

---

# HOW POWERFUL ARE GRAPH NEURAL NETWORKS?

**Keyulu Xu** <sup>\*†</sup>  
MIT  
keyulu@mit.edu

**Weihua Hu** <sup>\*‡</sup>  
Stanford University  
weihuahu@stanford.edu

**Jure Leskovec**  
Stanford University  
jure@cs.stanford.edu

**Stefanie Jegelka**  
MIT  
stefje@mit.edu

## ABSTRACT

Graph Neural Networks (GNNs) for representation learning of graphs broadly follow a neighborhood aggregation framework, where the representation vector of a node is computed by recursively aggregating and transforming feature vectors of its neighboring nodes. Many GNN variants have been proposed and have achieved state-of-the-art results on both node and graph classification tasks. However, despite GNNs revolutionizing graph representation learning, there is limited understanding of their representational properties and limitations. Here, we present a theoretical framework for analyzing the expressive power of GNNs in capturing different graph structures. Our results characterize the discriminative power of popular GNN variants, such as Graph Convolutional Networks and GraphSAGE, and show that they cannot learn to distinguish certain simple graph structures. We then develop a simple architecture that is provably the most expressive among the class of GNNs and is as powerful as the Weisfeiler-Lehman graph isomorphism test. We empirically validate our theoretical findings on a number of graph classification benchmarks, and demonstrate that our model achieves state-of-the-art performance.

## 1 INTRODUCTION

Learning with graph structured data, such as molecules, social, biological, and financial networks, requires effective representation of their graph structure (Hamilton et al., 2017b). Recently, there has been a surge of interest in Graph Neural Network (GNN) approaches for representation learning of graphs (Li et al., 2016; Hamilton et al., 2017a; Kipf & Welling, 2017; Velickovic et al., 2018; Xu et al., 2018). GNNs broadly follow a recursive neighborhood aggregation (or message passing) scheme, where each node aggregates feature vectors of its neighbors to compute its new feature vector (Gilmer et al., 2017; Xu et al., 2018). After  $k$  iterations of aggregation, a node is represented by its transformed feature vector, which captures the structural information within the node’s  $k$ -hop network neighborhood. The representation of an entire graph can then be obtained through pooling, for example, by summing the representation vectors of all nodes in the graph.

Many GNN variants with different neighborhood aggregation and graph-level pooling schemes have been proposed (Battaglia et al., 2016; Defferrard et al., 2016; Duvenaud et al., 2015; Hamilton et al., 2017a; Kearnes et al., 2016; Kipf & Welling, 2017; Li et al., 2016; Velickovic et al., 2018; Verma & Zhang, 2018; Ying et al., 2018; Zhang et al., 2018). Empirically, these GNNs have achieved state-of-the-art performance in many tasks such as node classification, link prediction, and graph classification. However, the design of new GNNs is mostly based on empirical intuition, heuristics, and experimental trial-and-error. There is little theoretical understanding of the properties and limitations of GNNs, and formal analysis of GNNs’ representational capacity is limited.

Here, we present a theoretical framework for analyzing the representational power of GNNs. We formally characterize how expressive different GNN variants are in learning to represent and distin-

---

<sup>\*</sup>Equal contribution.

<sup>†</sup>Work partially performed while in Tokyo, visiting Prof. Ken-ichi Kawarabayashi.

<sup>‡</sup>Work partially performed while at RIKEN AIP and University of Tokyo.

guish between different graph structures. Our framework is inspired by the close connection between GNNs and the Weisfeiler-Lehman (WL) graph isomorphism test (Weisfeiler & Lehman, 1968), a powerful test known to distinguish a broad class of graphs (Babai & Kucera, 1979). Similar to GNNs, the WL test iteratively updates a given node’s feature vector by aggregating feature vectors of its network neighbors. What makes the WL test so powerful is its injective aggregation update that maps different node neighborhoods to different feature vectors. Our key insight is that a GNN can have as large discriminative power as the WL test if the GNN’s aggregation scheme is highly expressive and can model injective functions.

To mathematically formalize the above insight, our framework first abstracts the feature vectors of a node’s neighbors as a *multiset*, *i.e.*, a set with possibly repeating elements. Then, the neighbor aggregation in GNNs can be abstracted as a *function over the multiset*. We rigorously study different variants of multiset functions and theoretically characterize their discriminative power, *i.e.*, how well different aggregation functions can distinguish different multisets. The more discriminative the multiset function is, the more powerful the representational power of the underlying GNN.

Our main results are summarized as follows:

- 1) We show that GNNs are *at most* as powerful as the WL test in distinguishing graph structures.
- 2) We establish conditions on the neighbor aggregation and graph pooling functions under which the resulting GNN is *as powerful as* the WL test.
- 3) We identify graph structures that cannot be distinguished by popular GNN variants, such as GCN (Kipf & Welling, 2017) and GraphSAGE (Hamilton et al., 2017a), and we precisely characterize the kinds of graph structures such GNN-based models can capture.
- 4) We develop a simple neural architecture, *Graph Isomorphism Network (GIN)*, and show that its discriminative/representational power is equal to the power of the WL test.

We validate our theory via experiments on graph classification datasets, where the expressive power of GNNs is crucial to capture graph structures. In particular, we compare the performance of GNNs with various aggregation functions. Our results confirm that the most powerful GNN (our Graph Isomorphism Network (GIN)) has high representational power as it almost perfectly fits the training data, whereas the less powerful GNN variants often severely underfit the training data. In addition, the representationally more powerful GNNs outperform the others by test set accuracy and achieve state-of-the-art performance on many graph classification benchmarks.

## 2 PRELIMINARIES

We begin by summarizing some of the most common GNN models and, along the way, introduce our notation. Let  $G = (V, E)$  denote a graph with node feature vectors  $X_v$  for  $v \in V$ . Then are two tasks of interest: (1) *Node classification*, where each node  $v \in V$  has an associated label  $y_v$  and the goal is to learn a representation vector  $h_v$  of  $v$  such that  $v$ ’s label can be predicted as  $y_v = f(h_v)$ ; (2) *Graph classification*, where, given a set of graphs  $\{G_1, \dots, G_N\} \subseteq \mathcal{G}$  and their labels  $\{y_1, \dots, y_N\} \subseteq \mathcal{Y}$ , we aim to learn a representation vector  $h_G$  that helps predict the label of an entire graph,  $y_G = g(h_G)$ .

**Graph Neural Networks.** GNNs use the graph structure and node features  $X_v$  to learn a representation vector of a node,  $h_v$ , or the entire graph,  $h_G$ . Modern GNNs follow a neighborhood aggregation strategy, where we iteratively update the representation of a node by aggregating representations of its neighbors. After  $k$  iterations of aggregation, a node’s representation captures the structural information within its  $k$ -hop network neighborhood. Formally, the  $k$ -th layer of a GNN is

$$a_v^{(k)} = \text{AGGREGATE}^{(k)} \left( \left\{ h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right), \quad h_v^{(k)} = \text{COMBINE}^{(k)} \left( h_v^{(k-1)}, a_v^{(k)} \right), \quad (2.1)$$

where  $h_v^{(k)}$  is the feature vector of node  $v$  at the  $k$ -th iteration/layer. We initialize  $h_v^{(0)} = X_v$ , and  $\mathcal{N}(v)$  is a set of nodes adjacent to  $v$ . The choice of  $\text{AGGREGATE}^{(k)}(\cdot)$  and  $\text{COMBINE}^{(k)}(\cdot)$  in GNNs is crucial. A number of architectures for AGGREGATE have been proposed. In Graph Convolutional Networks (GCN) (Kipf & Welling, 2017), AGGREGATE has been formulated as

$$a_v^{(k)} = \text{MEAN} \left( \left\{ \text{ReLU} \left( W \cdot h_u^{(k-1)} \right), \forall u \in \mathcal{N}(v) \right\} \right) \quad (2.2)$$



Figure 1: Subtree structures at the blue nodes in Weisfeiler-Lehman graph isomorphism test. Two WL iterations can capture and distinguish the structure of rooted subtrees of height 2.

where  $W$  is a learnable matrix. In the pooling variant of GraphSAGE (Hamilton et al., 2017a), the mean operation in Eq. 2.2 is replaced by an element-wise max-pooling. The COMBINE step could be a concatenation followed by a linear mapping  $W \cdot \left[ h_v^{(k-1)} \mid a_v^{(k)} \right]$  as in GraphSAGE. In GCN, the COMBINE step is omitted and the model simply aggregates node  $v$  with its neighbors as  $h_v^{(k)} = \text{AGGREGATE}(\{h_u^{(k-1)}, h_u^{(k-1)} : u \in \mathcal{N}(v)\})$ .

For node classification, the node representation  $h_v^{(K)}$  of the final iteration is used for prediction. For graph classification, the READOUT function aggregates node features from the final iteration to obtain the entire graph’s representation  $h_G$ :

$$h_G = \text{READOUT}(\{h_v^{(K)} \mid v \in G\}). \quad (2.3)$$

READOUT can be a simple permutation invariant function such as summation or a more sophisticated graph-level pooling function (Ying et al., 2018; Zhang et al., 2018).

**Weisfeiler-Lehman test.** The graph isomorphism problem asks whether two graphs are topologically identical. This is a challenging problem: no polynomial-time algorithm is known for it yet (Garey, 1979; Garey & Johnson, 2002; Babai, 2016). Despite some corner cases (Cai et al., 1992), the Weisfeiler-Lehman (WL) test of graph isomorphism (Weisfeiler & Lehman, 1968) is an effective and computationally efficient test that distinguishes a broad class of graphs (Babai & Kucera, 1979). Its 1-dimensional form, “naive vertex refinement”, is analogous to neighborhood aggregation in GNNs. Assuming each node has a categorical label<sup>1</sup>, the WL test iteratively (1) aggregates the labels of nodes and their neighborhoods, and (2) hashes the aggregated labels into *unique* new labels. The algorithm decides that two graphs are different if at some iteration their node labels are different.

Based on the WL test, Shervashidze et al. (2011) proposed the WL subtree kernel that measures the similarity between graphs. The kernel uses the counts of node labels at different iterations of the WL test as the feature vector of a graph. Intuitively, a node’s label at the  $k$ -th iteration of WL test represents a subtree structure of height  $k$  rooted at the node (Figure 1). Thus, the graph features considered by the WL subtree kernel are essentially counts of different rooted subtrees in the graph.

### 3 THEORETICAL FRAMEWORK: OVERVIEW

We start with an overview of our framework for analyzing the expressive power of GNNs. A GNN recursively updates each node’s feature vector to capture the network structure and features of other nodes around it, *i.e.*, its rooted subtree structures (Figure 1). For notational simplicity, we can assign each feature vector a unique label  $\in \{a, b, c, \dots\}$ . Then, feature vectors of a set of neighboring nodes form a *multiset*: the same element can appear multiple times since different nodes can have identical feature vectors.

**Definition 1 (Multiset).** A multiset is a generalized concept of a set that allows multiple instances for its elements. More formally, a multiset is a 2-tuple  $X = (S, m)$  where  $S$  is the *underlying set* of  $X$  that is formed from its *distinct elements*, and  $m : S \rightarrow \mathbb{N}_{\geq 1}$  gives the *multiplicity* of the elements.

In order to analyze the representational power of a GNN, we analyze when a GNN maps two nodes into the same location in the embedding space. Intuitively, the most powerful GNN maps two nodes to the same location *only if* they have identical subtree structures with identical features on the corresponding nodes. Since subtree structures are defined recursively via node neighborhoods (Figure 1), we can reduce our analysis recursively to the question when a GNN maps two neighborhoods to the same embedding. The most powerful GNN would *never* map two different neighborhoods, *i.e.*, multisets of feature vectors, to the same location. This means its aggregation scheme is *injective*. Thus, we

<sup>1</sup>In case each node has a feature vector, an injective function is used to map it to a categorical node label.

abstract a GNN’s aggregation scheme as a class of functions over multisets that its neural networks can represent, and analyze whether they are able to represent injective multiset functions. Next, we use this reasoning to develop a maximally powerful GNN. In Section 5, we study popular GNN variants and see that their aggregation schemes are inherently not injective and thus less powerful, but that they can capture other interesting properties of graphs.

## 4 GENERALIZING THE WL TEST WITH GRAPH NEURAL NETWORKS

Ideally, a representationally powerful GNN could distinguish different graphs by mapping them to different locations in the embedding space. This is, however, equivalent to solving graph isomorphism. In our analysis, we characterize the representational capacity of GNNs via a slightly weaker criterion: the *Weisfeiler-Lehman (WL) graph isomorphism test* that is known to work well in general, with some few exceptions. Proofs of all lemmas and theorems can be found in the appendix.

**Lemma 2.** *Let  $G_1$  and  $G_2$  be any non-isomorphic graphs. If a graph neural network  $A : \mathcal{G} \rightarrow \mathbb{R}^d$  following the neighborhood aggregation scheme maps  $G_1$  and  $G_2$  to different embeddings, the Weisfeiler-Lehman graph isomorphism test also decides  $G_1$  and  $G_2$  are not isomorphic.*

Hence, any aggregation-based GNN is at most as powerful as the WL test in distinguishing different graphs. A natural follow-up question is whether there exist GNNs that are, in principle, as powerful as the WL test? Our answer, in Theorem 3, is yes: if the neighbor aggregation and graph pooling functions are injective, then the resulting GNN is as powerful as the WL test.

**Theorem 3.** *Let  $A : \mathcal{G} \rightarrow \mathbb{R}^d$  be a GNN following the neighborhood aggregation scheme. With sufficient iterations,  $A$  maps any graphs  $G_1$  and  $G_2$  that the Weisfeiler-Lehman test of isomorphism decides as non-isomorphic, to different embeddings if the following conditions hold:*

a)  $A$  aggregates and updates node features iteratively with

$$h_v^{(k)} = \phi \left( h_v^{(k-1)}, f \left( \left\{ h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right) \right) \text{ or } h_v^{(k)} = f \left( \left\{ h_v^{(k-1)}, h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right)$$

where the functions  $f$ , which operates on multisets, and  $\phi$  are injective.

b)  $A$ ’s graph-level readout, which operates on the multiset of node features  $\{h_v^{(k)}\}$ , is injective.

We prove Theorem 3 in the appendix. Generally note that GNNs have an important benefit over the WL test: node feature vectors in the WL test are one-hot encodings and thus cannot capture the similarity between subtrees. In contrast, a GNN satisfying the criteria in Theorem 3 generalizes the WL test by *learning to embed* the subtrees to continuous space. This enables GNNs to not only discriminate different structures, but also to learn to map similar graph structures to similar embeddings and capture dependencies between graph structures. Such learned embeddings are particularly helpful for generalization when the co-occurrence of subtrees is sparse across different graphs or there are noisy edges (Yanardag & Vishwanathan, 2015).

### 4.1 GRAPH ISOMORPHISM NETWORK (GIN)

Next we develop a model that provably satisfies the conditions in Theorem 3 and thus generalizes the WL test. We name the resulting architecture *Graph Isomorphism Network (GIN)*.

To model injective multiset functions for the neighbor aggregation, we develop a theory of “deep multisets”, *i.e.*, parameterizing universal multiset functions with neural networks. Our next lemma states that sum aggregators can represent injective, in fact, universal functions over multisets.

**Lemma 4.** *Assume  $\mathcal{X}$  is countable. There exists a function  $f : \mathcal{X} \rightarrow \mathbb{R}^n$  so that  $h(X) = \sum_{x \in X} f(x)$  is unique for each finite multiset  $X \subset \mathcal{X}$ . Moreover, any multiset function  $g$  can be decomposed as  $g(X) = \phi \left( \sum_{x \in X} f(x) \right)$  for some function  $\phi$ .*

We prove Lemma 4 in the appendix. The proof extends the setting in (Zaheer et al., 2017) from sets to multisets. An important distinction between deep multisets and sets is that certain popular injective set functions, such as the mean aggregator, are not injective multiset functions. Thanks to the universal approximation theorem (Hornik et al., 1989; Hornik, 1991), we can use multi-layer perceptrons (MLPs) to model and learn  $f$  and  $\phi$  in Lemma 4 for universal injective embeddings. In practice, we model  $f^{(k+1)} \circ \phi^{(k)}$  with one MLP, because MLPs can represent the composition of




Figure 2: Ranking by expressive power for sum, mean and max-pooling aggregators over a multiset. Left panel shows the input multiset and the three panels illustrate the aspects of the multiset a given aggregator is able to capture: sum captures the full multiset, mean captures the proportion/distribution of elements of a given type, and the max aggregator ignores multiplicities (reduces the multiset to a simple set).




Figure 3: Examples of simple graph structures that mean and max-pooling aggregators fail to distinguish. Figure 2 gives reasoning about how different aggregators “compress” different graph structures/multisets.

functions. In the first iteration, we do not need MLPs before summation if input features are one-hot encodings as their summation alone is injective. GIN updates node representations as

$$h_v^{(k)} = \text{MLP}^{(k)}\left(h_v^{(k-1)} + \sum_{u \in \mathcal{N}(v)} h_u^{(k-1)}\right) \quad (4.1)$$

In contrast to GNNs, which combine a node’s feature with its aggregated neighborhood feature as in Eq. 2.1, GIN does not have the combine step and simply aggregates the node *along* with its neighbors. Although not intuitive, Theorem 3 suggests this simple scheme is *as powerful as* more complicated ones. Experimentally we observe that such simplicity improves performance.

#### 4.2 READOUT SUBTREE STRUCTURES OF DIFFERENT DEPTHS

An important aspect of the graph-level readout is that node representations, corresponding to subtree structures, get more refined and global as the number of iterations increases. A sufficient number of iterations is key to achieving good discriminative power. Yet, features from earlier iterations may sometimes generalize better. To consider all structural information, GIN uses information from all depths/iterations of the model. We achieve this by an architecture similar to Jumping Knowledge Networks (JK-Nets) (Xu et al., 2018), where we replace Eq. 2.3 with graph representations concatenated across *all iterations*:

$$h_G = \text{CONCAT}\left(\text{READOUT}\left(\left\{h_v^{(k)} \mid v \in G\right\}\right) \mid k = 0, 1, \dots, K\right). \quad (4.2)$$

By Theorem 3 and Lemma 4, if GIN replaces READOUT in Eq. 4.2 with summing all node features from the same iterations (we do not need an extra MLP before summation for the same reason as in Eq. 4.1), it provably generalizes the WL test and the WL subtree kernel.

### 5 LESS POWERFUL BUT STILL INTERESTING GNNs

Next we study GNNs that do not satisfy the conditions in Theorem 3, including GCN (Kipf & Welling, 2017) and GraphSAGE (Hamilton et al., 2017a). We conduct ablation studies on two aspects of the aggregator in Eq. 4.1: (1) 1-layer perceptrons instead of MLPs and (2) mean or max-pooling instead of the sum. We will see that these GNN variants get confused by surprisingly simple graphs and are less powerful than the WL test. Nonetheless, models with mean aggregators like GCN perform well for *node classification* tasks. To better understand this, we precisely characterize what different GNN variants can and cannot capture about a graph and discuss the implications for learning with graphs.

## 5.1 1-LAYER PERCEPTRON IS INSUFFICIENT FOR CAPTURING STRUCTURES

The function  $f$  in Lemma 4 helps map distinct multisets to unique embeddings. It can be parameterized by an MLP by the universal approximation theorem (Hornik, 1991). Nonetheless, many existing GNNs instead use a 1-layer perceptron  $\sigma \circ W$  (Duvenaud et al., 2015; Kipf & Welling, 2017; Zhang et al., 2018), a linear mapping followed by a non-linear activation function such as a ReLU. Such 1-layer mappings are examples of Generalized Linear Models (Nelder & Wedderburn, 1972). Therefore, we are interested in understanding whether 1-layer perceptrons are enough for graph learning. Lemma 5 suggests that there are indeed network neighborhoods (multisets) that models with 1-layer perceptrons can never distinguish.

**Lemma 5.** *There exist finite multisets  $X_1 \neq X_2$  so that for any linear mapping  $W$ ,  $\sum_{x \in X_1} \text{ReLU}(Wx) = \sum_{x \in X_2} \text{ReLU}(Wx)$ .*

The main idea of the proof for Lemma 5 is that 1-layer perceptrons can behave much like linear mappings, so the GNN layers degenerate into simply summing over neighborhood features. GNNs with 1-layer perceptrons lack representational capacity, and, as we will later see empirically, when applied to graph classification they may severely underfit, whereas GNNs with MLPs usually do not.

## 5.2 STRUCTURES THAT CONFUSE MEAN AND MAX-POOLING

What happens if we replace the sum in  $h(X) = \sum_{x \in X} f(x)$  with mean or max-pooling as in GCN and GraphSAGE? Mean and max-pooling aggregators are still well-defined multiset functions because they are permutation invariant. But, they are *not* injective. Figure 2 ranks the three aggregators by their representational power, and Figure 3 illustrates pairs of structures that the mean and max-pooling aggregators fail to distinguish. Here, node colors denote different node features, and we assume the GNNs aggregate neighbors first before combining them with the central node.

In Figure 3a, every node has the same feature  $a$  and  $f(a)$  is the same across all nodes (for any function  $f$ ). When performing neighborhood aggregation, the mean or maximum over  $f(a)$  remains  $f(a)$  and, by induction, we always obtain the same node representation everywhere. Thus, mean and max-pooling aggregators fail to capture any structural information. In contrast, a sum aggregator distinguishes the structures because  $2 \cdot f(a)$  and  $3 \cdot f(a)$  give different values. The same argument can be applied to any unlabeled graph. If node degrees instead of a constant value is used as node input features, in principle, mean can recover sum, but max-pooling cannot.

Fig. 3a suggests that mean and max have trouble distinguishing graphs with nodes that have repeating features. Let  $h_{\text{color}}$  ( $r$  for red,  $g$  for green) denote node features transformed by  $f$ . Fig. 3b shows that maximum over the neighborhood of the blue nodes yields  $\max(h_g, h_r)$  and  $\max(h_g, h_r, h_r)$ , which collapse to the same representation. Thus, max-pooling fails to distinguish them. In contrast, the sum aggregator still works because  $\frac{1}{2}(h_g + h_r)$  and  $\frac{1}{3}(h_g + h_r + h_r)$  are in general not equivalent. Similarly, in Fig. 3c, both mean and max fail as  $\frac{1}{2}(h_g + h_r) = \frac{1}{4}(h_g + h_g + h_r + h_r)$ .

## 5.3 MEAN LEARNS DISTRIBUTIONS

To characterize the class of multisets that the mean aggregator can distinguish, consider the example  $X_1 = (S, m)$  and  $X_2 = (S, k \cdot m)$ , where  $X_1$  and  $X_2$  have the same set of distinct elements, but  $X_2$  contains  $k$  copies of each element of  $X_1$ . Any mean aggregator maps  $X_1$  and  $X_2$  to the same embedding, because it simply takes averages over individual element features. Thus, the mean captures the *distribution* (proportions) of elements in a multiset, but not the *exact* multiset.

**Corollary 6.** *Assume  $\mathcal{X}$  is countable. There exists a function  $f : \mathcal{X} \rightarrow \mathbb{R}^n$  so that for  $h(X) = \frac{1}{|X|} \sum_{x \in X} f(x)$ ,  $h(X_1) = h(X_2)$  if and only if finite multisets  $X_1$  and  $X_2$  have the same distribution. That is, assuming  $|X_2| \geq |X_1|$ , we have  $X_1 = (S, m)$  and  $X_2 = (S, k \cdot m)$  for some  $k \in \mathbb{N}_{\geq 1}$ .*

The mean aggregator may perform well if, for the task, the statistical and distributional information in the graph is more important than the exact structure. Moreover, when node features are diverse and rarely repeat, the mean aggregator is as powerful as the sum aggregator. This may explain why, despite the limitations identified in Section 5.2, GNNs with mean aggregators are effective for node classification tasks, such as classifying article subjects and community detection, where node features are rich and the distribution of the neighborhood features provides a strong signal for the task.

---

## 5.4 MAX-POOLING LEARNS SETS WITH DISTINCT ELEMENTS

The examples in Figure 3 illustrate that max-pooling considers multiple nodes with the same feature as *only one* node (*i.e.*, treats a multiset as a set). Max-pooling captures neither the exact structure nor the distribution. However, it may be suitable for tasks where it is important to identify representative elements or the “skeleton”, rather than to distinguish the exact structure or distribution. Qi et al. (2017) empirically show that the max-pooling aggregator learns to identify the skeleton of a 3D point cloud and that it is robust to noise and outliers. For completeness, the next corollary shows that the max-pooling aggregator captures the underlying set of a multiset.

**Corollary 7.** *Assume  $\mathcal{X}$  is countable. Then there exists a function  $f : \mathcal{X} \rightarrow \mathbb{R}^\infty$  so that for  $h(X) = \max_{x \in X} f(x)$ ,  $h(X_1) = h(X_2)$  if and only if  $X_1$  and  $X_2$  have the same underlying set.*

## 6 EXPERIMENTS

We evaluate and compare the training and test performance of GIN and less powerful GNN variants.

**Datasets.** We use 9 graph classification benchmarks: 4 bioinformatics datasets (MUTAG, PTC, NCI1, PROTEINS) and 5 social network datasets (COLLAB, IMDB-BINARY, IMDB-MULTI, REDDIT-BINARY and REDDIT-MULTI5K) (Kersting et al., 2016). In the bioinformatic graphs, the nodes have categorical input features; in the social networks, they have no features. For the REDDIT datasets, we set all node feature vectors to be the same (thus, features here are uninformative); for the other social graphs, we use one-hot encodings of node degrees. Dataset statistics are summarized in Table 1, and more details of the data can be found in Appendix G.

**Models and configurations.** We evaluate GIN (Eqs. 4.1 and 4.2) and five less powerful GNN variants: We consider architectures that replace the sum in Eq. 4.1 with mean or max-pooling<sup>2</sup>, or replace MLPs with 1-layer perceptrons, *i.e.*, a linear mapping followed by ReLU. In Figure 4 and Table 1, a model is named by the aggregator/perceptron it uses. We apply the same graph-level readout (READOUT in Eq. 4.2) for GINs and all GNN variants, specifically, sum readout on bioinformatics datasets and mean readout on social datasets due to better test performance.

Following (Yanardag & Vishwanathan, 2015; Niepert et al., 2016), we perform 10-fold cross-validation with LIB-SVM (Chang & Lin, 2011), using 9 folds for training and 1 for testing. For all configurations, 5 GNN layers (including the input layer) are applied, and all MLPs have 2 layers. Batch normalization (Ioffe & Szegedy, 2015) is applied on every hidden layer. We use the Adam optimizer (Kingma & Ba, 2015) with initial learning rate 0.01 and decay the learning rate by 0.5 every 50 epochs. The hyper-parameters we tune for each dataset are: (1) The number of hidden units  $\in \{16, 32\}$  for bioinformatics graphs and 64 for social graphs; (2) batch size  $\in \{32, 128\}$ ; (3) dropout ratio  $\in \{0, 0.5\}$  after the dense layer (Srivastava et al., 2014); (4) the number of epochs.

**Baselines.** We compare the GNNs above with a number of state-of-the-art baselines for graph classification: (1) The WL subtree kernel (Shervashidze et al., 2011), where  $C$ -SVM (Chang & Lin, 2011) was used as a classifier. The hyper-parameters we tune are  $C$  in SVM and the number of WL iterations  $\in \{1, 2, \dots, 6\}$ ; (2) state-of-the-art deep learning architectures Diffusion-convolutional neural networks (DCNN) (Atwood & Towsley, 2016), PATCHY-SAN (Niepert et al., 2016) and Deep Graph CNN (DGCNN) (Zhang et al., 2018); (3) Anonymous Walk Embeddings (AWL) (Ivanov & Burnaev, 2018). For the deep learning methods and AWL, we report the accuracies reported in the original papers.

### 6.1 RESULTS

**Training set performance.** We validate our theoretical analysis of the representational power of GNNs by comparing their training accuracies. Figure 4 shows training curves of GINs and less powerful GNN variants with the same hyper-parameter settings. First, the theoretically most powerful GNN, *i.e.* GIN (Sum-MLP), is able to almost perfectly fit all training sets. In comparison, the less powerful GNN variants severely underfit on many datasets. In particular, the training accuracy patterns align with our ranking by the models’ representational power: GNN variants with MLPs

---

<sup>2</sup>For REDDIT-BINARY, REDDIT-MULTI5K, and COLLAB, we did not run experiments for max-pooling due to GPU memory constraints.

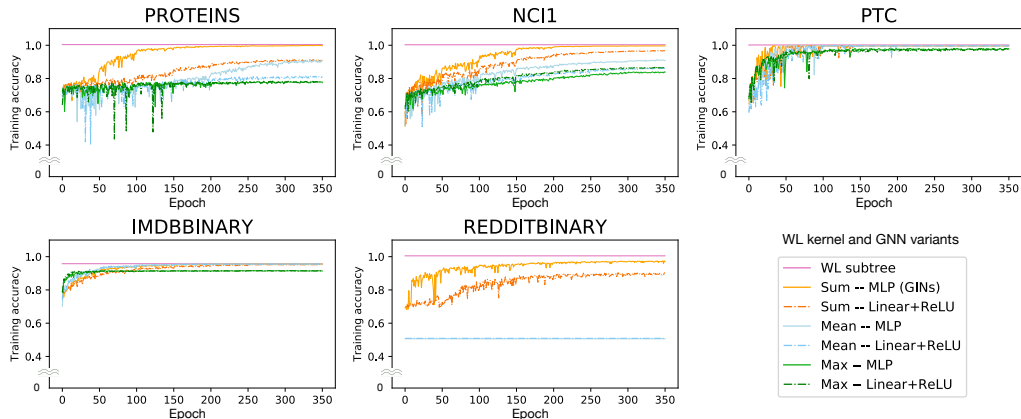


Figure 4: Training set performance of GINs, less powerful GNN variants, and the WL subtree kernel.

	IMDB-B	IMDB-M	RDT-B	RDT-M5K	COLLAB	MUTAG	PROTEINS	PTC	NC11
<b>Datasets</b>									
# graphs	1000	1500	2000	5000	5000	188	1113	344	4110
# classes	2	3	2	5	3	2	2	2	2
Avg # nodes	19.8	13.0	429.6	508.5	74.5	17.9	39.1	25.5	29.8
<b>Baselines</b>									
WL subtree	73.8	50.9	81.0	52.5	78.9	90.4	75.0	59.9	<b>86.0*</b>
DCNN	49.1	33.5	–	–	52.1	67.0	61.3	56.6	62.6
PATCHYSAN	71.0	45.2	86.3	49.1	72.6	<b>92.6*</b>	75.9	60.0	78.6
DGCNN	70.0	47.8	–	–	73.7	85.8	75.5	58.6	74.4
AWL	74.6	51.6	87.9	54.7	73.9	87.9	–	–	–
<b>GNN variants</b>									
GIN (SUM-MLP)	<b>75.1</b>	<b>52.3</b>	<b>92.4</b>	<b>57.5</b>	<b>80.2</b>	<b>89.4</b>	<b>76.2</b>	<b>64.6</b>	<b>82.7</b>
SUM-1-LAYER	74.1	52.2	90.0	55.1	<b>80.6</b>	<b>90.0</b>	<b>76.2</b>	63.1	82.0
MEAN-MLP	73.7	<b>52.3</b>	50.0 <sup>†</sup> (71.2)	20.0 <sup>†</sup> (41.3)	79.2	83.5	75.5	<b>66.6</b>	80.9
MEAN-1-LAYER	74.0	51.9	50.0 <sup>†</sup> (69.7)	20.0 <sup>†</sup> (39.7)	79.0	85.6	76.0	64.2	80.2
MAX-MLP	73.2	51.1	–	–	–	84.0	76.0	64.6	77.8
MAX-1-LAYER	72.3	50.9	–	–	–	85.1	75.9	63.9	77.7

Table 1: Classification accuracies (%). <sup>†</sup> indicate test accuracies (equal to chance rates) when all nodes have the same feature vector. We also report in the parentheses the test accuracies when the node degrees are used as input node features. The best-performing GNNs are highlighted with boldface. On datasets where GIN’s accuracy is not strictly the highest among GNN variants, GIN is comparable to the best because paired t-test at significance level 10% does not distinguish GIN from the best. If a baseline performs better than all GNNs, we highlight it with boldface and asterisk.

tend to have higher training accuracies than those with 1-layer perceptrons, and GNNs with sum aggregators tend to fit the training sets better than those with mean and max-pooling aggregators.

On our datasets training accuracies of the GNNs, however, never exceed those of the WL subtree kernel, which has the same discriminative power as the WL test. For example, on IMDBBINARY, none of the models can perfectly fit the training set, and the GNNs achieve at most the same training accuracy as the WL kernel. This pattern aligns with our result that the WL test provides an upper bound for the representational capacity of the aggregation-based GNNs. Our theoretical results focus on representational power and do not yet take into account optimization (*e.g.*, local minima). Nonetheless, the empirical results align very well with our theory.

**Test set performance.** Next, we compare test accuracies. Although our theoretical results do not directly speak about generalization ability of GNNs, it is reasonable to expect that GNNs with strong expressive power can accurately capture graph structures of interest and thus generalize well. Table 1 compares test accuracies of GINs (Sum-MLP), other GNN variants, as well as the state-of-the-art baselines.

First, GINs outperform (or achieve comparable performance as) the less powerful GNN variants on all the 9 datasets, achieving state-of-the-art performance. In particular, GINs shine on the social network datasets, which contain a relatively large number of training graphs. On the Reddit datasets, a random vector was used as a node feature. Here GINs and sum-aggregation GNNs accurately capture the graph structure (as predicted in Section 5.2) and significantly outperform other models. Mean-aggregation GNNs, however, fail to capture structural information and do not perform better



---

than random guessing. Even if node degrees are provided as input features, mean-based GNNs perform much worse than sum-based GNNs.

## 7 CONCLUSION

In this paper, we developed theoretical foundations for reasoning about expressive power of GNNs and proved tight bounds on the representational capacity of popular GNN variants. Along the way, we also designed a provably most powerful GNN under the neighborhood aggregation framework. An interesting direction for future work is to go beyond the neighborhood aggregation (or message passing) framework in order to pursue even more powerful architectures for learning with graphs. It would also be interesting to understand and improve the generalization properties of GNNs.

## ACKNOWLEDGMENTS

This research was supported by NSF CAREER award 1553284 and a DARPA D3M award. This research was also supported in part by NSF, ARO MURI, IARPA HFC, Boeing, Huawei, Stanford Data Science Initiative, and Chan Zuckerberg Biohub.

We *gratefully* thank Prof. Ken-ichi Kawarabayashi and Prof. Masashi Sugiyama for generously supporting this research with GPU computing resources, as well as giving us wonderful advices. We thank Tomohiro Sonobe and Kento Nozawa for their great management of the GPU servers used in this research. We thank Dr. Yasuo Tabei for inviting Keyulu for giving a talk at RIKEN AIP in Nihonbashi, Tokyo, which led to the collaboration of this research. We thank Jingling Li for initiating and arranging the collaboration of this research, as well as her great inspiration and discussions. We thank Rex Ying and William Hamilton for helpful reviews and positive comments on our paper. We thank Chengtao Li for helpful discussions on the title of the paper. Finally, we thank Simon S. Du for his very helpful discussions and positive comments on our work.

## REFERENCES

- James Atwood and Don Towsley. Diffusion-convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1993–2001, 2016.
- László Babai. Graph isomorphism in quasipolynomial time. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 684–697. ACM, 2016.
- László Babai and Ludik Kucera. Canonical labelling of graphs in linear average time. In *Foundations of Computer Science, 1979., 20th Annual Symposium on*, pp. 39–46. IEEE, 1979.
- Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 4502–4510, 2016.
- Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.
- Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):27, 2011.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3844–3852, 2016.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. pp. 2224–2232, 2015.
- Michael R Garey. A guide to the theory of np-completeness. *Computers and intractability*, 1979.
- Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.

- 
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, pp. 1273–1272, 2017.
- William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1025–1035, 2017a.
- William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *IEEE Data Engineering Bulletin*, 40(3):52–74, 2017b.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2): 251–257, 1991.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pp. 448–456, 2015.
- Sergey Ivanov and Evgeny Burnaev. Anonymous walk embeddings. In *International Conference on Machine Learning (ICML)*, pp. 2191–2200, 2018.
- Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8): 595–608, 2016.
- Kristian Kersting, Nils M. Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. Benchmark data sets for graph kernels, 2016. URL <http://graphkernels.cs.tu-dortmund.de>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- J. A. Nelder and R. W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society, Series A, General*, 135:370–384, 1972.
- Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International Conference on Machine Learning (ICML)*, pp. 2014–2023, 2016.
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proc. Computer Vision and Pattern Recognition (CVPR)*, *IEEE*, 1(2):4, 2017.
- Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(Sep): 2539–2561, 2011.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Saurabh Verma and Zhi-Li Zhang. Graph capsule convolutional neural networks. *arXiv preprint arXiv:1805.08090*, 2018.

- 
- Boris Weisfeiler and AA Lehman. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Tekhnicheskaya Informatsia*, 2(9):12–16, 1968.
- Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning (ICML)*, pp. 5453–5462, 2018.
- Pinar Yanardag and SVN Vishwanathan. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1365–1374. ACM, 2015.
- Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in Neural Information Processing Systems*, pp. 3391–3401, 2017.
- Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *AAAI Conference on Artificial Intelligence*, pp. 4438–4445, 2018.

## A PROOF FOR LEMMA 2

*Proof.* Suppose after  $k$  iterations, a graph neural network  $\mathcal{A}$  has  $\mathcal{A}(G_1) \neq \mathcal{A}(G_2)$  but the WL test cannot decide  $G_1$  and  $G_2$  are non-isomorphic. It follows that from iteration 0 to  $k$  in the WL test,  $G_1$  and  $G_2$  always have the same collection of node labels. In particular, because  $G_1$  and  $G_2$  have the same WL node labels for iteration  $i$  and  $i + 1$  for any  $i = 0, \dots, k - 1$ ,  $G_1$  and  $G_2$  have the same collection, i.e. multiset, of WL node labels  $\{l_v^{(i)}\}$  as well as the same collection of node neighborhoods  $\left\{ \left( l_v^{(i)}, \{l_u^{(i)} : u \in \mathcal{N}(v)\} \right) \right\}$ . Otherwise, the WL test would have obtained different collections of node labels at iteration  $i + 1$  for  $G_1$  and  $G_2$  as different multisets get unique new labels.

The WL test always relabels different multisets of neighboring nodes into different new labels. We show that on the same graph  $G = G_1$  or  $G_2$ , if WL node labels  $l_v^{(i)} = l_u^{(i)}$ , we always have GNN node features  $h_v^{(i)} = h_u^{(i)}$  for any iteration  $i$ . This apparently holds for  $i = 0$  because WL and GNN starts with the same node features. Suppose this holds for iteration  $j$ , if for any  $u, v$ ,  $l_v^{(j+1)} = l_u^{(j+1)}$ , then it must be the case that

$$\left( l_v^{(j)}, \{l_w^{(j)} : w \in \mathcal{N}(v)\} \right) = \left( l_u^{(j)}, \{l_w^{(j)} : w \in \mathcal{N}(u)\} \right)$$

By our assumption on iteration  $j$ , we must have

$$\left( h_v^{(j)}, \{h_w^{(j)} : w \in \mathcal{N}(v)\} \right) = \left( h_u^{(j)}, \{h_w^{(j)} : w \in \mathcal{N}(u)\} \right)$$

In the aggregation process of the GNN, the same AGGREGATE and COMBINE are applied. The same input, i.e. neighborhood features, generates the same output. Thus,  $h_v^{(j+1)} = h_u^{(j+1)}$ . By induction, if WL node labels  $l_v^{(i)} = l_u^{(i)}$ , we always have GNN node features  $h_v^{(i)} = h_u^{(i)}$  for any iteration  $i$ . This creates a valid mapping  $\phi$  such that  $h_v^{(i)} = \phi(l_v^{(i)})$  for any  $v \in G$ . It follows from  $G_1$  and  $G_2$  have the same multiset of WL neighborhood labels that  $G_1$  and  $G_2$  also have the same collection of GNN neighborhood features

$$\left\{ \left( h_v^{(i)}, \{h_u^{(i)} : u \in \mathcal{N}(v)\} \right) \right\} = \left\{ \left( \phi(l_v^{(i)}), \{ \phi(l_u^{(i)}) : u \in \mathcal{N}(v) \} \right) \right\}$$

Thus,  $\{h_v^{(i+1)}\}$  are the same. In particular, we have the same collection of GNN node features  $\{h_v^{(k)}\}$  for  $G_1$  and  $G_2$ . Because the graph level readout function is permutation invariant with respect to the collection of node features,  $\mathcal{A}(G_1) = \mathcal{A}(G_2)$ . Hence we have reached a contradiction.  $\square$

## B PROOF FOR THEOREM 3

*Proof.* Let  $\mathcal{A}$  be a graph neural network where a) and b) hold. Let  $G_1, G_2$  be any graphs which the WL test decides as non-isomorphic at iteration  $K$ . Because the graph-level readout function is injective, i.e. it maps distinct multiset of node features into unique embeddings, it suffices to show that  $\mathcal{A}$ 's neighborhood aggregation process, with sufficient iterations, embeds  $G_1$  and  $G_2$  into different multisets of node features. Let us first assume  $\mathcal{A}$  updates node representations as

$$h_v^{(k)} = \phi \left( h_v^{(k-1)}, f \left( \{h_u^{(k-1)} : u \in \mathcal{N}(v)\} \right) \right)$$

with injective functions  $f$  and  $\phi$ . The WL test applies a predetermined injective hash function  $g$  to update the WL node labels  $l_v^{(k)}$ .

$$l_v^{(k)} = g \left( l_v^{(k-1)}, \{l_u^{(k-1)} : u \in \mathcal{N}(v)\} \right)$$

We will show, by induction, that for any iteration  $k$ , there always exists an injective function  $\varphi$  such that  $h_v^{(k)} = \varphi(l_v^{(k)})$ . This apparently holds for  $k = 0$  because the initial node features are the same

for WL and GNN  $l_v^{(0)} = h_v^{(0)}$  for all  $v \in G_1, G_2$ . So  $\varphi$  could be the identity function for  $k = 0$ . Suppose this holds for iteration  $k - 1$ , we show that it also holds for  $k$ . Substituting  $h_v^{(k-1)}$  with  $\varphi(l_v^{(k-1)})$  gives us

$$h_v^{(k)} = \phi \left( \varphi \left( l_v^{(k-1)} \right), f \left( \left\{ \varphi \left( l_u^{(k-1)} \right) : u \in \mathcal{N}(v) \right\} \right) \right)$$

It follows from the fact composition of injective functions is injective that there exists some injective function  $\psi$  so that

$$h_v^{(k)} = \psi \left( l_v^{(k-1)}, \left\{ l_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right)$$

Then we have

$$h_v^{(k)} = \psi \circ g^{-1} g \left( l_v^{(k-1)}, \left\{ l_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right) = \psi \circ g^{-1} \left( l_v^{(k)} \right)$$

$\varphi = \psi \circ g^{-1}$  is injective because the composition of injective functions is injective. Hence for any iteration  $k$ , there always exists an injective function  $\varphi$  such that  $h_v^{(k)} = \varphi(l_v^{(k)})$ . At the  $K$ -th iteration, the WL test decides that  $G_1$  and  $G_2$  are non-isomorphic, that is the multisets  $\{l_v^{(K)}\}$  are different for  $G_1$  and  $G_2$ . The graph neural network  $\mathcal{A}$ 's node embeddings  $\{h_v^{(K)}\} = \{\varphi(l_v^{(K)})\}$  must also be different for  $G_1$  and  $G_2$  because of the injectivity of  $\varphi$ .

Now let us prove the theorem for the case where  $\mathcal{A}$  aggregates the central node along with its neighbors

$$h_v^{(k)} = f \left( \left\{ h_v^{(k-1)}, h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right)$$

The difficulty in proving this form of aggregation mainly lies in the fact that it does not immediately distinguish the root or central node from its neighbors. For example, let us consider the chain graphs  $a - b - b$  and  $b - a - b$ . If we aggregate once at the middle nodes of the chain graphs  $b$  and  $a$  respectively, we essentially get the same new representation  $(\{a, b, b\})$  for  $b$  and  $a$ , although the WL test would have represented these nodes as "rooted trees"  $(b, \{a, b\})$  and  $(a, \{b, b\})$  instead. The key insight into solving the problem described above is to notice that it is possible to distinguish the two structures (different roots, but the same neighborhood plus the central node) with two iterations of aggregation, unless the structures are symmetric, that is the two adjacent nodes (roots) in consideration are both adjacent to the same neighbors, in which case there is no need to distinguish because the multiset of node features will be the same. Before we formally prove the theorem, let us look at the chain graph example above again to get more intuition. This time let us apply the neighborhood aggregation again. The node features at the middle nodes  $b$  and  $a$  are essentially representative of  $\{\{a, b, b\}, \{a, b\}, \{b, b\}\}$  and  $\{\{a, b, b\}, \{a, b\}, \{a, b\}\}$ . This time we can successfully distinguish the two structures rooted at  $b$  and  $a$ . However, if we have complete graphs with nodes  $a, b, b$  again. Even if we apply the aggregation twice, the representations will stay the same for roots  $a$  and  $b$ , i.e.  $\{\{a, b, b\}, \{a, b, b\}, \{a, b, b\}\}$ . But that is fine because the two graphs are isomorphic and we will end up with the same collection of node features. In summary, unless two rooted subtrees of height 1 under consideration are not "symmetric", after two iterations, we can always recover the root and thus distinguish among different rooted trees. If they are symmetric, then we have the same multiset of node features and thus the correct representations. Next, we present a more formal proof.

Let  $v$  and  $u$  be adjacent nodes and we are interested in distinguishing the rooted tree structures  $v - \{\mathcal{N} \cup u\}$  and  $u - \{\mathcal{N} \cup v\}$  from graphs  $G_1$  and  $G_2$  respectively, where  $\mathcal{N}$  is a set of adjacent nodes. Both rooted trees have the same unrooted multiset representation  $\{u, v, \mathcal{N}\}$ . Here, we want our graph neural network  $\mathcal{A}$  to have a *unique* embedding for distinct structures and the goal is to recover the root with the unrooted multiset representations. Then we can apply the same argument as the "AGGREGATE and then COMBINE" aggregation process. Let us consider the node representations at  $u$  and  $v$  after two iterations of the aggregation below.

$$h_v^{(k)} = f \left( \left\{ h_v^{(k-1)}, h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right)$$

With some abuse of notation, let  $\mathcal{N}(\mathcal{N})$  denote the multiset of neighborhood multisets for each member from  $\mathcal{N}$ . Let  $\mathcal{N}_1(\cdot)$  denote the neighborhood of a node in  $G_1$ , i.e.  $v - \{\mathcal{N} \cup u\}$ . We can also define  $\mathcal{N}_2$  similarly.

Suppose it is the case that, for  $v - \{\mathcal{N} \cup u\}$  and  $u - \{\mathcal{N} \cup v\}$ ,  $\mathcal{N}_1(u) \cup u = \mathcal{N}_2(v) \cup v = \mathcal{N} \cup u \cup v$  and  $\mathcal{N}_1(\mathcal{N}) = \mathcal{N}_2(\mathcal{N})$ . That is, the two rooted trees from  $G_1$  and  $G_2$  are, in fact, symmetric in the sense that  $u, v, \mathcal{N}$  form the same triangle. Indeed, in this case, we cannot distinguish the subtrees at  $u$  and  $v$  after two iterations of aggregation. However, that is not a problem for distinguishing  $G_1$  and  $G_2$  because both  $G_1$  and  $G_2$  have the same structure of  $u, v, \mathcal{N}$ , as well as the same, thus correct, collection of node features  $\{h_u, h_v\}$ . Aggregating the central node along with neighbors would not reduce the discriminative power of the GNN in this case.

Now suppose we do not have the symmetric structure for  $v - \{\mathcal{N} \cup u\}$  and  $u - \{\mathcal{N} \cup v\}$ . Then either a) The neighborhood of  $u$  in  $v - \{\mathcal{N} \cup u\}$  and that of  $v$  in  $u - \{\mathcal{N} \cup v\}$  are not equivalent, i.e.  $\mathcal{N}_1(u) \cup u \neq \mathcal{N}_2(v) \cup v$ . b)  $\mathcal{N}_1(\mathcal{N}) \neq \mathcal{N}_2(\mathcal{N})$ . After two iterations of aggregation, the node representations for  $v$  in  $v - \{\mathcal{N} \cup u\}$  and  $u$  in  $u - \{\mathcal{N} \cup v\}$  will be  $\{\{u, v, \mathcal{N}\}, \{u, \mathcal{N}_1(u)\}, \{\mathcal{N}, \mathcal{N}_1(\mathcal{N})\}\}$  and  $\{\{v, u, \mathcal{N}\}, \{v, \mathcal{N}_2(v)\}, \{\mathcal{N}, \mathcal{N}_2(\mathcal{N})\}\}$  respectively. If any of a) or b) is true, i.e. non-symmetric case, then the node representations for  $u$  and  $v$  are different. Our graph neural network  $\mathcal{A}$  successfully distinguishes the embeddings of the rooted subtrees, despite only using unrooted representations.

In conclusion, the aggregation process in the following form

$$h_v^{(k)} = f \left( \left\{ h_v^{(k-1)}, h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right)$$

has the same discriminative power in embedding different graphs as the following aggregation form

$$h_v^{(k)} = \phi \left( h_v^{(k-1)}, f \left( \left\{ h_u^{(k-1)} : u \in \mathcal{N}(v) \right\} \right) \right)$$

if  $f, \phi$  are valid injective functions. Moreover, GNNs in both forms have the same discriminative power as the WL test. □

## C PROOF FOR LEMMA 4

*Proof.* We first prove that there exists a mapping  $f$  so that  $\sum_{x \in X} f(x)$  is unique for each finite multiset  $X$ . Because  $\mathcal{X}$  is countable, there exists a mapping  $Z : \mathcal{X} \rightarrow \mathbb{N}$  from  $x \in \mathcal{X}$  to natural numbers. Because the multisets  $X$  are finite, there exists a number  $N \in \mathbb{N}$  so that  $|X| < N$  for all  $X$ . Then an example of such  $f$  is  $f(x) = N^{-Z(x)}$ . This  $f$  can be viewed as a more compressed form of an one-hot vector or  $N$ -digit presentation. Thus,  $\tilde{f}(X) = \sum_{x \in X} f(x)$  is an injective function of multisets.

$\phi \left( \sum_{x \in X} f(x) \right)$  is permutation invariant so it is a well-defined multiset function. For any multiset function  $g$ , we can construct such  $\phi$  by letting  $\phi \left( \sum_{x \in X} f(x) \right) = g(X)$ . Note that such  $\phi$  is well-defined because  $\tilde{f}(X) = \sum_{x \in X} f(x)$  is injective. □

## D PROOF FOR LEMMA 5

*Proof.* Let us consider the example  $X_1 = \{1, 1, 1, 1, 1\}$  and  $X_2 = \{2, 3\}$ , i.e. two different multisets of positive numbers that sum up to the same value. We will be using the homogeneity of ReLU.

Let  $W$  be an arbitrary linear transform that maps  $x \in X_1, X_2$  into  $\mathbb{R}^n$ . It is clear that, at the same coordinates,  $Wx$  are either positive or negative for all  $x$  because all  $x$  in  $X_1$  and  $X_2$  are positive. It follows that  $\text{ReLU}(Wx)$  are either positive or 0 at the same coordinate for all  $x$  in  $X_1, X_2$ . For the coordinates where  $\text{ReLU}(Wx)$  are 0, we have  $\sum_{x \in X_1} \text{ReLU}(Wx) = \sum_{x \in X_2} \text{ReLU}(Wx)$ . For the coordinates where  $Wx$  are positive, linearity still holds. It follows from linearity that

$$\sum_{x \in X} \text{ReLU}(Wx) = \text{ReLU} \left( W \sum_{x \in X} x \right)$$

where  $X$  could be  $X_1$  or  $X_2$ . Because  $\sum_{x \in X_1} x = \sum_{x \in X_2} x$ , we have the following as desired.

$$\sum_{x \in X_1} \text{ReLU}(Wx) = \sum_{x \in X_2} \text{ReLU}(Wx)$$

□

## E PROOF FOR COROLLARY 6

*Proof.* Suppose multisets  $X_1$  and  $X_2$  have the same distribution, without loss of generality, let us assume  $X_1 = (S, m)$  and  $X_2 = (S, k \cdot m)$  for some  $k \in \mathbb{N}_{\geq 1}$ , i.e.  $X_1$  and  $X_2$  have the same underlying set and the multiplicity of each element in  $X_2$  is  $k$  times of that in  $X_1$ . Then we have  $|X_2| = k|X_1|$  and  $\sum_{x \in X_2} f(x) = k \cdot \sum_{x \in X_1} f(x)$ . Thus,

$$\frac{1}{|X_2|} \sum_{x \in X_2} f(x) = \frac{1}{k \cdot |X_1|} \cdot k \cdot \sum_{x \in X_1} f(x) = \frac{1}{|X_1|} \sum_{x \in X_1} f(x)$$

Now we show that there exists a function  $f$  so that  $\frac{1}{|X|} \sum_{x \in X} f(x)$  is unique for distributionally equivalent  $X$ . Because  $\mathcal{X}$  is countable, there exists a mapping  $Z : \mathcal{X} \rightarrow \mathbb{N}$  from  $x \in \mathcal{X}$  to natural numbers. Because the multisets  $X$  are finite, there exists a number  $N \in \mathbb{N}$  so that  $|X| < N$  for all  $X$ . Then an example of such  $f$  is  $f(x) = N^{-2Z(x)}$ . □

## F PROOF FOR COROLLARY 7

*Proof.* Suppose multisets  $X_1$  and  $X_2$  have the same underlying set  $S$ , then we have

$$\max_{x \in X_1} f(x) = \max_{x \in S} f(x) = \max_{x \in X_2} f(x)$$

Now we show that there exists a mapping  $f$  so that  $\max_{x \in X} f(x)$  is unique for  $X$ s with the same underlying set. Because  $\mathcal{X}$  is countable, there exists a mapping  $Z : \mathcal{X} \rightarrow \mathbb{N}$  from  $x \in \mathcal{X}$  to natural numbers. Then an example of such  $f : \mathcal{X} \rightarrow \mathbb{R}^\infty$  is defined as  $f_i(x) = 1$  for  $i = Z(x)$  and  $f_i(x) = 0$  otherwise, where  $f_i(x)$  is the  $i$ -th coordinate of  $f(x)$ . Such an  $f$  essentially maps a multiset to its one-hot embedding. □

## G DETAILS OF DATASETS

We give detailed descriptions of datasets used in our experiments. Further details can be found in (Yanardag & Vishwanathan, 2015).

**Social networks datasets.** IMDB-BINARY and IMDB-MULTI are movie collaboration datasets. Each graph corresponds to an ego-network for each actor/actress, where nodes correspond to actors/actresses and an edge is drawn between two actors/actresses if they appear in the same movie. Each graph is derived from a pre-specified genre of movies, and the task is to classify the genre graph it is derived from. REDDIT-BINARY and REDDIT-MULTI5K are balanced datasets where each graph corresponds to an online discussion thread and nodes correspond to users. An edge was drawn between two nodes if at least one of them responded to another’s comment. The task is to classify each graph to a community or a subreddit it belongs to. COLLAB is a scientific collaboration dataset, derived from 3 public collaboration datasets, namely, High Energy Physics, Condensed Matter Physics and Astro Physics. Each graph corresponds to an ego-network of different researchers from each field. The task is to classify each graph to a field the corresponding researcher belongs to.

**Bioinformatics datasets.** MUTAG is a dataset of 188 mutagenic aromatic and heteroaromatic nitro compounds with 7 discrete labels. PROTEINS is a dataset where nodes are secondary structure elements (SSEs) and there is an edge between two nodes if they are neighbors in the amino-acid sequence or in 3D space. It has 3 discrete labels, representing helix, sheet or turn. PTC is a dataset of 344 chemical compounds that reports the carcinogenicity for male and female rats and it has 19 discrete labels. NCI1 is a dataset made publicly available by the National Cancer Institute (NCI) and is a subset of balanced datasets of chemical compounds screened for ability to suppress or inhibit the growth of a panel of human tumor cell lines, having 37 discrete labels.