

Graph Classification with GNNs : Optimisation, Representation & Inductive Bias

P. Krishna Kumar^a and Harish G. Ramaswamy^{a,b}

^aDepartment of Computer Science & Engineering, IIT Madras

^bDepartment of Data Science & Artificial Intelligence, IIT Madras
{pkrishna, hariguru}@cse.iitm.ac.in

Abstract.

Theoretical studies on the representation power of GNNs have been centered around understanding the equivalence of GNNs, using WL-Tests for detecting graph isomorphism. In this paper, we argue that such equivalence ignores the accompanying optimization issues and does not provide a holistic view of the GNN learning process. We illustrate these gaps between representation and optimization with examples and experiments. We also explore the existence of an implicit inductive bias (e.g. fully connected networks prefer to learn low frequency functions in their input space) in GNNs, in the context of graph classification tasks. We further prove theoretically that the message-passing layers in the graph, have a tendency to search for either discriminative subgraphs, or a collection of discriminative nodes dispersed across the graph, depending on the different global pooling layers used. We empirically verify this bias through experiments over real-world and synthetic datasets. Finally, we show how our work can help in incorporating domain knowledge via attention based architectures, and can evince their capability to discriminate coherent subgraphs.

1 Introduction

Graph Neural Network's (GNNs) have enabled end-to-end learning over *relational* data due to differentiable loss functions, that can be trained with non-linear components like multi-layer perceptrons. Several real world applications such as fake news detection [22], physical simulation [2], traffic delay estimation [9], and fraudulent transactions prediction[7] have GNNs as a crucial component.

Graph classification is one of the most common downstream graph neural processing applications[11]. While different GNN operators update node level features via message-passing, the graph level predictions are done by pooling the member nodes into a single unified representation. This pooling is either done by coarsening functions that gradually reduce the size of the graph [4], or with the help of global pooling methods like average, max, sum [10].

The success of GNNs has also led to several attempts toward defining theoretical boundaries of what GNNs can and can not do. The strengths and weaknesses of graph neural networks have been extensively evaluated in terms of their *representation* capabilities. Most studies have focused on the capability of message-passing networks using Weisfeiler-Lehman test, which is constrained by its limitations in distinguishing isomorphic graphs [14]. The proposed architecture of Graph Isomorphism Network [32] implements an MLP to

model injective function, and be as powerful as 1-WL test. Higher order generalization of WL tests, using a combination of equivariant and invariant functions, is shown to surpass limitations of simple GNNs that use message-passing [24]. A disparity exists between the theoretical understandings and optimization practices within GNNs. For instance, previous works indicate that subgraph-based counting/classification is inherently unattainable for subgraphs exceeding a size of three [6]. However, there are empirical findings to challenge these constraints, primarily due to the relaxation of restrictive assumptions, such as fixed-node features derived from a countable set.

In this work, we introduce an implicit bias inherent in graph convolutional networks when integrated with self-attention-based global pooling functions. To our knowledge, this marks the initial investigation characterizing the graph-classification task with a focus on attention pooling layers. We empirically explore the full spectrum of attention, ranging from average pooling (equivalent to no attention) to max pooling (singular focus). Additionally, we substantiate our empirical findings with theoretical insights and conduct an in-depth analysis of attention-based global pooling. We summarize our main contributions as follows:

- **Develop an experimental setup to elicit GNN's implicit bias:** We introduce a dataset with unique characteristics of *occurrence* and *connectedness*, that underscores the ambiguity inherent in GNN models for graph classification tasks. Our approach provides empirical insights into deducing the implicit bias of GNNs by closely examining their nuanced behavior on synthetic grid-graph world.
- **Theoretically analyze of attention based global pooling:** We employ gradient-flow [8] to prove the existence and uniqueness of final vector, towards which our learned model parameters align. Our main result formalizes the biased nature of attention pooling, and shows it is preferring to use *closely-connected-substructures* as discriminative features, rather than collection of nodes that may be dispersed across the graph.
- **Empirical validation:** We study realworld datasets with planted ground-truth, to show the ambiguity caused by the inherent bias. We isolate the effect of the linear classifier from the message-passing layers. Lastly, our study underscores practical implications in tasks where domain knowledge informs graph classification. It aids in discerning whether classification hinges on the presence of coherent subgraphs or on merely the presence of nodes, irrespective of their neighbors.

2 Subgraph Based Graph Classification Task

Graph Neural Networks learn to infer graph labeling functions by leveraging latent representation of nodes, that extend to the entire graph structure. A trained classifier $h(X, A)$ operates on both node features and graph topology, producing a class label as output. However, the mechanisms driving classification are not fully understood. It remains unclear whether classifications arise from the presence or absence of specific structural components within the graph, such as nodes belonging to distinct classes forming unique subgraphs, or if predictions stem from numerical decision boundaries imposed over aggregated graph representations. Recent efforts in *explaining* GNN predictions aim to identify subgraphs or disjointed nodes that cause maximal change to the specific predictions [28, 33, 1]. Despite these endeavors, whether GNN architectures inherently learn to discern specific graph traits for prediction purposes, remains an open question. We aim to address such knowledge gaps through the subsequent sections of this paper.

We establish a common notation and terminology used throughout this paper. \mathcal{D} denotes a dataset of attributed graphs with N nodes in each graph. Each graph $G \in \mathcal{D}$ is represented by a tuple of its node-feature matrix and edge adjacency matrix $(X, A) \in (\mathbb{R}^{N \times d}, \{0, 1\}^{N \times N})$. For natural number $K \in \mathbb{N}$, $[K]$ denotes sequence $\{1, 2, 3 \dots K\}$. \mathcal{S} denote an ordered tuple $= (s_1, s_2, \dots s_M) \in [K]^M$. Let $G^* = (X^*, A^*) \in (\mathbb{R}^{M \times d}, \{0, 1\}^{M \times M})$ be a special graph with M nodes, where $M < N$. In our graph classification problem, the set of graphs in \mathcal{D} are labelled based on their relation to G^* , and our goal is to learn a model that can learn to map a graph to its correct label. Note that G^* is part of the data generation process, is not known to the learning algorithm, and is not necessarily required to be recovered by the learned graph classification model. We denote the node features and corresponding adjacency matrix of subgraph G^* by $X_{[S, :]} \in \mathbb{R}^{M \times d}$ and $A_{[S, S]} \in \{0, 1\}^{M \times M}$. The matrix subscripts corresponds to rows and columns of indices in the ordered tuple \mathcal{S} .

To illustrate this concept via a simple example, let's examine a dataset \mathcal{D} , that comprises of numerous graphs $G \in \mathcal{D}$. As per the law of dichotomy, each graph can either contain or lack a subgraph representing an orange-blue-green (O-B-G) chain (●●●), where colors denote the respective node features. Let us define a labeling function $y(G)$, such that $y(G) = +1$ if the O-B-G subgraph is present in G , and $y(G) = -1$ otherwise. Consequently, the dataset \mathcal{D} can be partitioned into two disjoint subsets, \mathcal{D}_1 and \mathcal{D}_{-1} , based on their respective class labels.

Can GNNs detect presence and absence of subgraphs?

To ascertain the capability of GNNs in discerning graphs based upon the existence or non-existence of a specified subgraph G^* , we have devised a simple yet informative experiment. We generated grid-graphs measuring 12×12 and randomly assigned one-hot embedding to its node features, corresponding to each node taking one of 4 colors. Subsequently, we delineated two distinct graph sets, denoted as \mathcal{D}_1 and \mathcal{D}_{-1} , contingent upon the presence of a subgraph comprising a special chain of length equal to three. The precise method used to generate \mathcal{D}_1 and \mathcal{D}_{-1} , and to ensure the presence or absence of this subgraph is detailed in Section 5.1. All graphs within \mathcal{D}_1 were assigned a ground-truth label of $y = 1$, while those in \mathcal{D}_{-1} were labeled with $y = 0$.

After training GNN models on the datasets \mathcal{D}_1 and \mathcal{D}_{-1} , we tested an independent set of grid-graphs, generated with dimensions 13×13

Table 1: Performance of GNN models that are trained on 12×12 with $\mathcal{D}_1, \mathcal{D}_{-1}$, and tested over unseen grid-graphs of varying sizes. Results show that GNNs are capable of learning to distinguish graphs based on presence/absence. (Sec.2)

Model	Pooling func	Training Loss	Train Accuracy	Test Acc 12 x 12	Test Acc 13 x 13
GCN	MAX	0.0766	0.987	0.979	0.991
	AVG	0.0451	1.000	0.998	0.517
	ATTN	0.0948	0.978	0.923	0.986
GAT	MAX	0.0715	0.991	0.968	0.994
	AVG	0.1213	1.000	0.998	0.675
	ATTN	0.0293	0.993	0.973	0.995

and 12×12 . Table 1 presents the empirical results, demonstrating the efficacy of popular GNN models, such as GCN and GAT, when coupled with various global pooling mechanisms. Notably, these models showcase remarkable generalization capabilities over previously unseen test data. Across all model-pooling combinations, the test accuracy consistently surpass the 90% threshold for both grid-graph sizes, except for the average pooling mechanism on the 13×13 grid. Given that all models achieve high accuracy and are optimized until attaining near-zero training loss (binary cross-entropy), we can say that GNNs possess the empirical capacity to differentiate between the presence and absence of the designated subgraph.

In the following sections, we will place this empirical observation in the context of theoretical results and attempt to reconcile multiple views.

3 Limits of Representation Power Arguments

There is an apparent difference in the theoretical representation pertaining to power of neural networks and their practical implication in real settings. For example, deeper neural networks have a definite advantage over shallow networks. It is known that networks with polynomial increase in depth can approximate functions that exponentially growing width networks cannot [29]. Simple architectures like deep feed-forward networks with piece-wise linear activation functions (like ReLU) can distinguish exponentially greater input regions due to the increased depth [23]. However, the mere ability of deeper networks to represent complex functions does not guarantee that they will be learned by gradient descent. For example, it has been observed that the number of distinct linear regions in ReLU based networks grows linearly along any single dimension. It highlights the gap between theoretical possibilities and empirical observations while training neural networks.

Similar to feed-forward networks, graph neural networks are shown to have restricted representation power, and their theoretical limitations have been an area of active research [15]. Message passing neural networks (MPNNs), encompassing different convolution operators for GNNs, are theoretically as powerful as 1-Weisfeiler Lehmann test for distinguishing non-isomorphic graphs [32]. In similar vein, higher order generalizations of k -WL test (for $k > 2$) are known to be strictly more powerful than 1-WL and capable of learning sub-graph patterns of maximally k -nodes [24]. In the context of subgraph-counting, MPNNs can only count a given subgraph (or induced-subgraph) if there are 3 or fewer nodes in the subgraph (barring exceptions of star-shaped subgraphs) [6]. As earlier, the higher order GNNs based on k -WL are shown to be powerful enough to count subgraphs of utmost k nodes.

The negative results from GNN's representation-theory are not as

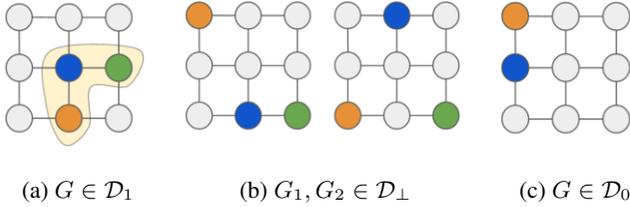


Figure 1: Sample graphs from each partition of data \mathcal{D} . (a) Subgraph G^* (orange-blue-green) is present in all $G \in \mathcal{D}_1$. (b) shows the presence of all nodes of G^* in $G \in \mathcal{D}_\perp$ but none of the nodes are mutually adjacent to each other. (c) among graphs $G \in \mathcal{D}_0$, subgraph G^* is only partially present (of size 1 or 2).

discouraging as they seem on the surface. For example, the negative result claiming that GNNs *cannot* detect presence/absence of a subgraph larger than size 4 [6], really corresponds to saying that there *exists* a distribution over graphs and labels for which the best GNN based classifier has low accuracy. We argue that such adversarial distributions are not indicative of real data. For example, as shown in Table 1, GNNs are able to learn to classify graphs based on presence/absence of subgraphs.

Despite these theoretic limitations, GNNs have proven to be empirically very successful in most graph based tasks, and achieved state-of-the-art performance in many real-world applications [22, 2, 9, 7]. In this paper we argue that studying GNN based graph classification algorithms in the context of optimisation and generalisation is still an open and interesting objective.

4 Inductive Bias of GNNs Under Gradient Descent

Table 2: Characteristics of the partitions of dataset \mathcal{D} with respect to nodes/edges of a subgraph G^*

Partition	Occurrence	Connected-ness
\mathcal{D}_1	✓	✓
\mathcal{D}_0	×	×
\mathcal{D}_\perp	✓	×

In this section, we explore if there exists a proclivity for the GNN to learn a particular category of functions. We study one such characteristic tendency of GNNs, where it exhibits distinct inclinations towards learning specific function types, when convolution operator is coupled with varied pooling mechanisms. Consider the graph classification task on dataset created around subgraph G^* as seen in the previous sections. For the purpose of illustration, let G^* be the O-B-G orange-blue-green chain and embedded in graphs of a fixed-size grid-graph dataset \mathcal{D} .¹ Let \mathcal{D}_1 be the partition of \mathcal{D} where subgraph G^* is *present*, for example Fig.1a contains O-B-G. Similarly, let \mathcal{D}_0 be the set of all graphs within \mathcal{D} where G^* is only *partially* present, as in at least one node of G^* is missing. For instance, Fig.1c does not contain the green node. All the remaining graphs together form \mathcal{D}_\perp , which implies the graphs in \mathcal{D}_\perp consist of all the nodes of G^* , (O,B,G), but do not contain all edge of G^* , for example Fig.1b shows that no two colored nodes are adjacent.

In the context of this example, let *occurrence* denote the presence of ‘all’ attributed-nodes of G^* and *connected-ness* denote the presence of ‘all’ edges of G^* . Table 2 characterizes the mutually exclusive partitions of $\mathcal{D}_1, \mathcal{D}_0$ and \mathcal{D}_\perp along these two properties and

¹ Arguments can trivially extend to graphs of varying sizes and multiple classes classification. There is nothing sacrosanct about the given setting.

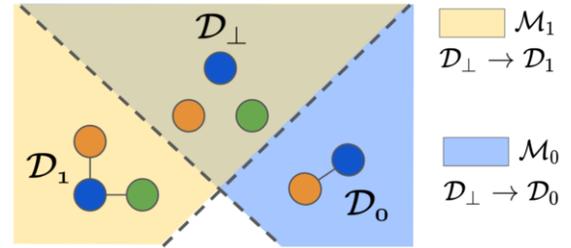


Figure 2: Multiple legal decision-boundaries (dotted lines) that can be learned while training GNNs on $\{\mathcal{D}_1 \cup \mathcal{D}_0\}$. \mathcal{D}_\perp can be classified along with either \mathcal{D}_1 or \mathcal{D}_0 . We show that (i) the graph convolution operator with global average pooling classifies \mathcal{D}_\perp as \mathcal{D}_1 and behaves as \mathcal{M}_1 family of functions denoted by the orange decision regions, while (ii) GCN with attention based pooling classifies \mathcal{D}_\perp as \mathcal{D}_0 and learns \mathcal{M}_0 denoted by blue decision region.

highlights their commonalities and differences. While \mathcal{D}_\perp shares its *occurrence and connected-ness* property with both the other partitions, it is not obvious whether it can be associated as more close to either \mathcal{D}_1 or \mathcal{D}_0 . To verify this empirically, we experiment with a GNN trained over $\{\mathcal{D}_1 \cup \mathcal{D}_0\}$ and further analyze its behavior upon \mathcal{D}_\perp . Say we assign label $y = 1$ to \mathcal{D}_1 and $y = 0$ to \mathcal{D}_0 , and train a GNN over set $\{\mathcal{D}_1 \cup \mathcal{D}_0\}$ to achieve perfect accuracy. It remains non-obvious to predict whether graphs of \mathcal{D}_\perp will be classified as label $y = 1$ or $y = 0$.

4.1 Ambiguity in Graph Classification: Figure 2

A GNN that has not seen graphs from the set \mathcal{D}_\perp while training on $\{\mathcal{D}_1 \cup \mathcal{D}_0\}$ could very well assign it either label 1 or 0. However, both these assignments indicate different inductive biases.

Case 1. ($\mathcal{D}_\perp \rightarrow \mathcal{D}_1$) Suppose, a trained GNN consistently classifies \mathcal{D}_\perp to be same as \mathcal{D}_1 , instead of \mathcal{D}_0 , it shows that the model looks for the *occurrence* of all nodes in G^* .

Case 2. ($\mathcal{D}_\perp \rightarrow \mathcal{D}_0$) Conversely, if a trained GNN aligns \mathcal{D}_\perp with \mathcal{D}_0 , it indicates that the model emphasizes on *connected-ness* of all nodes in G^* .

How to Deduce Implicit Bias? Let’s understand this bias using the example depicted in Figure 1, let us assume \mathcal{D} to be the universe of all fixed-size grid-graphs and $\mathcal{D}_1, \mathcal{D}_0$ to exhaustively contain all graphs satisfying the above discussed properties as in Table 2. Now, if a trained GNN classifies graph in *dataP* as \mathcal{D}_1 rather than as \mathcal{D}_0 , it implies that the GNN is looking for presence of all the three (orange,blue,green) nodes in each graph. Otherwise, there will exist $G_1, G_2 \in \mathcal{D}_0$ such that their nodes together will contain all $\{O, B, G\}$ but separated among the two graphs G_1 and G_2 . On the other hand, if a GNN classifies \mathcal{D}_\perp as \mathcal{D}_0 instead of \mathcal{D}_1 , it implies the GNN is looking for connection between the three (O-B-G) nodes. In other words, since \mathcal{D}_0 consists of all possible disintegrated subparts of G^* , the only reason \mathcal{D}_\perp can be classified as \mathcal{D}_0 is due to the lack of G^* being present as such.

4.2 Grid-Graphs Dataset

To empirically validate the capability of GNNs to distinguish *occurrence* versus *connectedness*, we start with a simple and controllable setting of grid-graphs with node-features, as discussed above. Each node in the graph is assigned a random one-hot feature vector, referred to as ‘node-color’, from a finite set of vectors. We further derive two partition sets, \mathcal{D}_1 and \mathcal{D}_0 , by selectively planting certain

subgraph G^* within the grid-graph. We assume our universe of data \mathcal{D} restricted to the grid-graphs \mathcal{G} of fixed size, and assign each node a random feature from the standard basis-vector set $\mathcal{C} = \{e_i\}$ for $i \in [K]$ classes. We further create a three-node graph \mathcal{G}^* , and assign each of its nodes distinct features, denoted by $\mathcal{C}' = \{e_j\}$. The partition \mathcal{D}_1 is generated by planting the subgraph \mathcal{G}^* in grid-graphs $G \in \mathcal{D}$. We randomly choose an anchor node $v_1 \in \mathcal{G}$ along with its two other neighbors $v_2, v_3 \in \mathcal{N}(v_1)$, and then alter their features such that \mathcal{G}^* is contained in \mathcal{G} .

Similarly, graphs in partition \mathcal{D}_0 are generated by *partially* planting \mathcal{G}^* , that is by taking a strict subset $\tilde{\mathcal{C}} \subset \mathcal{C}'$ of size 1 or 2 and replacing node features of random nodes in \mathcal{G} with $\tilde{\mathcal{C}}$. The graphs in \mathcal{D}_\perp are generated by randomly choosing $|\tilde{\mathcal{C}}|$ number of mutually non-adjacent nodes, and sequentially assigning them node features from $\tilde{\mathcal{C}}$. Note that the chosen universe of \mathcal{D} is small enough for complete enumeration and for exhaustive generation of \mathcal{D}_1 , i.e., any graph that satisfies the said characteristic properties will be present in \mathcal{D}_1 . The datasets size is restricted by number of nodes in the grid-graphs that we sequentially choose as anchor nodes to plant \mathcal{G}^* . To remove any class label-imbalance, we maintain parity between the sizes of \mathcal{D}_1 and \mathcal{D}_0 .

Labeling function: The graphs in the two partitions are assigned binary labels for the supervised training of the GNNs. Without loss of generality, \mathcal{D}_1 are assigned label $y = 1$, and \mathcal{D}_0 graphs with partially planted $\tilde{\mathcal{C}}$ are assigned label $y = 0$. Since \mathcal{D}_\perp is never used for training GNN models, it is not assigned any label. In case of 3-node \mathcal{G}^* , after 1-step of propagation, the central node’s representation will have a positive component along v^* ; the proof of its existence and uniqueness follows.

4.3 Model Baselines

We study the behaviors of two prominent models, the graph convolution network (GCN) [18] and the graph attention network (GAT) [30, 5].

Along with these message-passing layers, we use common global pooling functions, encompassing summation, averaging, and attention-based aggregation. For any graph $G = (X, A)$, GCN updates are given by $f_{\text{GCN}}(X, A; W) = \hat{A}XW$, where \hat{A} is the degree normalized adjacency matrix. After one step of propagation, the updated representation of node i can be given by $v_i = W \sum_{j \in \mathcal{N}(i)} x_j$. Further, all the nodes of the graph $\{v_1, v_2 \dots v_N\}$ are pooled using average, maximum or attention mechanism. For $(A, X) \in \mathbb{R}^{N \times N} \times \mathbb{R}^{N \times k}$, the end-to-end equation for graph classification using GCN and average global-pooling can be written as

$$f_{\text{GCN+AVG}}(X, A; \Theta, \mathbf{w}) = \mathbf{1}^\top \cdot \sigma(\hat{A}X\Theta) \cdot \mathbf{w} \quad (1)$$

where $\Theta \in \mathbb{R}^{k \times d}$ is the linear-transform-parameter of GCN, $\mathbf{1}$ is column of all 1’s, and $\mathbf{w} \in \mathbb{R}^d$ is the linear classifier. Similarly, the end-to-end model equation for GCN with global attention pooling can be written as

$$f_A(X, A; \Theta, \mathbf{a}, \mathbf{w}) = \text{sftmx} \left[\langle \sigma(\hat{A}X\Theta), \mathbf{a} \rangle \right] \sigma(\hat{A}X\Theta) \cdot \mathbf{w} \quad (2)$$

where $\mathbf{a} \in \mathbb{R}^d$ parameter is used to calculate the attention scores by taking *softmax* over the projection of nodes. Attention based global pooling is a special case of the self-attention pooling mechanism [21], where the representation of the entire graph is collapsed into one single vector.

4.4 Empirical Evidence for Inductive Bias

We empirically show that there exists a clear bias in GNNs, while classifying graphs from the above data partitions. We start with the most popular GNN architectures, graph convolution networks (GCN) and graph attention networks (GAT), along with most common global pooling methods like average, maximum and attention based readout function. We use the dataset described in previous Section 4.2, 12×12 grid graphs, partitioned into $\mathcal{D}_1, \mathcal{D}_0, \mathcal{D}_\perp$. All models, based on GCN/GAT layers coupled with MAX/AVG/ATTN global pooling layers, are trained to near perfect training accuracy (greater than 99%). While training, only \mathcal{D}_1 and \mathcal{D}_0 are exposed to the models and \mathcal{D}_\perp remains inaccessible. We then test each of these *trained* models on test data \mathcal{D}_\perp to analyze what labels $\hat{y} \in \{0, 1\}$ are assigned to graphs $G \in \mathcal{D}_\perp$.

Inductive bias from observations in Table 1: We observe that both GCN and GAT, with average pooling classify almost all unseen graphs as label $\hat{y} = 1$, while attention based pooling classify almost all unseen graphs as $\hat{y} = 0$. Assigning such skewed proportions of nodes with selective label $\hat{y} = 1$ or $\hat{y} = 0$, is a clear indicator that both models have preference to learn different functions. While this distinctive behavior shows an interesting biased behavior, we will see in later Section 5.1 that this behavior translates to model’s preference of subgraph selection.

5 Analyzing Attention Based Global Pooling

We analyze the behavior of GCN based models with attention pooling, and characterize it for grid-graph dataset. We start with formalizing notations for the three data partitions $\mathcal{D}_1, \mathcal{D}_0$ and \mathcal{D}_\perp below, along with precise definitions for *occurrence* and *connected-ness*.

Definition 1 (Data Partitions). *A graph dataset \mathcal{D} can be partitioned into three mutually exclusive sets on the basis of subgraph $G^* = (X^*, A^*)$ as follows:*

- \mathcal{D}_1 is the set of graphs $\{G \in \mathcal{D} \mid \exists \mathcal{S} = (s_1, s_2 \dots s_M) \in [N]^M\}$, such that $X_{[\mathcal{S},:]} = X^*$ and $A_{[\mathcal{S},\mathcal{S}]} \geq A^*$ for some ordered tuple \mathcal{S} (here matrix comparison is element-wise)
- \mathcal{D}_0 is the set of graphs $\{G \in \mathcal{D} \mid X_{\mathcal{S},:} \neq X^*, \forall \mathcal{S} \in [N]^M\}$.
- \mathcal{D}_\perp is the set of graphs $\{G \in \mathcal{D} \mid \exists \mathcal{S} \in [N]^M \text{ s.t. } X_{\mathcal{S},:} = X^* \text{ and } G \notin \mathcal{D}_1\}$

In other words, \mathcal{D}_1 contains G^* as its subgraph, \mathcal{D}_0 subgraphs do not contain at least one node from G^* , and lastly \mathcal{D}_\perp subgraphs have all the nodes of G^* , but does not contain G^* subgraph.

Definition 2 (Occurrence and Connection). *A subgraph $G^* = (X^*, A^*)$ is said to occur in $G = (X, A)$ iff there exists an ordered tuple $\mathcal{S} = (s_1, s_2 \dots s_M) \in [N]^M$, such that $X_{[\mathcal{S},:]} = X^*$. Also, G with G^* is said to be connected iff $A_{[\mathcal{S},\mathcal{S}]} \geq A^*$ for some tuple \mathcal{S} .*

Graph Convolutional Networks (GCN): Our theoretical analysis is restricted to a simpler setting of GCN after one step of message-propagation. GCN model updates the node-representations as $x_i = \frac{1}{\sqrt{d_i d_j}} W \sum_{j \in \mathcal{N}(i) \cup i} x_j$, where W is its learnable parameter and $d_i d_j$ is a degree based normalizing coefficient. The summation of neighboring nodes after one-hop message-passing can be denoted by matrix product $V = AX$. Ignoring the degree normalization, here each row $v_i = V_{i,:}$ is the updated feature vector of node i .

GCN with Global Attention Pooling (ATTN): While the graph convolution operator does the job of aggregating immediate neighbors, the global pooling functions, also known as READOUT functions, aggregate all node’s representation to get a single vector for

the entire graph. We analyze the behavior of attention based global pooling, which is calculated by the weighted summation over nodes as $h_G = \sum_{i=1}^N \alpha_i w (\sum_{j \in \mathcal{N}(i)} x_j) = \sum_{i=1}^N \alpha_i v_i = \hat{v}$, with two learnable parameters: $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{w} \in \mathbb{R}^d$.

The following lemma shows that for each graph G , there exists a node s (with its node representation v^*) that contains aggregated information from all the nodes of subgraph G^* after one-step of GCN aggregation.

Lemma 1 (Existence of v^*). *There exists $v^* \in \mathbb{R}^d$, such that for any graph $G = (X, A) \in \mathcal{D}_1$, there is an ordered tuple $\mathcal{S} \in [N]^M$ satisfying $\langle (A \cdot X)_{[\mathcal{S},:], v^*} \rangle > 0$. Here, the matrix-subscript $(A \cdot X)_{[\mathcal{S},:]}$ denotes selecting the rows indexed by elements of the set \mathcal{S} .*

Since G^* occurs in each graph of \mathcal{D}_1 , we can set $v^* = \sum_{j \in G^*} v_j$, and the lemma holds naturally by construction. We further note that there *cannot* exist any other node, apart from s , that contains aggregated information from all the nodes of subgraph G^* , after one-step of GCN aggregation.

Lemma 2 (Uniqueness of v^*). *For all $w \in \mathbb{R}^d$ perpendicular to v^* , $w \cdot v^* = 0$, there exists a graph $G = (X, A) \in \mathcal{D}_1$ such that $\forall \mathcal{S} \in [N]^M$, we have $\langle (A \cdot X)_{[\mathcal{S},:], w} \rangle \leq 0$.*

The above lemma highlights that G^* is the largest common-subgraph in \mathcal{D}_1 , because all the smaller subgraphs of G^* are also present in $G \in \mathcal{D}_1$ by the virtue of presence of G^* . Any other subgraph, apart from G^* , will find its aggregated representation w also in some $G \in \mathcal{D}_0$ graph. G^* being the largest common subgraph in \mathcal{D}_1 , v^* will not have any component along w .

For $G \in \mathcal{D}_1$, let the neighborhood of node s comprise of nodes of G^* in the set \mathcal{S} and other additional nodes $\{\mathcal{N}(s)/\mathcal{S}\}$. The aggregated representation of s can be split into those coming from neighbors in \mathcal{S} and its complement set $\{\mathcal{N}(s)/\mathcal{S}\}$. After attention based global pooling (readout), a graph $G \in \mathcal{D}_1$ can be represented as

$$h_G = \sum_{i \neq s, i=1}^N \alpha_i v_i + \alpha_s v_s = \sum_{i \neq s, i=1}^N \alpha_i v_i + \alpha_s w_s + \alpha_s v^* \quad (3)$$

Only the $\alpha_s v^*$ term is independent of the graph, as it is present in all graphs of \mathcal{D}_1 , hence it is the only vector to have largest dot-product component along v^* .

5.1 Loss Function and Gradient Flow of Parameters

The final GCN model's prediction $p \in [0, 1]$ for the graph label is calculated as $p = \sigma(\mathbf{w}^\top \hat{v})$, where $\sigma(x)$ is the sigmoid activation function. We characterize the evolution of the model parameters and their behavior by calculating their gradient flow [8]. For two class graph-classification, the model employs the binary cross entropy loss function $\mathbf{L}: -[y \log(p) + (1 - y) \log(1 - p)]$, where $y \in \{0, 1\}$ is the true label and $p \in [0, 1]$ is the model's predicted score. The binary cross entropy loss function is given by

$$\begin{aligned} \mathbf{L} &= -[y \log(p) + (1 - y) \log(1 - p)] \\ &= \log(1 + \exp(\mathbf{w}^\top \hat{v})) - y \mathbf{w}^\top \hat{v} \end{aligned}$$

Assumptions: In our theoretical analysis, we consider G^* as a sub-graph chain of radius one and derive results for a single-layer propagation of GCN, with attention based global pooling. During the course of optimization, we assume that the two partitions, \mathcal{D}_1 and \mathcal{D}_0 , are of approximately similar size and we are able to maintain

mutual label parity to avoid the class-imbalance problem. For ease of notation, we assume that the grid-graphs have orthogonal node features x_i , which get mixed after a single-step of message propagation and give v_i . Due to the symmetry of grid-graphs, we upper bound dot-product of nodes $v_i^\top v_j$ by θ . For bounded max-degree of the graphs, we represent $v_i^\top v_i = \theta_d$. By orthogonality among initial x_i , the relation $\theta_d \geq \theta$ follows. By construction, partitions $\mathcal{D}_1, \mathcal{D}_0$ are exhaustive datasets that contain all the grid-graphs satisfying the requisite property.

Over the course of optimization on the training data $\{\mathcal{D}_1 \cup \mathcal{D}_0\}$, we will next show that the weight parameter \mathbf{w} , as well as the attention parameter \mathbf{a} are most closely aligned along the direction of v^* . Proofs for the following results are deferred to the Appendix [19].

Lemma 3 (Orientation of \mathbf{w} parameter). *There exists a vector $v^* \in \mathbb{R}^d$, such that the component of $\frac{d\mathbf{w}}{dt}$ during optimization is always aligned along v^* and the weight parameter \mathbf{w} moves in the direction of v^* ,*

$$v^{*\top} \frac{d\mathbf{w}}{dt} > 0$$

We further characterize the gradient flow of attention parameter \mathbf{a} , and show that \mathbf{a} moves most along the direction v^* after $O(D)$ update iterations, where D is an upper bound on the maximum degree of graphs.

Lemma 4 (Orientation of \mathbf{a} parameter). *There exists a vector $v^* \in \mathbb{R}^d$, such that the component of $\frac{d\mathbf{a}}{dt}$ during optimization is always positive along v^* . After sufficient optimizer updates, the attention parameter \mathbf{a} moves most along the direction of v^**

$$v^{*\top} \frac{d\mathbf{a}}{dt} > 0$$

Based on the above lemmas, we can now state our main result regarding the biased behavior of graph convolution networks with attention based global pooling.

Theorem 5 (Implicit Bias of Global Attention Pooling Networks). *Consider the setting and assumptions of Section 5.1, and a 1-layer GCN coupled with attention based global pooling trained on \mathcal{D}_1 and \mathcal{D}_0 . The parameters $\{\mathbf{a}, \mathbf{w}\}$ at time t iterations of gradient descent, satisfy*

$$\langle \mathbf{a}, v^* \rangle > 0 \quad \text{and} \quad \langle \mathbf{w}, v^* \rangle > 0$$

for some v^* per Lemma 1 $\langle (A \cdot X)_{[\mathcal{S},:], v^*} \rangle > 0, \forall (A, X) \in \mathcal{D}_1$.

Proof sketch. It follows from Lemma 3 and 4 that parameters \mathbf{a} and \mathbf{w} are poised to align with v^* during optimization. Ensuring the existence of such v^* using Lemma 1 for training datasets \mathcal{D}_1 and \mathcal{D}_0 that are closed under completeness, it is easy to see that both the parameters will definitely have their component along v^* . \square

This is the main result of our work: it lays out an important property of GNN architectures that differentiate graphs based on the proximity of discriminative nodes found within the graph. Theorem 5 formalizes the biased nature of attention pooling and shows its preference to learn v^* , i.e. use *closely-connected-substructures* as discriminative features. Consequently, *average pooling*, which corresponds to uniform attention across all nodes, lacks the ability to selectively focus on specific nodes, thereby depending on aggregating nodes that may be distributed across the graph.

6 Experiments

To investigate the presence of inductive bias in diverse Graph Neural Network (GNN) architectures during gradient descent optimization, we pose the following pivotal question: Can two GNNs trained to achieve comparable accuracy on identical datasets ultimately learn characteristically different functions? We show this distinction by examining their behaviors on *unseen data*, which the models have not come across while training. Focusing specifically on graph classification, we narrow our scope to ascertain if GNNs develop classification capabilities based on the presence/absence of a designated subgraph G^* . To address this, we create both synthetic and semi-synthetic datasets, by embedding G^* within real-world datasets.

6.1 Datasets and Model Setting

Synthetic: The setting of grid-graphs, as seen in Section 4.2, provides a structured framework for investigating implicit biases within a controlled domain. The options for substituting and embedding G^* are inherently limited in this framework, due to the size of the grid that allows more control over its characteristics, including the balance between the positive and the negative graph instances. To elucidate the scenarios of extreme bias, it is imperative that the assumptions outlined in Section 5.1 remain valid, a condition readily attainable through these grid-graph datasets.

From Section 4.2 we know that the binary ground-truth labels for each partition are defined as $y_G = 1$, if $G \in \mathcal{D}_1$ and $y_G = 0$ if $G \in \mathcal{D}_0$.

Realworld (semi-synthetic): Due to the absence of precise ground truth node-information pertaining to the subgraph G^* , we resort to augmenting some of the well-established graph datasets, such as ZINC [13], Tox21 from MoleculeNet [31], and PROTEINS from TU Dataset [25], with G^* to create semi-synthetic datasets. This partial or complete planting of G^* yields distinct partitions \mathcal{D}_1 , \mathcal{D}_0 , and \mathcal{D}_\perp . Our full-version [19] shows some of the samples of real-world-derived graphs underscoring their random structural compositions. Importantly, this diversity serves to mitigate any unintended influences stemming from the inherent symmetry in the grid graphs. Furthermore, these datasets will help identify the presence of implicit biases within GNNs when applied to real-world graph data.

Planting G^* in realworld data \mathcal{D} : We adopt a parallel approach for realworld datasets akin to our treatment of grid-graphs, and embed the G^* subgraph either fully, partially, or sparsely. Details are deferred to the Appendix [19]. Our code is available [20].

Model baselines: In this study, we analyze GCN and GAT with standard global pooling (MAX,AVG,ATTN). Our focus lies in highlighting the phenomenon wherein GNNs that are trained on identical datasets, converge to different functional representations. We substantiate this bias through concrete examples showcasing divergent behaviors, achieved with minimal adjustments made to the architecture or training paradigms.

Table 3: Real datasets used to create partitions \mathcal{D}_1 , \mathcal{D}_0 , \mathcal{D}_\perp

Dataset	Average #nodes	Average #edges	Random colors	# Train graphs
ZINC [13]	23.2	49.8	24	10,000
Tox21 [31]	18.6	38.6	9	7,831
Proteins [25]	39.1	145.6	9	1,113

6.2 Performance on Real Datasets

To verify if the observations from Section 4.2 transcend to realworld scenarios, we perform experiments on four graph collection datasets, as described in Table 3. Appendix shows sample graphs used from these datasets while training [19]. The skewed proportions of label assignment is also seen in real datasets. Table 4 shows GCN with ATTN pooling assigns majority of the labels as 1, thus classifies on the basis of connections. This proportion is not as pronounced in AVG pooling because the classification is based on the presence of nodes as well. MAX pooling also has varied schemes of labeling among the three datasets. It is essential to acknowledge that these proportions are not as extreme as those observed in our toy dataset, owing to potential violations of assumptions such as data partition completeness. Nevertheless, the overarching trends persist, underscoring the robustness of our findings.

Table 4: Real world dataset results, GCN 1 layer, trained to 100% accuracy, averaged over multiple runs (standard deviation is order of magnitude smaller and so ignored)

Dataset	MAX		AVG		ATTN	
	$\hat{y} = 1$	$\hat{y} = 0$	$\hat{y} = 1$	$\hat{y} = 0$	$\hat{y} = 1$	$\hat{y} = 0$
ZINC	0	10000	8561	1439	3	9997
Protein	5531	34	1665	3900	6	5559
Tox21	0	23493	7576	15917	6	23487

6.3 Covering the Entire Attention-Spectrum

Table 5: Classification behavior of 1-layer GCN and ATTN pooling, with varying temperature coefficient of softmax operator. All baselines are uniformly run for 100 epochs of optimization with learning rate of 0.001, test data is of size 10k.

β Temperature	Train Accu.	$\hat{y} = 0$	$\hat{y} = 1$
1	0.9991	9806	194
4	0.9982	9704	296
10	0.993	9560	440
300	1.0000	3	9997

Attention pooling utilizes softmax to compute the weights for weighted summation. The softmax function can be modulated by a temperature coefficient, which governs the sharpness of the resulting distribution. $softmax_\beta = \frac{\exp(\mathbf{a}^T v_i / \beta)}{\sum_j \exp(\mathbf{a}^T v_j / \beta)}$. For β set to 1, the original softmax function is recovered. However, for smaller values of β approaching 0, the softmax tends towards behaving like a MAX pooling function. Conversely, as β tends towards infinity, the softmax function becomes increasingly smoothed out and resembles an AVG pooling, akin to $1/N(\sum v_i)$. We investigate different values of β to see the temperature point at which the inductive bias behavior switches. Table 5 shows that the flip in ATTN’s behavior happens at $\beta = 300$. Until then, ATTN pooling behaves much the same. Intuitively, this reinforces the separation between zero-attention (AVG) and non-zero attention (ATTN).

6.4 Delineating the Effect of Linear Layer

Equation.2 provides the formulation for end-to-end graph classification, employing attention-based global pooling within a single-layer GCN. It can appear that the biases discussed in preceding sections

Table 6: Effect of pruning out linear classifier and still observing consistent label assignment trends (Sec.6.4)

Pooling Method	With Linear Layer		Without Linear Layer	
	$\hat{y} = 0$	$\hat{y} = 1$	$\hat{y} = 0$	$\hat{y} = 1$
MAX	9988	12	9585	415
AVG	0	10000	2248	7752
ATTN	9940	60	9465	535

stem from the presence of the linear layer. While the GCN convolution layer facilitates message passing, the subsequent pooling layer condenses all nodes into a unified graph representation, followed by the linear layer assigning a logic for classification. To disentangle the influence of the linear layer from the GCN, we opt to prune out the linear classifier and allow the GCN to perform the necessary dimension reduction, resulting in $\Theta \in \mathbb{R}^{N \times 1}$. Subsequently, we look into the alterations in graph classification behavior resulting from the removal of the linear classifier.

Table 6 illustrates that the proportion of nodes classified as $\hat{Y} = 0$ remains consistent irrespective of the presence or absence of the linear layer. Upon closer examination, it becomes evident that the absence of the linear layer leads to elevated training loss, potentially indicating that the dimension of parameters $\Theta \in \mathbb{R}^{N \times 1}$ might be insufficient to effectively capture meaningful graph representations. Nonetheless, the initial bias persists in the distribution of nodes across each label. This observation underscores the intrinsic bias inherent in the layers of GNNs, even after pruning the linear classifier.

6.5 Effect of Hierarchical Pooling

Table 7: Classification behavior with hierarchical ASAP pooling.

Hierarchical Pooling	Global Pool	Training Accu.	$\hat{y} = 0$	$\hat{y} = 1$
GCN + ASAP Pooling	ATTN	0.9875	8736	1264
	AVG	0.9825	8519	1481
	MAX	0.9725	8491	1509

Global pooling methods operate simultaneously on all the nodes within a graph, thereby failing to encode its structural nuances. While this aggregation process condenses all nodes into a singular representation within one single time-step, in contrast, Hierarchical pooling methods progressively coarsen the graph across multiple iterations, thereby preserving its hierarchical structure. We scrutinize the behaviors of select hierarchical pooling operators, such as ASAP [27]. Intuitively, these operators are expected to identify and prioritize significant subgraphs for classification purposes. Table 7 shows a striking resemblance in the behaviors of hierarchical pooling operators to that of attention-based global pooling. A plausible explanation for this similarity lies in their shared objective of discerning connected patterns to effectively aggregate graph information, thereby exhibiting analogous behaviors.

6.6 Discussion

Our work has practical implications in tasks where domain knowledge guides graph classification. It can help distinguish whether classification depends on coherent subgraph presence or merely node presence, regardless of neighbors. For instance, in chemistry, chemical reaction likelihood may hinge on specific node configurations, while compound toxicity could relate to certain heavy metal atoms.

Thus, attention-based global pooling suits chemical reactions, while average pooling may be apt for detecting poisonous heavy metals.

We opted to analyze Attention pooling since it can be parameterized by adjusting the temperature coefficient of the sigmoid operator. Attention pooling can be analyzed across a spectrum ranging from MAX to AVG, with ATTN lying in between. The theoretical results hold for any graph G^* with diameter 2 (i.e., star-graphs). We have depicted using 3-node graphs for simpler illustrations, but the conclusions remain unchanged for larger star-graphs. Empirically we find that the same trends and conclusions hold for G^* with a larger diameter.

7 Related Works

Inductive biases are common in machine learning, and they shape the learned functions. Architecture-induced biases stem from design choices, enhancing performance on specific data types [15]. For instance, convolutional neural networks with max-pooling exhibit translation invariance over images, while graph neural networks are permutation invariant in node ordering over graphs [16, 17]. Optimizer-induced biases can arise in learned models due to factors such as the loss function’s nature, parameter initialization, and regularization [26]. In this work we characterize the nature of pooling mechanisms in learning different kinds of subgraph patterns.

Graph global pooling operators implement functions to incorporate node’s information into concise and reduced graph representation. After sufficient layers of message-passing, there are simple single-step operations such as `max`, `sum`, `avg` that take all nodes at once to derive graph’s representation [10, 3]. In similar spirit, global attention pooling can lay emphasis on certain *important* nodes and arrive at graph’s representation by asymmetric weighted addition [21]. There are other stepwise pooling methods, such as ASAP, Top-k [27, 12], that coarsen the graph structure in several steps. In our work, we focus on characterizing the implicit bias of single step global pooling mechanisms. To the best of our knowledge, there aren’t any previous works that evaluate implicit inductive bias of GNNs to preferentially learn certain functions for graph classification tasks. There is a large gap to be filled with practical considerations of optimization and learnability. MPNNs can count a context of subgraph-counting, for 3 or fewer nodes in the subgraph (except star-shaped subgraphs) [6]. While there is a great body of work on representation power of GNNs, there is a large gap with its practical optimization and learnability property.

8 Conclusion

We highlighted some of the gaps between the representation theory and optimization aspect of GNNs. We showed the existence of implicit inductive bias in GNNs, specifically we proved the bias of graph convolution networks with attention based global pooling. We showed theoretically and empirically the preference of attention based architectures, to look for a *closely-connected* patterns for graph classification. We discussed the implications of our work in incorporating domain knowledge.

The implicit bias herein represents just one facet, namely, the preferential treatment of nodes occurring together or dispersed across the graph. Other biases may exist, such as propensity to assign different labels to homophilic/heterophilic graphs versus treating them uniformly. Investigating these nuanced relationships between GNNs, optimizers, and real-world applications are avenues for future exploration.

Acknowledgements

We gratefully acknowledge the support of Robert Bosch Centre for Data Science & Artificial Intelligence (RBCDSAI) at Indian Institute of Technology Madras, as well as our anonymous reviewers for providing valuable feedback.

References

- [1] K. Amara, Z. Ying, Z. Zhang, Z. Han, Y. Zhao, Y. Shan, U. Brandes, S. Schemm, and C. Zhang. Graphframex: Towards systematic evaluation of explainability methods for graph neural networks. In *Learning on Graphs Conference, LoG 2022, 9-12 December 2022, Virtual Event*, volume 198 of *Proceedings of Machine Learning Research*, page 44. PMLR, 2022.
- [2] P. W. Battaglia, R. Pascanu, M. Lai, D. J. Rezende, and K. Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4502–4510, 2016.
- [3] F. M. Bianchi and V. Lachi. The expressive power of pooling in graph neural networks. In *Advances in Neural Information Processing Systems 36, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.
- [4] F. M. Bianchi, D. Grattarola, and C. Alippi. Spectral clustering with graph neural networks for graph pooling. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 874–883. PMLR, 2020.
- [5] S. Brody, U. Alon, and E. Yahav. How attentive are graph attention networks? In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [6] Z. Chen, L. Chen, S. Villar, and J. Bruna. Can graph neural networks count substructures? In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [7] D. Cheng, X. Wang, Y. Zhang, and L. Zhang. Graph neural network for fraud detection via spatial-temporal attention. *IEEE Trans. Knowl. Data Eng.*, 34(8):3800–3813, 2022.
- [8] L. Chizat and F. R. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3040–3050, 2018.
- [9] A. Derron-Pinon, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, and P. Velickovic. ETA prediction with graph neural networks in google maps. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, pages 3767–3776. ACM, 2021.
- [10] D. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2224–2232, 2015.
- [11] F. Errica, M. Podda, D. Bacciu, and A. Micheli. A fair comparison of graph neural networks for graph classification. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [12] H. Gao and S. Ji. Graph u-nets. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2083–2092. PMLR, 2019.
- [13] R. Gómez-Bombarelli, D. Duvenaud, J. M. Hernández-Lobato, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *CoRR*, abs/1610.02415, 2016.
- [14] M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, volume 47 of *Lecture Notes in Logic*. Cambridge University Press, 2017. ISBN 9781139028868.
- [15] S. Jegelka. Theory of graph neural networks: Representation and learning. *CoRR*, abs/2204.07697, 2022.
- [16] E. Kauderer-Abrams. Quantifying translation-invariance in convolutional neural networks. *CoRR*, abs/1801.01450, 2018.
- [17] N. Keriven and G. Peyré. Universal invariant and equivariant graph neural networks. In *Advances in Neural Information Processing Systems 32, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7090–7099, 2019.
- [18] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [19] P. K. Kumar and H. G. Ramaswamy. Graph classification with gnn: Optimisation, representation and inductive bias, 2024. URL <https://arxiv.org/abs/2408.09266>.
- [20] P. K. Kumar and H. G. Ramaswamy. Graph classification with gnn : Optimisation, representation and inductive bias. <https://github.com/pkrns4/gnn-inductive-bias/>, 2024.
- [21] J. Lee, I. Lee, and J. Kang. Self-attention graph pooling. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3734–3743. PMLR, 2019.
- [22] N. Mehta, M. L. Pacheco, and D. Goldwasser. Tackling fake news detection by continually improving social context representations using graph neural networks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1363–1380. Association for Computational Linguistics, 2022.
- [23] G. Montúfar, R. Pascanu, K. Cho, and Y. Bengio. On the number of linear regions of deep neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2924–2932, 2014.
- [24] C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. E. Lenssen, G. Rattan, and M. Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, 2019*, pages 4602–4609. AAAI Press, 2019.
- [25] C. Morris, N. M. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. *CoRR*, abs/2007.08663, 2020.
- [26] T. A. Poggio and Q. Liao. Explicit regularization and implicit bias in deep network classifiers trained with the square loss. *CoRR*, abs/2101.00072, 2021.
- [27] E. Ranjan, S. Sanyal, and P. P. Talukdar. ASAP: adaptive structure aware pooling for learning hierarchical graph representations. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 5470–5477. AAAI Press, 2020.
- [28] M. S. Schlichtkrull, N. D. Cao, and I. Titov. Interpreting graph neural networks for NLP with differentiable edge masking. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [29] M. Telgarsky. benefits of depth in neural networks. In V. Feldman, A. Rakhlin, and O. Shamir, editors, *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, volume 49 of *JMLR Workshop and Conference Proceedings*, pages 1517–1539. JMLR.org, 2016.
- [30] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [31] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. S. Pande. Moleculenet: A benchmark for molecular machine learning. *CoRR*, abs/1703.00564, 2017.
- [32] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [33] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 9240–9251, 2019.