i=1}^k 2^{2(k-n-i)+1} \right)}_{=: a_k} - \underbrace{\sum_{i=1}^{k-1} 2^{2(k-i)} x_i}_{=: b_k(\mathbf{x})} - x_k \quad (10)$$

Indeed, we have

$$\begin{aligned} 2^{2k} \rho(\mathbf{x}) &= \sum_{i=1}^n (2^{2(k-i)+1} + x_i 2^{2k-2i}) \\ &= \sum_{i=1}^{n+k} (2^{2(k-i)+1} + x_i 2^{2(k-i)}) - \sum_{i=n+1}^{n+k} 2^{2(k-i)+1} \\ &= \sum_{i=1}^k 2^{2(k-i)+1} - \sum_{i=n+1}^{n+k} 2^{2(k-i)+1} + \sum_{i=1}^k x_i 2^{2(k-i)} + \sum_{i=k+1}^{n+k} (2^{2(k-i)+1} + x_i 2^{2(k-i)}) \\ &= \sum_{i=1}^k 2^{2(k-i)+1} - \sum_{i=1}^k 2^{2(k-n-i)+1} + \sum_{i=1}^k x_i 2^{2(k-i)} + \sum_{i=1}^n (2^{-2i+1} + x_{k+i} 2^{-2i}) \\ &= \sum_{i=1}^k 2^{2(k-i)+1} - \sum_{i=1}^k 2^{2(k-n-i)+1} + \sum_{i=1}^{k-1} x_i 2^{2(k-i)} + x_k + \rho_k(\mathbf{x}) \\ &= a_k + b_k(\mathbf{x}) + x_k + \rho_k(\mathbf{x}), \end{aligned}$$

which proves (10). Now let

$$c_k(\mathbf{x}) := b_k(\mathbf{x}) + a_k + 1.$$

Then by (9) and (10), we have

$$x_k - \frac{1}{2} \leq 4^k \rho(\mathbf{x}) - c_k(\mathbf{x}) \leq x_k. \quad (11)$$

For $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ in $\{0, 1\}^n$, we write $\mathbf{x} \neq_k \mathbf{y}$ if $x_i \neq y_i$ for some $i < k$. Observe that $\mathbf{x} \neq_k \mathbf{y}$ implies $|b_k(\mathbf{x}) - b_k(\mathbf{y})| \geq 4$ and thus

$$|c_k(\mathbf{x}) - c_k(\mathbf{y})| \geq 4. \quad (12)$$

Let $A : \mathbb{R} \rightarrow \mathbb{R}$ be the continuous piecewise-linear function defined by

$$A(x) := \begin{cases} 0 & \text{if } x < 0, \\ 2x & \text{if } 0 \leq x < \frac{1}{2}, \\ 1 & \text{if } \frac{1}{2} \leq x < 1, \\ 3 - 2x & \text{if } 1 \leq x < \frac{3}{2}, \\ 0 & \text{if } \frac{3}{2} \leq x. \end{cases}$$

Since $x_k \in \{0, 1\}$, by (11) we have

$$x_k = A(4^k \rho(\mathbf{x}) - c_k(\mathbf{x})). \quad (13)$$

It follows from (12) that for \mathbf{y} with $\mathbf{y} \neq_k \mathbf{x}$ we have

$$A(4^k \rho(\mathbf{x}) - c_k(\mathbf{y})) = 0. \quad (14)$$

Let

$$\mathcal{C}_k := \{c_k(\mathbf{y}) \mid \mathbf{y} \in \{0, 1\}^n\}.$$

Then

$$x_k = \sum_{c \in \mathcal{C}_k} A(4^k \rho(\mathbf{x}) - c). \quad (15)$$

Note that the only dependence on \mathbf{x} of the right-hand side of (15) is in $\rho(\mathbf{x})$, because \mathcal{C}_k does not depend on \mathbf{x} .

Observe that $|\mathcal{C}_k| = 2^{k-1}$, because $c_k(\mathbf{y})$ only depends on $y_1, \dots, y_{k-1} \in \{0, 1\}$ and is distinct for distinct values of the y_i . We have

$$a_k = \sum_{i=1}^k 2^{2(k-i)+1} - \underbrace{\sum_{i=1}^k 2^{2(k-n-i)+1}}_{=: s \leq 1} = 2 \sum_{i=0}^{k-1} 4^i - s = \frac{2}{3}(4^k - 1) - s.$$

Thus

$$\frac{2}{3}(4^k - 1) - 1 \leq a_k \leq \frac{2}{3}(4^k - 1).$$

Furthermore,

$$0 \leq b_k(\mathbf{x}) \leq \sum_{i=1}^{k-1} 2^{2(k-i)} = \sum_{i=1}^{k-1} 4^{k-i} = 4 \sum_{i=0}^{k-2} 4^i = \frac{4}{3}(4^{k-1} - 1).$$

Thus

$$\frac{2}{3}(4^k - 1) \leq c \leq \frac{2}{3}(4^k - 1) + \frac{4}{3}(4^{k-1} - 1) + 1 = 4^k - 1. \quad (16)$$

Now for each ρ in \mathbb{R} we construct a 2-layer GNN \mathfrak{G}_ρ as follows:

- Initially, all nodes v carry the 1-dimensional feature $\mathbf{h}_v^{(0)} := 1$.
- The first layer computes the 2-dimensional feature $\begin{pmatrix} h_{v,1}^{(1)} \\ h_{v,2}^{(1)} \end{pmatrix}$ defined by

$$h_{v,1}^{(1)} := \sum_{w \in N(v)} \rho \cdot \mathbf{h}_w^{(0)} - \rho,$$

$$h_{v,2}^{(1)} := \sum_{w \in N(v)} \mathbf{h}_w^{(0)} - 1.$$

- The second layer computes the 1-dimensional feature $\mathbf{h}_v^{(2)}$ defined by

$$\mathbf{h}_v^{(2)} = A \left(h_{v,1}^{(1)} - \sum_{w \in N(v)} h_{w,2}^{(1)} \right).$$

- The readout functions just takes the sum of all the $\mathbf{h}_v^{(2)}$.

We define a graph F_k as follows. The graph F_k is a forest of height 2.

- F_k has a root node r_c for every $c \in \mathcal{C}_k$.

- Each r_c has a child s_c and 4^k additional children $t_{c,1}, \dots, t_{c,4^k}$.
- The $t_{c,i}$ are leaves.
- Each s_c has children $u_{c,1}, \dots, u_{c,c}$.
- The $u_{c,i}$ are leaves.

Now we run the GNN \mathfrak{G}_ρ on F_k with $\rho = \rho(\mathbf{x})$ for some $\mathbf{x} = (x_1, \dots, x_n)$ in $\{0, 1\}^n$.

- We have $\mathbf{h}_v^{(0)} = 1$ for all v in $V(F_k)$.
- We have

$$\begin{aligned}\mathbf{h}_{t_{c,i},1}^{(1)} &= \mathbf{h}_{u_{c,i},1}^{(1)}, \\ \mathbf{h}_{s_c,1}^{(1)} &= c\rho, \\ \mathbf{h}_{r_c,1}^{(1)} &= 4^k \rho,\end{aligned}$$

and

$$\begin{aligned}\mathbf{h}_{t_{c,i},2}^{(1)} &= \mathbf{h}_{u_{c,i},1}^{(1)} = 0, \\ \mathbf{h}_{s_c,2}^{(1)} &= c, \\ \mathbf{h}_{r_c,2}^{(1)} &= 4^k.\end{aligned}$$

- We have

$$\begin{aligned}\mathbf{h}_{t_{c,i}}^{(2)} &= A(-4^k) = 0 \\ \mathbf{h}_{u_{c,i}}^{(2)} &= A(-c) = 0 \\ \mathbf{h}_{s_c}^{(2)} &= A(c\rho - 4^k) = 0 \\ \mathbf{h}_{r_c}^{(2)} &= A(4^k \rho - c) = \begin{cases} x_k & \text{if } c = c_k(\mathbf{x}) \\ 0 & \text{otherwise} \end{cases} \quad \text{by (13) and (14)}.\end{aligned}$$

To see that the first three equalities hold, recall that $A(x) \neq 0$ only if $0 < x < \frac{3}{2}$. Thus $A(-4^k) = 0$. Moreover, by (16) we have $2 \leq c$ and thus $A(c) = 0$. Finally, $A(c\rho - 4^k) = 0$ because $\rho < 1$ and $c \leq 4^k - 1$ by (16) and therefore $c\rho - 4^k < 0$.

- As there is exactly one node r_c with $c = c_k(\mathbf{x})$, the readout is $\sum_{v \in V(F_k)} \mathbf{h}_v^{(2)} = x_k$.

Hence

$$\mathfrak{G}_{\rho(\mathbf{x})}(F_k) = x_i$$

Thus the GNNs $\mathfrak{G}_{\rho(\mathbf{x})}$ for $\mathbf{x} \in \{0, 1\}^n$ shatter the set $\{F_1, \dots, F_n\}$. Since the bitlength is upper bounded by $\mathcal{O}(b)$ and the number of parameters in the above construction is constant, the hypothesis set is finite, and the upper bound follows from standard learning-theoretic results; see, e.g., (Mohri et al., 2018). \square

F.3. Proof of Theorem 3.6

In the following, we outline the proof of Theorem 3.6. First, we define feedforward neural networks and show how simple GNNs can be interpreted as such.

Feedforward neural networks A *feedforward neural network* (FNN) is specified by a tuple $N = (\mathcal{N}, \beta, \gamma)$ where \mathcal{N} describes the underlying *architecture* and where β and γ define the *parameters* or *weights*. More specifically, $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, i_1, \dots, i_p, o_1, \dots, o_q, \alpha^{\mathcal{N}})$ where $(V^{\mathcal{N}}, E^{\mathcal{N}})$ is a *finite DAG* with p *input nodes* i_1, \dots, i_p of in-degree 0, and q *output nodes* o_1, \dots, o_q of out-degree 0. No other nodes have in- or out-degree zero. Moreover, $\alpha^{\mathcal{N}}$ is a function assigning to each node $v \in V^{\mathcal{N}} \setminus \{i_1, \dots, i_p\}$ an activation function $\alpha(v) : \mathbb{R} \rightarrow \mathbb{R}$. Furthermore, the function $\beta : V^{\mathcal{N}} \setminus \{i_1, \dots, i_p\} \rightarrow \mathbb{R}$

is a function assigning biases to nodes, and finally, the function $\gamma: E^{\mathcal{N}} \rightarrow \mathbb{R}$ assigns weights to edges. For an FNN N , we define its *size* s as the number of biases and weights, that is $s = |V^{\mathcal{N}}| - p + |E^{\mathcal{N}}|$.

Given an FNN $N = (\mathcal{N}, \beta, \gamma)$, we get a function $\text{fnn}_N: \mathbb{R}^p \rightarrow \mathbb{R}^q$ defined as follows. For all v in $V^{\mathcal{N}}$, we define a function $h_v^N: \mathbb{R}^p \rightarrow \mathbb{R}$ such for $\mathbf{a} = (a_1, \dots, a_p)$ in \mathbb{R}^p ,

$$h_v^N(\mathbf{a}) := \begin{cases} a_j & \text{if } v = i_j \text{ for } j \in [p], \\ \alpha^{\mathcal{N}}(v) \left(\sum_{u \in N^+(v)} \gamma(u, v) h_u^N(\mathbf{a}) + \beta(v) \right) & \text{otherwise.} \end{cases}$$

Finally, $\text{fnn}_N: \mathbb{R}^p \rightarrow \mathbb{R}^q$ is defined as $\mathbf{a} \mapsto \text{fnn}_N(\mathbf{a}) := (h_{o_1}^N(\mathbf{a}), \dots, h_{o_q}^N(\mathbf{a}))$.

Simple GNNs as FNNs We next connect simple GNNs in $\text{GNN}_{\text{slp}}(d, L)$ to FNNs. As described in Section E such GNNs are specified by $L + 1$ activation functions $\boldsymbol{\sigma} := (\sigma_1, \dots, \sigma_{L+1})$ and a weight vector Θ in $\mathbb{R}^{d(2dL+L+1)+1}$ describing weight matrices and bias vectors in all the layers. We show that for any attributed graph of order at most n $\mathbf{G} = (G, \mathbf{L})$ in $\mathcal{G}_{n,d}$ with $G = (V(G), E(G))$ and \mathbf{L} in $\mathbb{R}^{n \times d}$ there exists an architecture $\mathcal{N}_G(\boldsymbol{\sigma})$ such that for any weight assignment Θ in $\mathbb{R}^{d(2dL+L+1)+1}$ of the GNN, there exists $\beta_\Theta: V^{\mathcal{N}_G} \rightarrow \mathbb{R}$ and $\gamma_\Theta: E^{\mathcal{N}_G} \rightarrow \mathbb{R}$, satisfying

$$\text{gnn}_{\Theta, \boldsymbol{\sigma}}(G, \mathbf{L}) = \text{fnn}_{N=(\mathcal{N}_G(\boldsymbol{\sigma}), \beta_\Theta, \gamma_\Theta)}(\mathbf{L}'), \quad (17)$$

where \mathbf{L}' in \mathbb{R}^{nd} is the (column-wise) concatenation of the rows of the matrix \mathbf{L} . Moreover, $\mathcal{N}_G(\boldsymbol{\sigma})$ is of polynomial size in the number of vertices and edges in G , the feature dimension d , and the number of layers L . Furthermore, $\mathcal{N}_G(\boldsymbol{\sigma})$ has only a single output node o .

The idea behind the construction of $\mathcal{N}_G(\boldsymbol{\sigma})$ is to consider the tree unraveling or unrolling, see, e.g., (Morris et al., 2020b), of the computation of $\text{gnn}_{\Theta, \boldsymbol{\sigma}}(G, \mathbf{L})$ but instead of a tree we represent the computation more concisely as a DAG. The DAG $\mathcal{N}_G(\boldsymbol{\sigma})$ is defined as follows.

- The node set $V^{\mathcal{N}_G}$ consists of the following nodes.
 - We have input nodes $i_{v,j}$ for v in $V(G)$ and j in $[d]$ which will take the vertex labels L_{vj} in \mathbb{R} as value.
 - For each t in $[L]$, we include the following nodes: $n_{v,j}^{(t)}$ for v in $V(G)$, j in $[d]$.
 - Finally, we have a single output node o .

We thus have $d(L+1)|V(G)| + 1$ nodes in total.

- The edge set $E^{\mathcal{N}_G}$ consists of the following edges.
 - We have edges encoding the adjacency structure of the graph G in every layer. More specifically, we have an edge $e_{u,j,v,k,t} := (n_{u,j}^{(t-1)}, n_{v,k}^{(t)})$ whenever u in $N(v) \cup \{v\}$ and where u and v are in $V(G)$, j and k in $[d]$, and t in $\{2, \dots, L\}$.
 - We also have edges from the input nodes $i_{u,j}$ to $n_{v,k}^{(1)}$ for all u in $N_G(v) \cup \{v\}$ and where u and v are in $V(G)$ and j and k in $[d]$.
 - Finally, we have edges connecting the last layer nodes to the output, i.e., edges $e_{v,j} := (n_{v,j}^{(L)}, o)$ for all v in $V(G)$ and j in $[d]$.

We thus have $d|V(G)| + d^2((L-1)(E(G) + V(G)) + (E(G) + V(G)))$ edges in total.

- Finally, we define the activation functions.
 - $\alpha^{\mathcal{N}}(n_{v,j}^{(t)}) := \sigma_t$ for all v in $V(G)$, j in $[d]$ and t in $[L]$, and $\alpha^{\mathcal{N}}(o) := \sigma_{L+1}$.

This fixes the architecture $\mathcal{N}_G(\boldsymbol{\sigma})$. We next verify Equation (17). Let Θ in $\mathbb{R}^{d(2dL+L+1)+1}$ and $\mathbf{G} = (G, \mathbf{L})$ in $\mathcal{G}_{n,d}$. Let $\mathcal{N}_G(\boldsymbol{\sigma})$ be the architecture defined above for the graph G . We define β_Θ and γ_Θ , as follows.

- $\beta_\Theta: V^{\mathcal{N}_G} \rightarrow \mathbb{R}$ is such that $\beta_\Theta(n_{v,j}^{(t)}) = b_j^{(t)}$ for all v in $V(G)$, j in $[d]$ and t in $[L]$. We also set $\beta_\Theta(o) = b$.
- $\gamma_\Theta: E^{\mathcal{N}_G} \rightarrow \mathbb{R}$ is such that $\gamma_\Theta(e_{u,j,v,k,t}) := W_{jk}^{(2,t)}$ if $u \neq v$ and $\gamma_\Theta(e_{u,j,v,k,t}) := W_{jk}^{(1,t)}$ otherwise, and $\gamma_\Theta(e_{v,j}) = w_j$, for u and v in $V(G)$, j and k in $[d]$, and t in $[L]$.

Note that we share weights across edges that correspond to the same edge in the underlying graph.

Now, if we denote by $\mathbf{f}_v^{(t)}$ the feature vector in \mathbb{R}^d computed in the t th layer by the GNN $\text{gnn}_{\Theta, \sigma}(G, \mathbf{L})$, then it is readily verified, by induction on the layers, that for $N = (\mathcal{N}_G, \alpha_\sigma, \beta_\Theta, \gamma_\Theta)$:

$$h_{n_{v,j}}^N = \mathbf{f}_{v,j}^{(t)} \text{ and thus } h_o^N := \sigma_{L+1} \left(\sum_{v \in V(G)} \sum_{j \in [d]} w_j \mathbf{f}_{v,j}^{(L)} + b \right),$$

from which Equation (17) follows.

We next expand the construction by obtaining an FNN that simulates GNNs on multiple input graphs. More specifically, consider a set \mathfrak{G} consisting of m graphs $\mathbf{G}_1 = (G_1, \mathbf{L}_1), \dots, \mathbf{G}_m = (G_m, \mathbf{L}_m)$ in $\mathcal{G}_{n,d}$ and a GNN in $\text{GNN}_{\text{slp}}(d, L)$ using activation functions $\sigma = (\sigma_1, \dots, \sigma_{L+1})$ in its layers. We first construct an FNN architecture $\mathcal{N}_{G_i}(\sigma)$ for each graph separately, as explained above, such that for every Θ in \mathbb{R}^P , there exists β_Θ and γ_Θ such that

$$\text{gnn}_{\Theta, \sigma}(G_i, \mathbf{L}_i) = \text{fnn}_{N_{G_i} := (\mathcal{N}_{G_i}(\sigma), \beta_\Theta, \gamma_\Theta)}(\mathbf{L}'_i),$$

with \mathbf{L}'_i is the concatenation of rows in \mathbf{L}_i , as before.

Then, let $\mathcal{N}_{\mathfrak{G}}(\sigma)$ be the FNN architecture obtained as the disjoint union of $\mathcal{N}_{G_1}(\sigma), \dots, \mathcal{N}_{G_m}(\sigma)$. If we denote by o_i the output node of $\mathcal{N}_{G_i}(\sigma)$ in $\mathcal{N}_{\mathfrak{G}}(\sigma)$, then we have again that for every Θ in \mathbb{R}^P , there exists β_Θ and γ_Θ such that

$$\text{gnn}_{\Theta, \sigma}(\mathbf{G}_i) = h_{o_i}^{N_{\mathfrak{G}} := (\mathcal{N}_{\mathfrak{G}}(\sigma), \beta_\Theta, \gamma_\Theta)}(\mathbf{L}')$$

for all i in $[m]$, where $\mathbf{L}' := (\mathbf{L}'_1, \dots, \mathbf{L}'_m)$.

We recall that, for t in $[L]$, the nodes in $\mathcal{N}_{\mathfrak{G}}(\sigma)$ are of the form $\nu_{v,j}^{(t),g}$ for v in V_{G_g} , j in $[d]$ and g in $[m]$. In layer $L+1$, we have the output nodes o_1, \dots, o_m . If the order of the graphs in \mathfrak{G} is at most n , then every layer, except the last one, has ndm nodes. The last layer only has m nodes.

Piece-wise polynomial activation functions A *piece-wise polynomial activation function* $\sigma_{p,\delta} : \mathbb{R} \rightarrow \mathbb{R}$ is specified by a partition of \mathbb{R} into p intervals I_j and corresponding polynomials $p_j(x)$ of degree at most δ , for j in $[p]$. That is, $\sigma_{p,\delta}(x) = p_j(x)$ if x in I_j . Examples of $\sigma_{p,\delta}(x)$ are: $\text{sign}(x) : \mathbb{R} \rightarrow \mathbb{R} : x \mapsto \mathbf{1}_{x \geq 0}$ for which $p = 2$ and $\delta = 0$, $\text{relu}(x) : \mathbb{R} \rightarrow \mathbb{R} : x \mapsto \max(0, x)$ for which $p = 2$ and $\delta = 1$, and $\text{lsig}(x) : \mathbb{R} \rightarrow \mathbb{R} : x \mapsto \max(0, \min(1, x))$ for which $p = 3$ and $\delta = 1$. *Piece-wise linear activation functions* are of the form $\sigma_{p,1}$, i.e., they are defined in terms of linear polynomials. The parameters of an activation function $\sigma_{p,\delta}$ consist of the coefficients of the polynomials involved and the boundary points (numbers) of the intervals in the partition of \mathbb{R} .

Proof of Theorem 3.6 We next derive upper bounds on the VC dimension of GNNs by the approach used in Bartlett et al. (2019a), where they used it for bounding the VC dimension of FNNs using piecewise polynomial activation functions. Their approach allows for recovering known bounds on the VC dimension of FNNs in a unified manner. As we will see, the bounds by Bartlett et al. (2019a) for FNNs naturally translate to bounds for GNNs.

Assume d and L in \mathbb{N} . In this section, we will consider the subclass of GNNs in $\text{GNN}_{\text{slp}}(d, L)$ that use *piece-wise polynomial activation functions* with $p > 0$ pieces and degree $\delta \geq 0$. As explained in Section E, $d(2dL + L + 1) + 1$ is the total number of (learnable) parameters for our GNNs in $\text{GNN}_{\text{slp}}(d, L)$. As shorthand notation, we define $P := d(2dL + L + 1) + 1$. We first bound $\text{VC-dim}_{\mathcal{G}_{d,n}}(\text{GNN}_{\text{slp}}(d, L))$ and then use this bound to obtain a bound for $\text{VC-dim}_{\mathcal{G}_{d,\leq n}}(\text{GNN}_{\text{slp}}(d, L))$.

We take \mathfrak{G} consisting of m graphs $\mathbf{G}_1 = (G_1, \mathbf{L}_1), \dots, \mathbf{G}_m = (G_m, \mathbf{L}_m)$ in $\mathcal{G}_{n,d}$ and consider the FNN architecture $\mathcal{N}_{\mathfrak{G}}(\sigma)$ define above with output nodes o_1, \dots, o_m . Let $\text{tresh} : \mathbb{R} \rightarrow \mathbb{R}$ such that $\text{tresh}(x) = 1$ if $x \geq 2/3$ and $\text{tresh}(x) = 0$ if $x \leq 1/3$. We will bound

$$K' := \left| \left\{ \left(\text{tresh}(h_{o_1}^{N_{\mathfrak{G}}}(\mathbf{L}')), \dots, \text{tresh}(h_{o_m}^{N_{\mathfrak{G}}}(\mathbf{L}')) \right) : N_{\mathfrak{G}} := (\mathcal{N}_{\mathfrak{G}}(\sigma), \beta_\Theta, \gamma_\Theta), \Theta \in \mathbb{R}^P \right\} \right|,$$

as this number describes how many 0/1 patterns can occur when Θ ranges over \mathbb{R}^P . These 0/1 patterns correspond, by the construction of $N_{\mathfrak{G}}$ and the semantics of its output nodes, to how many of the input graphs in \mathfrak{G} can be shattered. To bound K' using the approach in Bartlett et al. (2019a) we need to slightly change the activation function σ_{L+1} used in the FNN

architecture. The reason is that Bartlett et al. use the sign function to turn a real-valued function into a 0/1-valued function. In contrast, we use the tresh function described above.

Let $\sigma' := (\sigma_1, \dots, \sigma_L, \sigma_{L+1} - 1/3)$. We will bound K' by bounding

$$K := \left| \left\{ \left(\text{sign}(h_{o_1}^{N_{\mathfrak{G}}}(\mathbf{L}')), \dots, \text{sign}(h_{o_m}^{N_{\mathfrak{G}}}(\mathbf{L}')) \right) : N_{\mathfrak{G}} := (\mathcal{N}_{\mathfrak{G}}(\sigma'), \beta_{\Theta}, \gamma_{\Theta}), \Theta \in \mathbb{R}^P \right\} \right|.$$

Note that $K' \leq K$ because if $\text{tresh}(\sigma_{L+1})(\mathbf{x}) = 1$ then $\sigma_{L+1}(\mathbf{x}) \geq 2/3$ and hence $\sigma_{L+1}(\mathbf{x}) - 1/3 > 0$ and hence $\text{sign}(\sigma_{L+1}(\mathbf{x}) - 1/3) = 1$. Similarly, $\text{tresh}(\sigma_{L+1})(\mathbf{x}) = 0$ then $\sigma_{L+1}(\mathbf{x}) \leq 1/3$ and hence $\sigma_{L+1}(\mathbf{x}) - 1/3 \leq 0$ and hence $\text{sign}(\sigma_{L+1}(\mathbf{x}) - 1/3) = 0$.

Then, if $\text{VC-dim}_{\mathcal{G}_{d,n}}(\text{GNN}_{\text{slp}}(d, L)) = m$ then $K \geq 2^m$. We thus bound K in terms of a function κ in m and then use $2^m \leq \kappa(m)$ to find an upper bound for m , i.e., an upper bound for $\text{VC-dim}_{\mathcal{G}_{d,n}}(\text{GNN}_{\text{slp}}(d, L))$. To bound K we can now use the approach of Bartlett et al. (2019a). In a nutshell, the entire parameter space \mathbb{R}^P is partitioned into pieces S_1, \dots, S_{ℓ} such that whenever Θ and Θ' belong to the same piece (i) they incur the same sign pattern in $\{0, 1\}^m$; and (ii) each $h_{o_1}^{N_{\mathfrak{G}}}(\mathbf{L}')$ is a polynomial of degree at most $1 + L\delta^L$. For $\delta = 0$, these are polynomials in $d + 1$ variables, for $\delta > 0$, the number of variables is P . Crucial in Bartlett's approach is the following lemma.

Lemma F.4 (Lemma 17 in Bartlett et al. (2019a)). *Let $p_1(\mathbf{x}), \dots, p_r(\mathbf{x})$ be polynomials of degree at most δ and in variables \mathbf{x} satisfying $|\mathbf{x}| \leq r$, where $|\cdot|$ denotes the number of components of a vector. Then*

$$\left| \left\{ \left(\text{sign}(p_1(\Theta)), \dots, \text{sign}(p_r(\Theta)) \right) \mid \Theta \in \mathbb{R}^{|\mathbf{x}|} \right\} \right| \leq 2 \left(\frac{2er\delta}{|\mathbf{x}|} \right)^{|\mathbf{x}|}.$$

Given property (ii) of the pieces S_1, \dots, S_{ℓ} , we can apply the above lemma to the polynomials $h_{o_1}^{N_{\mathfrak{G}}}(\mathbf{L}'), \dots, h_{o_m}^{N_{\mathfrak{G}}}(\mathbf{L}')$ and, provided that the number of variables is at most m , obtain a bound for K by $\ell 2 \left(\frac{2em}{d+1} \right)^{d+1}$, when $\delta = 0$, and $K \leq \ell 2 \left(\frac{2em(1+L\delta^L)}{P} \right)^P$ for $\delta > 0$.

It then remains to bound the number of parts ℓ . Bartlett et al. show how to do this inductively (on the number of layers), again using Lemma F.4. More precisely, every node in the FNN architecture is associated with a number of polynomials. In layer t we have nmd nodes (number of computation nodes), and we associate with each node p polynomials (number of breakpoints of activation function) of degree at most $1 + (t-1)\delta^{t-1}$ and have $(2d+1)d$ variables for $\delta = 0$ and $(2d+1)dt$ variables for $\delta > 0$. We then get, for $\delta = 0$,

$$K \leq 2^L \left(\left(\frac{2edmp}{(2d+1)d} \right)^{(2d+1)d} \right)^L 2 \left(\frac{2em}{d+1} \right)^{d+1}, \quad (18)$$

and for $\delta > 0$,

$$K \leq \prod_{t=1}^L 2 \left(\frac{2edmp(1+(t-1)\delta^{t-1})}{(2d+1)dt} \right)^{(2d+1)dt} 2 \left(\frac{2em(1+L\delta^L)}{P} \right)^P. \quad (19)$$

These are precisely the bounds given in Bartlett et al. (2019a) applied to our FNN. It is important, however, to note that this upper bound is only valid when Lemma F.4 can be applied, and hence, the number of variables must be smaller than the number of polynomials. For t in $[L]$, we must have that the number of variables is less than the number of polynomials. We have $nmdp$ polynomials, and up to layer t we have $(2d+1)dt$ parameters (variables). Hence, we must have $(2d+1)dt \leq nmdp$ or $(2d+1)t \leq nmp$, and also $D \leq m$ or $P \leq nmdp$ for $\delta > 0$. For $\delta = 0$, we need $(2d+1)d \leq nmdp$ (or $(2d+1) \leq nmp$) and $d+1 \leq m$. The following conditions are sufficient

$$P \leq m \text{ for } \delta > 0, \text{ and } 2d+1 \leq m \text{ for } \delta = 0. \quad (\dagger)$$

FNN size reduction based on 1-WL We next bring in 1-WL into consideration by collapsing computation nodes in $N_{\mathfrak{G}}$ in each layer based on their equivalence with regards to 1-WL. In other words, if we assume that the graphs to be shattered have at most u vertex colors, then we have at most u —rather than n —computation nodes per graph. This implies that the parameter n in the above expression can be replaced by u .

As a consequence, following [Bartlett et al. \(2019a\)](#), using the weighted AM–GM inequality on the right-hand side of the inequalities (18) and (19), we obtain a bound for the VC dimension by finding maximal m satisfying, for $\delta = 0$,

$$2^m \leq 2^{L+1} \left(\frac{2ep(udL+1)m}{P} \right)^P$$

and for $\delta > 0$,

$$2^m \leq 2^{L+1} \left(\frac{2ep(ud \sum_{t=1}^L (1+(t-1)\delta^{t-1}) + (1+L\delta^L))m}{\frac{L(L+1)}{2}(2d+1)d+P} \right)^{\frac{L(L+1)}{2}(2d+1)d+P}.$$

Such m are found in ([Bartlett et al., 2019a](#)), resulting in the following bounds.

Proposition F.5 (([Bartlett et al., 2019a](#)) modified for our FNN $N_{\mathfrak{G}}$).

$$\text{VC-dim}_{\mathcal{G}_{d,\leq u}}(\text{GNN}_{\text{slp}}(d, L)) \leq \begin{cases} \mathcal{O}(Ld^2 \log(p(udL+1))) & \text{if } \delta = 0 \\ \mathcal{O}(L^2 d^2 \log(p(udL+1))) & \text{if } \delta = 1 \\ \mathcal{O}(L^2 d^2 \log(p(udL+1)) + L^3 d^2 \log(\delta)) & \text{if } \delta > 1. \end{cases}$$

We can simplify this to $\mathcal{O}(P \log(puP))$, $\mathcal{O}(LP \log(puP))$, and $\mathcal{O}(LP \log(puP) + L^2 P \log(\delta))$, respectively. We note that since these are larger than P , the condition (\dagger) is satisfied.

G. Additional experimental data and results

Here, we report on additional experimental details, results, and dataset generation.

G.1. Simple GNN layer used for Q1 and Q2

The simple GNN layer used in **Q1** and **Q2** updates the feature of vertex v at layer t via

$$\mathbf{f}_v^{(t)} := \sigma \left(\mathbf{f}_v^{(t-1)} \mathbf{W}_1^{(t)} + \sum_{u \in N(v)} \mathbf{f}_u^{(t-1)} \mathbf{W}_2^{(t)} \right) \in \mathbb{R}^d, \quad (20)$$

where $\mathbf{W}_1^{(t)}$ and $\mathbf{W}_2^{(t)} \in \mathbb{R}^{d \times d}$ are parameter matrices. In the experiments, we used ReLU activation functions.

G.2. GNN architecture used for Q3

For the experiments on **Q3** we extend the simple GNN layer from Equation (20) used in **Q1** and **Q2**. We update the feature of vertex v at layer t via

$$\mathbf{f}_v^{(t)} := \sigma \left(\text{BN} \left(\mathbf{f}_v^{(t-1)} \mathbf{W}^{(t)} + \sum_{u \in N(v)} \text{mlp}^{(t)}(\mathbf{f}_u^{(t-1)}) \right) \right) \in \mathbb{R}^d, \quad (21)$$

where BN is a batch normalization module ([Ioffe & Szegedy, 2015](#)) and $\text{mlp}^{(t)}$ is a two-layer perceptron with architecture,

$$\text{Linear} \rightarrow \text{BN} \rightarrow \text{ReLU} \rightarrow \text{Linear}.$$

We, therefore, use a normalized 2-layer MLP to generate messages in each layer. We found this change necessary to ensure smooth convergence on the challenging synthetic task posed by **Q3**, where the GNN has to memorize an arbitrary binary graph labeling. Moreover, in the experiments, we used ReLU activation functions.

G.3. Synthetic dataset generation

To address **Q3**, we aim to empirically estimate how well GNNs of different sizes can fit arbitrary binary labelings of graphs. We construct a synthetic dataset that focuses on a simple class of trees. Formally, for two natural numbers m and n in $\mathbb{N}_{\geq 0}$, we define the graph $T_{m,n} = (V_{m,n}, E_{m,n})$ as a directed tree with vertex set $V = \{v_0, \dots, v_{m+n+3}\}$. The root v_0 has two children v_1 and v_2 . The remaining $m+n$ vertices form the leaves such that vertex v_1 has m children and v_2 has n children. Figure 4 provides a visual example.

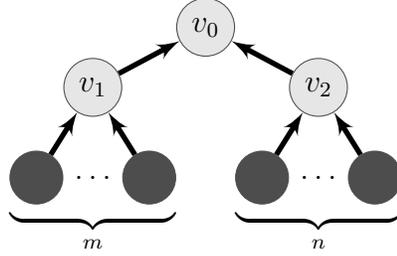

 Figure 4: A visualization of a $T_{m,n}$.

Table 3: Dataset statistics and properties.

Dataset	Properties					
	Number of graphs	Number of classes/targets	∅ Number of nodes	∅ Number of edges	Node labels	Edge labels
ENZYMES	600	6	32.6	62.1	✓	✗
MCF-7	27 770	2	26.4	28.5	✓	✓
MCF-7H	27 770	2	47.3	49.4	✓	✓
MUTAGENICITY	4 337	2	30.3	30.8	✓	✓
NCI1	4 110	2	29.9	32.3	✓	✗
NCI109	4 127	2	29.7	32.1	✓	✗

For a chosen k in \mathbb{N} , $k \geq 4$, we define:

$$\mathcal{T}_k = \left\{ T_{m,n} \mid 0 \leq m \leq \left\lfloor \frac{(k-3)}{2} \right\rfloor, n = k - 3 - m \right\}.$$

Therefore, \mathcal{T}_k contains all distinct graphs $T_{m,n}$ with $|V_{m,n}(T_{m,n})| = k$. In particular, we observe $|\mathcal{T}_k| = \left\lfloor \frac{(k-3)}{2} \right\rfloor$.

For $k \in \{10, 20, \dots, 90\}$, we aim to test how well a GNN with a given feature dimension d in $\{4, 16, 64, 256\}$ can learn binary labelings $y: \mathcal{T}_k \rightarrow \{0, 1\}$ of \mathcal{T}_k . The labeling y is obtained by sampling the label $y(T)$ uniformly at random for all T in \mathcal{T}_k . We then train a GNN model with stochastic gradient descent to minimize the binary cross entropy on the dataset (\mathcal{T}_k, y) . For each combination of k and d , we repeat the experiment 50 times. We resample a new labeling y and a new random initialization of the GNN model in each repetition.

G.4. Simulating bitlength via higher feature dimension

Here, we outline how to simulate a higher bitlength via a higher feature dimension. Assume the simple GNN layer of Equation (20). Clearly, we can express the matrices $\mathbf{W}_1^{(t)}$ and $\mathbf{W}_2^{(t)}$ as the sum of k matrix with smaller bitlength, e.g.,

$$\mathbf{W}_2^{(t)} = \mathbf{W}_2^{(1,t)} + \dots + \mathbf{W}_2^{(k,t)}.$$

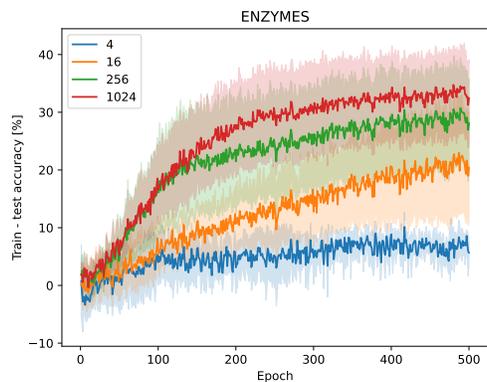
Hence, we can re-write the aggregation in Equation (20) as

$$\sum_{u \in N(v)} \mathbf{f}_u^{(t-1)} [\mathbf{W}_2^{(1,t)}, \dots, \mathbf{W}_2^{(k,t)}] \cdot \mathbf{M} \in \mathbb{R}^d,$$

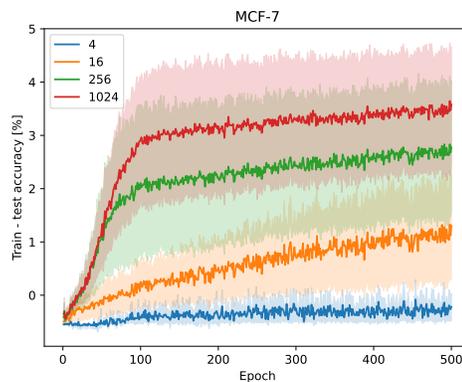
where $[\dots]$ denotes column-wise matrix concatenation and \mathbf{M} in $\{0, 1\}^{kd \times d}$ is a matrix such that

$$\sigma \left(\mathbf{f}_v^{(t-1)} \mathbf{W}_1^{(t)} + \sum_{u \in N(v)} \mathbf{f}_u^{(t-1)} [\mathbf{W}_2^{(1,t)}, \dots, \mathbf{W}_2^{(k,t)}] \cdot \mathbf{M} \right) = \sigma \left(\mathbf{f}_v^{(t-1)} \mathbf{W}_1^{(t)} + \sum_{u \in N(v)} \mathbf{f}_u^{(t-1)} \mathbf{W}_2^{(t)} \right),$$

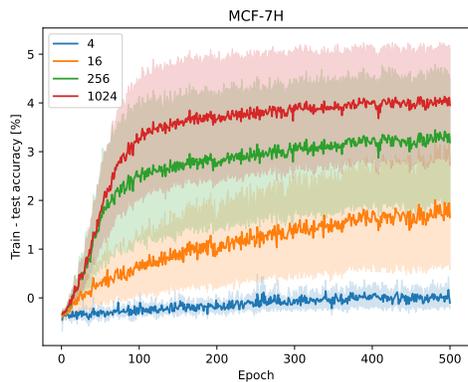
i.e., the matrix \mathbf{M} sums together columns of the aggregated features such that they have feature dimension d .



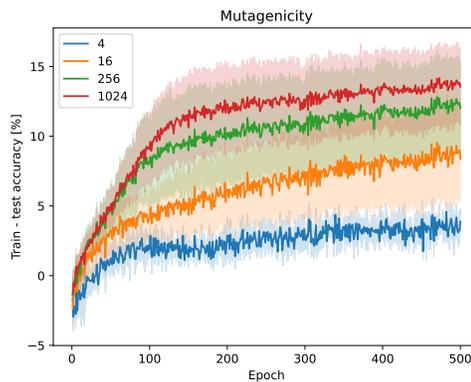
(a) ENZYMES



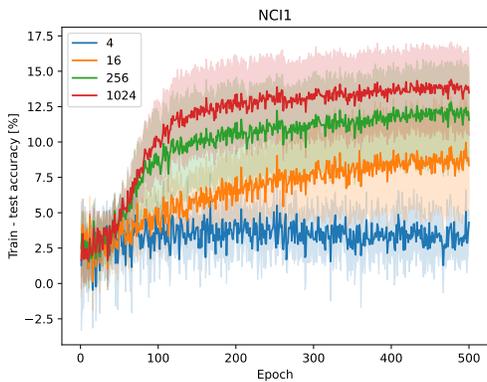
(b) MCF-7



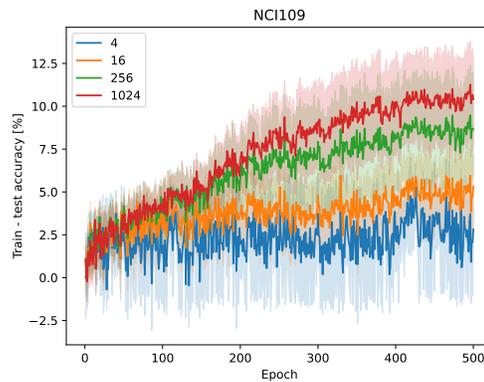
(c) MCF-7H



(d) MUTAGENICITY



(e) NCI1



(f) NCI109

Figure 5: Difference between train and test accuracy for different feature dimensions in $\{4, 16, 256, 1024\}$.