
Can Graph Neural Networks Count Substructures?

Zhengdao Chen
New York University
zc1216@nyu.edu

Lei Chen
New York University
lc3909@nyu.edu

Soledad Villar
Johns Hopkins University
soledad.villar@jhu.edu

Joan Bruna
New York University
bruna@cims.nyu.edu

Abstract

The ability to detect and count certain substructures in graphs is important for solving many tasks on graph-structured data, especially in the contexts of computational chemistry and biology as well as social network analysis. Inspired by this, we propose to study the expressive power of graph neural networks (GNNs) via their ability to count attributed graph substructures, extending recent works that examine their power in graph isomorphism testing and function approximation. We distinguish between two types of substructure counting: induced-subgraph-count and subgraph-count, and establish both positive and negative answers for popular GNN architectures. Specifically, we prove that Message Passing Neural Networks (MPNNs), 2-Weisfeiler-Lehman (2-WL) and 2-Invariant Graph Networks (2-IGNs) cannot perform induced-subgraph-count of any connected substructure consisting of 3 or more nodes, while they can perform subgraph-count of star-shaped substructures. As an intermediary step, we prove that 2-WL and 2-IGNs are equivalent in distinguishing non-isomorphic graphs, partly answering an open problem raised in [38]. We also prove positive results for k -WL and k -IGNs as well as negative results for k -WL with a finite number of iterations. We then conduct experiments that support the theoretical results for MPNNs and 2-IGNs. Moreover, motivated by substructure counting and inspired by [45], we propose the Local Relational Pooling model and demonstrate that it is not only effective for substructure counting but also able to achieve competitive performance on molecular prediction tasks.

1 Introduction

In recent years, graph neural networks (GNNs) have achieved empirical success on processing data from various fields such as social networks, quantum chemistry, particle physics, knowledge graphs and combinatorial optimization [4, 5, 9, 10, 11, 13, 14, 31, 47, 55, 58, 66, 68, 69, 70, 71, 73, 74]. Thanks to such progress, there have been growing interests in studying the expressive power of GNNs. One line of work does so by studying their ability to distinguish non-isomorphic graphs. In this regard, Xu et al. [64] and Morris et al. [44] show that GNNs based on neighborhood-aggregation schemes are at most as powerful as the classical Weisfeiler-Lehman (WL) test [60] and propose GNN architectures that can achieve such level of power. While graph isomorphism testing is very interesting from a theoretical viewpoint, one may naturally wonder how relevant it is to real-world tasks on graph-structured data. Moreover, WL is powerful enough to distinguish almost all pairs of non-isomorphic graphs except for rare counterexamples [3]. Hence, from the viewpoint of graph isomorphism testing, existing GNNs are in some sense already not far from being maximally powerful, which could make the pursuit of more powerful GNNs appear unnecessary.

Another perspective is the ability of GNNs to approximate permutation-invariant functions on graphs. For instance, Maron et al. [41] and Keriven and Peyré [29] propose architectures that achieve universal approximation of permutation-invariant functions on graphs, though such models involve tensors with order growing in the size of the graph and are therefore impractical. Chen et al. [8] establishes an equivalence between the ability to distinguish any pair of non-isomorphic graphs and the ability to approximate arbitrary permutation-invariant functions on graphs. Nonetheless, for GNNs used in practice, which are not universally approximating, more efforts are needed to characterize *which* functions they can or cannot express. For example, Loukas [37] shows that GNNs under assumptions are Turing universal but lose power when their depth and width are limited, though the arguments rely on the nodes all having distinct features and the focus is on the asymptotic depth-width tradeoff. Concurrently to our work, Garg et al. [17] provide impossibility results of several classes of GNNs to decide graph properties including girth, circumference, diameter, radius, conjoint cycle, total number of cycles, and k -cliques. Despite these interesting results, we still need a perspective for understanding the expressive power of different classes of GNNs in a way that is intuitive, relevant to goals in practice, and potentially helpful in guiding the search for more powerful architectures.

Meanwhile, graph substructures (also referred to by various names including *graphlets*, *motifs*, *subgraphs* and *graph fragments*) are well-studied and relevant for graph-related tasks in computational chemistry [12, 14, 26, 27, 28, 46], computational biology [32] and social network studies [25]. In organic chemistry, for example, certain patterns of atoms called functional groups are usually considered indicative of the molecules’ properties [33, 49]. In the literature of molecular chemistry, substructure counts have been used to generate molecular fingerprints [43, 48] and compute similarities between molecules [1, 53]. In addition, for general graphs, substructure counts have been used to create graph kernels [56] and compute spectral information [51]. The connection between GNNs and graph substructures is explored empirically by Ying et al. [67] to interpret the predictions made by GNNs. Thus, the ability of GNN architectures to count graph substructures not only serves as an intuitive theoretical measure of their expressive power but also is highly relevant to practical tasks.

In this work, we propose to understand the expressive power of GNN architectures via their ability to count attributed substructures, that is, counting the number of times a given pattern (with node and edge features) appears as a *subgraph* or *induced subgraph* in the graph. We formalize this question based on a rigorous framework, prove several results that partially answer the question for Message Passing Neural Networks (MPNNs) and Invariant Graph Networks (IGNs), and finally propose a new model inspired by substructure counting. In more detail, our main contributions are:

1. We prove that neither MPNNs [18] nor 2nd-order Invariant Graph Networks (2-IGNs) [41] can count *induced subgraphs* for any connected pattern of 3 or more nodes. For any such pattern, we prove this by constructing a pair of graphs that provably cannot be distinguished by any MPNN or 2-IGN but with different induced-subgraph-counts of the given pattern. This result points at an important class of simple-looking tasks that are provably hard for classical GNN architectures.
2. We prove that MPNNs and 2-IGNs can count *subgraphs* for star-shaped patterns, thus generalizing the results in Arvind et al. [2] to incorporate node and edge features. We also show that k -WL and k -IGNs can count *subgraphs* and *induced subgraphs* for patterns of size k , which provides an intuitive understanding of the hierarchy of k -WL’s in terms of increasing power in counting substructures.
3. We prove that T iterations of k -WL is unable to count *induced subgraphs* for *path* patterns of $(k + 1)2^T$ or more nodes. The result is relevant since real-life GNNs are often shallow, and also demonstrates an interplay between k and depth.
4. Since substructures present themselves in local neighborhoods, we propose a novel GNN architecture called *Local Relation Pooling (LRP)*¹, with inspirations from Murphy et al. [45]. We empirically demonstrate that it can count both *subgraphs* and *induced subgraphs* on random synthetic graphs while also achieving competitive performances on molecular datasets. While variants of GNNs have been proposed to better utilize substructure information [34, 35, 42], often they rely on handcrafting rather than learning such information. By contrast, LRP is not only powerful enough to count substructures but also able to learn from data *which* substructures are relevant.

¹Code available at <https://github.com/leichen2018/GNN-Substructure-Counting>.




Figure 1: Illustration of the two types substructure-counts of the patterns $G^{[P_1]}$ and $G^{[P_2]}$ in the graph G , as defined in Section 2.1. The node features are indicated by colors. We have $C_I(G; G^{[P_1]}) = C_S(G; G^{[P_1]}) = 1$, and $C_I(G; G^{[P_2]}) = 0$ while $C_S(G; G^{[P_2]}) = 1$.

2 Framework

2.1 Attributed graphs, (induced) subgraphs and two types of counting

We define an *attributed graph* as $G = (V, E, x, e)$, where $V = [n] := \{1, \dots, n\}$ is the set of vertices, $E \subset V \times V$ is the set of edges, $x_i \in \mathcal{X}$ represents the feature of node i , and $e_{i,j} \in \mathcal{Y}$ represent the feature of the edge (i, j) if $(i, j) \in E$. The adjacency matrix $A \in \mathbb{R}^{n \times n}$ is defined by $A_{i,j} = 1$ if $(i, j) \in E$ and 0 otherwise. We let $D_i = \sum_{j \in V} A_{i,j}$ denote the degree of node i . For simplicity, we only consider undirected graphs without self-connections or multi-edges. Note that an unattributed graph $G = (V, E)$ can be viewed as an attributed graph with identical node and edge features.

Unlike the node and edge features, the indices of the nodes are not inherent properties of the graph. Rather, different ways of ordering the nodes result in different representations of the same underlying graph. This is characterized by the definition of *graph isomorphism*: Two attributed graphs $G^{[1]} = (V^{[1]}, E^{[1]}, x^{[1]}, e^{[1]})$ and $G^{[2]} = (V^{[2]}, E^{[2]}, x^{[2]}, e^{[2]})$ are *isomorphic* if there exists a bijection $\pi : V^{[1]} \rightarrow V^{[2]}$ such that (1) $(i, j) \in E^{[1]}$ if and only if $(\pi(i), \pi(j)) \in E^{[2]}$, (2) $x_i^{[1]} = x_{\pi(i)}^{[2]}$ for all i in $V^{[1]}$, and (3) $e_{i,j}^{[1]} = e_{\pi(i), \pi(j)}^{[2]}$ for all $(i, j) \in E^{[1]}$.

For $G = (V, E, x, e)$, a *subgraph* of G is any graph $G^{[S]} = (V^{[S]}, E^{[S]}, x, e)$ with $V^{[S]} \subseteq V$ and $E^{[S]} \subseteq E$. An *induced subgraph* of G is any graph $G^{[S']} = (V^{[S']}, E^{[S']}, x, e)$ with $V^{[S']} \subseteq V$ and $E^{[S']} = E \cap (V^{[S']})^2$. In words, the edge set of an induced subgraph needs to include all edges in E that have both end points belonging to $V^{[S']}$. Thus, an induced subgraph of G is also its subgraph, but the converse is not true.

We now define two types of counting associated with *subgraphs* and *induced subgraphs*, as illustrated in Figure 1. Let $G^{[P]} = (V^{[P]}, E^{[P]}, x^{[P]}, e^{[P]})$ be a (typically smaller) graph that we refer to as a *pattern* or *substructure*. We define $C_S(G, G^{[P]})$, called the *subgraph-count* of $G^{[P]}$ in G , to be the number of *subgraphs* of G that are isomorphic to $G^{[P]}$. We define $C_I(G; G^{[P]})$, called the *induced-subgraph-count* of $G^{[P]}$ in G , to be the number of *induced subgraphs* of G that are isomorphic to $G^{[P]}$. Since all induced subgraphs are subgraphs, we always have $C_I(G; G^{[P]}) \leq C_S(G; G^{[P]})$.

Below, we formally define the ability for certain function classes to count substructures as the ability to distinguish graphs with different subgraph or induced-subgraph counts of a given substructure.

Definition 2.1. Let \mathcal{G} be a space of graphs, and \mathcal{F} be a family of functions on \mathcal{G} . We say \mathcal{F} is able to perform *subgraph-count* (or *induced-subgraph-count*) of a pattern $G^{[P]}$ on \mathcal{G} if for all $G^{[1]}, G^{[2]} \in \mathcal{G}$ such that $C_S(G^{[1]}, G^{[P]}) \neq C_S(G^{[2]}, G^{[P]})$ (or $C_I(G^{[1]}, G^{[P]}) \neq C_I(G^{[2]}, G^{[P]})$), there exists $f \in \mathcal{F}$ that returns different outputs when applied to $G^{[1]}$ and $G^{[2]}$.

In Appendix A, we prove an equivalence between Definition 2.1 and the notion of approximating *subgraph-count* and *induced-subgraph-count* functions on the graph space. Definition 2.1 also naturally allows us to define the ability of graph isomorphism tests to count substructures. A graph isomorphism test, such as the Weisfeiler-Lehman (WL) test, takes as input a pair of graphs and returns whether or not they are judged to be isomorphic. Typically, the test will return true if the two graphs are indeed isomorphic but does not necessarily return false for every pair of non-isomorphic graphs. Given such a graph isomorphism test, we say it is able to perform *induced-subgraph-count* (or *subgraph-count*) of a pattern $G^{[P]}$ on \mathcal{G} if $\forall G^{[1]}, G^{[2]} \in \mathcal{G}$ such that $C_I(G^{[1]}, G^{[P]}) \neq C_I(G^{[2]}, G^{[P]})$ (or $C_S(G^{[1]}, G^{[P]}) \neq C_S(G^{[2]}, G^{[P]})$), the test can distinguish these two graphs.

3 Message Passing Neural Networks and k -Weisfeiler-Lehman tests

The Message Passing Neural Network (MPNN) is a generic model that incorporates many popular architectures, and it is based on learning local aggregations of information in the graph [18]. When applied to an undirected graph $G = (V, E, x, e)$, an MPNN with T layers is defined iteratively as follows. For $t < T$, to compute the message $m_i^{(t+1)}$ and the hidden state $h_i^{(t+1)}$ for each node $i \in V$ at the $(t + 1)$ th layer, we apply the following update rule:

$$m_i^{(t+1)} = \sum_{\mathcal{N}(i)} M_t(h_i^{(t)}, h_j^{(t)}, e_{i,j}), \quad h_i^{(t+1)} = U_t(h_i^{(t)}, m_i^{(t+1)}),$$

where $\mathcal{N}(i)$ is the neighborhood of node i in G , M_t is the message function at layer t and U_t is the vertex update function at layer t . Finally, a graph-level prediction is computed as $\hat{y} = R(\{h_i^{(T)} : i \in V\})$, where R is the readout function. Typically, the hidden states at the first layer are set as $h_i^{(0)} = x_i$. Learnable parameters can appear in the functions M_t, U_t (for all $t \in [T]$) and R .

Xu et al. [64] and Morris et al. [44] show that, when the graphs' edges are unweighted, such models are at most as powerful as the Weisfeiler-Lehman (WL) test in distinguishing non-isomorphic graphs. Below, we will first prove an extension of this result that incorporates edge features. To do so, we first introduce the hierarchy of k -Weisfeiler-Lehman (k -WL) tests. The k -WL test takes a pair of graphs $G^{[1]}$ and $G^{[2]}$ and attempts to determine whether they are isomorphic. In a nutshell, for each of the graphs, the test assigns an initial color in some color space to every k -tuple in V^k according to its *isomorphism type*, and then it updates the colorings iteratively by aggregating information among neighboring k -tuples. The test will terminate and return the judgement that the two graphs are not isomorphic if and only if at some iteration t , the coloring multisets differ. We refer the reader to Appendix C for a rigorous definition.

Remark 3.1. *For graphs with unweighted edges, 1-WL and 2-WL are known to have the same discriminative power [39]. For $k \geq 2$, it is known that $(k + 1)$ -WL is strictly more powerful than k -WL, in the sense that there exist pairs of graph distinguishable by the former but not the latter [6]. Thus, with growing k , the set of k -WL tests forms a hierarchy with increasing discriminative power. Note that there has been an different definition of WL in the literature, sometimes known as Folklore Weisfeiler-Lehman (FWL), with different properties [39, 44].²*

Our first result is an extension of Morris et al. [44], Xu et al. [64] to incorporate edge features.

Theorem 3.2. *Two attributed graphs that are indistinguishable by 2-WL cannot be distinguished by any MPNN.*

The theorem is proven in Appendix D. Thus, it motivates us to first study what patterns 2-WL can or cannot count.

3.1 Substructure counting by 2-WL and MPNNs

It turns out that whether or not 2-WL can perform *induced-subgraph-count* of a pattern is completely characterized by the number of nodes in the pattern. Any connected pattern with 1 or 2 nodes (i.e., representing a node or an edge) can be easily counted by an MPNN with 0 and 1 layer of message-passing, respectively, or by 2-WL with 0 iterations³. In contrast, for all larger connected patterns, we have the following negative result, which we prove in Appendix E.

Theorem 3.3. *2-WL cannot induced-subgraph-count any connected pattern with 3 or more nodes.*

The intuition behind this result is that, given any connected pattern with 3 or more nodes, we can construct a pair of graphs that have different *induced-subgraph-counts* of the pattern but cannot be distinguished from each other by 2-WL, as illustrated in Figure 2. Thus, together with Theorem 3.2, we have

Corollary 3.4. *MPNNs cannot induced-subgraph-count any connected pattern with 3 or more nodes.*

For *subgraph-count*, if both nodes and edges are unweighted, Arvind et al. [2] show that the only patterns 1-WL (and equivalently 2-WL) can count are either star-shaped patterns and pairs of disjoint

²When ‘‘WL test’’ is used in the literature without specifying ‘‘ k ’’, it usually refers to 1-WL, 2-WL or 1-FWL.

³In fact, this result is a special case of Theorem 3.7.




Figure 2: Illustration of the construction in the proof of Theorem 3.3 for the pattern $G^{[P]}$ on the left. Note that $C_I(G^{[1]}; G^{[P]}) = 0$ and $C_I(G^{[2]}; G^{[P]}) = 2$, and the graphs $G^{[1]}$ and $G^{[2]}$ cannot be distinguished by 2-WL, MPNNs or 2-IGNs.

edges. We prove the positive result that MPNNs can count star-shaped patterns even when node and edge features are allowed, utilizing a result in Xu et al. [64] that the message functions are able to approximate any function on multisets.

Theorem 3.5. *MPNNs can perform subgraph-count of star-shaped patterns.*

By Theorem 3.2, this implies that

Corollary 3.6. *2-WL can perform subgraph-count of star-shaped patterns.*

3.2 Substructure counting by k -WL

There have been efforts to extend the power of GNNs by going after k -WL for higher k , such as Morris et al. [44]. Thus, it is also interesting to study the patterns that k -WL can and cannot count. Since k -tuples are assigned initial colors based on their isomorphism types, the following is easily seen, and we provide a proof in Appendix G.

Theorem 3.7. *k -WL, at initialization, is able to perform both induced-subgraph-count and subgraph-count of patterns consisting of at most k nodes.*

This establishes a potential hierarchy of increasing power in terms of substructure counting by k -WL. However, tighter results can be much harder to achieve. For example, to show that 2-FWL (and therefore 3-WL) cannot count cycles of length 8, Fürer [16] has to rely on computers for counting cycles in the classical Cai-Fürer-Immerman counterexamples to k -WL [6]. We leave the pursuit of general and tighter characterizations of k -WL’s substructure counting power for future research, but we are nevertheless able to prove a partial negative result concerning finite iterations of k -WL.

Definition 3.8. *A path pattern of size m , denoted by H_m , is an unattributed graph, $H_m = (V^{[H_m]}, E^{[H_m]})$, where $V^{[H_m]} = [m]$, and $E^{[H_m]} = \{(i, i + 1) : 1 \leq i < m\} \cup \{(i + 1, i) : 1 \leq i < m\}$.*

Theorem 3.9. *Running T iterations of k -WL cannot perform induced-subgraph-count of any path pattern of $(k + 1)2^T$ or more nodes.*

The proof is given in Appendix H. This bound grows quickly when T becomes large. However, since in practice, many if not most GNN models are designed to be shallow [62, 74], this result is still relevant for studying finite-depth GNNs that are based on k -WL.

4 Invariant Graph Networks

Recently, diverging from the strategy of local aggregation of information as adopted by MPNNs and k -WLs, an alternative family of GNN models called *Invariant Graph Networks (IGNs)* was introduced in Maron et al. [39, 40, 41]. Here, we restate its definition. First, note that if the node and edge features are vectors of dimension d_n and d_e , respectively, then an input graph can be represented by a *second-order tensor* $\mathbf{B} \in \mathbb{R}^{n \times n \times (d+1)}$, where $d = \max(d_n, d_e)$, defined by

$$\begin{aligned} \mathbf{B}_{i,i,1:d_n} &= x_i, \quad \forall i \in V = [n], \\ \mathbf{B}_{i,j,1:d_e} &= e_{i,j}, \quad \forall (i,j) \in E, \\ \mathbf{B}_{1:n,1:n,d+1} &= A. \end{aligned} \tag{1}$$

If the nodes and edges do not have features, \mathbf{B} simply reduces to the adjacency matrix. Thus, GNN models can be alternatively defined as functions on such second-order tensors. More generally, with graphs represented by k th-order tensors, we can define:

Definition 4.1. *A k th-order Invariant Graph Network (k -IGN) is a function $F : \mathbb{R}^{n^k \times d_0} \rightarrow \mathbb{R}$ that can be decomposed in the following way:*

$$F = m \circ h \circ L^{(T)} \circ \sigma \circ \dots \circ \sigma \circ L^{(1)},$$

where each $L^{(t)}$ is a linear equivariant layer [40] from $\mathbb{R}^{n^k \times d_{t-1}}$ to $\mathbb{R}^{n^k \times d_t}$, σ is a pointwise activation function, h is a linear invariant layer from $\mathbb{R}^{n^k \times d_T}$ to \mathbb{R} , and m is an MLP.

Maron et al. [41] show that if k is allowed to grow as a function of the size of the graphs, then k -IGNs can achieve universal approximation of permutation-invariant functions on graphs. Nonetheless, due to the quick growth of computational complexity and implementation difficulty as k increases, in practice it is hard to have $k > 2$. If $k = 2$, on one hand, it is known that 2-IGNs are at least as powerful as 2-WL [38]; on the other hand, 2-IGNs are not universal [8]. However, it remains open to establish a strict upper bound on the expressive power of 2-IGNs in terms of the WL tests as well as to characterize concretely their limitations. Here, we first answer the former question by proving that 2-IGNs are no more powerful than 2-WL:

Lemma 4.2. *If two graphs are indistinguishable by 2-WL, then no 2-IGN can distinguish them either.*

We give the full proof of Lemma 4.2 in Appendix I. As a consequence, we then have

Corollary 4.3. *2-IGNs are exactly as powerful as 2-WL.*

Thanks to this equivalence, the following results on the ability of 2-IGNs to count substructures are immediate corollaries of Theorem 3.3 and Corollary 3.6 (though we also provide a direct proof of Corollary 4.4 in Appendix J):

Corollary 4.4. *2-IGNs cannot perform induced-subgraph-count of any connected pattern with 3 or more nodes.*

Corollary 4.5. *2-IGNs can perform subgraph-count of star-shaped patterns.*

In addition, as k -IGNs are no less powerful than k -WL [39], as a corollary of Theorem 3.7, we have

Corollary 4.6. *k -IGNs can perform both induced-subgraph-count and subgraph-count of patterns consisting of at most k nodes.*

5 Local Relational Pooling

While deep MPNNs and 2-IGNs are able to aggregate information from multi-hop neighborhoods, our results show that they are unable to preserve information such as the induced-subgraph-counts of nontrivial patterns. To bypass such limitations, we suggest going beyond the strategy of iteratively aggregating information in an equivariant way, which underlies both MPNNs and IGNs. One helpful observation is that, if a pattern is present in the graph, it can always be found in a sufficiently large local neighborhood, or *egonet*, of some node in the graph [50]. An *egonet* of depth l centered at a node i is the induced subgraph consisting of i and all nodes within distance l from it. Note that any pattern with radius r is contained in some *egonet* of depth $l = r$. Hence, we can obtain a model capable of counting patterns by applying a powerful local model to each *egonet* separately and then aggregating the outputs across all *egonets*, as we will introduce below.

For such a local model, we adopt the Relational Pooling (RP) approach from Murphy et al. [45]. In summary, it creates a powerful permutation-invariant model by symmetrizing a powerful permutation-sensitive model, where the symmetrization is performed by averaging or summing over all permutations of the nodes' ordering. Formally, let $\mathbf{B} \in \mathbb{R}^{n \times n \times d}$ be a permutation-sensitive second-order tensor representation of the graph G , such as the one defined in (1). Then, an RP model is defined by

$$f_{\text{RP}}(G) = \frac{1}{|S_n|} \sum_{\pi \in S_n} f(\pi \star \mathbf{B}),$$

where f is some function that is not necessarily permutation-invariant, such as a general multi-layer perceptron (MLP) applied to the vectorization of its tensorial input, S_n is the set of permutations on n nodes, and $\pi \star \mathbf{B}$ is \mathbf{B} transformed by permuting its first two dimensions according to π , i.e., $(\pi \star \mathbf{B})_{j_1, j_2, p} = \mathbf{B}_{\pi(j_1), \pi(j_2), p}$. For choices of f that are sufficiently expressive, such f_{RP} 's are shown to be an universal approximator of permutation-invariant functions [45]. However, the summation quickly becomes intractable once n is large, and hence approximation methods have been introduced. In comparison, since we apply this model to small *egonets*, it is tractable to compute the model exactly. Moreover, as *egonets* are rooted graphs, we can reduce the symmetrization over all permutations in S_n to the subset $S_n^{\text{BFS}} \subseteq S_n$ of permutations which order the nodes in a way that

is compatible with breath-first-search (BFS), as suggested in Murphy et al. [45] to further reduce the complexity. Defining $G_{i,l}^{[\text{ego}]}$ as the egonet centered at node i of depth l , $\mathbf{B}_{i,l}^{[\text{ego}]}$ as the tensor representation of $G_{i,l}^{[\text{ego}]}$ and $n_{i,l}$ as the number of nodes in $G_{i,l}^{[\text{ego}]}$, we consider models of the form

$$f_{\text{LRP}}^l(G) = \sum_{i \in V} H_i, \quad H_i = \frac{1}{|S_{n_{i,l}}^{\text{BFS}}|} \sum_{\pi \in S_{n_{i,l}}^{\text{BFS}}} f\left(\pi \star \mathbf{B}_{i,l}^{[\text{ego}]}\right), \quad (2)$$

To further improve efficiency, we propose to only consider *ordered subsets* of the nodes in each egonet that are compatible with k -truncated-BFS rather than all orderings of the full node set of the egonet, where we define k -truncated-BFS to be a BFS-like procedure that only adds at most k children of every node to the priority queue for future visits and uses zero padding when fewer than k children have not been visited. We let $\tilde{S}_{i,l}^{k\text{-BFS}}$ denote the set of order subsets of the nodes in $G_{i,l}^{[\text{ego}]}$ that are compatible with k -truncated-BFS. Each $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$ can be written as the ordered list $[\tilde{\pi}(1), \dots, \tilde{\pi}(|\tilde{\pi}|)]$, where $|\tilde{\pi}|$ is the length of $\tilde{\pi}$, and for $i \in [|\tilde{\pi}|]$, each $\tilde{\pi}(i)$ is the index of a distinct node in $G_{i,l}^{[\text{ego}]}$. In addition, for each $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$, we introduce a learnable normalization factor, $\alpha_{\tilde{\pi}}$, which can depend on the degrees of the nodes that appear in $\tilde{\pi}$, to adjust for the effect that adding irrelevant edges can alter the fraction of permutations in which a substructure of interest appears. It is a vector whose dimension matches the output dimension of f . More detail on this factor will be given below. Using \odot to denote the element-wise product between vectors, our model becomes

$$f_{\text{LRP}}^{l,k}(G) = \sum_{i \in V} H_i, \quad H_i = \frac{1}{|\tilde{S}_{i,l}^{k\text{-BFS}}|} \sum_{\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}} \alpha_{\tilde{\pi}} \odot f\left(\tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}\right). \quad (3)$$

We call this model depth- l size- k *Local Relational Pooling (LRP- l - k)*. Depending on the task, the summation over all nodes can be replaced by taking average in the definition of $f_{\text{LRP}}^l(G)$. In this work, we choose f to be an MLP applied to the vectorization of its tensorial input. For fixed l and k , if the node degrees are upper-bounded, the time complexity of the model grows linearly in n .

In the experiments below, we focus on two particular variants, where either $l = 1$ or $k = 1$. When $l = 1$, we let $\alpha_{\tilde{\pi}}$ be the output of an MLP applied to the degree of root node i . When $k = 1$, note that each $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$ consists of nodes on a path of length at most $(l + 1)$ starting from node i , and we let $\alpha_{\tilde{\pi}}$ be the output of an MLP applied to the concatenation of the degrees of all nodes on the path. More details on the implementations are discussed in Appendix M.1.

Furthermore, the LRP procedure can be applied iteratively in order to utilize multi-scale information. We define a *Deep LRP- l - k* model of T layers as follows. For $t \in [T]$, we iteratively compute

$$H_i^{(t)} = \frac{1}{|\tilde{S}_{i,l}^{k\text{-BFS}}|} \sum_{\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}} \alpha_{\tilde{\pi}}^{(t)} \odot f^{(t)}\left(\tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}\left(H^{(t-1)}\right)\right), \quad (4)$$

where for an $H \in \mathbb{R}^{n \times d}$, $\mathbf{B}_{i,l}^{[\text{ego}]}(H)$ is the subtensor of $\mathbf{B}(H) \in \mathbb{R}^{n \times n \times d}$ corresponding to the subset of nodes in the egonet $G_{i,l}^{[\text{ego}]}$, and $\mathbf{B}(H)$ is defined by replacing each x_i by H_i in (1). The dimensions of $\alpha_{\tilde{\pi}}^{(t)}$ and the output of $f^{(t)}$ are both $d^{(t)}$, and we set $H_i^{(0)} = x_i$. Finally, we define the graph-level output to be, depending on the task,

$$f_{\text{DLRP}}^{l,k,T}(G) = \sum_{i \in V} H_i^{(T)} \quad \text{or} \quad \frac{1}{|V|} \sum_{i \in V} H_i^{(T)}. \quad (5)$$

The efficiency in practice can be greatly improved by leveraging a pre-computation of the set of maps $H \mapsto \tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}(H)$ as well as sparse tensor operations, which we describe in Appendix K.

6 Experiments

6.1 Counting substructures in random graphs

We first complement our theoretical results with numerical experiments on counting the five substructures illustrated in Figure 3 in synthetic random graphs, including the *subgraph-count* of 3-stars and

the *induced-subgraph-counts* of triangles, tailed triangles, chordal cycles and attributed triangles. By Theorem 3.3 and Corollary 3.4, MPNNs and 2-IGNs cannot exactly solve the *induced-subgraph-count* tasks; while by Theorem 3.5 and Corollary 4.5, they are able to express the *subgraph-count* of 3-stars.




Figure 3: Substructures considered in the experiments: (a) 3-star (b) triangle (c) tailed triangle (d) chordal cycle (e) attributed triangle.

Datasets. We create two datasets of random graphs, one consisting of Erdős-Renyi graphs and the other random regular graphs. Further details on the generation of these datasets are described in Appendix M.2.1. As the target labels, we compute the ground-truth counts of these unattributed and attributed patterns in each graph with a counting algorithm proposed by Shervashidze et al. [56].

Models. We consider LRP, GraphSAGE (using full 1-hop neighborhood) [19], GIN [64], GCN [31], 2-IGN [40], PPGN [39] and spectral GNN (sGNN) [7], with GIN and GCN under the category of MPNNs. Details of the model architectures are given in Appendix M.1. We use mean squared error (MSE) for regression loss. Each model is trained on 1080ti five times with different random seeds.

Results. The results on the *subgraph-count* of 3-stars and the *induced-subgraph-count* of triangles are shown in Table 1, measured by the MSE on the test set divided by the variance of the ground truth counts of the pattern computed over all graphs in the dataset, while the results on the other counting tasks are shown in Appendix M.2.2. Firstly, the almost-negligible errors of LRP on all the tasks support our theory that LRP exploiting only egonets of depth 1 is powerful enough for counting patterns with radius 1. Moreover, GIN, 2-IGN and sGNN yield small error for the 3-star task compared to the variance of the ground truth counts, which is consistent with their theoretical power to perform subgraph-count of star-shaped patterns. Relative to the variance of the ground truth counts, GraphSAGE, GIN and 2-IGN have worse top performance on the triangle task than on the 3-star task, which is also expected from the theory (see Appendix L for a discussion on GraphSAGE). PPGN [39] with provable 3-WL discrimination power also performs well on both counting triangles and 3-stars. Moreover, the results provide interesting insights into the average-case performance in the substructure counting tasks, which are beyond what our theory can predict.

Table 1: Performance of different GNNs on learning the induced-subgraph-count of triangles and the subgraph-count of 3-stars on the two datasets, measured by test MSE divided by variance of the ground truth counts. Shown here are the best and the median (i.e., third-best) performances of each model over five runs with different random seeds. Values below 1E-3 are emboldened. Note that we select the best out of four variants for each of GCN, GIN, sGNN and GraphSAGE, and the better out of two variants for 2-IGN. Details of the GNN architectures and the results on the other counting tasks can be found in Appendices M.1 and M.2.2.

	Erdős-Renyi				Random Regular			
	triangle		3-star		triangle		3-star	
	best	median	best	median	best	median	best	median
GCN	6.78E-1	8.27E-1	4.36E-1	4.55E-1	1.82	2.05	2.63	2.80
GIN	1.23E-1	1.25E-1	1.62E-4	3.44E-4	4.70E-1	4.74E-1	3.73E-4	4.65E-4
GraphSAGE	1.31E-1	1.48E-1	2.40E-10	1.96E-5	3.62E-1	5.21E-1	8.70E-8	4.61E-6
sGNN	9.25E-2	1.13E-1	2.36E-3	7.73E-3	3.92E-1	4.43E-1	2.37E-2	1.41E-1
2-IGN	9.83E-2	9.85E-1	5.40E-4	5.12E-2	2.62E-1	5.96E-1	1.19E-2	3.28E-1
PPGN	5.08E-8	2.51E-7	4.00E-5	6.01E-5	1.40E-6	3.71E-5	8.49E-5	9.50E-5
LRP-1-3	1.56E-4	2.49E-4	2.17E-5	5.23E-5	2.47E-4	3.83E-4	1.88E-6	2.81E-6
Deep LRP-1-3	2.81E-5	4.77E-5	1.12E-5	3.78E-5	1.30E-6	5.16E-6	2.07E-6	4.97E-6

6.2 Molecular prediction tasks

We evaluate LRP on the molecular prediction datasets ogbg-molhiv [63], QM9 [54] and ZINC [15]. More details of the setup can be found in Appendix M.3.

Results. The results on the ogbg-molhiv, QM9 and ZINC are shown in Tables 2 - 4, where for each task or target, the top performance is colored red and the second best colored violet. On ogbg-molhiv, Deep LRP-1-3 with early stopping (see Appendix M.3 for details) achieves higher testing ROC-AUC than the baseline models. On QM9, Deep LRP-1-3 and Deep-LRP-5-1 consistently outperform MPNN and achieve comparable performances with baseline models that are more powerful than

1-WL, including 123-gnn and PPGN. In particular, Deep LRP-5-1 attains the lowest test error on three targets. On ZINC, Deep LRP-7-1 achieves the best performance among the models that do not use feature augmentation (the top baseline model, GateGCN-E-PE, additionally augments the node features with the top Laplacian eigenvectors).

Table 2: Performances on ogbg-molhiv measured by ROC-AUC (%). †: Reported on the OGB leaderboard [21]. ‡: Reported in [22].

Model	Training	Validation	Testing
GIN†	88.64±2.54	82.32±0.90	75.58±1.40
GIN + VN†	92.73±3.80	84.79±0.68	77.07±1.49
GCN†	88.54±2.19	82.04±1.41	76.06±0.97
GCN + VN†	90.07±4.69	83.84±0.91	75.99±1.19
GAT [59]‡	-	-	72.9±1.8
GraphSAGE [19]‡	-	-	74.4±0.7
Deep LRP-1-3	89.81±2.90	81.31±0.88	76.87±1.80
Deep LRP-1-3 (ES)	87.56±2.11	82.09±1.16	77.19±1.40

Table 3: Performances on ZINC measured by the Mean Absolute Error (MAE). †: Reported in Dwivedi et al. (2020).

Model	Training	Testing	Time / Ep
GraphSAGE†	0.081 ± 0.009	0.398 ± 0.002	16.61s
GIN†	0.319 ± 0.015	0.387 ± 0.015	2.29s
MoNet†	0.093 ± 0.014	0.292 ± 0.006	10.82s
GatedGCN-E†	0.074 ± 0.016	0.282 ± 0.015	20.50s
GatedGCN-E-PE†	0.067 ± 0.019	0.214 ± 0.013	10.70s
PPGN†	0.140 ± 0.044	0.256 ± 0.054	334.69s
Deep LRP-7-1	0.028 ± 0.004	0.223 ± 0.008	72s
Deep LRP-5-1	0.020 ± 0.006	0.256 ± 0.033	42s

Table 4: Performances on QM9 measured by the testing Mean Absolute Error. All baseline results are from [39], including DTNN [63] and 123-gnn [44]. The loss value on the last row is defined in Appendix M.3.

Target	DTNN	MPNN	123-gnn	PPGN	Deep LRP-1-3	Deep LRP-5-1
μ	0.244	0.358	0.476	0.231	0.399	0.364
α	0.95	0.89	0.27	0.382	0.337	0.298
ϵ_{homo}	0.00388	0.00541	0.00337	0.00276	0.00287	0.00254
ϵ_{lumo}	0.00512	0.00623	0.00351	0.00287	0.00309	0.00277
Δ_ϵ	0.0112	0.0066	0.0048	0.00406	0.00396	0.00353
$\langle R^2 \rangle$	17	28.5	22.9	16.07	20.4	19.3
ZPVE	0.00172	0.00216	0.00019	0.00064	0.00067	0.00055
U_0	2.43	2.05	0.0427	0.234	0.590	0.413
U	2.43	2	0.111	0.234	0.588	0.413
H	2.43	2.02	0.0419	0.229	0.587	0.413
G	2.43	2.02	0.0469	0.238	0.591	0.413
C_v	0.27	0.42	0.0944	0.184	0.149	0.129
Loss	0.1014	0.1108	0.0657	0.0512	0.0641	0.0567

7 Conclusions

We propose a theoretical framework to study the expressive power of classes of GNNs based on their ability to count substructures. We distinguish two kinds of counting: subgraph-count and induced-subgraph-count. We prove that neither MPNNs nor 2-IGNs can induced-subgraph-count any connected structure with 3 or more nodes; k -IGNs and k -WL can subgraph-count and induced-subgraph-count any pattern of size k . We also provide an upper bound on the size of “path-shaped” substructures that finite iterations of k -WL can induced-subgraph-count. To establish these results, we prove an equivalence between approximating graph functions and discriminating graphs. Also, as intermediary results, we prove that MPNNs are no more powerful than 2-WL on attributed graphs, and that 2-IGNs are equivalent to 2-WL in distinguishing non-isomorphic graphs, which partly answers an open problem raised in Maron et al. [38]. In addition, we perform numerical experiments that support our theoretical results and show that the Local Relational Pooling approach inspired by Murphy et al. [45] can successfully count certain substructures. In summary, we build the foundation for using substructure counting as an intuitive and relevant measure of the expressive power of GNNs, and our concrete results for existing GNNs motivate the search for more powerful designs of GNNs.

One limitation of our theory is that it is only concerned with the expressive power of GNNs and not their optimization or generalization. Our theoretical results are also worst-case in nature and cannot predict average-case performance. Many interesting theoretical questions remain, including better characterizing the ability to count substructures of general k -WL and k -IGNs as well as other architectures such as spectral GNNs [7] and polynomial IGNs [38]. On the practical side, we hope our framework can help guide the search for more powerful GNNs by considering substructure counting as a criterion. It will be interesting to quantify the relevance of substructure counting in empirical tasks, perhaps following the work of Ying et al. [67], and also to consider tasks where substructure counting is explicitly relevant, such as subgraph matching [36].

Broader impact

In this work we propose to understand the power of GNN architectures via the substructures that they can and cannot count. Our work is motivated by the relevance of detecting and counting *graph substructures* in applications, and the current trend on using deep learning – in particular, graph neural networks – in such scientific fields. The ability of different GNN architectures to count graph substructures not only serves as an intuitive theoretical measure of their expressive power but also is highly relevant to real-world scenarios. Our results show that some widely used GNN architectures are not able to count substructures. Such knowledge may indicate that some widely-used graph neural network architectures are actually not the right tool for certain scientific problems. On the other hand, we propose a GNN model that not only has the ability to count substructures but also can learn from data what the relevant substructures are.

Acknowledgements

We are grateful to Haggai Maron, Jiaxuan You, Ryoma Sato and Christopher Morris for helpful conversations. This work is partially supported by the Alfred P. Sloan Foundation, NSF RI-1816753, NSF CAREER CIF 1845360, NSF CHS-1901091, Samsung Electronics, and the Institute for Advanced Study. SV is supported by NSF DMS 2044349, EOARD FA9550-18-1-7007, and the NSF-Simons Research Collaboration on the Mathematical and Scientific Foundations of Deep Learning (MoDL) (NSF DMS 2031985).

References

- [1] Alon, N., Dao, P., Hajirasouliha, I., Hormozdiari, F., and Sahinalp, S. C. (2008). Biomolecular network motif counting and discovery by color coding. *Bioinformatics*, 24(13):i241–i249.
- [2] Arvind, V., Fuhlbrück, F., Köbler, J., and Verbitsky, O. (2018). On weisfeiler-leman invariance: Subgraph counts and related graph properties. *arXiv preprint arXiv:1811.04801*.
- [3] Babai, L., Erdos, P., and Selkow, S. M. (1980). Random graph isomorphism. *SIAM Journal on computing*, 9(3):628–635.
- [4] Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42.
- [5] Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. (2013). Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*.
- [6] Cai, J.-Y., Fürer, M., and Immerman, N. (1992). An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410.
- [7] Chen, Z., Li, L., and Bruna, J. (2019a). Supervised community detection with line graph neural networks. *International Conference on Learning Representations*.
- [8] Chen, Z., Villar, S., Chen, L., and Bruna, J. (2019b). On the equivalence between graph isomorphism testing and function approximation with gnns. In *Advances in Neural Information Processing Systems*, pages 15868–15876.
- [9] Choma, N., Monti, F., Gerhardt, L., Palczewski, T., Ronaghi, Z., Prabhat, P., Bhimji, W., Bronstein, M., Klein, S., and Bruna, J. (2018). Graph neural networks for icecube signal classification. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 386–391. IEEE.
- [10] Dai, H., Khalil, E. B., Zhang, Y., Dilkina, B., and Song, L. (2017). Learning combinatorial optimization algorithms over graphs. *arXiv preprint arXiv: 1704.01665*.
- [11] Defferrard, M., Bresson, X., and Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852.
- [12] Deshpande, M., Kuramochi, M., and Karypis, G. (2002). Automated approaches for classifying structures. Technical report, Minnesota University Minneapolis Department of Computer Science.
- [13] Ding, M., Zhou, C., Chen, Q., Yang, H., and Tang, J. (2019). Cognitive graph for multi-hop reading comprehension at scale. In *Proceedings of the 57th Annual Meeting of the Association*

- for *Computational Linguistics*, pages 2694–2703, Florence, Italy. Association for Computational Linguistics.
- [14] Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. (2015). Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pages 2224–2232.
- [15] Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., and Bresson, X. (2020). Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982*.
- [16] Fürer, M. (2017). On the combinatorial power of the Weisfeiler-Lehman algorithm. *arXiv preprint arXiv:1704.01023*.
- [17] Garg, V. K., Jegelka, S., and Jaakkola, T. (2020). Generalization and representational limits of graph neural networks.
- [18] Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org.
- [19] Hamilton, W., Ying, Z., and Leskovec, J. (2017). Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034.
- [20] Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- [21] Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. (2020). Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*.
- [22] Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. (2019). Strategies for pre-training graph neural networks. In *International Conference on Learning Representations*.
- [23] Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- [24] Irwin, J. J., Sterling, T., Mysinger, M. M., Bolstad, E. S., and Coleman, R. G. (2012). Zinc: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):1757–1768.
- [25] Jiang, C., Coenen, F., and Zito, M. (2010). Finding frequent subgraphs in longitudinal social network data using a weighted graph mining approach. In *International Conference on Advanced Data Mining and Applications*, pages 405–416. Springer.
- [26] Jin, W., Barzilay, R., and Jaakkola, T. (2019). Hierarchical graph-to-graph translation for molecules.
- [27] Jin, W., Barzilay, R., and Jaakkola, T. (2020). Composing molecules with multiple property constraints. *arXiv preprint arXiv:2002.03244*.
- [28] Jin, W., Barzilay, R., and Jaakkola, T. S. (2018). Junction tree variational autoencoder for molecular graph generation. *CoRR*, abs/1802.04364.
- [29] Keriven, N. and Peyré, G. (2019). Universal invariant and equivariant graph neural networks. *arXiv preprint arXiv:1905.04943*.
- [30] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [31] Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- [32] Koyutürk, M., Grama, A., and Szpankowski, W. (2004). An efficient algorithm for detecting frequent subgraphs in biological networks. *Bioinformatics*, 20(suppl 1):i200–i207.
- [33] Lemke, T. L. (2003). *Review of organic functional groups: introduction to medicinal organic chemistry*. Lippincott Williams & Wilkins.
- [34] Liu, S., Chandereng, T., and Liang, Y. (2018). N-gram graph, A novel molecule representation. *arXiv preprint arXiv:1806.09206*.
- [35] Liu, X., Pan, H., He, M., Song, Y., and Jiang, X. (2019). Neural subgraph isomorphism counting.

- [36] Lou, Z., You, J., Wen, C., Canedo, A., Leskovec, J., et al. (2020). Neural subgraph matching. *arXiv preprint arXiv:2007.03092*.
- [37] Loukas, A. (2019). What graph neural networks cannot learn: depth vs width. *arXiv preprint arXiv:1907.03199*.
- [38] Maron, H., Ben-Hamu, H., and Lipman, Y. (2019a). Open problems: Approximation power of invariant graph networks.
- [39] Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y. (2019b). Provably powerful graph networks. In *Advances in Neural Information Processing Systems*, pages 2153–2164.
- [40] Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. (2018). Invariant and equivariant graph networks.
- [41] Maron, H., Fetaya, E., Segol, N., and Lipman, Y. (2019c). On the universality of invariant networks. *arXiv preprint arXiv:1901.09342*.
- [42] Monti, F., Otness, K., and Bronstein, M. M. (2018). Motifnet: a motif-based graph convolutional network for directed graphs. *CoRR*, abs/1802.01572.
- [43] Morgan, H. L. (1965). The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. *Journal of Chemical Documentation*, 5(2):107–113.
- [44] Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen, J. E., Rattan, G., and Grohe, M. (2019). Weisfeiler and leman go neural: Higher-order graph neural networks. *Association for the Advancement of Artificial Intelligence*.
- [45] Murphy, R. L., Srinivasan, B., Rao, V., and Ribeiro, B. (2019). Relational pooling for graph representations. *arXiv preprint arXiv:1903.02541*.
- [46] Murray, C. W. and Rees, D. C. (2009). The rise of fragment-based drug discovery. *Nature chemistry*, 1(3):187.
- [47] Nowak, A., Villar, S., Bandeira, A. S., and Bruna, J. (2017). A note on learning algorithms for quadratic assignment with graph neural networks. *arXiv preprint arXiv:1706.07450*.
- [48] O’Boyle, N. M. and Sayle, R. A. (2016). Comparing structural fingerprints using a literature-based similarity benchmark. *Journal of cheminformatics*, 8(1):1–14.
- [49] Pope, P., Kolouri, S., Rostrami, M., Martin, C., and Hoffmann, H. (2018). Discovering molecular functional groups using graph convolutional neural networks. *arXiv preprint arXiv:1812.00265*.
- [50] Preciado, V. M., Draief, M., and Jadbabaie, A. (2012). Structural analysis of viral spreading processes in social and communication networks using egonets.
- [51] Preciado, V. M. and Jadbabaie, A. (2010). From local measurements to network spectral properties: Beyond degree distributions. In *49th IEEE Conference on Decision and Control (CDC)*, pages 2686–2691. IEEE.
- [52] Puny, O., Ben-Hamu, H., and Lipman, Y. (2020). From graph low-rank global attention to 2-fwl approximation. *arXiv preprint arXiv:2006.07846*.
- [53] Rahman, S. A., Bashton, M., Holliday, G. L., Schrader, R., and Thornton, J. M. (2009). Small molecule subgraph detector (smsd) toolkit. *Journal of cheminformatics*, 1(1):12.
- [54] Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. (2014). Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1:140022.
- [55] Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2008). The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80.
- [56] Shervashidze, N., Vishwanathan, S., Petri, T., Mehlhorn, K., and Borgwardt, K. (2009). Efficient graphlet kernels for large graph comparison. In *Artificial Intelligence and Statistics*, pages 488–495.
- [57] Steger, A. and Wormald, N. C. (1999). Generating random regular graphs quickly. *Combinatorics, Probability and Computing*, 8(4):377–396.
- [58] Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackerman, Z., et al. (2020). A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702.

- [59] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2017). Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- [60] Weisfeiler, B. and Leman, A. (1968). The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Tekhnicheskaya Informatsia*, 2(9):12-16.
- [61] Weisstein, E. W. (2020). Quartic graph. From MathWorld—A Wolfram Web Resource. <https://mathworld.wolfram.com/QuarticGraph.html>.
- [62] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. (2019). A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*.
- [63] Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. (2018). Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530.
- [64] Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2018a). How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- [65] Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.-i., and Jegelka, S. (2018b). Representation learning on graphs with jumping knowledge networks. *arXiv preprint arXiv:1806.03536*.
- [66] Yao, W., Bandeira, A. S., and Villar, S. (2019). Experimental performance of graph neural networks on random instances of max-cut. In *Wavelets and Sparsity XVIII*, volume 11138, page 111380S. International Society for Optics and Photonics.
- [67] Ying, R., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. (2019). Gnn explainer: A tool for post-hoc explanation of graph neural networks. *arXiv preprint arXiv:1903.03894*.
- [68] Ying, R., You, J., Morris, C., Ren, X., Hamilton, W. L., and Leskovec, J. (2018). Hierarchical graph representation learning with differentiable pooling. *CoRR*, abs/1806.08804.
- [69] You, J., Liu, B., Ying, Z., Pande, V., and Leskovec, J. (2018a). Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in neural information processing systems*, pages 6410–6421.
- [70] You, J., Wu, H., Barrett, C., Ramanujan, R., and Leskovec, J. (2019). G2sat: Learning to generate sat formulas. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 10553–10564. Curran Associates, Inc.
- [71] You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. (2018b). Graphrnn: A deep generative model for graphs. *CoRR*, abs/1802.08773.
- [72] Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. (2017). Deep sets. In *Advances in neural information processing systems*, pages 3391–3401.
- [73] Zhang, M. and Chen, Y. (2018). Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pages 5165–5175.
- [74] Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M. (2018). Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*.

A Function approximation perspective of substructure counting

On a space of graphs \mathcal{G} , we call $C_I(\cdot; G^{[P]})$ the *induced-subgraph-count function* of the pattern $G^{[P]}$, and $C_S(\cdot; G^{[P]})$ the *subgraph-count function* of $G^{[P]}$. To formalize the probe into whether certain GNN architectures can count different substructures, a natural question to study is whether they are able to approximate the induced-subgraph-count and the subgraph-count functions arbitrarily well. Formally, given a target function $g : \mathcal{G} \rightarrow \mathbb{R}$, and family of functions, \mathcal{F} , which in our case is typically the family of functions that a GNN architecture can represent, we say \mathcal{F} is able to approximate g on \mathcal{G} if for all $\epsilon > 0$ there exists $f \in \mathcal{F}$ such that $|g(G) - f(G)| < \epsilon$, for all $G \in \mathcal{G}$.

However, such criterion based on function approximation is hard to work with directly when we look at concrete examples later on. For this reason, below we will look for an alternative and equivalent definition from the perspective of graph discrimination.

A.1 From function approximation to graph discrimination

Say \mathcal{G} is a space of graphs, and \mathcal{F} is a family of functions from \mathcal{G} to \mathbb{R} . Given two graphs $G^{[1]}, G^{[2]} \in \mathcal{G}$, we say \mathcal{F} is able to distinguish them if there exists $f \in \mathcal{F}$ such that $f(G^{[1]}) \neq f(G^{[2]})$. Such a perspective has been explored in Chen et al. [8], for instance, to build an equivalence between function approximation and graph isomorphism testing by GNNs. In the context of substructure counting, it is clear that the ability to approximate the count functions entails the ability to distinguish graphs in the following sense:

Observation 1. *If \mathcal{F} is able to approximate the induced-subgraph-count (or subgraph-count) function of a pattern $G^{[P]}$ on the space \mathcal{G} , then for all $G^{[1]}, G^{[2]} \in \mathcal{G}$ such that $C_I(G^{[1]}, G^{[P]}) \neq C_I(G^{[2]}, G^{[P]})$ (or $C_S(G^{[1]}, G^{[P]}) \neq C_S(G^{[2]}, G^{[P]})$), they can be distinguished by \mathcal{F} .*

What about the converse? When the space \mathcal{G} is finite, such as if the graphs have bounded numbers of nodes and the node as well as edge features belong to finite alphabets, we can show a slightly weaker statement than the exact converse. Following Chen et al. [8], we define an augmentation of families of functions using feed-forward neural networks as follows:

Definition A.1. *Given \mathcal{F} , a family of functions from a space \mathcal{X} to \mathbb{R} , we consider an augmented family of functions also from \mathcal{X} to \mathbb{R} consisting of all functions of the following form*

$$x \mapsto h_{\mathcal{NN}}([f_1(x), \dots, f_d(x)]),$$

where $d \in \mathbb{N}$, $f_1, \dots, f_d \in \mathcal{F}$, and $h_{\mathcal{NN}}$ is a feed-forward neural network / multi-layer perceptron. When \mathcal{NN} is restricted to have L layers at most, we denote this augmented family by \mathcal{F}^{+L} .

Lemma A.2. *Suppose \mathcal{X} is a finite space, g is a finite function on \mathcal{X} , and \mathcal{F} is a family of functions on \mathcal{X} . Then, \mathcal{F}^{+1} is able to approximate f on \mathcal{G} if $\forall x_1, x_2 \in \mathcal{X}$ with $g(x_1) \neq g(x_2)$, $\exists f \in \mathcal{F}$ such that $f(x_1) \neq f(x_2)$.*

Proof. Since \mathcal{X} is a finite space, for some large enough integer d , \exists a collection of d functions, $f_1, \dots, f_d \in \mathcal{F}$ such that, if we define the function $\mathbf{f}(x) = (f_1(x), \dots, f_d(x)) \in \mathbb{R}^d$, then it holds that $\forall x_1, x_2 \in \mathcal{X}, \mathbf{f}(x_1) = \mathbf{f}(x_2) \Rightarrow g(x_1) = g(x_2)$. (In fact, we can choose $d \leq \frac{|\mathcal{X}| \cdot (|\mathcal{X}| - 1)}{2}$, since in the worst case we need one f_i per pair of $x_1, x_2 \in \mathcal{X}$ with $x_1 \neq x_2$.) Then, \exists a well-defined function h from \mathbb{R}^d to \mathbb{R} such that $\forall x \in \mathcal{X}, g(x) = h(\mathbf{f}(x))$. By the universal approximation power of neural networks, h can then be approximated arbitrarily well by some neural network $h_{\mathcal{NN}}$. \square

Thus, in the context of substructure counting, we have the following observation.

Observation 2. *Suppose \mathcal{G} is a finite space. If $\forall G^{[1]}, G^{[2]} \in \mathcal{G}$ with $C_I(G^{[1]}, G^{[P]}) \neq C_I(G^{[2]}, G^{[P]})$ (or $C_S(G^{[1]}, G^{[P]}) \neq C_S(G^{[2]}, G^{[P]})$), \mathcal{F} is able to distinguish $G^{[1]}$ and $G^{[2]}$, then \mathcal{F}^{+1} is able to approximate the induced-subgraph-count (or subgraph-count) function of the pattern $G^{[P]}$ on \mathcal{G} .*

For many GNN families, \mathcal{F}^{+1} in fact has the same expressive power as \mathcal{F} . For example, consider $\mathcal{F}_{\text{MPNN}}$, the family of all Message Passing Neural Networks on \mathcal{G} . $\mathcal{F}_{\text{MPNN}}^{+1}$ consists of functions that run several MPNNs on the input graph in parallel and stack their outputs to pass through an MLP. However, running several MPNNs in parallel is equivalent to running one MPNN with larger dimensions of hidden states and messages, and moreover the additional MLP at the end can be merged into the readout function. Similar holds for the family of all k -Invariant Graph Functions (k -IGNs). Hence, for such GNN families, we have an exact equivalence on finite graph spaces \mathcal{G} .

B Additional notations

For two positive integers a and b , we define $\text{MOD}_a(b)$ to be a if a divides b and the number c such that $b \equiv c \pmod{a}$ otherwise. Hence the value ranges from 1 to a as we vary $b \in \mathbb{N}^*$.

For a positive integer c , let $[c]$ denote the set $\{1, \dots, c\}$.

Two k -tuples, $(i_1, \dots, i_k), (j_1, \dots, j_k) \in V^k$ are said to be in the same *equivalent class* if \exists a permutation π on V such that $(\pi(i_1), \dots, \pi(i_k)) = (j_1, \dots, j_k)$. Note that belonging to the same equivalence class is a weaker condition than having the same isomorphism type, as will be defined in Appendix C, which has to do with what the graphs look like.

For any k -tuple, $s = (i_1, \dots, i_k)$, and for $w \in [k]$, use $I_w(s)$ to denote the w th entry of s , i_w .

C Definition of k -WL on attributed graphs

In this section, we introduce the general k -WL test for $k \in \mathbb{N}^*$ applied to a pair of graphs, $G^{[1]}$ and $G^{[2]}$. Assume that the two graphs have the same number of vertices, since otherwise they can be told apart easily. Without loss of generality, we assume that they share the same set of vertex indices, V (but can differ in E , x or e). For each of the graphs, at iteration 0, the test assigns an initial color in some color space to every k -tuple in V^k according to its isomorphism type (we define isomorphism types rigorously in Section C.1), and then updates the coloring in every iteration. For any k -tuple $s = (i_1, \dots, i_k) \in V^k$, we let $\mathbf{c}_k^{(t)}(s)$ denote the color of s in $G^{[1]}$ assigned at t th iteration, and let $\mathbf{c}'_k^{(t)}(s)$ denote the color it receives in $G^{[2]}$. $\mathbf{c}_k^{(t)}(s)$ and $\mathbf{c}'_k^{(t)}(s)$ are updated iteratively as follows. For each $w \in [k]$, define the neighborhood

$$N_w(s) = \{(i_1, \dots, i_{w-1}, j, i_{w+1}, \dots, i_k) : j \in V\}$$

Given $\mathbf{c}_k^{(t-1)}$ and $\mathbf{c}'_k^{(t-1)}$, define

$$C_w^{(t)}(s) = \text{HASH}_{t,1} \left(\{ \mathbf{c}_k^{(t-1)}(\tilde{s}) : \tilde{s} \in N_w(s) \} \right)$$

$$C'_w{}^{(t)}(s) = \text{HASH}_{t,1} \left(\{ \mathbf{c}'_k{}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_w(s) \} \right)$$

with “ $\{ \}$ ” representing a multiset, and $\text{HASH}_{t,1}$ being some hash function that maps injectively from the space of multisets of colors to some intermediate space. Then let

$$\begin{aligned} \mathbf{c}_k^{(t)}(s) &= \text{HASH}_{t,2} \left(\left(\mathbf{c}_k^{(t-1)}(s), \left(C_1^{(t)}(s), \dots, C_k^{(t)}(s) \right) \right) \right) \\ \mathbf{c}'_k{}^{(t)}(s) &= \text{HASH}_{t,2} \left(\left(\mathbf{c}'_k{}^{(t-1)}(s), \left(C'_1{}^{(t)}(s), \dots, C'_k{}^{(t)}(s) \right) \right) \right) \end{aligned}$$

where $\text{HASH}_{t,2}$ maps injectively from its input space to the space of colors. The test will terminate and return the result that the two graphs are not isomorphic if at some iteration t , the following two multisets differ:

$$\{ \mathbf{c}_k^{(t)}(s) : s \in V^k \} \neq \{ \mathbf{c}'_k{}^{(t)}(s) : s \in V^k \}$$

C.1 Isomorphism types of k -tuples in k -WL for attributed graphs

Say $G^{[1]} = (V^{[1]}, E^{[1]}, x^{[1]}, e^{[1]})$, $G^{[2]} = (V^{[2]}, E^{[2]}, x^{[2]}, e^{[2]})$.

a) $\forall s = (i_1, \dots, i_k), s' = (i'_1, \dots, i'_k) \in (V^{[1]})^k$, s and s' are said to have the same isomorphism type if

1. $\forall \alpha, \beta \in [k], i_\alpha = i_\beta \Leftrightarrow i'_\alpha = i'_\beta$
2. $\forall \alpha \in [k], x_{i_\alpha}^{[1]} = x_{i'_\alpha}^{[1]}$

3. $\forall \alpha, \beta \in [k], (i_\alpha, i_\beta) \in E^{[1]} \Leftrightarrow (i'_\alpha, i'_\beta) \in E^{[1]}$, and moreover, if either side is true, then $e_{i_\alpha, i_\beta}^{[1]} = e_{i'_\alpha, i'_\beta}^{[1]}$

b) Similar if both $s, s' \in (V^{[2]})^k$.

c) $\forall s = (i_1, \dots, i_k) \in (V^{[1]})^k, s' = (i'_1, \dots, i'_k) \in (V^{[2]})^k$, s and s' are said to have the same isomorphism type if

1. $\forall \alpha, \beta \in [k], i_\alpha = i_\beta \Leftrightarrow i'_\alpha = i'_\beta$
2. $\forall \alpha \in [k], x_{i_\alpha}^{[1]} = x_{i'_\alpha}^{[2]}$
3. $\forall \alpha, \beta \in [k], (i_\alpha, i_\beta) \in E^{[1]} \Leftrightarrow (i'_\alpha, i'_\beta) \in E^{[2]}$, and moreover, if either side is true, then $e_{i_\alpha, i_\beta}^{[1]} = e_{i'_\alpha, i'_\beta}^{[2]}$

In k -WL tests, two k -tuples s and s' in either $(V^{[1]})^k$ or $(V^{[2]})^k$ are assigned the same color at iteration 0 if and only if they have the same isomorphism type.

For a reference, see Maron et al. [39].

D Proof of Theorem 3.2 (MPNNs are no more powerful than 2-WL)

Proof. Suppose for contradiction that there exists an MPNN with T_0 layers that can distinguish the two graphs. Let $m^{(t)}$ and $h^{(t)}$, $m'^{(t)}$ and $h'^{(t)}$ be the messages and hidden states at layer t obtained by applying the MPNN on the two graphs, respectively. Define

$$\tilde{h}_{i,j}^{(t)} = \begin{cases} h_i^{(t)} & \text{if } i = j \\ (h_i^{(t)}, h_j^{(t)}, a_{i,j}, e_{i,j}) & \text{otherwise} \end{cases}$$

$$\tilde{h}'_{i,j}^{(t)} = \begin{cases} h'_i{}^{(t)} & \text{if } i = j \\ (h'_i{}^{(t)}, h'_j{}^{(t)}, a'_{i,j}, e'_{i,j}) & \text{otherwise,} \end{cases}$$

where $a_{i,j} = 1$ if $(i, j) \in E^{[1]}$ and 0 otherwise, $e_{i,j} = e_{i,j}^{[1]}$ is the edge feature of the first graph, and a', e' are defined similarly for the second graph.

Since the two graphs cannot be distinguished by 2-WL, then for the T_0 th iteration, there is

$$\{\mathbf{c}_2^{(T_0)}(s) : s \in V^2\} = \{\mathbf{c}'_2{}^{(T_0)}(s) : s \in V^2\},$$

which implies that there exists a permutation on V^2 , which we can call η_0 , such that $\forall s \in V^2$, there is $\mathbf{c}_2^{(T_0)}(s) = \mathbf{c}'_2{}^{(T_0)}(\eta_0(s))$. To take advantage of this condition, we introduce the following lemma, which is central to the proof.

Lemma D.1. $\forall t \leq T_0, \forall i, j, i', j' \in V$, if $\mathbf{c}_2^{(t)}((i, j)) = \mathbf{c}'_2{}^{(t)}((i', j'))$, then

1. $i = j \Leftrightarrow i' = j'$.
2. $\tilde{h}_{i,j}^{(t)} = \tilde{h}'_{i',j'}{}^{(t)}$

Proof of Lemma D.1: First, we state the following simple observation without proof, which is immediate given the update rule of k -WL:

Lemma D.2. For k -WL, $\forall s, s' \in V^k$, if for some t_0 , $\mathbf{c}_k^{(t_0)}(s) = \mathbf{c}'_k{}^{(t_0)}(s')$, then $\forall t \in [0, t_0]$, $\mathbf{c}_k^{(t)}(s) = \mathbf{c}'_k{}^{(t)}(s')$.

For the first condition, assuming $\mathbf{c}_2^{(t)}((i, j)) = \mathbf{c}'_2{}^{(t)}((i', j'))$, Lemma D.2 then tells us that $\mathbf{c}_2^{(0)}((i, j)) = \mathbf{c}'_2{}^{(0)}((i', j'))$. Since the colors in 2-WL are initialized by the isomorphism type of the node pair, it has to be that $i = j \Leftrightarrow i' = j'$.

We will prove the second condition by induction on t . For the base case, $t = 0$, we want to show that $\forall i, j, i', j' \in V$, if $\mathbf{c}_2^{(0)}((i, j)) = \mathbf{c}'_2{}^{(0)}((i', j'))$ then $\tilde{h}_{i,j}^{(0)} = \tilde{h}'_{i',j'}{}^{(0)}$. If $i = j$, then $\mathbf{c}_2^{(0)}((i, i)) = \mathbf{c}'_2{}^{(0)}((i', i'))$ if and only if $x_i = x_{i'}$, which is equivalent to $h_i^{(0)} = h_{i'}^{(0)}$, and hence $\tilde{h}_i^{(0)} = \tilde{h}'_{i'}{}^{(0)}$. If $i \neq j$, then by the definition of isomorphism types given in Appendix C, $\mathbf{c}_2^{(0)}((i, j)) = \mathbf{c}'_2{}^{(0)}((i', j'))$ implies that

$$\begin{aligned} x_i = x_{i'} &\Rightarrow h_i^{(0)} = h_{i'}^{(0)} \\ x_j = x_{j'} &\Rightarrow h_j^{(0)} = h_{j'}^{(0)} \\ a_{i,j} &= a'_{i',j'} \\ e_{i,j} &= e'_{i',j'} \end{aligned}$$

which yields $\tilde{h}_{i,j}^{(0)} = \tilde{h}'_{i',j'}{}^{(0)}$.

Next, to prove the inductive step, assume that for some $T \in [T_0]$, the statement in Lemma D.1 holds for all $t \leq T - 1$, and consider $\forall i, j, i', j' \in V$ such that $\mathbf{c}_2^{(T)}((i, j)) = \mathbf{c}'_2{}^{(T)}((i', j'))$. By the update rule of 2-WL, this implies that

$$\begin{aligned} \mathbf{c}_2^{(T-1)}((i, j)) &= \mathbf{c}'_2{}^{(T-1)}((i', j')) \\ \{\mathbf{c}_2^{(T-1)}((k, j)) : k \in V\} &= \{\mathbf{c}'_2{}^{(T-1)}((k, j')) : k \in V\} \\ \{\mathbf{c}_2^{(T-1)}((i, k)) : k \in V\} &= \{\mathbf{c}'_2{}^{(T-1)}((i', k)) : k \in V\} \end{aligned} \quad (6)$$

The first condition, thanks to the inductive hypothesis, implies that $\tilde{h}_{i,j}^{(T-1)} = \tilde{h}'_{i',j'}{}^{(T-1)}$. In particular, if $i \neq j$, then we have

$$\begin{aligned} a_{i,j} &= a'_{i',j'} \\ e_{i,j} &= e'_{i',j'} \end{aligned} \quad (7)$$

The third condition implies that \exists a permutation on V , which we can call $\xi_{i,i'}$, such that $\forall k \in V$,

$$\mathbf{c}_2^{(T-1)}((i, k)) = \mathbf{c}'_2{}^{(T-1)}((i', \xi_{i,i'}(k)))$$

By the inductive hypothesis, there is $\forall k \in V$,

$$\tilde{h}_{i,k}^{(T-1)} = \tilde{h}'_{i',\xi_{i,i'}(k)}{}^{(T-1)}$$

and moreover, $\xi_{i,i'}(k) = i'$ if and only if $k = i$. For $k \neq i$, we thus have

$$\begin{aligned} h_i^{(T-1)} &= h_{i'}^{(T-1)} \\ h_k^{(T-1)} &= h_{\xi_{i,i'}(k)}^{(T-1)} \\ a_{i,k} &= a'_{i',\xi_{i,i'}(k)} \\ e_{i,k} &= e'_{i',\xi_{i,i'}(k)} \end{aligned}$$

Now, looking at the update rule at the T th layer of the MPNN,

$$\begin{aligned}
m_i^{(T)} &= \sum_{k \in \mathcal{N}(i)} M_T(h_i^{(T-1)}, h_k^{(T-1)}, e_{i,k}) \\
&= \sum_{k \in V} a_{i,k} \cdot M_T(h_i^{(T-1)}, h_k^{(T-1)}, e_{i,k}) \\
&= \sum_{k \in V} a'_{i', \xi_{i,i'}(k)} \cdot M_T(h'_{i'}^{(T-1)}, h'_{\xi_{i,i'}(k)}^{(T-1)}, e'_{i', \xi_{i,i'}(k)}) \\
&= \sum_{k' \in V} a'_{i', k'} \cdot M_T(h'_{i'}^{(T-1)}, h'_{k'}^{(T-1)}, e'_{i', k'}) \\
&= m'_{i'}^{(T)}
\end{aligned}$$

where between the third and the fourth line we made the substitution $k' = \xi_{i,i'}(k)$. Therefore,

$$\begin{aligned}
h_i^{(T)} &= U_t(h_i^{(T-1)}, m_i^{(T)}) \\
&= U_t(h'_{i'}^{(T-1)}, m'_{i'}^{(T)}) \\
&= h'_{i'}^{(T)}
\end{aligned}$$

By the symmetry between i and j , we can also show that $h_j^{(T)} = h'_{j'}^{(T)}$. Hence, together with 7, we can conclude that

$$\tilde{h}_{i,j}^{(T)} = \tilde{h}'_{i',j'}^{(T)},$$

which proves the lemma. \square

Thus, the second result of this lemma tells us that $\forall i, j \in V^2$, $\tilde{h}_{i,j}^{(T_0)} = \tilde{h}'_{\eta_0(i,j)}^{(T_0)}$. Moreover, by the first result, \exists a permutation on V , which we can call τ_0 , such that $\forall i \in V$, $\eta((i, i)) = (\tau_0(i), \tau_0(i))$. Combining the two, we have that $\forall i \in V$, $h_i^{(T_0)} = h'_{\tau_0(i)}^{(T_0)}$, and hence

$$\{h_i^{(T_0)} : i \in V\} = \{h'_{i'}^{(T_0)} : i' \in V\} \quad (8)$$

Therefore, $\hat{y} = \hat{y}'$, meaning that the MPNN returns identical outputs on the two graphs. \square

E Proof of Theorem 3.3 (2-WL is unable to induced-subgraph-count patterns of 3 or more nodes)

Proof Intuition. Given any connected pattern of at least 3 nodes, such as the one in the left of Figure 2, we can construct a pair of graphs, such as the pair in the center and the right of Figure 2. They that have different induced-subgraph-counts of the pattern, and we can show that 2-WL cannot distinguish them. but cannot be distinguished from each other by 2-WL. For instance, if we run 2-WL on the pair of graphs in Figure 2, then there will be $c_2^{(t)}((1, 3)) = c'^{(t)}_2((1, 3))$, $c_2^{(t)}((1, 2)) = c'^{(t)}_2((1, 6))$, $c_2^{(t)}((1, 6)) = c'^{(t)}_2((1, 2))$, and so on. We can in fact show that $\{c_2^{(t)}(s) : s \in V^2\} = \{c'^{(t)}_2(s) : s \in V^2\}$, $\forall t$, which implies that 2-WL cannot distinguish the two graphs.

Proof. Say $G^{[P]} = (V^{[P]}, E^{[P]}, x^{[P]}, e^{[P]})$ is a connected pattern of m nodes, where $m > 2$, and thus $V^{[P]} = [m]$.

First, if $G^{[P]}$ is not a clique, then by definition, there exists two distinct nodes $i, j \in V^{[P]}$ such that i and j are not connected by an edge. Assume without loss of generality that $i = 1$ and $j = 2$. Now, construct two graphs $G^{[1]} = (V = [2m], E^{[1]}, x^{[1]}, e^{[1]})$, $G^{[2]} = (V = [2m], E^{[2]}, x^{[2]}, e^{[2]})$ both with $2m$ nodes. For $G^{[1]}$, let $E^{[1]} = \{(i, j) : i, j \leq m, (i, j) \in E^{[P]}\} \cup \{(i+m, j+m) : i, j \leq m, (i, j) \in E^{[P]}\} \cup \{(1, 2), (2, 1), (1+m, 2+m), (2+m, 1+m)\}$; $\forall i \leq m$, $x_i^{[1]} = x_{i+m}^{[1]} = x_i^{[P]}$; $\forall (i, j) \in E^{[P]}$, $e_{i,j}^{[1]} = e_{i+m, j+m}^{[1]} = e_{i,j}^{[P]}$, and moreover we can randomly choose a value of edge

feature for $e_{1,2}^{[1]} = e_{2,1}^{[1]} = e_{1+m,2+m}^{[1]} = e_{2+m,1+m}^{[1]}$. For $G^{[2]}$, let $E^{[2]} = \{(i, j) : i, j \leq m, (i, j) \in E^{[P]}\} \cup \{(i+m, j+m) : i, j \leq m, (i, j) \in E^{[P]}\} \cup \{(1, 2+m), (2+m, 1), (1+m, 2), (2, 1+m)\}$; $\forall i \leq m, x_i^{[2]} = x_{i+m}^{[2]} = x_i^{[P]}$; $\forall (i, j) \in E^{[P]}, e_{i,j+m}^{[2]} = e_{i+m,j}^{[2]} = e_{i,j}^{[P]}$, and moreover we let $e_{1,2+m}^{[2]} = e_{2+m,1}^{[2]} = e_{1+m,2}^{[2]} = e_{2,1+m}^{[2]} = e_{1,2}^{[1]}$. In words, both $G^{[1]}$ and $G^{[2]}$ are constructed based on two copies of $G^{[P]}$, and the difference is that, $G^{[1]}$ adds the edges $\{(1, 2), (2, 1), (1+m, 2+m), (2+m, 1+m)\}$, whereas $G^{[2]}$ adds the edges $\{(1, 2+m), (2+m, 1), (1+m, 2), (2, 1+m)\}$, all with the same edge feature.

On one hand, by construction, 2-WL will not be able to distinguish $G^{[1]}$ from $G^{[2]}$. This is intuitive if we compare the rooted subtrees in the two graphs, as there exists a bijection from $V^{[1]}$ to $V^{[2]}$ that preserves the rooted subtree structure. A rigorous proof is given at the end of this section. In addition, we note that this is also consequence of the direct proof of Corollary 4.4 given in Appendix J, in which we will show that the same pair of graphs cannot be distinguished by 2-IGNs. Since 2-IGNs are no less powerful than 2-WL [39], this implies that 2-WL cannot distinguish them either.

On the other hand, $G^{[1]}$ and $G^{[2]}$ has different matching-count of the pattern. $G^{[1]}$ contains no subgraph isomorphic to $G^{[P]}$. Intuitively this is obvious; to be rigorous, note that firstly, neither the subgraph induced by the nodes $\{1, \dots, m\}$ nor the subgraph induced by the nodes $\{1+m, \dots, 2m\}$ is isomorphic to $G^{[P]}$, and secondly, the subgraph induced by any other set of m nodes is not connected, whereas $G^{[P]}$ is connected. $G^{[2]}$, however, has at least two induced subgraphs isomorphic to $G^{[P]}$, one induced by the nodes $\{1, \dots, m\}$, and the other induced by the nodes $\{1+m, \dots, 2m\}$.

If $G^{[P]}$ is a clique, then we also first construct $G^{[1]}, G^{[2]}$ from $G^{[P]}$ as two copies of $G^{[P]}$. Then, for $G^{[1]}$, we pick two distinct nodes $1, 2 \in V^{[P]}$ and remove the edges $(1, 2), (2, 1), (1+m, 2+m)$ and $(2+m, 1+m)$ from $V^{[1]}$, while adding edges $(1, 2+m), (2+m, 1), (1+m, 2), (2, 1+m)$ with the same edge features. Then, $G^{[1]}$ contains no subgraph isomorphic to $G^{[P]}$, while $G^{[2]}$ contains two. Note that the pair of graphs is the same as the counterexample pair of graphs that could have been constructed in the non-clique case for the pattern that is a clique with one edge deleted. Hence 2-WL still can't distinguish $G^{[1]}$ from $G^{[2]}$. □

Proof of 2-WL failing to distinguish $G^{[1]}$ and $G^{[2]}$:

To show that 2-WL cannot distinguish $G^{[1]}$ from $G^{[2]}$, we need to show that if we run 2-WL on the two graphs, then $\forall T, \{\mathbf{c}^{(T)}((i, j)) : i, j \in V\} = \{\mathbf{c}'^{(T)}((i, j)) : i, j \in V\}$. For this to hold, it is sufficient to find a bijective map $\eta : V^2 \rightarrow V^2$ such that $\mathbf{c}^{(T)}((i, j)) = \mathbf{c}'^{(T)}(\eta((i, j)))$, $\forall i, j \in V$. First, we define a set $S = \{(1, 2), (2, 1), (1+m, 2+m), (2+m, 1+m), (1, 2+m), (2+m, 1), (1+m, 2), (2, 1+m)\}$, which represents the ‘‘special’’ pairs of nodes that capture the difference between $G^{[1]}$ and $G^{[2]}$. Then we can define $\eta : V^2 \rightarrow V^2$ as

$$\eta((i, j)) = \begin{cases} (i, j), & \text{if } (i, j) \notin S \\ (i, \text{MOD}_{2m}(j+m)), & \text{if } (i, j) \in S \end{cases}$$

Note that η is a bijective. It is easy to verify that η is a color-preserving map between node pairs in $G^{[1]}$ and node pairs in $G^{[2]}$ at initialization, i.e. $\mathbf{c}^{(0)}((i, j)) = \mathbf{c}'^{(0)}(\eta((i, j)))$, $\forall i, j \in V$. We will prove by induction that in fact it remains such a color-preserving map at any iteration T . The inductive step that we need to prove is,

Lemma E.1. *For any positive integer t , supposing that $\mathbf{c}^{(t-1)}((i, j)) = \mathbf{c}'^{(t-1)}(\eta((i, j)))$, $\forall i, j \in V$, then we also have $\mathbf{c}^{(t)}((i, j)) = \mathbf{c}'^{(t)}(\eta((i, j)))$, $\forall i, j \in V$.*

Proof of Lemma E.1: By the update rule of 2-WL, $\forall i, j \in V$, to show that $\mathbf{c}^{(t)}((i, j)) = \mathbf{c}'^{(t)}(\eta((i, j)))$, we need to establish three conditions:

$$\mathbf{c}^{(t-1)}((i, j)) = \mathbf{c}'^{(t-1)}(\eta((i, j))) \tag{9}$$

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j))\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1(\eta((i, j)))\} \quad (10)$$

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j))\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2(\eta((i, j)))\} \quad (11)$$

The first condition is already guaranteed by the inductive hypothesis. Now we prove the last two conditions by examining different cases separately below.

Case 1 $i, j \notin \{1, 2, 1+m, 2+m\}$

Then $\eta((i, j)) = (i, j)$, and $N_1((i, j)) \cap S = \emptyset$, $N_2((i, j)) \cap S = \emptyset$. Therefore, η restricted to $N_1((i, j))$ or $N_2((i, j))$ is the identity map, and thus

$$\begin{aligned} \{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j))\} &= \{\mathbf{c}'^{(t-1)}(\eta(\tilde{s})) : \tilde{s} \in N_1((i, j))\} \\ &= \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1(\eta((i, j)))\}, \end{aligned}$$

thanks to the inductive hypothesis. Similar for the condition (11).

Case 2 $i \in \{1, 1+m\}, j \notin \{1, 2, 1+m, 2+m\}$

Then $\eta((i, j)) = (i, j)$, $N_2((i, j)) \cap S = \{(i, 2), (i, 2+m)\}$, and $N_1((i, j)) \cap S = \emptyset$. To show condition (11), note that η is the identity map when restricted to $N_2((i, j)) \setminus \{(i, 2), (i, 2+m)\}$, and hence

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j)) \setminus \{(i, 2), (i, 2+m)\}\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j)) \setminus \{(i, 2), (i, 2+m)\}\}$$

Moreover, $\eta((i, 2)) = (i, 2+m)$ and $\eta((i, 2+m)) = (i, 2)$. Hence, by the inductive hypothesis, $\mathbf{c}^{(t-1)}((i, 2)) = \mathbf{c}'^{(t-1)}((i, 2+m))$ and $\mathbf{c}^{(t-1)}((i, 2+m)) = \mathbf{c}'^{(t-1)}((i, 2))$. Therefore,

$$\begin{aligned} \{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j))\} &= \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j))\} \\ &= \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2(\eta((i, j)))\}, \end{aligned}$$

which shows condition (11). Condition (10) is easily seen as η restricted to $N_1((i, j))$ is the identity map.

Case 3 $j \in \{1, 1+m\}, i \notin \{1, 2, 1+m, 2+m\}$

There is $\eta((i, j)) = (i, j)$, $N_1((i, j)) \cap S = \{(2, j), (2+m, j)\}$, and $N_2((i, j)) \cap S = \emptyset$. Hence the proof can be carried out analogously to case 2.

Case 4 $i \in \{2, 2+m\}, j \notin \{1, 2, 1+m, 2+m\}$

There is $\eta((i, j)) = (i, j)$, $N_2((i, j)) \cap S = \{(i, 1), (i, 1+m)\}$, and $N_1((i, j)) \cap S = \emptyset$. Hence the proof can be carried out analogously to case 2.

Case 5 $j \in \{2, 2+m\}, i \notin \{1, 2, 1+m, 2+m\}$

There is $\eta((i, j)) = (i, j)$, $N_1((i, j)) \cap S = \{(1, j), (1+m, j)\}$, and $N_2((i, j)) \cap S = \emptyset$. Hence the proof can be carried out analogously to case 2.

Case 6 $(i, j) \in S$

There is $\eta((i, j)) = (i, \text{MOD}_{2m}(j))$, $N_1((i, j)) \cap S = \{(i, j), (\text{MOD}_{2m}(i), j)\}$, $N_2((i, j)) \cap S = \{(i, j), (i, \text{MOD}_{2m}(j))\}$. Thus, $N_1(\eta((i, j))) = N_1((i, \text{MOD}_{2m}(j)))$, $N_2(\eta((i, j))) = N_2((i, \text{MOD}_{2m}(j))) = N_2((i, j))$. Once again, η is the identity map when restricted to $N_1((i, j)) \setminus S$ or $N_2((i, j)) \setminus S$. Hence, by the inductive hypothesis, there is

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j)) \setminus \{(i, j), (\text{MOD}_{2m}(i), j)\}\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j)) \setminus \{(i, j), (\text{MOD}_{2m}(i), j)\}\}$$

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j)) \setminus \{(i, j), (i, \text{MOD}_{2m}(j))\}\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j)) \setminus \{(i, j), (i, \text{MOD}_{2m}(j))\}\}$$

Also from the inductive hypothesis, we have

$$\begin{aligned} \mathbf{c}^{(t-1)}((i, j)) &= \mathbf{c}'^{(t-1)}(\eta((i, j))) \\ &= \mathbf{c}'^{(t-1)}((i, \text{MOD}_{2m}(j))), \end{aligned} \quad (12)$$

$$\begin{aligned}
\mathbf{c}^{(t-1)}((i, j)) &= \mathbf{c}^{(t-1)}((j, i)) \\
&= \mathbf{c}'^{(t-1)}(\eta((j, i))) \\
&= \mathbf{c}'^{(t-1)}((j, \text{MOD}_{2m}(i))) \\
&= \mathbf{c}'^{(t-1)}((\text{MOD}_{2m}(i), j)),
\end{aligned} \tag{13}$$

$$\begin{aligned}
\mathbf{c}^{(t-1)}((i, \text{MOD}_{2m}(j))) &= \mathbf{c}'^{(t-1)}(\eta((i, \text{MOD}_{2m}(j)))) \\
&= \mathbf{c}'^{(t-1)}((i, \text{MOD}_{2m}(\text{MOD}_{2m}(j)))) \\
&= \mathbf{c}'^{(t-1)}((i, j)),
\end{aligned} \tag{14}$$

$$\begin{aligned}
\mathbf{c}^{(t-1)}((\text{MOD}_{2m}(i), j)) &= \mathbf{c}^{(t-1)}((j, \text{MOD}_{2m}(i))) \\
&= \mathbf{c}'^{(t-1)}(\eta((j, \text{MOD}_{2m}(i)))) \\
&= \mathbf{c}'^{(t-1)}((j, \text{MOD}_{2m}(\text{MOD}_{2m}(i)))) \\
&= \mathbf{c}'^{(t-1)}((j, i)) \\
&= \mathbf{c}'^{(t-1)}((i, j)),
\end{aligned} \tag{15}$$

where in (13) and (15), the first and the last equalities are thanks to the symmetry of the coloring between any pair of nodes (i', j') and its “reversed” version (j', i') , which persists throughout all iterations, as well as the fact that if $(i', j') \in S$, then $(j', i') \in S$. Therefore, we now have

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j))\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j))\} \tag{16}$$

$$\{\mathbf{c}^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j))\} = \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_2((i, j))\} \tag{17}$$

Since $\eta((i, j)) = (i, \text{MOD}_{2m}(j))$, we have

$$\begin{aligned}
N_1(\eta((i, j))) &= \{(k, \text{MOD}_{2m}(j)) : k \in V\} \\
&= \{(k, \text{MOD}_{2m}(j)) : (\text{MOD}_{2m}(k), j) \in N_1((i, j))\} \\
&= \{(\text{MOD}_{2m}(k), \text{MOD}_{2m}(j)) : (k, j) \in N_1((i, j))\}
\end{aligned}$$

Thanks to the symmetry of the coloring under the map $(i', j') \rightarrow (\text{MOD}_{2m}(i'), \text{MOD}_{2m}(j'))$, we then have

$$\begin{aligned}
\{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1(\eta((i, j)))\} &= \{\mathbf{c}'^{(t-1)}((\text{MOD}_{2m}(k), \text{MOD}_{2m}(j))) : (k, j) \in N_1((i, j))\} \\
&= \{\mathbf{c}'^{(t-1)}((k, j)) : (k, j) \in N_1((i, j))\} \\
&= \{\mathbf{c}'^{(t-1)}(\tilde{s}) : \tilde{s} \in N_1((i, j))\}
\end{aligned}$$

Therefore, combined with (16), we see that (10) is proved. (11) is a straightforward consequence of (17), since $N_2((i, j)) = N_2(\eta((i, j)))$.

Case 7 $i, j \in \{1, 1+m\}$

There is $\eta((i, j)) = (i, j)$, $N_2((i, j)) \cap S = \{(i, 2), (i, 2+m)\}$, and $N_1((i, j)) \cap S = \{(2, j), (2+m, j)\}$. Thus, both (10) and (11) can be proved analogously to how (11) is proved for case 2.

Case 8 $i, j \in \{2, 2+m\}$

There is $\eta((i, j)) = (i, j)$, $N_2((i, j)) \cap S = \{(i, 1), (i, 1+m)\}$, and $N_1((i, j)) \cap S = \{(1, j), (1+m, j)\}$. Thus, both (10) and (11) can be proved analogously to how (11) is proved for case 2.

With conditions (10) and (11) shown for all pairs of $(i, j) \in V^2$, we know that by the update rules of 2-WL, there is $\mathbf{c}^{(t)}((i, j)) = \mathbf{c}'^{(t)}(\eta((i, j)))$, $\forall i, j \in V$.

□

With Lemma E.1 justifying the inductive step, we see that for any positive integer T , there is $\mathbf{c}^{(T)}((i, j)) = \mathbf{c}'^{(T)}(\eta((i, j))), \forall i, j \in V$. Hence, we can conclude that $\forall T, \{\mathbf{c}^{(T)}((i, j)) : i, j \in V\} = \{\mathbf{c}'^{(T)}((i, j)) : i, j \in V\}$, which implies that the two graphs cannot be distinguished by 2-WL.

□

F Proof of Theorem 3.5 (MPNNs are able to subgraph-count star-shaped patterns)

(See Section 2.1 of Arvind et al. [2] for a proof for the case where all nodes have identical features.)

Proof. Without loss of generality, we represent a star-shaped pattern by $G^{[P]} = (V^{[P]}, E^{[P]}, x^{[P]}, e^{[P]})$, where $V^{[P]} = [m]$ (with node 1 representing the center) and $E^{[P]} = \{(1, i) : 2 \leq i \leq m\} \cup \{(i, 1) : 2 \leq i \leq m\}$.

Given a graph G , for each of its node j , we define $N(j)$ as the set of its neighbors in the graph. Then the neighborhood centered at j contributes to $C_S(G, G^{[P]})$ if and only if $x_j = x_1^{[P]}$ and $\exists S \subseteq N(j)$ such that the multiset $\{(x_k, e_{jk}) : k \in S\}$ equals the multiset $\{(x_k^{[P]}, e_{1k}^{[P]}) : 2 \leq k \leq m\}$. Moreover, the contribution to the number $C_S(G, G^{[P]})$ equals the number of all such subsets $S \subseteq N(j)$. Hence, we have the following decomposition

$$C_S(G, G^{[P]}) = \sum_{j \in V} f^{[P]}(x_j, \{(x_k, e_{jk}) : k \in N(j)\}),$$

where $f^{[P]}$, is defined for every 2-tuple consisting of a node feature and a multiset of pairs of node feature and edge feature (i.e., objects of the form

$$(x, M = \{(x_\alpha, e_\alpha) : \alpha \in K\})$$

where K is a finite set of indices) as

$$f^{[P]}(x, M) = \begin{cases} 0 & \text{if } x \neq x_1^{[P]} \\ \#_M^{[P]} & \text{if } x = x_1^{[P]} \end{cases}$$

where $\#_M^{[P]}$ denotes the number of sub-multisets of M that equals the multiset $\{(x_k^{[P]}, e_{1k}^{[P]}) : 2 \leq k \leq m\}$.

Thanks to Corollary 6 of Xu et al. [64] based on Zaheer et al. [72], we know that $f^{[P]}$ can be expressed by some message-passing function in an MPNN. Thus, together with summation as the readout function, MPNN is able to express $C_S(G, G^{[P]})$. □

G Proof of Theorem 3.7 (k -WL is able to count patterns of k or fewer nodes)

Proof. Suppose we run k -WL on two graphs, $G^{[1]}$ and $G^{[2]}$. In k -WL, the colorings of the k -tuples are initialized according to their isomorphism types as defined in Appendix C. Thus, if for some pattern of no more than k nodes, $G^{[1]}$ and $G^{[2]}$ have different matching-count or containment-count, then there exists an isomorphism type of k -tuples such that $G^{[1]}$ and $G^{[2]}$ differ in the number of k -tuples under this type. This implies that $\{\mathbf{c}_k^{(0)}(s) : s \in (V^{[1]})^k\} \neq \{\mathbf{c}'_k^{(0)}(s') : s' \in (V^{[2]})^k\}$, and hence the two graphs can be distinguished at the 0th iteration of k -WL. □

H Proof of Theorem 3.9 (T iterations of k -WL cannot induced-subgraph-count path patterns of size $(k + 1)2^T$ or more)

Proof. For any integer $m \geq (k + 1)2^T$, we will construct two graphs $G^{[1]} = (V^{[1]} = [2m], E^{[1]}, x^{[1]}, e^{[1]})$ and $G^{[2]} = (V^{[2]} = [2m], E^{[2]}, x^{[2]}, e^{[2]})$, both with $2m$ nodes but with

different matching-counts of H_m , and show that k -WL cannot distinguish them. Define $E_{double} = \{(i, i+1) : 1 \leq i < m\} \cup \{(i+1, i) : 1 \leq i < m\} \cup \{(i+m, i+m+1) : 1 \leq i < m\} \cup \{(i+m+1, i+m) : 1 \leq i < m\}$, which is the edge set of a graph that is exactly two disconnected copies of H_m . For $G^{[1]}$, let $E^{[1]} = E_{double} \cup \{(1, m), (m, 1), (1+m, 2m), (2m, 1+m)\}$; $\forall i \leq m, x_i^{[1]} = x_{i+m}^{[1]} = x_i^{[H_m]}$; $\forall (i, j) \in E^{[H_m]}, e_{i,j}^{[1]} = e_{j,i}^{[1]} = e_{i+m,j+m}^{[1]} = e_{j+m,i+m}^{[1]} = e_{i,j}^{[H_m]}$, and moreover, we can randomly choose a value of edge feature for $e_{1,m}^{[1]} = e_{m,1}^{[1]} = e_{1+m,2m}^{[1]} = e_{2m,1+m}^{[1]}$. For $G^{[2]}$, let $E^{[2]} = E_{double} \cup \{(1, 2m), (2m, 1), (m, 1+m), (1+m, 2m)\}$; $\forall i \leq m, x_i^{[2]} = x_{i+m}^{[2]} = x_i^{[H_m]}$; $\forall (i, j) \in E^{[H_m]}, e_{i,j}^{[1]} = e_{j,i}^{[1]} = e_{i+m,j+m}^{[1]} = e_{j+m,i+m}^{[1]} = e_{i,j}^{[H_m]}$, and moreover, set $e_{1,2m}^{[2]} = e_{2m,1}^{[2]} = e_{m,1+m}^{[2]} = e_{1+m,m}^{[2]} = e_{1,m}^{[1]}$. In words, both $G^{[1]}$ and $G^{[2]}$ are constructed based on two copies of H_m , and the difference is that, $G^{[1]}$ adds the edges $\{(1, m), (m, 1), (1+m, 2m), (2m, 1+m)\}$, whereas $G^{[2]}$ adds the edges $\{(1, 2m), (2m, 1), (m, 1+m), (1+m, m)\}$, all with the same edge feature. For the case $k=3, m=8, T=1$, for example, the constructed graphs are illustrated in Figure 4.

Can $G^{[1]}$ and $G^{[2]}$ be distinguished by k -WL? Let $c_k^{(t)}, c'_k{}^{(t)}$ be the coloring functions of k -tuples for $G^{[1]}$ and $G^{[2]}$, respectively, obtained after running k -WL on the two graphs simultaneously for t iterations. To show that the answer is negative, we want to prove that

$$\{c_k^{(T)}(s) : s \in [2m]^k\} = \{c'_k{}^{(T)}(s) : s \in [2m]^k\} \quad (18)$$

To show this, it is sufficient to find a permutation $\eta : [2m]^k \rightarrow [2m]^k$ such that $\forall k$ -tuple $s \in [2m]^k, c_k^{(T)}(s) = c'_k{}^{(T)}(\eta(s))$. Before defining such an η , we need the following lemma.

Lemma H.1. *Let p be a positive integer. If $m \geq (k+1)p$, then $\forall s \in [2m]^k, \exists i \in [m]$ such that $\{i, i+1, \dots, i+p-1\} \cap \{\text{MOD}_m(j) : j \in s\} = \emptyset$.*

Proof of Lemma H.1: We can use a simple counting argument to show this. For $u \in [k+1]$, define $A_u = \{up, up+1, \dots, (u+1)p-1\} \cup \{up+m, up+1+m, \dots, (u+1)p-1+m\}$. Then $|A_u| = 2p$, $A_u \cap A_{u'} = \emptyset$ if $u \neq u'$, and

$$[2m] \supseteq \bigcup_{u \in [k+1]} A_u, \quad (19)$$

since $m \geq (k+1)p$. Suppose that the claim is not true, then each A_i contains at least one node in s , and therefore

$$s \supseteq (s \cap [2m]) \supseteq \bigcup_{u \in [k+1]} (s \cap A_u),$$

which contains at least $k+1$ nodes, which is contradictory. \square

With this lemma, we see that $\forall s \in [2m]^k, \exists i \in [m]$ such that $\forall j \in s, \text{MOD}_m(j)$ either $< i$ or $\geq i + 2^{T+1} - 1$. Thus, we can first define the mapping $\chi : [2m]^k \rightarrow [m]$ from a k -tuple s to the smallest such node index $i \in [m]$. Next, $\forall i \in [m]$, we define a mapping τ_i from $[2m]$ to $[2m]$ as

$$\tau_i(j) = \begin{cases} j, & \text{if } \text{MOD}_m(j) \leq i \\ \text{MOD}_{2m}(j+m), & \text{otherwise} \end{cases} \quad (20)$$

τ_i is a permutation on $[2m]$. For $\forall i \in [m]$, this allows us to define a mapping ζ_i from $[2m]^k \rightarrow [2m]^k$ as, $\forall s = (i_1, \dots, i_k) \in [2m]^k$,

$$\zeta_i(s) = (\tau_i(i_1), \dots, \tau_i(i_k)). \quad (21)$$

Finally, we define a mapping η from $[2m]^k \rightarrow [2m]^k$ as,

$$\eta(s) = \zeta_{\chi(s)}(s) \quad (22)$$

The maps χ, τ and η are illustrated in Figure 4.




Figure 4: Illustration of the construction in the proof of Theorem 3.9 in Appendix H. In this particular case, $k = 3$, $m = 8$, $T = 1$. If we consider $s = (1, 12, 8)$ as an example, where the corresponding nodes are marked by blue squares in $G^{[1]}$, there is $\chi(s) = 2$, and thus $\eta(s) = \zeta_2(s) = (1, 4, 16)$, which are marked by blue squares in $G^{[2]}$. Similarly, if we consider $s = (3, 14, 15)$, then $\chi(s) = 4$, and thus $\eta(s) = \zeta_4(s) = (3, 6, 7)$. In both cases, we see that the isomorphism type of s in $G^{[1]}$ equals the isomorphism type of $\eta(s)$ in $G^{[2]}$. In the end, we will show that $\mathbf{c}_k^{(T)}(s) = \mathbf{c}'_k^{(T)}(\eta(s))$.

To fulfill the proof, there are two things we need to show about η . First, we want it to be a permutation on $[2m]^k$. To see this, observe that $\chi(s) = \chi(\eta(s))$, and hence $\forall s \in [2m]^k$, $(\eta \circ \eta)(s) = (\zeta_{\chi(\eta(s))} \circ \zeta_{\chi(s)})(s) = s$, since $\forall i \in [m]$, $\tau_i \circ \tau_i$ is the identity map on $[2m]$.

Second, we need to show that $\forall s \in [2m]^k$, $\mathbf{c}_k^{(T)}(s) = \mathbf{c}'_k^{(T)}(\eta(s))$. This will be a consequence of the following lemma.

Lemma H.2. *At iteration t , $\forall s \in [2m]^k$, $\forall i$ such that $\forall j \in s$, either $\text{MOD}_m(j) < i$ or $\text{MOD}_m(j) \geq i + 2^t$, there is*

$$\mathbf{c}_k^{(t)}(s) = \mathbf{c}'_k^{(t)}(\zeta_i(s)) \quad (23)$$

Remark: This statement allows i to depend on s , as will be the case when we apply this lemma to $\eta(s) = \zeta_{\chi(s)}(s)$, where we set i to be $\chi(s)$.

Proof of Lemma H.2: Notation-wise, for any k -tuple, $s = (i_1, \dots, i_k)$, and for $w \in [k]$, use $\mathbb{I}_w(s)$ to denote the w th entry of s , i_w .

The lemma can be shown by using induction on t . Before looking at the base case $t = 0$, we will first show the inductive step, which is:

$$\begin{aligned} \forall \bar{T}, \text{ suppose the lemma holds for all } t \leq \bar{T} - 1, \\ \text{ then it also holds for } t = \bar{T}. \end{aligned} \quad (24)$$

Inductive step:

Fix a \bar{T} and suppose the lemma holds for all $t \leq \bar{T} - 1$. Under the condition that $\forall j \in s$, either $\text{MOD}_m(j) < i$ or $\text{MOD}_m(j) \geq i + 2^{\bar{T}}$, to show $\mathbf{c}_k^{(\bar{T})}(s) = \mathbf{c}'_k^{(\bar{T})}(\zeta_i(s))$, we need two things to hold:

1. $\mathbf{c}_k^{(\bar{T}-1)}(s) = \mathbf{c}'_k^{(\bar{T}-1)}(\zeta_i(s))$
2. $\forall w \in [k], \{\mathbf{c}_k^{(\bar{T}-1)}(\tilde{s}) : \tilde{s} \in N_w(s)\} = \{\mathbf{c}'_k^{(\bar{T}-1)}(\tilde{s}) : \tilde{s} \in N_w(\zeta_i(s))\}$

The first condition is a consequence of the inductive hypothesis, as $i + 2^{\bar{T}} > i + 2^{(\bar{T}-1)}$. For the second condition, it is sufficient to find for all $w \in [k]$, a bijective mapping ξ from $N_w(s)$ to $N_w(\zeta_i(s))$ such that $\forall \tilde{s} \in N_w(s)$, $\mathbf{c}_k^{(\bar{T}-1)}(\tilde{s}) = \mathbf{c}'_k^{(\bar{T}-1)}(\xi(\tilde{s}))$.

We then define $\beta(i, \tilde{s}) =$

$$\begin{cases} \text{MOD}_m(\mathbb{I}_w(\tilde{s})) + 1, & \text{if } i \leq \text{MOD}_m(\mathbb{I}_w(\tilde{s})) < i + 2^{\bar{T}-1} \\ i, & \text{otherwise} \end{cases} \quad (25)$$

Now, consider any $\tilde{s} \in N_w(s)$. Note that \tilde{s} and s differ only in the w th entry of the k -tuple.

- If $i \leq \text{MOD}_m(\mathbb{I}_w(\tilde{s})) < i + 2^{\bar{T}-1}$, then $\forall j \in \tilde{s}$,
 - either $j \in s$, in which case either $\text{MOD}_m(j) < i < \text{MOD}_m(\mathbb{I}_w(\tilde{s})) + 1 = \beta(i, \tilde{s})$ or $\text{MOD}_m(j) \geq i + 2^{\bar{T}} \geq \text{MOD}_m(\mathbb{I}_w(\tilde{s})) + 1 + 2^{\bar{T}-1} = \beta(i, \tilde{s}) + 2^{\bar{T}-1}$,
 - or $j = \mathbb{I}_w(\tilde{s})$, in which case $\text{MOD}_m(j) < \text{MOD}_m(\mathbb{I}_w(\tilde{s})) + 1 = \beta(i, \tilde{s})$.
- If $\text{MOD}_m(\mathbb{I}_w(\tilde{s})) < i$ or $\text{MOD}_m(\mathbb{I}_w(\tilde{s})) \geq i + 2^{\bar{T}-1}$, then $\forall j \in \tilde{s}$,
 - either $j \in s$, in which case either $\text{MOD}_m(j) < i = \beta(i, \tilde{s})$ or $\text{MOD}_m(j) \geq i + 2^{\bar{T}} \geq \beta(i, \tilde{s}) + 2^{\bar{T}-1}$,
 - or $j = \mathbb{I}_w(\tilde{s})$, in which case either $\text{MOD}_m(j) < i = \beta(i, \tilde{s})$ or $\text{MOD}_m(j) \geq i + 2^{\bar{T}-1} \geq \beta(i, \tilde{s}) + 2^{\bar{T}-1}$.

Thus, in all cases, there is $\forall j \in \tilde{s}$, either $\text{MOD}_m(j) < \beta(i, \tilde{s})$, or $\text{MOD}_m(j) \geq i + 2^{\bar{T}-1}$. Hence, by the inductive hypothesis, we have $\mathbf{c}_k^{(\bar{T}-1)}(\tilde{s}) = \mathbf{c}'_k^{(\bar{T}-1)}(\zeta_{\beta(i, \tilde{s})}(\tilde{s}))$. This inspires us to define, for $\forall w \in [k], \forall \tilde{s} \in N_w(s)$,

$$\xi(\tilde{s}) = \zeta_{\beta(i, \tilde{s})}(\tilde{s}) \quad (26)$$

Additionally, we still need to prove that, firstly, ξ maps $N_w(s)$ to $N_w(\zeta_i(s))$, and secondly, ξ is a bijection. For the first statement, note that $\forall \tilde{s} \in N_w(s)$, $\zeta_{\beta(i, \tilde{s})}(s) = \zeta_i(s)$ because s contains no entry between i and $\beta(i, \tilde{s})$, with the latter being less than $i + 2^{\bar{T}}$. Hence, if $\tilde{s} \in N_w(s)$, then $\forall w' \in [k]$ with $w' \neq w$, there is $\mathbb{I}_{w'}(\tilde{s}) = \mathbb{I}_{w'}(s)$, and therefore $\mathbb{I}_{w'}(\xi(\tilde{s})) = \mathbb{I}_{w'}(\zeta_{\beta(i, \tilde{s})}(\tilde{s})) = \tau_{\beta(i, \tilde{s})}(\mathbb{I}_{w'}(\tilde{s})) = \tau_{\beta(i, \tilde{s})}(\mathbb{I}_{w'}(s)) = \mathbb{I}_{w'}(\zeta_{\beta(i, \tilde{s})}(s)) = \mathbb{I}_{w'}(\zeta_i(s))$, which ultimately implies that $\xi(\tilde{s}) \in N_w(\zeta_i(s))$.

For the second statement, note that since $\mathbb{I}_w(\xi(\tilde{s})) = \tau_{\beta(i, \tilde{s})}(\mathbb{I}_w(\tilde{s}))$ (by the definition of ζ), there is $\text{MOD}_m(\mathbb{I}_w(\xi(\tilde{s}))) = \text{MOD}_m(\tau_{\beta(i, \tilde{s})}(\mathbb{I}_w(\tilde{s}))) = \text{MOD}_m(\mathbb{I}_w(\tilde{s}))$, and therefore $\beta(i, \xi(\tilde{s})) = \beta(i, \tilde{s})$. Thus, we know that $(\xi \circ \xi)(\tilde{s}) = (\zeta_{\beta(i, \xi(\tilde{s}))} \circ \zeta_{\beta(i, \tilde{s})})(\tilde{s}) = (\zeta_{\beta(i, \tilde{s})} \circ \zeta_{\beta(i, \tilde{s})})(\tilde{s}) = \tilde{s}$. This implies that ξ is a bijection from $N_w(s)$ to $N_w(\zeta_i(s))$.

This concludes the proof of the inductive step.

Base case:

We need to show that

$$\begin{aligned} \forall s \in [2m]^k, \forall i^* \text{ such that } \forall j \in s, \text{ either } \text{MOD}_m(j) < i^* \\ \text{or } \text{MOD}_m(j) \geq i^* + 1, \text{ there is } \mathbf{c}_k^{(0)}(s) = \mathbf{c}'_k^{(0)}(\zeta_{i^*}(s)) \end{aligned} \quad (27)$$

Due to the way in which the colorings of the k -tuples are initialized in k -WL, the statement above is equivalent to showing that s in $G^{[1]}$ and $\zeta_{i^*}(s)$ in $G^{[2]}$ have the same isomorphism type, for which we need the following to hold.

Lemma H.3. *Say $s = (i_1, \dots, i_k)$, in which case $\zeta_{i^*}(s) = (\tau_{i^*}(i_1), \dots, \tau_{i^*}(i_k))$. Then*

1. $\forall i_\alpha, i_\beta \in s, i_\alpha = i_\beta \Leftrightarrow \tau_{i^*}(i_\alpha) = \tau_{i^*}(i_\beta)$
2. $\forall i_\alpha \in s, x_{i_\alpha}^{[1]} = x_{\tau_{i^*}(i_\alpha)}^{[2]}$
3. $\forall i_\alpha, i_\beta \in s, (i_\alpha, i_\beta) \in E^{[1]} \Leftrightarrow (\tau_{i^*}(i_\alpha), \tau_{i^*}(i_\beta)) \in E^{[2]}$, and moreover, if either is true, $e_{i_\alpha, i_\beta}^{[1]} = e_{\tau_{i^*}(i_\alpha), \tau_{i^*}(i_\beta)}^{[2]}$

Proof of Lemma H.3:

1. This is true since τ_{i^*} is a permutation on $[2m]$.
2. This is true because by the construction of the two graphs, $\forall i \in [2m], x_i^{[1]} = x_i^{[2]}$, and moreover $x_i^{[1]} = x_{i+m}^{[1]}$ if $i \leq m$.
3. Define $S = \{(1, m), (m, 1), (1+m, 2m), (2m, 1+m), (1, 2m), (2m, 1), (m, 1+m), (1+m, 2m)\}$, which is the set of ‘‘special’’ pairs of nodes in which $G^{[1]}$ and $G^{[2]}$ differ. Note that $\forall (i_\alpha, i_\beta) \in [2m]^2, (i_\alpha, i_\beta) \in S$ if and only if the sets $\{\text{MOD}_m(i_\alpha), \text{MOD}_m(i_\beta)\} = \{1, m\}$.
By the assumption on i^* in (27), we know that $i_\alpha, i_\beta \notin \{i^*, i^* + m\}$. Now we look at 16 different cases separately, which comes from 4 possibilities for each of i_α and i_β : i_α (or i_β) belonging to $\{1, \dots, i^* - 1\}, \{i^* + 1, \dots, m\}, \{1+m, \dots, i^* - 1+m\}$, or $\{i^* + 1+m, \dots, 2m\}$

- Case 1** $1 \leq i_\alpha, i_\beta < i^*$
Then $\tau_{i^*}(i_\alpha) = i_\alpha, \tau_{i^*}(i_\beta) = i_\beta$. In addition, as $\text{MOD}_m(i_\alpha), \text{MOD}_m(i_\beta) \neq m$, there is $(i_\alpha, i_\beta) \notin S$. Thus, if $(i_\alpha, i_\beta) \in E^{[1]}$, then $(i_\alpha, i_\beta) \in E_{\text{double}} \subset E^{[2]}$, and moreover, $e_{i_\alpha, i_\beta}^{[1]} = e_{i_\alpha, i_\beta}^{[H_m]} = e_{i_\alpha, i_\beta}^{[2]} = e_{\tau_{i^*}(i_\alpha), \tau_{i^*}(i_\beta)}^{[2]}$. Same for the other direction.
- Case 2** $1+m \leq i_\alpha, i_\beta < i^* + m$
Similar to case 1.
- Case 3** $i^* + 1 \leq i_\alpha, i_\beta \leq m$
Then $\tau_{i^*}(i_\alpha) = i_\alpha + m, \tau_{i^*}(i_\beta) = i_\beta + m$. In addition, as $\text{MOD}_m(i_\alpha), \text{MOD}_m(i_\beta) \neq 1$, there is $(i_\alpha, i_\beta) \notin S$. Thus, if $(i_\alpha, i_\beta) \in E^{[1]}$, then $(i_\alpha, i_\beta) \in E_{\text{double}}$, and hence $(i_\alpha + m, i_\beta + m) \in E_{\text{double}} \subset E^{[2]}$, and moreover, $e_{i_\alpha, i_\beta}^{[1]} = e_{i_\alpha, i_\beta}^{[H_m]} = e_{i_\alpha + m, i_\beta + m}^{[2]} = e_{\tau_{i^*}(i_\alpha), \tau_{i^*}(i_\beta)}^{[2]}$.
- Case 4** $i^* + 1 + m \leq i_\alpha, i_\beta \leq 2m$
Similar to case 3.
- Case 5** $1 \leq i_\alpha < i^*, i^* + 1 \leq i_\beta \leq m$
If $i_\alpha \neq 1$ or $i_\beta \neq m$, then since H_m is a path and $i_\alpha < i^* \leq i_\beta - 1$, $(i_\alpha, i_\beta) \notin E^{[1]}$ or $E^{[2]}$. Now we consider the case where $i_\alpha = 1, i_\beta = m$. As $1 \leq i^* < m$, by the definition of τ , there is $\tau_{i^*}(1) = 1$, and $\tau_{i^*}(m) = 2m$. Note that both $(1, m) \in E^{[1]}$ and $(1, 2m) \in E^{[2]}$ are true, and moreover, $e_{1, m}^{[1]} = e_{1, 2m}^{[2]}$.
- Case 6** $1 \leq i_\beta < i^*, i^* + 1 \leq i_\alpha \leq m$
Similar to case 5.
- Case 7** $1+m \leq i_\alpha < i^* + m, i^* + 1 + m \leq i_\beta \leq 2m$
Similar to case 5.
- Case 8** $1+m \leq i_\beta < i^* + m, i^* + 1 + m \leq i_\alpha \leq 2m$
Similar to case 5.
- Case 9** $1 \leq i_\alpha < i^*$ and $1+m \leq i_\beta < i^* + m$
Then $\tau_s(i_\alpha) = i_\alpha, \tau_s(i_\beta) = i_\beta$, and $(i_\alpha, i_\beta) \notin E^{[1]}$ or $E^{[2]}$.
- Case 10** $1 \leq i_\beta < i^*$ and $1+m \leq i_\alpha < i^* + m$
Similar to case 9.
- Case 11** $i^* + 1 \leq i_\alpha < m$ and $i^* + 1 + m \leq i_\beta \leq 2m$
 $(i_\alpha, i_\beta) \notin E^{[1]}$. $\tau_s(i_\alpha) = i_\alpha + m, \tau_s(i_\beta) = i_\beta - m$. Hence $(\tau_s(i_\alpha), \tau_s(i_\beta)) \notin E^{[2]}$ either.
- Case 12** $i^* + 1 \leq i_\beta \leq m$ and $i^* + 1 + m \leq i_\alpha \leq 2m$
Similar to case 11.
- Case 13** $1 \leq i_\alpha < i^*$ and $i^* + 1 + m \leq i_\beta \leq 2m$
 $(i_\alpha, i_\beta) \notin E^{[1]}$ obviously. We also have $\tau_s(i_\alpha) = i_\alpha \in [1, i^*], \tau_s(i_\beta) = i_\beta - 1 \in [i^* + 1, m]$, and hence $(\tau_s(i_\alpha), \tau_s(i_\beta)) \notin E^{[2]}$.
- Case 14** $1 \leq i_\beta < i^*$ and $i^* + 1 + m \leq i_\alpha \leq 2m$
Similar to case 13.
- Case 15** $1+m \leq i_\alpha < i^* + m$ and $i^* + 1 \leq i_\beta \leq m$
Similar to case 13.

Case 16 $1 + m \leq i_\beta < i^* + m$ and $i^* + 1 \leq i_\alpha \leq m$
 Similar to case 13.

This concludes the proof of Lemma H.3. □

Lemma H.3 completes the proof of the base case, and hence the induction argument for Lemma H.2. □

$\forall s \in [2m]^k$, since $\eta(s) = \zeta_{\chi(s)}(s)$, and $\chi(s)$ satisfies $\forall j \in s$, either $\text{MOD}_m(j) < i$ or $\text{MOD}_m(j) \geq i + 2^T$, Lemma H.2 implies that at iteration T , we have $\mathbf{c}_k^{(T)}(s) = \mathbf{c}'_k^{(T)}(\zeta_{\chi(s)}(s)) = \mathbf{c}'_k^{(T)}(\eta(s))$. Since we have shown that η is a permutation on $[2m]^k$, this let's us conclude that

$$\{\mathbf{c}_k^{(T)}(s) : s \in [2m]^k\} = \{\mathbf{c}'_k^{(T)}(s) : s \in [2m]^k\}, \quad (28)$$

and therefore k -WL cannot distinguish between the two graphs in T iterations. □

I Proof of Theorem 4.2 (2-IGNs are no more powerful than 2-WL)

Note that a 2-IGN takes as input a third-order tensor, $\mathbf{B}^{(0)}$, defined as in (1). If we use $\mathbf{B}^{(t)}$ to denote the output of the t th layer of the 2-IGN, then they are obtained iteratively by

$$\mathbf{B}^{(t+1)} = \sigma(L^{(t)}(\mathbf{B}^{(t)})) \quad (29)$$

Proof. For simplicity of notations, we assume $d_t = 1$ in every layer of a 2-IGN. The general case can be proved by adding more subscripts. For 2-WL, we use the definition in Appendix C except for omitting the subscript k in $\mathbf{c}_k^{(t)}$.

To start, it is straightforward to show (and we will prove it at the end) that the theorem can be deduced from the following lemma:

Lemma I.1. *Say $G^{[1]}$ and $G^{[2]}$ cannot be distinguished by the 2-WL. Then $\forall t \in \mathbb{N}$, it holds that*

$$\forall s, s' \in V^2, \text{ if } \mathbf{c}^{(t)}(s) = \mathbf{c}'^{(t)}(s'), \text{ then } \mathbf{B}_s^{(t)} = \mathbf{B}'_{s'} \quad (30)$$

This lemma can be shown by induction. To see this, first note that the lemma is equivalent to the statement that

$$\forall T \in \mathbb{N}, \forall t \leq T, (30) \text{ holds.}$$

This allows us to carry out an induction in $T \in \mathbb{N}$. For the base case $t = T = 0$, this is true because $\mathbf{c}^{(0)}$ and $\mathbf{c}'^{(0)}$ in WL and $\mathbf{B}^{(0)}$ and $\mathbf{B}'^{(0)}$ in 2-IGN are both initialized in the same way according to the subgraph isomorphism. To be precise, $\mathbf{c}^{(0)}(s) = \mathbf{c}'^{(0)}(s')$ if and only if the subgraph in $G^{[1]}$ induced by the pair of nodes s is isomorphic to the subgraph in $G^{[2]}$ induced by the pair of nodes s' , which is also true if and only if $\mathbf{B}_s^{(0)} = \mathbf{B}'_{s'}^{(0)}$.

Next, to show that the induction step holds, we need to prove the following statement:

$$\begin{aligned} \forall T \in \mathbb{N}, \text{ if } \forall t \leq T - 1, (30) \text{ holds,} \\ \text{then (30) also holds for } t = T. \end{aligned}$$

To prove the consequent, we assume that for some $s, s' \in V^2$, there is $\mathbf{c}^{(T)}(s) = \mathbf{c}'^{(T)}(s')$, and then attempt to show that $\mathbf{B}_s^{(T)} = \mathbf{B}'_{s'}^{(T)}$. By the update rules of k -WL, the statement $\mathbf{c}^{(T)}(s) = \mathbf{c}'^{(T)}(s')$ implies that

$$\begin{cases} \mathbf{c}^{(T-1)}(s) = \mathbf{c}'^{(T-1)}(s') \\ \{\mathbf{c}^{(T-1)}(\tilde{s}) : \tilde{s} \in N_1(s)\} = \{\mathbf{c}'^{(T-1)}(\tilde{s}) : \tilde{s} \in N_1(s')\} \\ \{\mathbf{c}^{(T-1)}(\tilde{s}) : \tilde{s} \in N_2(s)\} = \{\mathbf{c}'^{(T-1)}(\tilde{s}) : \tilde{s} \in N_2(s')\} \end{cases} \quad (31)$$

Case 1: $s = (i, j) \in V^2$ with $i \neq j$

Let's first consider the case where $s = (i, j) \in V^2$ with $i \neq j$. In this case, we can also write $s' = (i', j') \in V^2$ with $i' \neq j'$, thanks to Lemma D.1. Then, note that V^2 can be written as the union of 9 disjoint sets that are defined depending on s :

$$V^2 = \bigcup_{w=1}^9 A_{s,w},$$

where we define $A_{s,1} = \{(i, j)\}$, $A_{s,2} = \{(i, i)\}$, $A_{s,3} = \{(j, j)\}$, $A_{s,4} = \{(i, k) : k \neq i \text{ or } j\}$, $A_{s,5} = \{(k, i) : k \neq i \text{ or } j\}$, $A_{s,6} = \{(j, k) : k \neq i \text{ or } j\}$, $A_{s,7} = \{(k, j) : k \neq i \text{ or } j\}$, $A_{s,8} = \{(k, l) : k \neq l \text{ and } \{k, l\} \cap \{i, j\} = \emptyset\}$, and $A_{s,9} = \{(k, k) : k \notin \{i, j\}\}$. In this way, we partition V^2 into 9 different subsets, each of which consisting of pairs (k, l) that yield a particular equivalence class of the 4-tuple (i, j, k, l) . Similarly, we can define $A_{s',w}$ for $w \in [9]$, which will also give us

$$V^2 = \bigcup_{w=1}^9 A_{s',w}$$

Moreover, note that

$$\begin{aligned} N_1(s) &= \bigcup_{w=1,3,7} A_{s,w} \\ N_2(s) &= \bigcup_{w=1,2,4} A_{s,w} \\ N_1(s') &= \bigcup_{w=1,3,7} A_{s',w} \\ N_2(s') &= \bigcup_{w=1,2,4} A_{s',w} \end{aligned}$$

Before proceeding, we make the following definition to simplify notations:

$$\begin{aligned} \mathfrak{C}_{s,w} &= \{\mathbf{c}^{(T-1)}(\tilde{s}) : \tilde{s} \in A_{s,w}\} \\ \mathfrak{C}'_{s',w} &= \{\mathbf{c}'^{(T-1)}(\tilde{s}) : \tilde{s} \in A_{s',w}\} \end{aligned}$$

This allows us to rewrite (31) as

$$\mathfrak{C}_{s,1} = \mathfrak{C}'_{s',1} \tag{32}$$

$$\bigcup_{w=1,3,7} \mathfrak{C}_{s,w} = \bigcup_{w=1,3,7} \mathfrak{C}'_{s',w} \tag{33}$$

$$\bigcup_{w=1,2,4} \mathfrak{C}_{s,w} = \bigcup_{w=1,2,4} \mathfrak{C}'_{s',w} \tag{34}$$

Combining (32) and (33), we obtain

$$\bigcup_{w=3,7} \mathfrak{C}_{s,w} = \bigcup_{w=3,7} \mathfrak{C}'_{s',w} \tag{35}$$

Combining (32) and (34), we obtain

$$\bigcup_{w=2,4} \mathfrak{C}_{s,w} = \bigcup_{w=2,4} \mathfrak{C}'_{s',w} \tag{36}$$

Note that V^2 can also be partitioned into two disjoint subsets:

$$V^2 = \left(\bigcup_{w=1,4,5,6,7,8} A_{s,w} \right) \cap \left(\bigcup_{w=2,3,9} A_{s,w} \right),$$

where the first subset represent the edges: $\{(i, j) \in V^2 : i \neq j\}$ and the second subset represent the nodes: $\{(i, i) : i \in V\}$. Similarly,

$$V^2 = \left(\bigcup_{w=1,4,5,6,7,8} A_{s',w} \right) \cap \left(\bigcup_{w=2,3,9} A_{s',w} \right),$$

As shown in Lemma D.1, pairs of nodes that represent edges cannot share the same color with pairs of nodes the represent nodes in any iteration of 2-WL. Thus, we have

$$\left(\bigcup_{w=1,4,5,6,7,8} \mathfrak{C}_{s,w} \right) \cap \left(\bigcup_{w=2,3,9} \mathfrak{C}'_{s',w} \right) = \emptyset \quad (37)$$

$$\left(\bigcup_{w=1,4,5,6,7,8} \mathfrak{C}'_{s',w} \right) \cap \left(\bigcup_{w=2,3,9} \mathfrak{C}_{s,w} \right) = \emptyset \quad (38)$$

Combining (35) and (37) or (38), we get

$$\mathfrak{C}_{s,3} = \mathfrak{C}'_{s',3} \quad (39)$$

$$\mathfrak{C}_{s,7} = \mathfrak{C}'_{s',7} \quad (40)$$

Combining (36) and (37) or (38), we get

$$\mathfrak{C}_{s,2} = \mathfrak{C}'_{s',2} \quad (41)$$

$$\mathfrak{C}_{s,4} = \mathfrak{C}'_{s',4} \quad (42)$$

Thanks to symmetry between (i, j) and (j, i) , as we work with undirected graphs, there is

$$\mathfrak{C}_{s,5} = \mathfrak{C}_{s,4} = \mathfrak{C}'_{s',4} = \mathfrak{C}'_{s',5} \quad (43)$$

$$\mathfrak{C}_{s,6} = \mathfrak{C}_{s,7} = \mathfrak{C}'_{s',7} = \mathfrak{C}'_{s',6} \quad (44)$$

In addition, since we assume that $G^{[1]}$ and $G^{[2]}$ cannot be distinguished by 2-WL, there has to be

$$\bigcup_{w=1}^9 \mathfrak{C}_{s,w} = \bigcup_{w=1}^9 \mathfrak{C}'_{s',w}$$

Combining this with (37) or (38), we get

$$\bigcup_{w=1,4,5,6,7,8} \mathfrak{C}_{s,w} = \bigcup_{w=1,4,5,6,7,8} \mathfrak{C}'_{s',w} \quad (45)$$

$$\bigcup_{w=2,3,9} \mathfrak{C}_{s,w} = \bigcup_{w=2,3,9} \mathfrak{C}'_{s',w} \quad (46)$$

Combining (45) with (32), (42), (43), (44), (40), we get

$$\mathfrak{C}_{s,8} = \mathfrak{C}'_{s',8} \quad (47)$$

Combining (46) with (41) and (39), we get

$$\mathfrak{C}_{s,9} = \mathfrak{C}'_{s',9} \quad (48)$$

Hence, in conclusion, we have that $\forall w \in [9]$,

$$\mathfrak{C}_{s,w} = \mathfrak{C}'_{s',w} \quad (49)$$

By the inductive hypothesis, this implies that $\forall w \in [9]$,

$$\{\mathbf{B}_{\tilde{s}}^{(T-1)} : \tilde{s} \in A_{s,w}\} = \{\mathbf{B}'_{\tilde{s}}^{(T-1)} : \tilde{s} \in A_{s',w}\} \quad (50)$$

Let us show how (50) may be leveraged. First, to prove that $\mathbf{B}_s^{(T)} = \mathbf{B}'_{s'}^{(T)}$, recall that

$$\begin{aligned} \mathbf{B}^{(T)} &= \sigma(L^{(T)}(\mathbf{B}^{(T-1)})) \\ \mathbf{B}'^{(T)} &= \sigma(L^{(T)}(\mathbf{B}'^{(T-1)})) \end{aligned} \quad (51)$$

Therefore, it is sufficient to show that for all linear equivariant layer L , we have

$$L(\mathbf{B}^{(T-1)})_{i,j} = L(\mathbf{B}'^{(T-1)})_{i',j'} \quad (52)$$

Also, recall that

$$\begin{aligned} L(\mathbf{B}^{(T-1)})_{i,j} &= \sum_{(k,l) \in V^2} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\ L(\mathbf{B}'^{(T-1)})_{i',j'} &= \sum_{(k',l') \in V^2} T_{i',j',k',l'} \mathbf{B}'_{k',l'} + Y_{i',j'} \end{aligned} \quad (53)$$

By the definition of the $A_{s,w}$'s and $A_{s',w}$'s, there is $\forall w \in [9], \forall (k, l) \in A_{s,w}, \forall (k', l') \in A_{s',w}$, we have the 4-tuples $(i, j, k, l) \sim (i', j', k', l')$, i.e., \exists a permutation π on V such that $(i, j, k, l) = (\pi(i'), \pi(j'), \pi(k'), \pi(l'))$, which implies that $T_{i,j,k,l} = T_{i',j',k',l'}$. Therefore, together with (50), we have the following:

$$\begin{aligned}
L(\mathbf{B}^{(T-1)})_{i,j} &= \sum_{(k,l) \in V^2} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\
&= \sum_{w=1}^9 \sum_{(k,l) \in A_{s,w}} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\
&= \sum_{w=1}^9 \sum_{(k',l') \in A_{s',w}} T_{i',j',k',l'} \mathbf{B}'_{k',l'} + Y_{i',j'} \\
&= L(\mathbf{B}'^{(T-1)})_{i',j'}
\end{aligned} \tag{54}$$

and hence $\mathbf{B}_{i,j}^{(T)} = \mathbf{B}'_{i',j'}^{(T)}$, which concludes the proof for the case that $s = (i, j)$ for $i \neq j$.

Case 2: $s = (i, i) \in V^2$

Next, consider the case $s = (i, i) \in V^2$. In this case, $s' = (i', i')$ for some $i' \in V$. This time, we write V^2 as the union of 5 disjoint sets that depend on s (or s'):

$$V^2 = \bigcup_{w=1}^5 A_{s,w},$$

where we define $A_{s,1} = \{(i, i)\}$, $A_{s,2} = \{(i, j) : j \neq i\}$, $A_{s,3} = \{(j, i) : j \neq i\}$, $A_{s,4} = \{(j, k) : j, k \neq i \text{ and } j \neq k\}$, and $A_{s,5} = \{(j, j) : j \neq i\}$. Similar for s' . We can also define $\mathfrak{C}_{s,w}$ and $\mathfrak{C}'_{s',w}$ as above. Note that

$$N_1(s) = \bigcup_{w=1,3} A_{s,w}$$

$$N_2(s) = \bigcup_{w=1,2} A_{s,w}$$

$$N_1(s') = \bigcup_{w=1,3} A_{s',w}$$

$$N_2(s') = \bigcup_{w=1,2} A_{s',w}$$

Hence, we can rewrite (31) as

$$\mathfrak{C}_{s,1} = \mathfrak{C}'_{s',1} \tag{55}$$

$$\bigcup_{w=1,3} \mathfrak{C}_{s,w} = \bigcup_{w=1,3} \mathfrak{C}'_{s',w} \tag{56}$$

$$\bigcup_{w=1,2} \mathfrak{C}_{s,w} = \bigcup_{w=1,2} \mathfrak{C}'_{s',w} \tag{57}$$

Combining (55) with (56), we get

$$\mathfrak{C}_{s,3} = \mathfrak{C}'_{s',3} \tag{58}$$

Combining (55) with (57), we get

$$\mathfrak{C}_{s,2} = \mathfrak{C}'_{s',2} \tag{59}$$

Moreover, since we can decompose V^2 as

$$\begin{aligned}
V^2 &= \left(\bigcup_{w=1,5} A_{s,w} \right) \bigcup \left(\bigcup_{w=2,3,4} A_{s,w} \right) \\
&= \left(\bigcup_{w=1,5} A_{s',w} \right) \bigcup \left(\bigcup_{w=2,3,4} A_{s',w} \right)
\end{aligned}$$

with $\bigcup_{w=1,5} A_{s,w} = \bigcup_{w=1,5} A_{s',w}$ representing the nodes and $\bigcup_{w=2,3,4} A_{s,w} = \bigcup_{w=2,3,4} A_{s',w}$ representing the edges, we have

$$\left(\bigcup_{w=1,5} \mathfrak{C}_{s,w} \right) \cap \left(\bigcup_{w=2,3,4} \mathfrak{C}'_{s',w} \right) = \emptyset \quad (60)$$

$$\left(\bigcup_{w=1,5} \mathfrak{C}'_{s',w} \right) \cap \left(\bigcup_{w=2,3,4} \mathfrak{C}_{s,w} \right) = \emptyset \quad (61)$$

Since $G^{[1]}$ and $G^{[2]}$ cannot be distinguished by 2-WL, there is

$$\bigcup_{w=1}^5 \mathfrak{C}_{s,w} = \bigcup_{w=1}^5 \mathfrak{C}'_{s',w}$$

Therefore, combining this with (60) or (61), we obtain

$$\bigcup_{w=1,5} \mathfrak{C}_{s,w} = \bigcup_{w=1,5} \mathfrak{C}'_{s',w} \quad (62)$$

$$\bigcup_{w=2,3,4} \mathfrak{C}_{s,w} = \bigcup_{w=2,3,4} \mathfrak{C}'_{s',w} \quad (63)$$

Combining (62) with (55), we get

$$\mathfrak{C}_{s,5} = \mathfrak{C}'_{s',5} \quad (64)$$

Combining (63) with (59) and (58), we get

$$\mathfrak{C}_{s,4} = \mathfrak{C}'_{s',4} \quad (65)$$

Hence, in conclusion, we have that $\forall w \in [5]$,

$$\mathfrak{C}_{s,w} = \mathfrak{C}'_{s',w} \quad (66)$$

By the inductive hypothesis, this implies that $\forall w \in [5]$,

$$\{\mathbf{B}_{\tilde{s}}^{(T-1)} : \tilde{s} \in A_{s,w}\} = \{\mathbf{B}'_{\tilde{s}}^{(T-1)} : \tilde{s} \in A_{s',w}\} \quad (67)$$

Thus,

$$\begin{aligned} L(\mathbf{B}^{(T-1)})_{i,i} &= \sum_{(k,l) \in V^2} T_{i,i,k,l} \mathbf{B}_{k,l} + Y_{i,i} \\ &= \sum_{w=1}^5 \sum_{(k,l) \in A_{s,w}} T_{i,i,k,l} \mathbf{B}_{k,l} + Y_{i,i} \\ &= \sum_{w=1}^5 \sum_{(k',l') \in A_{s',w}} T_{i',i',k',l'} \mathbf{B}'_{k',l'} + Y_{i',i'} \\ &= L(\mathbf{B}'^{(T-1)})_{i',i'} \end{aligned}$$

and hence $\mathbf{B}_{i,j}^{(T)} = \mathbf{B}'_{i',j'}^{(T)}$, which concludes the proof for the case that $s = (i, i)$ for $i \in V$. \square

Now, suppose we are given any 2-IGN with T layers. Since $G^{[1]}$ and $G^{[2]}$ cannot be distinguished by 2-WL, together with Lemma D.1, there is

$$\{\mathbf{c}^{(T)}((i, j)) : i, j \in V, i \neq j\} = \{\mathbf{c}'^{(T)}((i', j')) : i', j' \in V, i' \neq j'\}$$

and

$$\{\mathbf{c}^{(T)}((i, i)) : i \in V\} = \{\mathbf{c}'^{(T)}((i', i')) : i' \in V\}$$

Hence, by the lemma, we have

$$\{\mathbf{B}_{(i,j)}^{(T)} : i, j \in V, i \neq j\} = \{\mathbf{B}'_{(i',j')}^{(T)} : i', j' \in V, i' \neq j'\}$$

and

$$\{\mathbf{B}_{(i,i)}^{(T)} : i \in V\} = \{\mathbf{B}'_{(i',i')}^{(T)} : i' \in V\}$$

Then, since the second-last layer h in the 2-IGN can be written as

$$h(\mathbf{B}) = \alpha \sum_{i,j \in V, i \neq j} \mathbf{B}_{i,j} + \beta \sum_{i \in V} \mathbf{B}_{i,i} \quad (68)$$

there is

$$h(\mathbf{B}^{(T)}) = h(\mathbf{B}'^{(T)}) \quad (69)$$

and finally

$$m \circ h(\mathbf{B}^{(T)}) = m \circ h(\mathbf{B}'^{(T)}) \quad (70)$$

which means the 2-IGN yields identical outputs on the two graphs.

J Direct proof of Corollary 4.4 (2-IGNs are unable to induced-subgraph-count patterns of 3 or more nodes)

Proof. The same counterexample as in the proof of Theorem 3.3 given in Appendix E applies here, as we are going to show below. Note that we only need to consider the non-clique case, since the set of counterexample graphs for the non-clique case is a superset of the set of counterexample graphs for the clique case.

Let \mathbf{B} be the input tensor corresponding to $G^{[1]}$, and \mathbf{B}' corresponding to $G^{[2]}$. For simplicity, we assume in the proof below that $d_0, \dots, d_T = 1$. The general case can be proved in the same way but with more subscripts. (In particular, for our counterexamples, (74) can be shown to hold for each of the d_0 feature dimensions.)

Define a set $S = \{(1, 2), (2, 1), (1 + m, 2 + m), (2 + m, 1 + m), (1, 2 + m), (2 + m, 1), (1 + m, 2), (2, 1 + m)\}$, which represents the “special” edges that capture the difference between $G^{[1]}$ and $G^{[2]}$. We aim to show something like this:

$\forall t,$

$$\left\{ \begin{array}{l} \mathbf{B}_{i,j}^{(t)} = \mathbf{B}'_{i,j}^{(t)}, \forall (i, j) \notin S \\ \mathbf{B}_{1,2}^{(t)} = \mathbf{B}'_{1+m,2}^{(t)} \\ \mathbf{B}_{2,1}^{(t)} = \mathbf{B}'_{2,1+m}^{(t)} \\ \mathbf{B}_{1+m,2+m}^{(t)} = \mathbf{B}'_{1,2+m}^{(t)} \\ \mathbf{B}_{2+m,1+m}^{(t)} = \mathbf{B}'_{2+m,1}^{(t)} \\ \mathbf{B}_{1,2+m}^{(t)} = \mathbf{B}'_{1+m,2+m}^{(t)} \\ \mathbf{B}_{2+m,1}^{(t)} = \mathbf{B}'_{2+m,1+m}^{(t)} \\ \mathbf{B}_{1+m,2}^{(t)} = \mathbf{B}'_{1,2}^{(t)} \\ \mathbf{B}_{2,1+m}^{(t)} = \mathbf{B}'_{2,1}^{(t)} \end{array} \right. \quad (71)$$

If this is true, then it is not hard to show that the 2-IGN returns identical outputs on \mathbf{B} and \mathbf{B}' , which we will leave to the very end. To represent the different cases above compactly, we define a permutation η_1 on $V \times V$ in the following way. First, define the following permutations on V :

$$\kappa_1(i) = \begin{cases} \text{MOD}_{2m}(1 + m), & \text{if } i \in \{1, 1 + m\} \\ i, & \text{otherwise} \end{cases}$$

Next, define the permutation τ_1 on $V \times V$:

$$\tau_1((i, j)) = (\kappa_1(i), \kappa_1(j))$$

and then η_1 as the restriction of τ_1 on the set $S \subset V \times V$:

$$\eta_1((i, j)) = \begin{cases} \tau_1((i, j)), & \text{if } (i, j) \in S \\ (i, j), & \text{otherwise} \end{cases}$$

Thus, (71) can be rewritten as

$$\forall t, \mathbf{B}_{i,j}^{(t)} = \mathbf{B}'_{\eta_1((i,j))}^{(t)} \quad (72)$$

Before trying to prove (72), let's define κ_2, τ_2 and η_2 analogously:

$$\begin{aligned} \kappa_2(i) &= \begin{cases} \text{MOD}_{2m}(2+m), & \text{if } i \in \{2, 2+m\} \\ i, & \text{otherwise} \end{cases} \\ \tau_2((i,j)) &= (\kappa_2(i), \kappa_2(j)) \\ \eta_2((i,j)) &= \begin{cases} \tau_2((i,j)), & \text{if } (i,j) \in S \\ (i,j), & \text{otherwise} \end{cases} \end{aligned}$$

Thus, by symmetry, (72) is equivalent to

$$\forall t, \mathbf{B}_{i,j}^{(t)} = \mathbf{B}'_{\eta_1((i,j))}^{(t)} = \mathbf{B}'_{\eta_2((i,j))}^{(t)} \quad (73)$$

Because of the recursive relation (29), we will show (73) by induction on t . For the base case, it can be verified that

$$\mathbf{B}_{i,j}^{(0)} = \mathbf{B}'_{\eta_1((i,j))}^{(0)} = \mathbf{B}'_{\eta_2((i,j))}^{(0)} \quad (74)$$

thanks to the construction of $G^{[1]}$ and $G^{[2]}$. Moreover, if we define another permutation $V \times V, \zeta_1$:

$$\zeta_1((i,j)) = \begin{cases} (\text{MOD}_{2m}(i+m), \text{MOD}_{2m}(j+m)), \\ \quad \text{if } j \in \{1, 1+m\}, i \notin \{2, 2+m\} \\ \quad \text{or } i \in \{1, 1+m\}, j \notin \{2, 2+m\} \\ (i,j), \text{ otherwise} \end{cases} \quad (75)$$

then thanks to the symmetry between (i,j) and $(i+m, j+m)$, there is

$$\mathbf{B}_{i,j}^{(0)} = \mathbf{B}_{\zeta_1((i,j))}^{(0)}, \mathbf{B}'_{i,j}^{(0)} = \mathbf{B}'_{\zeta_1((i,j))}^{(0)}$$

Thus, for the induction to hold, and since σ applies entry-wise, it is sufficient to show that

Lemma J.1. *If*

$$\mathbf{B}_{i,j} = \mathbf{B}_{\zeta_1((i,j))}, \mathbf{B}'_{i,j} = \mathbf{B}'_{\zeta_1((i,j))} \quad (76)$$

$$\mathbf{B}_{i,j} = \mathbf{B}'_{\eta_1((i,j))} = \mathbf{B}'_{\eta_2((i,j))}, \quad (77)$$

then

$$L(\mathbf{B})_{i,j} = L(\mathbf{B})_{\zeta_1((i,j))}, L(\mathbf{B}')_{i,j} = L(\mathbf{B}')_{\zeta_1((i,j))} \quad (78)$$

$$L(\mathbf{B})_{i,j} = L(\mathbf{B}')_{\eta_1((i,j))} = L(\mathbf{B}')_{\eta_2((i,j))}, \quad (79)$$

Proof of Lemma J.1: Again, by symmetry between (i,j) and $(i+m, j+m)$, (78) can be easily shown.

For (79), because of the symmetry between η_1 and η_2 , we will only prove the first equality. By Maron et al. [40], we can express the linear equivariant layer L by

$$L(\mathbf{B})_{i,j} = \sum_{(k,l)=(1,1)}^{(2m,2m)} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j}$$

where crucially, $T_{i,j,k,l}$ depends only on the equivalence class of the 4-tuple (i, j, k, l) .

We consider eight different cases separately.

Case 1 $i, j \notin \{1, 2, 1+m, 2+m\}$

There is $\eta_1((i, j)) = (i, j)$, and $(i, j, k, l) \sim (i, j, \eta_1((k, l)))$, and thus $T_{i,j,k,l} = T_{i,j,\eta_1((k,l))}$. Therefore,

$$\begin{aligned}
L(\mathbf{B}')_{\eta_1((i,j))} &= L(\mathbf{B}')_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}'_{k,l} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{\eta_1((k,l))=(1,1)} T_{i,j,\eta_1((k,l))} \mathbf{B}'_{\eta_1((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,\eta_1((k,l))} \mathbf{B}'_{\eta_1((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}'_{\eta_1((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\
&= \mathbf{B}_{i,j}
\end{aligned}$$

Case 2 $i \in \{1, 1+m\}, j \notin \{1, 2, 1+m, 2+m\}$

There is $\eta_1((i, j)) = (i, j)$, and $(i, j, k, l) \sim (i, j, \eta_2((k, l)))$, because η_2 only involves permutation between nodes 2 and $2+m$, while i and $j \notin \{2, 2+m\}$. Thus, $T_{i,j,k,l} = T_{i,j,\eta_2((k,l))}$. Therefore,

$$\begin{aligned}
L(\mathbf{B}')_{\eta_1((i,j))} &= L(\mathbf{B}')_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}'_{k,l} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{\eta_2((k,l))=(1,1)} T_{i,j,\eta_2((k,l))} \mathbf{B}'_{\eta_2((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,\eta_2((k,l))} \mathbf{B}'_{\eta_2((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}'_{\eta_2((k,l))} + Y_{i,j} \\
&\stackrel{(2m,2m)}{=} \sum_{(k,l)=(1,1)} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\
&= \mathbf{B}_{i,j}
\end{aligned}$$

Case 3 $j \in \{1, 1+m\}, i \notin \{1, 2, 1+m, 2+m\}$

Analogous to case 2.

Case 4 $i \in \{2, 2+m\}, j \notin \{1, 2, 1+m, 2+m\}$

There is $\eta_1((i, j)) = (i, j)$, and $(i, j, k, l) \sim (i, j, \eta_1((k, l)))$, because η_1 only involves permutation between nodes 1 and $1+m$, while i and $j \notin \{1, 1+m\}$. Thus, $T_{i,j,k,l} = T_{i,j,\eta_1((k,l))}$. Therefore, we can apply the same proof as for case 2 here except for changing η_2 's to η_1 's.

Case 5 $j \in \{2, 2 + m\}, i \notin \{1, 2, 1 + m, 2 + m\}$
Analogous to case 4.

Case 6 $(i, j) \in S$
Define one other permutation on $V \times V$, ξ_1 , as $\xi_1((i, j)) =$

$$\begin{cases} (\text{MOD}_{2m}(i + m), j), & \text{if } \text{MOD}_m(j) = 1, \text{MOD}_m(i) \neq 1 \text{ or } 2 \\ (i, \text{MOD}_{2m}(j + m)), & \text{if } \text{MOD}_m(i) = 1, \text{MOD}_m(j) \neq 1 \text{ or } 2 \\ (i, j), & \text{otherwise} \end{cases}$$

It can be verified that

$$\xi_1 \circ \tau_1 = \eta_1 \circ \zeta_1$$

Moreover, it has the property that if $(i, j) \in S$, then

$$(i, j, k, l) \sim (i, j, \xi_1(k, l))$$

because ξ_1 only involves permutations among nodes not in $\{1, 2, 1 + m, 2 + m\}$ while $i, j \in \{1, 2, 1 + m, 2 + m\}$. Thus, we have

$$\begin{aligned} (i, j, k, l) &\sim (\kappa_1(i), \kappa_1(j), \kappa_1(k), \kappa_1(l)) \\ &= (\tau_1(i, j), \tau_1(k, l)) \\ &= (\eta_1(i, j), \tau_1(k, l)) \\ &\sim (\eta_1(i, j), \xi_1 \circ \tau_1(k, l)) \\ &= (\eta_1(i, j), \eta_1 \circ \zeta_1(k, l)), \end{aligned}$$

implying that $T_{i,j,k,l} = T_{\eta_1(i,j), \eta_1 \circ \zeta_1(k,l)}$. In addition, as $\eta_1((i, j)) \sim (i, j)$, there is $Y_{\eta_1((i,j))} = Y_{i,j}$. Moreover, by (76),

$$\mathbf{B}'_{\eta_1 \circ \zeta_1((k,l))} = \mathbf{B}'_{\eta_1((k,l))} = \mathbf{B}_{k,l}$$

Therefore,

$$\begin{aligned} L(\mathbf{B}')_{\eta_1((i,j))} &= \sum_{(k,l)=(1,1)}^{(2m,2m)} T_{\eta_1((i,j)),k,l} \mathbf{B}'_{k,l} + Y_{\eta_1((i,j))} \\ &= \sum_{\eta_1 \circ \zeta_1((k,l))=(1,1)}^{(2m,2m)} T_{\eta_1((i,j)), \eta_1 \circ \zeta_1((k,l))} \mathbf{B}'_{\eta_1 \circ \zeta_1((k,l))} + Y_{\eta_1((i,j))} \\ &= \sum_{(k,l)=(1,1)}^{(2m,2m)} T_{\eta_1((i,j)), \eta_1 \circ \zeta_1((k,l))} \mathbf{B}'_{\eta_1 \circ \zeta_1((k,l))} + Y_{\eta_1((i,j))} \\ &= \sum_{(k,l)=(1,1)}^{(2m,2m)} T_{i,j,k,l} \mathbf{B}_{k,l} + Y_{i,j} \\ &= \mathbf{B}_{i,j} \end{aligned}$$

Case 7 $i, j \in \{1, 1 + m\}$

There is $\eta_1(i, j) = (i, j)$ and $(i, j, k, l) \sim (i, j, \eta_2((k, l)))$. Thus, $T_{i,j,k,l} = T_{i,j,\eta_2((k,l))}$, and the rest of the proof proceeds as for case 2.

Case 8 $i, j \notin \{1, 1 + m\}$

There is $\eta_1(i, j) = (i, j)$ and $(i, j, k, l) \sim (i, j, \eta_1((k, l)))$. Thus, $T_{i,j,k,l} = T_{i,j,\eta_1((k,l))}$, and the rest of the proof proceeds as for case 4.

□

With the lemma above, (72) can be shown by induction as a consequence. Thus,

$$\mathbf{B}_{i,j}^{(T)} = \mathbf{B}_{\eta_1(i,j)}^{(T)}$$

Maron et al. [40] show that the space of linear invariant functions on $\mathbb{R}^{n \times n}$ is two-dimensional, and so for example, the second-last layer h in the 2-IGN can be written as

$$h(\mathbf{B}) = \alpha \sum_{i,j=(1,1)}^{(2m,2m)} \mathbf{B}_{i,j} + \beta \sum_{i=1}^{2m} \mathbf{B}_{i,i}$$

for some $\alpha, \beta \in \mathbb{R}$. Then since η_1 is a permutation on $V \times V$ and also is the identity map when restricted to $\{(i, i) : i \in V\}$, we have

$$\begin{aligned} h(\mathbf{B}'^{(T)}) &= \alpha \sum_{(i,j)=(1,1)}^{(2m,2m)} \mathbf{B}'_{i,j}{}^{(T)} + \beta \sum_{i=1}^{2m} \mathbf{B}'_{i,i}{}^{(T)} \\ &= \alpha \sum_{(i,j)=(1,1)}^{(2m,2m)} \mathbf{B}'_{\eta_1((i,j))}{}^{(T)} + \beta \sum_{i=1}^{2m} \mathbf{B}'_{\eta_1((i,i))}{}^{(T)} \\ &= \alpha \sum_{(i,j)=(1,1)}^{(2m,2m)} \mathbf{B}_{i,j}{}^{(T)} + \beta \sum_{i=1}^{2m} \mathbf{B}_{i,i}{}^{(T)} \\ &= h(\mathbf{B}^{(T)}) \end{aligned}$$

Therefore, finally,

$$m \circ h(\mathbf{B}^{(T)}) = m \circ h(\mathbf{B}'^{(T)})$$

□

K Leveraging sparse tensor operations for LRP

Following our definition of Deep LRP in (4), in each layer, for each egonet $G_{i,l}^{[\text{ego}]}$ and each ordered subset $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$ of nodes in $G_{i,l}^{[\text{ego}]}$, we need to compute the tensor $\tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}(H^{(t-1)})$ out of the hidden node states of the previous layer, $H^{(t-1)}$. This is computationally challenging for stacking multiple layers. Moreover, the tensor operations involved in (4) are dense. In particular, if we batch multiple graphs together, the computational complexity grows quadratically in the number of graphs in a batch, whereas a more reasonable cost would be linear with respect to batch size. In this section, we outline an approach to improve efficiency in implementation via pre-computation and sparse tensor operations. Specifically, we propose to represent the mapping from an H to the set of all $\tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}(H)$'s as a sparse matrix, which can be pre-computed and then applied in every layer. We will also define a similar procedure for the edge features.

The first step is to translate the *local* definitions of $\tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}$ in (4) to a *global* definition. The difference lies in the fact that $\mathbf{B}_{i,l}^{[\text{ego}]}$ implicitly defines a local node index for each node in the egonet, $G_{i,l}^{[\text{ego}]}$ – e.g., $(\mathbf{B}_{i,l}^{[\text{ego}]})_{j,j}$ gives the node feature of the j th node in $G_{i,l}^{[\text{ego}]}$ according to this local index, which is not necessarily the j th node in the whole graph, G . To deal with this notational subtlety, for each ordered subset $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$, we associate with it an ordered subset $\Pi[\tilde{\pi}]$ with elements in V , such that the $(\tilde{\pi}(j))$ th node in $G_{i,l}^{[\text{ego}]}$ according to the local index is indexed to be the $(\Pi[\tilde{\pi}](j))$ th node in the whole graph. Thus, by this definition, we have $\Pi[\tilde{\pi}] \star \mathbf{B} = \tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}$.

Our next task is to efficiently implement the mapping from an H to each $\Pi[\tilde{\pi}] \star \mathbf{B}(H)$. We propose to represent this mapping as a sparse matrix. To illustrate, below we consider the example of Deep LRP- l - k with $l = 1$, and Figure 5 illustrates each step in a layer of Deep LRP-1-3 in particular. For Deep LRP-1- k , each ordered subset $\tilde{\pi} \in \tilde{S}_{i,1}^{k\text{-BFS}}$ consists of $(k + 1)$ nodes, and therefore the first two dimensions of $\Pi[\tilde{\pi}] \star \mathbf{B} = \tilde{\pi} \star \mathbf{B}_{i,l}^{[\text{ego}]}$ are $(k + 1) \times (k + 1)$. We use the following definition of \mathbf{B} , which is slightly simpler than (1) by neglecting the adjacency matrix (whose information is already contained in the edge features): $\mathbf{B} \in \mathbb{R}^{n \times n \times d}$, with $d = \max(d_n, d_e)$, and

$$\begin{aligned} \mathbf{B}_{i,i,1:d_n} &= x_i, \quad \forall i \in V = [n], \\ \mathbf{B}_{i,j,1:d_e} &= e_{i,j}, \quad \forall (i, j) \in E. \end{aligned} \tag{80}$$

Similarly, for $H \in \mathbb{R}^{n \times d'}$, $\mathbf{B}(H)$ is defined to be an element of $\mathbb{R}^{n \times n \times \max(d', d_e)}$, with

$$\begin{aligned} \mathbf{B}_{i,i,1:d'} &= H_i, \quad \forall i \in V = [n], \\ \mathbf{B}_{i,j,1:d_e} &= e_{i,j}, \quad \forall (i,j) \in E. \end{aligned} \quad (81)$$

Below, we assume for the simplicity of presentation that $d_n = d_e = d'$. We let $|E|$ denote the number of edges in G . Define $Y \in \mathbb{R}^{|E| \times d_e}$ to be the matrix of edge features, where Y_q is the feature vector of the q th edge in the graph according to some ordering of the edges. Let P_i be the cardinality of $\tilde{S}_{i,l}^{k\text{-BFS}}$, and define $P = \sum_{i \in [n]} P_i$, where the summation is over all nodes in the graph. Note that these definitions can be generalized to the case where we have a batch of graphs.

We define *Node_to_perm*, denoted by $\mathbf{N2P}$, which is a matrix of size $((k+1)^2 P) \times N$ with each entry being 0 or 1. The first dimension corresponds to the flattening of the first two dimension of $\Pi[\tilde{\pi}] \star \mathbf{B}$ for all legitimate choices of $\tilde{\pi}$. Hence, each row corresponds to one of the $(k+1) \times (k+1)$ ‘‘slots’’ of the first two dimension of $\Pi[\tilde{\pi}] \star \mathbf{B}$ for some $\tilde{\pi}$. In addition, each column of $\mathbf{N2P}$ corresponds to a node in G . Thus, each entry (m, j) of $\mathbf{N2P}$ is 1 if and only if the ‘‘slot’’ indexed by m is filled by the H_j . By the definition of $\mathbf{B}(H)$, $\mathbf{N2P}$ is a sparse matrix. For the edge features, we similarly define *Edge_to_perm*, denoted by $\mathbf{E2P}$, with size $((k+1)^2 P) \times |E|$. Similar to $\mathbf{N2P}$, each entry (m, q) of $\mathbf{E2P}$ is 1 if and only if the ‘‘slot’’ indexed by m is filled by the e_{j_1, j_2} , where (j_1, j_2) is the q th edge. Hence, by these definitions, the list of the vectorizations of all $\tilde{\pi} \star \mathbf{B}_{i,l}^{\text{egol}}(H)$ can be obtained by

$$\text{RESHAPE}\left(\mathbf{N2P} \cdot H + \mathbf{E2P} \cdot Y\right) \in \mathbb{R}^{P \times ((k+1)^2 d)}, \quad (82)$$

where RESHAPE is a tensor-reshapping operation that splits the first dimension from $(k+1)^2 P$ to $P \times (k+1)^2$. Hence, with our choice of f to be an MLP on the vectorization of its tensorial input, the list of all $f(\tilde{\pi} \star \mathbf{B}_{i,l}^{\text{egol}}(H))$ is obtained by

$$\text{MLP1}\left(\text{RESHAPE}\left(\mathbf{N2P} \cdot H + \mathbf{E2P} \cdot Y\right)\right), \quad (83)$$

where MLP1 acts on the second dimension.

Next, we define the $\alpha_{\tilde{\pi}}$ factor as the output of an MLP applied to the relevant node degrees. When $l = 1$, we implement it as $\text{MLP2}(D_i)$, the output of an MLP applied to the degree of the root node i of the egonet. When $k = 1$, since each $\tilde{\pi} \in \tilde{S}_{i,l}^{k\text{-BFS}}$ consists of nodes on a path of length at most $(l+1)$ starting from node i , we let $\alpha_{\tilde{\pi}}$ be the output of an MLP applied to the concatenation of the degrees of all nodes on the path, as discussed in the main text. This step is also compatible with sparse operations similar to (83), in which we substitute H with the degree vector $D \in \mathbb{R}^{n \times 1}$ and neglect Y . In the $l = 1$ case, the list of all the list of all $\alpha_{\tilde{\pi}} f(\tilde{\pi} \star \mathbf{B}_{i,l}^{\text{egol}}(H))$ is obtained by

$$\text{MLP1}\left(\text{RESHAPE}\left(\mathbf{N2P} \cdot H + \mathbf{E2P} \cdot Y\right)\right) \odot \text{MLP2}(D_i) \quad (84)$$

Note that the output dimensions of MLP1 and MLP2 are chosen to be the same, and \odot denotes the element-wise product between vectors.

The final step is to define *Permutation_pooling*, denoted by \mathbf{PPL} , which is a sparse matrix in $\mathbb{R}^{N \times P}$. Each non-zero entry at position (j, p) means that the p -th $\tilde{\pi}$ among all P of them for the whole graph (or a batch of graphs) contributes to the representation of node i in the next layer. In particular, sum-pooling corresponds to setting all non-zero entries in \mathbf{PPL} as 1, while average-pooling corresponds to first setting all non-zero entries in \mathbf{PPL} as 1 and then normalizing it for every row, which is equivalent to having the factor $\frac{1}{|\tilde{S}_{i,1}^{k\text{-BFS}}|}$ in (4).

Therefore, we can now write the update rule (4) for LRP-1- k as

$$H_i^{(t)} = \mathbf{PPL} \cdot \left[\text{MLP1}^{(t)}\left(\text{RESHAPE}\left(\mathbf{N2P} \cdot H^{(t-1)} + \mathbf{E2P} \cdot Y\right)\right) \odot \text{MLP2}^{(t)}(D_i) \right], \quad (85)$$

L Theoretical limitations of GraphSAGE in substructure counting

In order for substructure counting to be well-defined, we do not consider random node sampling and only consider GraphSAGE with aggregation over a full neighborhood. If only 1-hop neighborhood is




Figure 5: Illustration of the t -th local relational pooling layer in Deep LRP-1-3. Rounded rectangles denote representations (Rep) after each operation (denoted as arrows).




Figure 6: A pair of non-isomorphic attributed graphs that GraphSAGE cannot distinguish.

used for aggregation in each iteration, its expressive power is upper-bounded by that of WL, just like MPNNs. If multi-hop neighborhood is used for aggregation, the question becomes more interesting. Compared to LRP, however, GraphSAGE aggregates neighborhood information as a set or sequence rather than a tensor, which results in a loss of the information of the edge features and high-order structures. In particular,

1. The original GraphSAGE does not consider edge features. Even if we allow it to incorporate edge feature information via augmenting the node features by applying an invariant function to the features of its immediate edges (e.g. summing or averaging), GraphSAGE cannot distinguish the pairs of graphs shown in Figure 6, for example, while LRP-1-2 can.
2. GraphSAGE cannot distinguish the pair of 12-circular graphs $C_{12}(1, 3)$ and $C_{12}(1, 5)$ (see [61]), no matter the hop-size being used, because $\forall k$, the k -hop neighborhood of every node

in the two graphs has the same size. This means GraphSAGE cannot count the number of 4-cycles as either *subgraphs* or *induced subgraphs*, whereas LRP-2-2 is able to.

Further, Table 1 shows the performance on the synthetic tasks of GraphSAGE + LSTM using full 1-hop neighborhood for aggregation. We see that it can count stars but not triangles, consistent with the limitation of the information in the 1-hop neighborhood, in the same way as MPNNs.

M Additional details of the numerical experiments

M.1 Models

As reported in Section 6, we run experiments on synthetic and real datasets using different GNN models. Below are some details regarding their architecture and implementation:

- **LRP- l - k** : Local Relational Pooling with egonet depth l and k -truncated BFS, as described in the main text. For LRP-1-3, for example, with d being the dimension of the initial tensor representation, \mathbf{B} , we define

$$\tilde{f}_{\text{LRP}}^{1,3}(G) = \mathbf{W}_1 \sum_{i \in V} \sigma \left[\frac{\text{MLP}(D_i)}{|\tilde{S}_{i,1}^{3\text{-BFS}}|} \odot \sum_{\tilde{\pi} \in \tilde{S}_{i,1}^{3\text{-BFS}}} f_*(\tilde{\pi} \circ \mathbf{B}_{i,1}^{\text{[ego]}}) \right], \quad (86)$$

where D_i is the degree of node i , σ is ReLU, MLP maps from \mathbb{R} to \mathbb{R}^H , where H is the hidden dimension, $\mathbf{W}_1 \in \mathbb{R}^{1 \times H}$ and $\forall p \in [H]$, $(f_*(\mathbf{X}))_p = \tanh(\sum \mathbf{W}_{2,p} \odot \mathbf{X}) \in \mathbb{R}$ with $\mathbf{W}_{2,p} \in \mathbb{R}^{4 \times 4 \times d}$. Note that each $\tilde{\pi} \in \tilde{S}_{i,1}^{3\text{-BFS}}$ is an ordered set of 4 nodes that begin with node i , and $\tilde{\pi} \circ \mathbf{B}_{i,1}^{\text{[ego]}}$ is a $4 \times 4 \times d$ tensor such that $(\tilde{\pi} \circ \mathbf{B}_{i,1}^{\text{[ego]}})_{j,j,:} = (\mathbf{B}_{i,1}^{\text{[ego]}})_{\tilde{\pi}(j),\tilde{\pi}(j),:}$.

As discussed in the main text, $\text{MLP}(D_i)$ plays the role of $\alpha_{\tilde{\pi}}$, which adaptively learns an invariant function over permutation, such as summing and averaging.

The nonlinear activation functions are chosen between ReLU and tanh by hand. The models are trained using the Adam optimizer [30] with learning rate 0.1. The number of hidden dimensions is searched in $\{1, 8, 16, 64, 128\}$.

- **Deep LRP- l - k** : The nonlinear activation functions are ReLU. For synthetic experiments, we set the depth of the model as 1. The number of hidden dimensions is searched in $\{64, 128\}$. We use summation for the final graph-level aggregation function. For real experiments, we search the depth of the model in $\{4, 5, 6, 7, 8, 10, 12, 20, 24\}$. The number of hidden dimensions is searched in $\{8, 16, 32, 50, 100, 128, 150, 200, 256, 300, 512\}$. The final graph-level aggregation function is average. We involve Batch Normalization [23] and Jumping Knowledge [65]. On ogbg-molhiv, we utilize AtomEncoder and BondEncoder following the official implementation of GIN [64] on the OGB leaderboard [21]. The models are trained using the Adam optimizer [30] with learning rate searched in $\{0.01, 0.005, 0.001, 0.0001\}$.
- **2-IGN**: The 2nd-order Invariant Graph Networks proposed by Maron et al. [40]. In our synthetic experiments, we chose 8 hidden dimensions for the invariant layers and 16 hidden dimensions for the output MLP. The models are trained using the Adam optimizer with learning rate 0.1. The numbers of hidden dimensions are searched in $\{(16, 32), (8, 16), (64, 64)\}$.
- **PPGN**: The Provably Powerful Graph Network model proposed in Maron et al. [39]. In our synthetic experiments, we choose the depth of the model to be 4 and select the hidden dimension in $\{16, 64\}$. The models are trained using the Adam optimizer [30] with learning rate searched in $\{0.01, 0.001, 0.0001, 0.00001\}$. The depth of each MLP involved in the model is 2.
- **GCN**: The Graph Convolutional Network proposed by Kipf and Welling [31]. In our experiments, we adopt a 4-layer GCN with 128 hidden dimensions. The models are trained using the Adam optimizer with learning rate 0.01. The number of hidden dimensions is searched in $\{8, 32, 128\}$. The depth is searched in $\{2, 3, 4, 5\}$.
- **GIN**: The Graph Isomorphism Network proposed by Xu et al. [64]. In our experiments, we adopt a 4-layer GIN with 32 hidden dimensions. The models are trained using the Adam optimizer with learning rate 0.01. The number of hidden dimensions is searched in $\{8, 16, 32, 128\}$.

- **sGNN**: Spectral GNN with operators from family $\{\mathbf{I}, \mathbf{A}, \min(\mathbf{A}^2, 1)\}$. In our experiments, we adopt a 4-layer sGNN with 128 hidden dimensions. The models are trained using the Adam optimizer with learning rate 0.01. The number of hidden dimensions is searched in $\{8, 128\}$.
- **GraphSAGE**: GraphSAGE [19] using LSTM [20] for aggregation over the full 1-hop neighborhood. In our experiments, we adopt a 5-layer GraphSAGE with 16 hidden dimensions. The models are trained using the Adam optimizer with learning rate 0.1.

For the experiments on substructure counting in random graphs, for GCN, GIN and sGNN, we always train four variants for each architecture, depending on whether Jump Knowledge [65] or Batch Normalization [23] is included or not. All models are trained for 100 epochs. Learning rates are searched in $\{1, 0.1, 0.05, 0.01\}$. We pick the best model with the lowest MSE loss on validation set to generate results.

M.2 Counting substructures in random graphs

M.2.1 Dataset generation

We generate two synthetic datasets of random unattributed graphs. The first one is a set of 5000 Erdős-Renyi random graphs denoted as $ER(m, p)$, where $m = 10$ is the number of nodes in each graph and $p = 0.3$ is the probability that an edge exists. The second one is a set of 5000 random regular graphs [57] denoted as $RG(m, d)$, where m is the number of nodes in each graph and d is the node degree. We uniformly sample (m, d) from $\{(10, 6), (15, 6), (20, 5), (30, 5)\}$. We also randomly delete m edges in each graph from the second dataset. For both datasets, we randomly split them into training-validation-test sets with percentages 30%-20%-50%. For the attributed task, we mark nodes with even indices as red and nodes with odd indices as blue, and set the color as node feature using 1-hot encoding.

M.2.2 Additional results

For the synthetic experiments, we design five substructure-counting tasks with patterns illustrated in Figure 3. In Section 6, we show the results for the subgraph-count of 3-stars and the induced-subgraph-count of triangles. In this section, we give results for the the remaining patterns: tailed triangles, chordal cycles and attributed triangles. As we see in Table M.2.2, while Deep LRP-1-3 achieves the best overall performance, all three models perform well in learning the induced-subgraph-count of each of these three patterns on at least one of the two synthetic datasets.

Table 5: Performance of the different models on learning the induced-subgraph-count of tailed triangles, chordal cycles and attributed triangles on the two datasets, measured by test MSE divided by variance of the ground truth counts (given in Table 6). Shown here are the best and the median performances of each model over five runs.

	Erdős-Renyi				Random Regular			
	Tailed Triangle		Chordal Cycle		Tailed Triangle		Chordal Cycle	
	top 1	top 3	top 1	top 3	top 1	top 3	top 1	top 3
LRP-1-3	7.61E-5	1.94E-4	5.97E-4	7.73E-4	9.80E-5	2.01E-4	8.19E-5	1.63E-4
Deep LRP-1-3	3.00E-6	1.25E-5	8.03E-6	9.65E-5	1.37E-7	2.25E-5	7.54E-13	3.22E-7
PPGN	7.11E-3	2.03E-2	2.14E-2	1.31E-1	2.29E-3	6.88E-3	5.90E-4	3.12E-2

	Erdős-Renyi		Random Regular	
	Attributed Triangle		Attributed Triangle	
	top 1	top 3	top 1	top 3
LRP-1-3	9.23E-4	2.12E-3	4.50E-1*	4.72E-1*
Deep LRP-1-3	1.48E-4	1.35E-3	9.06E-5	5.05E-4
PPGN	2.58E-5	8.02E-5	4.30E-1*	4.33E-1*

M.3 Molecular prediction tasks

M.3.1 ogbg-molhiv

The molecular dataset ogbg-molhiv from the Open Graph Benchmark (OGB) contains 41127 graphs, with 25.5 nodes and 27.5 edges per graph on average, and the task is to predict 1 target graph-level

Table 6: Variance of the ground truth labels for each synthetic task.

Task	Erdős-Renyi	Random Regular
3-star	311.17	316.13
triangle	7.3441	9.4249
tailed triangle	607.78	1472.93
chordal cycle	86.48	102.58
attributed triangle	2.11	2.71

label. Each graph represents a molecule, where the nodes represent atoms and the edges represent chemical bonds. We use binary cross entropy as the loss function, and we utilize the official APIs including an evaluator provided by OGB (version 1.1.1) [21].

ogbg-molhiv adopts the *scaffold splitting* procedure that splits the data based on their two-dimensional structural frameworks. Because of this, more training epochs might lead to overfitting and therefore worse performance on the test set. Hence, we report the results of LRP-1-3 trained different number of epochs: “LRP-1-3” is trained for 100 epochs, same as other models reported on the OGB leaderboard, and “LRP-1-3 (ES)” is trained for 20 epochs only. To ensure the reproducibility of our results, LRP-1-3 (ES) is run with 35 random seeds, from 0 to 34.

We report the average training time of Deep LRP-1-3 on ogbg-molhiv in Table 7. We can see that Deep LRP-1-3 approximately takes 5-8 \times time as much as GIN. However, the ratio goes down to 3-5 \times when we utilize more numbers of workers to load the data, because the dataloader involves batching operations as defined in Appendix K. We also split the training time for one epoch into several components in Table 8. It turns out the operations **N2P**, **E2P** and **Ppl** account for most of the forward running time, which indicates a possible direction to optimize the current implementation.

Table 7: Training time per epoch for different GNNs on ogbg-molhiv with batch size of 64. All results are generated from a computing node with a GTX 1080Ti, 4 CPUs and 32GB RAM. “#workers” stands for the number of workers in Dataloader of PyTorch. “Ours-2” is the model reported as “Deep LRP-1-3” in Table 2 while “Ours-3” is the model reported as “Deep LRP-1-3 (ES)”.

model	time/epoch (sec)	#params	#workers
GIN	26	189K	0
GIN	26	189K	4
Ours-1	133	166K	0
Ours-1	82	166K	4
Ours-2	136	98K	0
Ours-2	83	98K	4
Ours-3	194	630K	0
Ours-3	122	630K	4

Table 8: Components of training time in an epoch. The setting is the same as that in Table 7.

model	#workers	total time	forward	N2P&E2P	Ppl	backward
Ours-1	4	81.8	39.3	16.4	17.2	13.2

M.3.2 QM9

QM9 has 134K graphs and 12 graph-level target labels for regression. The data is randomly split into 80% for training, 10% for validation and 10% for testing. For training loss, we use the 12-target average of the normalized Mean Absolute Error, where normalization means dividing by the standard deviation of all training labels in the dataset for each of the 12 targets. We report this averaged normalized MAE as “Loss” in the last row Table 4.

M.3.3 ZINC

ZINC [24] is a real-world molecular dataset of 250K graphs. We follow the setting of [15] that selects 12K graphs for regression out of the entire dataset. The dataset is split into 10K/1K/1K for training/validation/testing. We use Mean Absolute Error as the loss for training, validation and testing. Baselines in Table 3 are picked as the best results from [15] regardless of numbers of parameters. Here we also list results with numbers of parameters in Table 9. It turns out our models outperforms all other baselines with the same level of numbers of parameters.

Following [15], we train Deep LRPs with a learning rate scheduler, in which the learning rate decay factor is 0.5 and the patience value for validation loss is 10. The stopping criterion is whether the current learning rate is smaller than 1 percent of the initial learning rate.

Table 9: Additional ZINC test results measured by Mean Absolute Error (MAE). All baselines are taken from [15, 52]. †: Also reported in Table 3.

Model	#Params	Testing MAE	#Params	Testing MAE
MLP	106970	0.681±0.005	2289351	0.704 ±0.003
GCN	103077	0.469±0.002	2189531	0.479±0.007
GraphSAGE	105031	0.410±0.005	2176751	0.439±0.006
GIN	103079	0.408±0.008	2028508	0.382±0.008
DiffPool	110561	0.466±0.006	2291521	0.448±0.005
GAT	102385	0.463±0.002	2080881	0.471±0.005
MoNet	106002	0.407±0.007	2244343	0.372±0.01
GatedGCN	105875	0.363±0.009	2134081	0.338±0.003
LRGA + GatedGCN	94457	0.367±0.008	1989730	0.285±0.01
Deep LRP-7-1	92073	0.317±0.031	1695137	0.244±0.012
Deep LRP-5-1†	-	-	6590593	0.256±0.033
Deep LRP-7-1†	-	-	11183233	0.223±0.008