

On the Expressiveness and Generalization of Hypergraph Neural Networks

by

Zhezheng Luo

B.S. in Computer Science and Engineering and in Mathematics,
Massachusetts Institute of Technology (2021)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2023

© Massachusetts Institute of Technology 2023. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
Dec 12, 2022

Certified by.....
Leslie Pack Kaelbling
Panasonic Professor of Computer Science and Engineering
Thesis Supervisor

Accepted by
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

On the Expressiveness and Generalization of Hypergraph Neural Networks

by

Zhezheng Luo

Submitted to the Department of Electrical Engineering and Computer Science
on Dec 12, 2022, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Graph Neural Networks have demonstrated their success on many applications, including analyzing molecules and social networks. Although these graph neural networks can effectively determine pairwise connections between nodes, the data structure in reality sometimes goes beyond pairwise relations and can be complicated, involving multiple nodes. This requires the graph neural networks to be extended to hypergraphs to deal with higher-order relations. It is critical to understand what type of problems these hypergraph neural networks can solve and effectively learn from data.

In this thesis, we describe how we use Neural Logical Machines as a unified framework for analyzing the expressiveness, learning, and (structural) generalization of hypergraph neural networks (HyperGNNs). Specifically, we focus on how HyperGNNs can learn from finite datasets and generalize structurally to graph reasoning problems of arbitrary input sizes. Our first contribution is a fine-grained analysis of the expressiveness of HyperGNNs, that is, the set of functions that they can realize. Our result is a hierarchy of problems they can solve, defined in terms of various hyperparameters such as depths and edge arities. Next, we analyze the learning properties of these neural networks, especially focusing on how they can be trained on a finite set of small graphs and generalize to larger graphs, which we term structural generalization. Our theoretical results are further supported by the empirical results.

Thesis Supervisor: Leslie Pack Kaelbling

Title: Panasonic Professor of Computer Science and Engineering

Acknowledgments

First and foremost, I would like to express my heartfelt gratitude to my advisor, Professor Leslie Kaelbling. Throughout my research, Leslie has been a constant source of inspiration and support. She has guided me in finding research topics that interest me and has provided valuable insights on my project from a variety of perspectives. Additionally, Leslie is the most hands-on professor I have ever worked with, carefully reviewing every detail of my research from the stories, principles, and flows to every figure, sentence, and connection. I have learned so much from Leslie, including the importance of making my research educational, thinking about problems in a systematic and principled way, and being open-minded and critical while conducting research.

My gratitude also extends to Jiayuan Mao, who has been my mentor and friend throughout my research since my junior year. I have learned so much from Jiayuan about research, including the crucial role that problem formulation plays in research and the importance of presentation in collaboration and communication. Jiayuan has also been a great source of enthusiasm and inspiration for me, and we have had many stimulating discussions about various ideas. I am also grateful to Professor Jiajun Wu, who introduced me to Jiayuan and guided me on how to conduct research in a systematic way, as well as taught me the mindset for research. I would also like to thank Chuang Gan for introducing me to my first research project and providing valuable advice on both research and life.

I also want to express my thanks to Josh Tenenbaum, Professor Tomer Ullman and Professor Tomás Lozano-Pérez for their guidance and support in my projects.

I am also very grateful to Professor Aleksander Mądry, who has been my undergraduate advisor and has supported me in planning both my academic and personal lives. I also want to thank Professor Michael Sipser and Professor Piotr Indyk for giving me the opportunity to TA their classes. Being a TA has been a fun and rewarding experience, and I have learned a lot about teaching and organizing large classes.

Then, I would like to thank my friends for the games, trips, and parties we have

shared, as well as their company, support, and encouragement during the tough times. I want to give a special shout-out to Lijie Chen for his help, suggestions, and conversations about academia, career, and life decisions.

I cannot express enough gratitude to Yunyun Wang for the love and companionship that she has given me over the past six years. Together, we have crossed half the earth, taking in the grandeur of life and overcoming its obstacles hand in hand. She has always been there for me, offering me her unwavering faith, encouragement, and support.

Finally, I thank my parents, for their unconditional love.

Contents

1	Introduction	13
2	Preliminary: Hypergraph Reasoning Problems and Hypergraph Neural Networks	15
2.1	Hypergraph representation and functions	15
2.2	Hypergraph Neural Networks	16
2.2.1	Higher-order Graph Neural Networks	16
2.2.2	Neural Logic Machines	18
2.3	Equivalence of Expressiveness between Hypergraph Models	20
3	Expressiveness of Relational Neural Networks	23
3.1	Arity Hierarchy	24
3.2	Depth Hierarchy	25
4	Learning and Generalization in Relational Neural Networks	29
4.1	Structural Generalization under Complete Presentation	30
5	Experimental Results	33
5.1	Experiment Setup	34
5.2	Substructure Detection	34
5.2.1	Relational Reasoning	34
5.2.2	Model Implementation Details	35
5.3	Results and Analysis	35

6	Related Work	39
7	Conclusion	41
A	Hypergraph Neural Networks	43
A.1	Expressiveness Equivalence of Relational Neural Networks	43
A.2	Expressiveness of hypergraph convolution and attention	47
A.3	The Time and Space Complexity of NLMs	48
B	Arity and Depth Hierarchy: Proofs and Analysis	51
B.1	Proof of Theorem 3.1: Arity Hierarchy.	51
B.2	Upper Depth Bound for Unbounded-Precision NLM.	52
B.3	Graph Problems	52

List of Figures

2-1	The overall architecture of our Neural Logic Machines (NLMs). It follows the computation graph of NLM [11] and can be applied to hypergraphs.	19
5-1	How the performance of models drop when generalizing to larger graphs on the problem connectivity-4 (trained on graphs with size 10). . . .	37

List of Tables

3.1	The minimum depth and arity of NLMs for solving graph classification and regression tasks.	24
5.1	Overall accuracy on structure detecting problems.	36
5.2	Overall accuracy on relational reasoning problems.	36

Chapter 1

Introduction

Graph-like structures are very common in daily life. Road maps, social media networks and the internet can all be represented by graphs. Reasoning over these graph structures is an important task in many applications, including route planning, molecule analysis, social network modeling, and knowledge graph reasoning. In recent years, Graph Neural Networks have demonstrated their success in many applications [14, 23, 30], showing that graph neural networks can effectively determine the pairwise connection between nodes. However, some of the graph-structured data is beyond pairwise relations and can be complicated involving multiple nodes, such as triangles in social networks and molecular groups. These require the graph neural networks to be extended to hypergraphs to deal with higher-order relations.

While we have seen great success of various relational neural networks that can extend to hypergraphs, such as Graph Neural Networks [GNNs; 29] and Neural Logical Machines [NLM; 11] in a variety of applications [4, 25, 31], we do not yet have a full understanding of how different design parameters, such as the depth of the neural network, affects the expressiveness of these models, or how effectively these models generalize from limited data.

This thesis analyzes the *expressiveness* and *generalization* of relational neural networks applied to *hypergraphs*, which are graphs with edges connecting more than two nodes. Literature has shown that even when the inputs and outputs of models have only unary and binary relations, allowing intermediate hyperedge representations

increases the expressiveness [1, 5]. In this thesis, we further formally show the “if and only if” conditions for the expressive power with respect to the edge arity. That is, k -ary hyper-graph neural networks are sufficient and necessary for realizing FOC- k , a fragment of first-order logic with counting quantification which involves at most k variables. This is a helpful result because now we can determine whether a specific hypergraph neural network can solve a problem by understanding what form of logic formula can represent the solution to this problem. Next, we formally describe the relationship between expressiveness and non-constant-depth networks. We state a conjecture about the “depth hierarchy,” and connect the potential proof of this conjecture to the distributed computing literature.

Furthermore, we prove, under certain assumptions, it is possible to train a hyper-graph neural network on a finite set of small graphs, and it will generalize to arbitrarily large graphs. This ability results from the weight-sharing nature of hypergraph neural networks.

We hope our work can serve as a foundation for designing hypergraph neural networks: to solve a specific problem, what arity do you need? What depth do you need? Will my model have structural generalization (i.e., to larger graphs)?

Our theoretical results on learning are further supported by experiments, for empirical demonstration of the theorems and for exploring the tightness of the sample-complexity bounds and effectiveness of different aggregation functions. Specifically, we show that NLMs with sufficient expressiveness power can be trained using gradient descent to realize the desired function. Meanwhile, when the training data distribution is uniform, NLMs do demonstrate strong generalization to larger graphs than those seen during training. We also empirically show the generalization of NLMs with different aggregation functions.

Chapter 2

Preliminary: Hypergraph Reasoning Problems and Hypergraph Neural Networks

2.1 Hypergraph representation and functions

A *hypergraph representation* G is a tuple (V, X) , where V is a set of entities (nodes), and X is a set of *hypergraph representation functions*. Specifically, $X = \{X_0, X_1, X_2, \dots, X_k\}$, where $X_j : (v_1, v_2, \dots, v_j) \rightarrow \mathcal{S}$ is a function mapping every tuple of j nodes to a value. We call j the *arity* of the hyperedge and k is the max arity of input hyperedges. The range \mathcal{S} can be any set of discrete labels that describes relation type, or a scalar number (e.g., the length of an edge), or a vector. We will use the arity 0 representation $X_0(\emptyset) \rightarrow \mathcal{S}$ to represent any global properties of the graph.

A *graph reasoning function* f is a mapping from a hypergraph representation $G = (V, X)$ to another hyperedge representation function Y on V . As concrete examples, asking whether a graph is fully connected is a graph classification problem, where the output $Y = \{Y_0\}$ and $Y_0(\emptyset) \rightarrow \mathcal{S}' = \{0, 1\}$ is a global label; finding the set of disconnected subgraphs of size k is a k -ary hyperedge classification problem, where the output $Y = \{Y_k\}$ is a label for each k -ary hyperedges.

Example (undirected graph connectivity). The input is $G = (V, X)$, where $X = \{X_0, X_1, X_2\}$, in which X_2 is mapping from each pair of nodes $v_1, v_2 \in V$ to $\{0, 1\}$, indicating whether (v_1, v_2) are connected. The input is an undirected graph; thus, $X_2(v_1, v_2) = X_2(v_2, v_1)$ for every v_1, v_2 . Note that X_0 and X_1 are constant functions and not used for describing the undirected graph. The output is $Y = \{Y_0\}$, where $Y_0(\emptyset) \rightarrow \{0, 1\}$ is a binary classification label. The desired label is 1 when the graph is connected, and 0 otherwise.

2.2 Hypergraph Neural Networks

There are two main motivations and constructions of a neural network applied to graph reasoning problems: message-passing-based and first-order-logic-inspired. Both approaches construct the computation graph layer by layer. The input is the features of nodes and hyperedges, while the output is the per-node or per-edge prediction of desired properties, depending on the task.

2.2.1 Higher-order Graph Neural Networks

In a nutshell, within each layer, *message-passing-based* hypergraph neural networks, Higher-Order GNNs [26], perform message passing between each hyperedge and its neighbours. Specifically, Higher-order Graph Neural Networks [HO-GNNs; 26] are Graph Neural Networks (GNNs) that apply to hypergraphs. A GNN is usually defined based on two message passing operations.

- Edge update: the feature of each edge is updated by features of its ends.
- Node update: the feature of each node is updated by features of all edges adjacent to it.

However, computing only node-wise and edge-wise features does not handle higher-order relations, such as triangles in the graph. In order to obtain more expressive power, GNNs have be extend to hypergraphs of higher arity [26]. Specifically, HO-GNNs on B -ary hypergraph maintains features for all B -tuple of nodes, and the neighborhood

is extended to B -tuples accordingly: the feature of tuple (v_1, v_2, \dots, v_B) is updated by the $|V|$ element multiset (contain $|V|$ elements for each $u \in V$) of B -tuples of features

$$(H_i[u, v_2, \dots, v_B], H_{i-1}[v_1, u, v_2, \dots, v_B], \dots, H_{i-1}[v_1, \dots, v_{B-1}, u]) \quad (2.2.1)$$

where $H_{i-1}[\mathbf{v}]$ is the feature of tuple \mathbf{v} from the previous iteration.

We now introduce the formal definition of the high-dimensional message passing. We denote \mathbf{v} as a B -tuple of nodes (v_1, v_2, \dots, v_B) , and generalize the neighborhood to a higher dimension by defining the neighborhood of \mathbf{v} as all node tuples that differ from \mathbf{v} at one position.

$$\begin{aligned} \text{Neighbors}(\mathbf{v}, u) = & ((u, v_2, \dots, v_B), \\ & (v_1, u, v_3, \dots, v_B), \\ & \dots, \\ & (v_1, \dots, v_{B-1}, u)) \end{aligned} \quad (2.2.2)$$

$$N(\mathbf{v}) = \{\text{Neighbors}(\mathbf{v}, u) | u \in V\} \quad (2.2.3)$$

Then message passing scheme naturally generalizes to high-dimensional features using the high-dimensional neighborhood.

$$\text{Received}_i[\mathbf{v}] = \sum_u (\text{NN}_1 (H_{i-1}[\mathbf{v}]; \text{CONCAT}_{\mathbf{v}' \in \text{neighbors}(\mathbf{v}, u)} H_{i-1}[\mathbf{v}'])) \quad (2.2.4)$$

»We now introduce two important hypergraph neural network implementations that can be trained to solve graph reasoning problems: Higher-order Graph Neural Networks [HO-GNN; 26] and Neural Logic Machines [NLM; 11]. They are equivalent to each other in terms of expressiveness. Showing this equivalence allows us to focus the rest of the paper on analyzing a single model type, with the understanding that the conclusions generalize to a broader class of hypergraph neural networks.

2.2.2 Neural Logic Machines

On the other hand, *first-order-logic-inspired* hypergraph neural networks consider building neural networks that can emulate first logic formulas. Neural Logic Machines [NLM; 11] are defined in terms of a set of input hyperedges; each hyperedge of arity k is represented by a vector of (possibly real) values obtained by applying all of the k -ary predicates in the domain to the tuple of vertices it connects. Each layer in an NLM learns to apply a linear transformation with nonlinear activation and quantification operators (analogous to the for all \forall and exists \exists quantifiers in first-order logic), on these values. It is easy to prove, by construction, that given a sufficient number of layers and maximum arity, NLMs can learn to realize any first-order-logic formula.

In practice, a NLM is a multi-layer neural network that operates on hypergraph representations, in which the hypergraph representation functions are represented as tensors. The input is a hypergraph representation (V, X) . There are then several computational layers, each of which produces a hypergraph representation with nodes V and a new set of representation functions. Specifically, a B -ary NLM produces hypergraph representation functions with arities from 0 up to a maximum hyperedge arity of B . We let $T_{i,j}$ denote the tensor representation for the output at layer i and arity j . Each entry in the tensor is a mapping from a set of node indices (v_1, v_2, \dots, v_j) to a vector in a latent space \mathbb{R}^W . Thus, $T_{i,j}$ is a tensor of $j + 1$ dimensions, with the first j dimensions corresponding to j -tuple of nodes, and the last feature dimension. For convenience, we write $h_{0,\cdot}$ for the input hypergraph representation and $h_{D,\cdot}$ for the output of the NLM.

Fig. 2-1a shows the overall architecture of NLMs. It has $D \times B$ computation blocks, namely relational reasoning layers (RRLs). Each block $\text{RRL}_{i,j}$, illustrated in Fig. 2-1b, takes the output from neighboring arities in the previous layer, $T_{i-1,j-1}$, $T_{i-1,j}$ and $T_{i-1,j+1}$, and produces $T_{i,j}$. Below we show the computation of each primitive operation in an RRL.

The *expand* operation takes tensor $T_{i-1,j-1}$ (arity $j - 1$) and produces a new tensor $T_{i-1,j-1}^E$ of arity j . The *reduce* operation takes tensor $T_{i-1,j+1}$ (arity $j + 1$) and produces

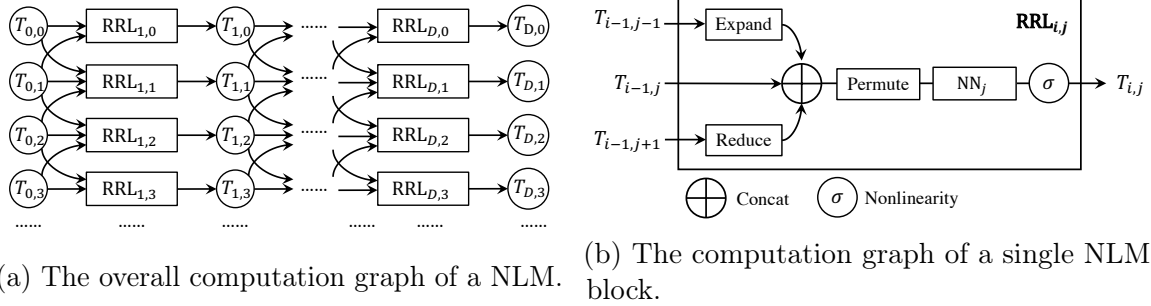


Figure 2-1: The overall architecture of our Neural Logic Machines (NLMs). It follows the computation graph of NLM [11] and can be applied to hypergraphs.

a new tensor $T_{i-1,j+1}^R$ of arity $j + 1$. Mathematically,

$$T_{i-1,j-1}^E[v_1, v_2, \dots, v_j] = T_{i-1,j-1}[v_1, v_2, \dots, v_{j-1}];$$

$$T_{i-1,j+1}^R[v_1, v_2, \dots, v_j] = \text{Agg}_{v_{j+1}} \{T_{i-1,j+1}[v_1, v_2, \dots, v_j, v_{j+1}]\}.$$

Here, Agg is called the aggregation function of a NLM. For example, a sum aggregation function takes the summation along the dimension $j + 1$ of the tensor, and a max aggregation function takes the max along that dimension.

The *concat* (concatenate) operation \oplus is applied at the “vector representation” dimension. The *permute* operation generates a new tensor of the same arity, but it fuses the representations of hyperedges that share the same set of entities but in different order, such as (v_1, v_2) and (v_2, v_1) . Mathematically, for tensor X of arity j , if $Y = \text{permute}(X)$ then

$$Y[v_1, v_2, \dots, v_j] = \text{Concat}_{\sigma \in S_j} \{X[v_{\sigma_1}, v_{\sigma_2}, \dots, v_{\sigma_j}]\},$$

where $\sigma \in S_j$ iterates over all permutations of $\{1, 2, \dots, j\}$. NN_j is a multi-layer perceptron (MLP) applied to each entry in the tensor produced after permutation, with nonlinearity σ (e.g., ReLU).

It is important to note that we intentionally name the MLPs NN_j instead of $\text{NN}_{i,j}$.

In generalized relational neural networks, for a given arity j , all MLPs across all layers i are shared. It is straightforward to see that this “weight-shared” model can realize a “non-weight-shared” NLM that uses different weights for MLPs at different layers when the number of layers is a constant. With a sufficiently large length of the representation vector, we can simulate the computation of applying different transformations by constructing block matrix weights. (A more formal proof is in Appendix A) The advantage of this weight sharing is that the network can be easily extended to a “recurrent” model. For example, we can apply the NLM for a number of layers that is a function of n , where n is the the number of nodes in the input graph. Thus, we will use the term *layers* and *iterations* interchangeably.

Handling high-arity features and using deeper models usually increase the computational cost. In Appendix A.3, we show that the time and space complexity of NLM $[D, B]$ is $O(Dn^B)$.

Note that even when hyperparameters such as the maximum arity and the number of iterations are fixed, a NLM is still a *model family* \mathcal{M} : the weights for MLPs will be trained on some data. Furthermore, each model $M \in \mathcal{M}$ is a NLM with a specific set of MLP weights.

2.3 Equivalence of Expressiveness between Hypergraph Models

In this section, we will prove the equivalence between HO-GNNs and NLMs. We first describe a framework for quantifying if two hypergraph neural network models are equally expressive on regression tasks (which is more general than classification problems). The framework view the expressiveness from the perspective of computation. Specifically, we will prove the expressiveness equivalence between models by showing that their computation can be aligned.

In complexity, we usually show a problem is at least as hard as the other one by showing a reduction from the other problem to the problem. Similarly, on the

expressiveness of NLMs, we can construct reduction from model family \mathcal{A} to model family \mathcal{B} to show that \mathcal{B} can realize all computation that \mathcal{A} does, or even more. Formally, we have the following definition.

Definition 2.1 (Expressiveness reduction). For two model families \mathcal{A} and \mathcal{B} , we say \mathcal{A} can be **reduced** to \mathcal{B} if and only if there is a function $r : \mathcal{A} \rightarrow \mathcal{B}$ such that for each model instance $A \in \mathcal{A}$, $r(A) \in \mathcal{B}$ and A have the same outputs on all inputs. In this case, we say \mathcal{B} is at least as expressive as \mathcal{A} .

Definition 2.2 (Expressiveness equivalence). For two model families \mathcal{A} and \mathcal{B} , if \mathcal{A} and \mathcal{B} can be reduced to each other, then \mathcal{A} and \mathcal{B} are equally expressive. Note that this definition of expressiveness equivalence generalizes to both classification and regression tasks.

Theorem 2.1. HO-GNNs [26] are equivalent to NLMs in terms of expressiveness. Specifically, a B -ary HO-GNN is equivalent to an NLM applied to $B+1$ -ary hyperedges. Proofs are in Appendix A.1.

Given Theorem 2.1, we can focus on just one single type of hypergraph neural network. Specifically, we will focus on Neural Logic Machines [NLM; 11] because its architecture naturally aligns with first-order logic formula structures, which will aid some of our analysis. An NLM is characterized by hyperparameters D (depth), and B maximum arity. We are going to assume that B is a constant, but D can be dependent on the size of the input graph. We will use $\text{NLM}[D, B]$ to denote an NLM family with depth D and max arity B . Other parameters such as the width of neural networks affects the precise details of what functions can be realized, as it does in a regular neural network, but does not affect the analyses in this thesis. Furthermore, we will be focusing on neural networks with bounded precision, and briefly discuss how our results generalize to unbounded precision cases.

Chapter 3

Expressiveness of Relational Neural Networks

A NLM is characterized by hyperparameters D (depth), and B maximum arity. In this section, we will focus on D and B , quantifying how they affect the expressiveness of a model; i.e., the set of graph reasoning functions a model family can realize. We will use the notation $\text{NLM}[D, B]$ to represent a NLM family with depth D and max arity B . A summary of some concrete results is in Table 3.1 and discussed in detail in Appendix B.3.

We will provide definition of the arity hierarchy and prove it, showing that NLMs with higher arity are always more expressive. We will also show an example of the arity hierarchy on the FOL realization bounds. Then, we raise a conjecture on the depth hierarchy and discuss it under different scenarios of precisions. We start from a formal definition of hypergraph neural network expressiveness.

Definition 3.1 (Expressiveness). We say a model family \mathcal{M}_1 is *at least expressive as* \mathcal{M}_2 , written as $\mathcal{M}_1 \succcurlyeq \mathcal{M}_2$, if for all $M_2 \in \mathcal{M}_2$, there exists $M_1 \in \mathcal{M}_1$ such that M_1 can realize M_2 . A model family \mathcal{M}_1 is *more expressive than* \mathcal{M}_2 , written as $\mathcal{M}_1 \succ \mathcal{M}_2$, if $\mathcal{M}_1 \succcurlyeq \mathcal{M}_2$ and $\exists M_1 \in \mathcal{M}_1, \forall M_2 \in \mathcal{M}_2, M_2$ can not realize M_1 .

$B = 4$	4-Clique Detection NLM[$O(1), 4$]	4-Clique Count NLM[$O(1), 4$]
$B = 3$	Triangle Detection NLM[$O(1), 3$]	All-Pair Distance NLM[$O(\log n), 3$]*
	Bipartiteness NLM[$O(\log n), 3$]*	
	All-Pair Connectivity NLM[$O(\log n), 3$]*	
	All-Pair Connectivity- k NLM[$O(\log k), 3$]*	
$B = 2$	FOC ₂ Realization NLM[$\cdot, 2$] [3]	S-T Distance NLM[$O(n), 2$]
	3/4-Link Detection NLM[$O(1), 2$]	Max Degree NLM[$O(1), 2$]
	S-T Connectivity NLM[$O(n), 2$]	Max Flow NLM[$O(n^3), 2$]*
	S-T Connectivity- k NLM[$O(k), 2$]	
$B = 1$	Node Color Majority: NLM[$O(1), 1$]	Count Red Nodes: NLM[$O(1), 1$]
	Classification Tasks	Regression Tasks

Table 3.1: The minimum depth and arity of NLMs for solving graph classification and regression tasks. The * symbol indicates that these are conjectured lower bounds.

3.1 Arity Hierarchy

We first aim to quantify how the maximum arity B of the network’s representation affects its expressiveness and find that, in short, even if the inputs and outputs of neural networks are of low arity, the higher the maximum arity for intermediate layers, the more expressive the NLM is.

Corollary 3.1 (Arity Hierarchy). For any maximum arity B , there exists a depth D^* such that: $\forall D \geq D^*$, NLM[$D, B + 1$] is more expressive than NLM[D, B]. This theorem applies to both fixed-precision* and unbounded-precision networks.

Proof sketch: Our proof slightly extends the proof of Morris et al. [26]. First, the set of graphs distinguishable by NLM[D, B] is bounded by graphs distinguishable by a D -round order- B Weisfeiler-Leman test [21]. If models in NLM[D, B] cannot generate different outputs for two distinct hypergraphs G_1 and G_2 , but there exists $M \in$ NLM[$D, B + 1$] that *can* generate different outputs for G_1 and G_2 , then we can construct a graph classification function f that NLM[$D, B + 1$] (with some fixed

*Fixed precision: the results of intermediate layers (tensors) are constant-sized (e.g., W bits per entry). Practical GNNs are all fixed-precision because real number types in modern computers have finite precision.

precision) can realize but $\text{NLM}[D, B]$ (even with unbounded precision) cannot.[†] The full proof is described in Appendix B.1.

It is also important to quantify the minimum arity for realizing certain graph reasoning functions.

Corollary 3.2 (FOL realization bounds). Let FOC_B denote a fragment of first order logic with at most B variables, extended with counting quantifiers of the form $\exists^{\geq n} \phi$, which state that there are at least n nodes satisfying formula ϕ [7].

- (Upper Bound) Any function f in FOC_B can be realized by $\text{NLM}[D, B]$ for some D .
- (Lower Bound) There exists a function $f \in \text{FOC}_B$ such that for all D , f cannot be realized by $\text{NLM}[D, B - 1]$.

Proof: The upper bound part of the claim has been proved by Barceló et al. [3] for $B = 2$. The results generalize easily to arbitrary B because the counting quantifiers can be realized by sum aggregation. The lower bound part can be proved by applying Section 5 of [7], in which they show that FOC_B is equivalent to a $(B - 1)$ -dimensional WL test in distinguishing non-isomorphic graphs. Given that $\text{NLM}[D, B - 1]$ is equivalent to the $(B - 2)$ -dimensional WL test of graph isomorphism, there must be an FOL_B formula that distinguishes two non-isomorphic graphs that $\text{NLM}[D, B - 1]$ cannot. Hence, FOL_B cannot be realized by $\text{NLM}[\cdot, B - 1]$.

3.2 Depth Hierarchy

We now study the dependence of the expressiveness of NLMs on depth D . Neural networks are generally defined to have a fixed depth, but allowing them to have a depth that is dependent on the number of nodes $n = |V|$ in the graph, in many cases, can substantially increase their expressive power [8, see also Theorem 3.4 and Appendix B for examples]. In the following, we define a *depth hierarchy* by analogy to the *time hierarchy* in computational complexity theory [16], and we extend our

[†]Note that the arity hierarchy is applied to fixed-precision and unbounded-precision separately. For example, $\text{NLM}[D, B]$ with unbounded precision is incomparable with $\text{NLM}[D, B + 1]$ with fixed precision.

notation to let $\text{NLM}[O(f(n)), B]$ denote the class of adaptive-depth NLMs in which the growth-rate of depth D is bounded by $O(f(n))$.

Conjecture 3.3 (Depth hierarchy). For any maximum arity B , for any two functions f and g , if $g(n) = o(f(n)/\log n)$, that is, f grows logarithmically more quickly than g , then fixed-precision $\text{NLM}[O(f(n)), B]$ is more expressive than fixed-precision $\text{NLM}[O(g(n)), B]$.

There is a closely related result for the *congested clique* model in distributed computing, where [20] proved that $\text{CLIQUE}(g(n)) \subsetneq \text{CLIQUE}(f(n))$ if $g(n) = o(f(n))$. This result does not have the $\log n$ gap because the congested clique model allows $\log n$ bits to transmit between nodes at each iteration, while fixed-precision NLM allows only a constant number of bits. The reason why the result on congested clique can not be applied to fixed-precision NLMs is that congested clique assumes unbounded precision representation for each individual node.

However, Conjecture 3.3 is not true for NLMs with unbounded precision, because there is an upper bound depth $O(n^{B-1})$ for a model’s expressiveness power.[‡] That is, an unbounded-precision NLM can not achieve stronger expressiveness by increasing its depth beyond $O(n^{B-1})$.

It is important to point out that, to realize a specific graph reasoning function, NLMs with different maximum arity B may require different depth D . Fürer [13] provides a general construction for problems that higher-dimensional NLMs can solve in asymptotically smaller depth than lower-dimensional NLMs. In the following we give a concrete example for computing *S-T Connectivity-k*, which asks whether there is a path of nodes from S and T in a graph, with length $\leq k$.

Theorem 3.4 (S-T Connectivity- k with Different Max Arity). For any function $f(k)$, if $f(k) = o(k)$, $\text{NLM}[O(f(k)), 2]$ cannot realize S-T Connectivity- k . That is, S-T Connectivity- k requires depth at least $O(k)$ for a relational neural network with an maximum arity of $B = 2$. However, S-T Connectivity- k can be realized by $\text{NLM}[O(\log k), 3]$.

[‡]See Appendix B.2 for a formal statement and the proof.

Proof sketch. For any integer k , we can construct a graph with two chains of length k , so that if we mark two of the four ends as S or T , any NLM[$k - 1, 2$] cannot tell whether S and T are on the same chain. The full proof is described in Appendix B.3.

There are many important graph reasoning tasks that do not have known depth lower bounds, including all-pair connectivity and shortest distance [18, 27]. In Appendix B.3, we discuss the concrete complexity bounds for a series of graph reasoning problems.

Chapter 4

Learning and Generalization in Relational Neural Networks

Given our understanding of what functions can be realized by NLMs, we move on to the problems of learning them: Can we effectively learn a NLMs to solve a desired task given a sufficient number of input-output examples? In this paper, we show that applying *enumerative training* with examples up to some fixed graph size can ensure that the trained neural network will generalize to all graphs *larger* than those appearing in the training set.

A critical determinant of the generalization ability for NLMs is the aggregation function. Specifically, Xu et al. [32] have shown that using *sum* as the aggregation function provides maximum expressiveness for graph neural networks. However, sum aggregation cannot be implemented in fixed-precision models, because as the graph size n increases, the range of the sum aggregation also increases.

Definition 4.1 (Fixed-precision aggregation function). An aggregation function is *fixed precision* if it maps from any finite *set* of inputs with values drawn from *finite domains* to a *fixed finite* set of possible output values; that is, the cardinality of the range of the function cannot grow with the number of elements in the input set.

Two useful fixed-precision aggregation functions are *max*, which computes the dimension-wise maximum over the set of input values, and *fixed-precision mean*, which

approximates the dimension-wise mean to a fixed decimal place.

Definition 4.2 (Fixed-precision NLM). A NLM is *fixed precision* if its aggregation function is fixed precision and all of the component MLPs, NN_j , are also fixed precision, in the same sense as for aggregation functions above. This guarantees that the cardinality of the range of the hypergraph representation functions cannot grow with the number of nodes in the hypergraph.

4.1 Structural Generalization under Complete Presentation

In order to focus on structural generalization in this section, we consider an *enumerative* training paradigm. When the input hypergraph representation domain \mathcal{S} is a finite set, we can enumerate the set $\mathcal{G}_{\leq N}$ of all possible input hypergraph representations of size bounded by N . We first enumerate all graph sizes $n \leq N$; for each n , we enumerate all possible values assigned to the hyperedges in the input. Given training size N , we enumerate all inputs in $\mathcal{G}_{\leq N}$, associate with each one the corresponding ground-truth output representation, and train the model with these input-output pairs.

This has much stronger data requirements than the standard sampling-based training mechanisms in machine learning. In practice, this can be approximated well when the input domain \mathcal{S} is small and the input data distribution is approximately uniformly distributed. The enumerative learning setting is studied by the *language identification in the limit* community [15], in which it is called *complete presentation*. This is an interesting learning setting because even if the domain for each individual hyperedge representation is finite, as the graph size can go arbitrarily large, the number of possible inputs is enumerable but unbounded.

Theorem 4.1 (Fixed-precision generalization under complete presentation). For any hypergraph reasoning function f , if it can be realized by a fixed-precision relational neural network model \mathcal{M} , then there exists an integer N , such that if we train the model with complete presentation on all input hypergraph representations with size

smaller than N , $\mathcal{G}_{\leq N}$, then for all $M \in \mathcal{M}$,

$$\sum_{G \in \mathcal{G}_{\leq N}} 1[M(G) \neq f(G)] = 0 \implies \forall G \in \mathcal{G}_{\infty} : M(G) = f(G).$$

That is, as long as M fits all training examples, it will generalize to all possible hypergraphs in \mathcal{G}_{∞} .

Proof. The key observation is that for any fixed vector representation length W , there are only a finite number of distinctive models in a fixed-precision NLM family, *independent of the graph size n* . Let W_b be the number of bits in each intermediate representation of a fixed-precision NLM. There are at most $(2^{W_b})^{2^{W_b}}$ different mappings from inputs to outputs. Hence, if N is sufficiently large to enumerate all input hypergraphs, we can always identify the correct model in the hypothesis space.

Our results are related to the *algorithmic alignment* approach [33, 34]. In contrast to their Probably Approximately Correct (PAC) Learning bounds for sample efficiency, our expressiveness results directly quantifies whether a hypergraph neural network can be trained to realize a specific function.

Chapter 5

Experimental Results

We now study how our theoretical results on model expressiveness and learning apply to relational neural networks trained with gradient descent on practically meaningful problems. We begin by describing two synthetic benchmarks: graph substructure detection and relational reasoning.

In the graph substructure detection dataset, there are several tasks of predicting whether there input graph contained a sub-graph with specific structure. The tasks are: *3-link* (length-3 path), *4-link*, *triangle*, and *4-clique*. These are important graph properties with many potential applications.

The relational reasoning dataset is composed of two family-relationship prediction tasks and two connectivity-prediction tasks. They are all *binary edge classification* tasks. In the family-relationship prediction task, the input contains the *mother* and *father* relationships, and the task is to predict the *grandparent* and *uncle* relationships between all pairs of entities. In the connectivity-prediction tasks, the input is the edges in an undirected graph and the task is to predict, for all pairs of nodes, whether they are connected with a path of length ≤ 4 (*connectivity-4*) and whether they are connected with a path of arbitrary length (*connectivity*).

5.1 Experiment Setup

For all problems, we have 800 training samples, 100 validation samples, and 300 test samples for each different n we are testing the models on.

We then provide the details on how we synthesize the data. For most of the problems, we generate the graph by randomly selecting from all potential edges i.e. the Erdős–Rényi model. We sample the number of edges around n , $2n$, $n \log n$ and $n^2/2$. For all problems, with 50% probability the graph will first be divided into 2, 3, 4 or 5 parts with equal number of components, where we use the first generated component to fill the edges for rest of the components. Some random edges are added afterwards. This make the data contain more isomorphic sub-graphs, which we found challenging empirically.

5.2 Substructure Detection

To generate a graph that does not contain a certain substructure, we randomly add edges when reaching a maximal graph not containing the substructure or reaching the edge limit. For generating a graph that does contain the substructure, we first generate one that does not contain, and then randomly replace present edges with missing edges until we detect the substructure in the graph. This aim to change the label from “No” to “Yes” while minimizing the change to the graph, and we found that data generated using edge replacing is much more difficult for neural networks compared to random generated graphs from scratch.

5.2.1 Relational Reasoning

Family tree. We generate the family trees using the algorithm modified from [11]. We add people to the family one by one. When a person is added, with probability p we will try to find a single woman and a single man, get them married and let the new children be their child, and otherwise the new person is introduced as a non-related person. Every new person is marked as single and set the gender with a coin flip.

We adjust p based on the ratio of single population: $p = 0.7$ when more than 40%

of the population are single, and $p = 0.3$ when less than 20% of the population are single, and $p = 0.5$ otherwise.

Connectivity. For connectivity problems, we use the similar generation method as the substructure detection. We sample the query pairs so that the labels are balanced.

5.2.2 Model Implementation Details

For all models, we use a hidden dimension 128 except for 3-dimensional HO-GNN and 4-dimensional NLM where we use hidden dimension 64.

All model have 4 layers that each has its own parameters, except for connectivity where we use the recurrent models that apply the second layer k times, where k is sampled from integers in $[2 \log n, 3 \log n]$. The depths are proven to be sufficient for solving these problems (unless the model itself can not solve).

All models are trained for 100 epochs using adam optimizer with learning rate 3×10^{-4} decaying at epoch 50 and 80.

We have varied the depth, the hidden dimension, and the activation function of different models. We select sufficient hidden dimension and depth for every model and problem (i.e., we stop when increasing depth or hidden dimension doesn't increase the accuracy). We tried linear, ReLU, and Sigmoid activation functions, and ReLU performed the best overall combinations of models and tasks.

5.3 Results and Analysis

Our main results on all datasets are shown in Table 5.1 and Table 5.2. We empirically compare relational neural networks with different maximum arity B , different model architecture (GNN and NLM), and different aggregation functions (max and sum). All models use sigmoidal activation for all MLPs. For each task on both datasets we train on a set of small graphs ($n = 10$) and test the trained model on both small graphs and large graphs ($n = 10$ and $n = 30$). We summarize the findings below.

Model	Agg.	3-link		4-link		triangle		4-clique	
		$n = 10$	$n = 30$	$n = 10$	$n = 30$	$n = 10$	$n = 30$	$n = 10$	$n = 30$
1-ary GNN	Max	70.0 \pm 0.0	82.7 \pm 0.0	92.0 \pm 0.0	91.7 \pm 0.0	73.7 \pm 3.2	50.2 \pm 1.8	55.3 \pm 4.0	46.2 \pm 1.3
	Sum	100.0 \pm 0.0	89.4 \pm 0.4	100.0 \pm 0.0	86.1 \pm 1.2	77.7 \pm 8.5	48.6 \pm 1.6	53.7 \pm 0.6	55.2 \pm 0.8
2-ary NLM	Max	65.3 \pm 0.6	54.0 \pm 0.6	93.0 \pm 0.0	95.7 \pm 0.0	51.0 \pm 1.7	49.2 \pm 0.4	55.0 \pm 0.0	45.7 \pm 0.0
	Sum	100.0 \pm 0.0	88.3 \pm 0.0	100.0 \pm 0.0	67.4 \pm 16.4	82.0 \pm 2.6	48.3 \pm 0.0	53.0 \pm 0.0	54.4 \pm 1.5
2-ary GNN	Max	78.7 \pm 0.6	76.0 \pm 17.3	97.7 \pm 4.0	98.6 \pm 2.5	100.0 \pm 0.0	100.0 \pm 0.0	55.0 \pm 0.0	45.7 \pm 0.0
	Sum	100.0 \pm 0.0	51.2 \pm 7.9	100.0 \pm 0.0	45.7 \pm 7.6	100.0 \pm 0.0	49.2 \pm 1.0	61.0 \pm 5.6	54.3 \pm 0.0
3-ary NLM	Max	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	59.0 \pm 6.9	45.9 \pm 0.4
	Sum	100.0 \pm 0.0	87.6 \pm 11.0	100.0 \pm 0.0	65.4 \pm 14.3	100.0 \pm 0.0	80.6 \pm 8.8	73.7 \pm 13.8	53.3 \pm 8.8
3-ary GNN	Max	79.0 \pm 0.0	86.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	84.0 \pm 0.0	93.3 \pm 0.0
	Sum	100.0 \pm 0.0	84.1 \pm 18.6	100.0 \pm 0.0	61.1 \pm 15.0	100.0 \pm 0.0	95.1 \pm 7.3	80.5 \pm 0.7	66.2 \pm 19.6
4-ary NLM	Max	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	82.0 \pm 1.7	93.1 \pm 0.2
	Sum	100.0 \pm 0.0	59.1 \pm 5.3	100.0 \pm 0.0	67.7 \pm 24.1	100.0 \pm 0.0	82.1 \pm 12.8	84.0 \pm 0.0	67.0 \pm 18.9

Table 5.1: Overall accuracy on structure detecting problems. All models are trained on $n = 10$, and tested on $n = 30$. The standard error of all values are computed based on three random seeds.

Model	Agg.	grand parent		uncle		connectivity-4*		connectivity	
		$n = 20$	$n = 80$	$n = 20$	$n = 80$	$n = 10$	$n = 80$	$n = 10$	$n = 80$
1-ary GNN	Max	84.0 \pm 0.3	64.8 \pm 0.0	93.6 \pm 0.3	66.1 \pm 0.0	72.6 \pm 3.6	67.5 \pm 0.5	85.6 \pm 0.3	75.1 \pm 1.9
	Sum	84.7 \pm 0.1	64.4 \pm 0.0	94.3 \pm 0.2	66.2 \pm 0.0	79.6 \pm 0.1	68.3 \pm 0.1	87.1 \pm 0.3	75.0 \pm 0.2
2-ary NLM	Max	82.3 \pm 0.5	65.6 \pm 0.1	93.1 \pm 0.0	66.6 \pm 0.0	91.2 \pm 0.2	51.0 \pm 0.6	88.9 \pm 2.6	67.1 \pm 4.8
	Sum	82.9 \pm 0.1	64.6 \pm 0.1	93.4 \pm 0.0	66.7 \pm 0.2	96.0 \pm 0.4	68.3 \pm 0.5	84.0 \pm 0.0	71.9 \pm 0.0
2-ary GNN	Max	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	84.0 \pm 0.0	71.9 \pm 0.0
	Sum	100.0 \pm 0.0	35.7 \pm 0.0	100.0 \pm 0.0	33.9 \pm 0.0	100.0 \pm 0.0	51.3 \pm 5.3	84.0 \pm 0.0	71.9 \pm 0.0
3-ary NLM	Max	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.1
	Sum	100.0 \pm 0.0	35.7 \pm 0.0	100.0 \pm 0.0	50.8 \pm 29.4	100.0 \pm 0.0	77.8 \pm 11.8	100.0 \pm 0.0	88.2 \pm 8.0
3-ary NLM _{HE}	Max	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	N/A	N/A	N/A	N/A
	Sum	100.0 \pm 0.0	35.7 \pm 0.0	100.0 \pm 0.0	33.8 \pm 29.4	N/A	N/A	N/A	N/A

Table 5.2: Overall accuracy on relational reasoning problems. Models for family-relationship prediction are trained on $n = 20$, while models for connectivity problems are trained on $n = 10$. All model are tested on $n = 80$. The standard error of all values are computed based on three random seeds. The 3-ary NLMs marked with “HE” have hyperedges in inputs, where each family is represented by a 3-ary hyperedge instead of two parent-child edges, and the results are similar to binary edges.

Expressiveness. We have seen a theoretical equal expressiveness between GNNs and NLMs applied to hypergraphs. That is, a GNN applied to B -ary hyperedges is equivalent to a $(B + 1)$ -ary NLM. Table 5.1 and Table 5.2 further suggest their similar performance on tasks when trained with gradient descent.

Formally, triangle detection requires NLMs with at least $B = 3$ to solve. Thus, we see that all NLMs with arity $B = 2$ fail on this task, but models with $B = 3$ perform

well. Formally, 4-clique is realizable by NLMs with maximum arity $B = 4$, but we failed to reliably train models to reach perfect accuracy on this problem. It is not yet clear what the cause of this behavior is.




Figure 5-1: How the performance of models drop when generalizing to larger graphs on the problem connectivity-4 (trained on graphs with size 10).

Structural generalization. We discussed the structural generalization properties of NLMs in Chapter 4, in a learning setting based on fixed-precision networks and enumerative training. This setting can be *approximated* by training NLMs with max aggregation and sigmoidal activation on sufficient data.

We run a case study on the problem connectivity-4 about how the generalization performance changes when the test graph size gradually becomes larger. Fig. 5-1 show how these models generalize to gradually larger graphs with size increasing from 10 to 80. From the curves we can see that only models with sufficient expressiveness can get 100% accuracy on the same size graphs, and among them the models using max aggregation generalize to larger graphs with no performance drop. 2-ary GNN and 3-ary NLM that use max aggregation have sufficient expressiveness and better generalization property. They achieve 100% accuracy on the original graph size and generalize perfectly to larger graphs.

Chapter 6

Related Work

Solving problems on graphs of arbitrary size is studied in many fields. NLMs can be viewed as circuit families with constrained architecture. In distributed computation, the congested clique model can be viewed as 2-arity NLMs, where nodes have identities as extra information. Common graph problems including sub-structure detection[22, 28] and connectivity[18] are studied for lower bounds in terms of depth, width and communication. This has been connected to GNNs for deriving expressiveness bounds [24].

Studies have been conducted on the expressiveness of GNNs and their variants. Xu et al. [32] provide an illuminating characterization of GNN expressiveness in terms of the WL graph isomorphism test. Azizian and Lelarge [1] analyze the expressiveness of higher-order Folklore GNNs by connecting them with high-dimensional WL-tests. We have the similar results in the arity hierarchy. Barceló et al. [3] reviewed GNNs from the logical perspective and rigorously refined their logical expressiveness with respect to fragments of first-order logic. Dong et al. [11] proposed Neural Logical Machines (NLMs) to reason about higher-order relations, and showed that increasing order increases expressiveness. It is also possible to gain expressiveness using unbounded computation time, as shown by the work of Dehghani et al. [8] on dynamic halting in transformers.

It is interesting that GNNs may generalize to larger graphs. Xu et al. [33, 34] have studied the notion of *algorithmic alignment* to quantify such structural generalization.

Dong et al. [11] provided empirical results showing that NLMs generalize to much larger graphs on certain tasks. Buffelli et al. [6] introduced a regularization technique to improve GNNs' generalization to larger graphs and demonstrated its effectiveness empirically. In Xu et al. [33], they analyzed and compared the sample complexity of Graph Neural Networks. This is different from our notion of expressiveness for realizing functions. In Xu et al. [34], they showed empirically on some problems (e.g., Max-Degree, Shortest Path, and n-body problem) that algorithm alignment helps GNNs to extrapolate, and theoretically proved the improvement by algorithm alignment on the Max-Degree problem. In this extended abstract, instead of focusing on computing specific graph problems, we analyzed how GNNs can extrapolate to larger graphs in a general case, based on the assumption of fixed precision computation.

Chapter 7

Conclusion

In this thesis, we have shown that Neural Logical Machines [11] as a framework that unifies several powerful classes of hypergraph neural networks in terms of expressiveness, NLMs also characterize aspects of the representational and generalization power of hypergraph neural networks, both theoretically and empirically.

In particular, we have shown the substantial increase of expressive power due to higher-arity relations and increasing depth, and have characterized very powerful structural generalization from training on small graphs to performance on larger ones. We further discuss the relationship between these results and existing results in Chapter 6. All theoretical results are further supported by the empirical results, discussed in Chapter 5. Although many questions remain open about the overall generalization capacity of these models in continuous and noisy domains, we believe this work has shed some light on their utility and potential for application in a variety of problems. Specifically, this work suggest the future directions:

First, the relation between the expressiveness and depth of the neural network remains unclear. Although it can be implied that more layers bring more expressiveness to constant depth NLMs, we still lack theoretical support for the necessity of additional layers when the depth is a function of the input size or the termination is determined by the neural network itself.

Second, we have proved the existence of the upper bound of the training graph sizes under enumerating training. However, there is a large gap between the theoretical

bound and the actual graph size needed in our experiments (i.e., some model achieved 100% accuracy on larger graphs that are not seen during training, which indicates that the models learned to solve the graph problem precisely rather than approximating it.) It is important to study the sample complexity for learning graph problems. In particular, among the problems that require the same expressive power, which of them requires a larger complexity? Does easier graph problems (e.g. simulating a simple formula in first-order logic) indeed easier to learn?

Third, our experiment results on different aggregation functions have shown that the aggregation function can effectively affect the learning and generalization of hypergraph neural networks. It would be interesting to examine and analyze common aggregation functions on their learning performance. On what type of problems does one aggregation function learn or generalize better than others? The algorithm alignment framework [33, 34] studied common aggregation functions on a variety of graph problems, and we believe more theoretical and empirical results can be produced in this direction that will help designing large graph neural networks in practice. Also, there are aggregation functions that are not commonly used, such as XOR and SoftMax, and exploring their performance on graph problems may bring us better understanding of the learning process.

Appendix A

Hypergraph Neural Networks

A.1 Expressiveness Equivalence of Relational Neural Networks

Since we are going to study both constant-depth and adaptive-depth graph neural networks, we first prove the following lemma (for general multi-layer neural networks), which helps us simplify the analysis.

Lemma A.1. A neural network with representation width W that has D different layers $\text{NN}_1, \dots, \text{NN}_D$ can be realized by a neural network that applies a single layer NN' for D iterations with width $(D + 1)(W + 1)$.

Proof. The representation for NN' can be partitioned into $D + 1$ segments each of length $W + 1$. Each segment consist of a “flag” element and a W -element representation, which are all 0 initially, except for the first segment, where the flag is set to 1, and the representation is the input.

NN' has the weights for all $\text{NN}_1, \dots, \text{NN}_D$, where weights NN_i are used to compute the representation in segment $i + 1$ from the representation in segment i . Additionally, at each iteration, segment $i + 1$ can only be computed if the flag in segment i is 1, in which case the flag of segment $i + 1$ is set to 1. Clearly, after D iterations, the output of NN'_k should be the representation in segment $D + 1$. \square

Due to Lemma A.1, we consider the neural networks that recurrently apply the same layer because a) they are as expressive as those using layers of different weights, b) it is easier to analyze a single neural network layer than D layers, and c) they naturally generalize to neural networks that runs for adaptive number of iterations (e.g. GNNs that run $O(\log n)$ iterations where n is the size of the input graph).

Equivalence between HO-GNNs and NLMs. We will prove the equivalence between HO-GNNs and NLMs by making reductions in both directions.

Lemma A.2. A B -ary HO-GNN with depth D can be realized by a NLM with maximum arity $B + 1$ and depth $2D$.

Proof. We prove Lemma A.2 by showing that one layer of GNNs on B -ary hypergraphs can be realized by two NLM with maximum arity $B + 1$.

Firstly, a GNN layer maintain features of B -tuples, which are stored in correspondingly in an NLM layer at dimension B . Then we will realize the message passing scheme using the NLM features of dimension B and $B + 1$ in two steps.

Recall the message passing scheme generalized to high dimensions (to distinguish, we use H for HO-GNN features and T for NLM features.)

$$\text{Received}_i(\mathbf{v}) = \sum_u (\text{NN}_1 (H_{i-1,B}[\mathbf{v}]; \text{CONCAT}_{\mathbf{v}' \in \text{neighbors}(\mathbf{v},u)} H_{i-1}[\mathbf{v}'])) \quad (\text{A.1.1})$$

At the first step, the Expand operation first raise the dimension to $B + 1$ by expanding a non-related variable u to the end, and the Permute operation can then swap u with each of the elements (or no swap). Particularly, $T_{i,B}[v_1, v_2, \dots, v_B]$ will be expand to

$$T_{i+1,B+1}[u, v_2, v_3, \dots, v_B, v_1], T_{i+1,B+1}[v_1, u, v_3, \dots, v_B, v_2], \dots, \\ T_{i+1,B+1}[v_1, v_2, \dots, v_{B-1}, u, v_B], \text{and } T_{i+1,B+1}[v_1, v_2, \dots, v_{B-1}, v_B, u]$$

Hence, $T_{i+1,B+1}[v_1, v_2, v_3, \dots, v_B, u]$ receives the features from

$$T_{i,B}[v_1, v_2, \dots, v_B], T_{i,B}[u, v_2, v_3, \dots, v_B], T_{i,B}[v_1, u, v_3, \dots, v_B], \dots, T_{i,B}[v_1, v_2, \dots, v_{B-1}, u]$$

These features matches the input of NN_1 in Equation A.1.1, and in this layer NN_1 can be applied to compute things inside the summation.

Then at the second step, the last element is reduced to get what tuple \mathbf{v} should receive, so \mathbf{v} can be updated. Since each HO-GNN layer can be realized by such two NLM layers, each B -ary HO-GNN with depth D can be realized by a NLM of maximum arity $(B + 1)$ and depth $2D$. \square

To complete the proof we need to find a reduction from NLMs of maximum arity $B + 1$ to B -ary HO-GNNs. The key observation here is that the features of $(B + 1)$ -tuples in NLMs can only be expanded from sub-tuples, and the expansion and reduction involving $(B + 1)$ -tuples can be simulated by the message passing process.

Lemma A.3. The features of $(B + 1)$ -tuples feature $T_{i,B+1}[v_1, v_2, \dots, v_{B+1}]$ can be computed from the following tuples

$$(T_{i,B}[v_2, v_3, \dots, v_{B+1}], T_{i,B}[v_1, v_3, \dots, v_{B+1}], \dots, T_{i,B}[v_1, v_2, \dots, v_B]).$$

Proof. Lemma A.3 is true because $(B + 1)$ -dimensional representations can either be computed from themselves at the previous iteration, or expanded from B -dimensional representations. Since representations at all previous iterations $j < i$ can be contained in $T_{i,B}$, it is sufficient to compute $T_{i,B+1}[v_1, v_2, \dots, v_{B+1}]$ from all its B -ary sub-tuples. \square

Then let's construct the HO-GNN for given NLM to show the existence of the reduction.

Lemma A.4. A NLM of maximum arity $B + 1$ and depth D can be realized by a B -ary HO-GNN with no more than D iterations.

Proof. We can realize the Expand and Reduce operation with only the B -dimensional features using the broadcast message passing scheme. Note that Expand and Reduce between B -dimensional features and $(B + 1)$ -dimensional features in the NLM is a special case where Lemma A.3 is applied.

Let's start with Expand and Reduce operations between features of dimension B or lower. For the b -dimensional feature in the NLM, we keep $n^b n^{B-b*}$ copies of it and store them the representation of every B -tuple who has a sub-tuple[†] that is a permutation of the b -tuple. That is, for each B -tuple in the B -ary HO-GNN, for its every sub-tuple of length b , we store $b!$ representations corresponding to every permutation of the b -tuple in the NLM. Keeping representation for all sub-tuple permutations make it possible to realize the Permute operation. Also, it is easy to notice that Expand operation is realized already, as all features with dimension lower than B are naturally expanded to B dimension by filling in all possible combinations of the rest elements. Finally, the Reduce operation can be realized using a broadcast casting message passing on certain position of the tuple.

Now let's move to the special case – the Expand and Reduce operation between features of dimensions B and $B + 1$. Lemma A.3 suggests how the $(B + 1)$ -dimensional features are stored in B -dimensional representations in GNNs, and we now show how the Reduce can be realized by message passing.

We first bring in Lemma A.3 to the HO-GNN message passing, where we have $\text{Received}_i[\mathbf{v}]$ to be

$$\sum_u \left(\text{NN}_1 \left(T_{i-1,B}[v_2, v_3, \dots, v_B, u], T_{i-1,B}[v_1, v_3, \dots, v_B, u], \dots, T_{(i-1),B}[v_1, v_2, \dots, v_B] \right) \right)$$

Note that the last term $T_{i-1,B}[v_1, v_2, \dots, v_B]$ is contained in $H_{i-1}(v)$ in Equation A.1.1, and other terms are contained in $H_{i-1}(v')$ for $v' \in \text{neighbors}(\mathbf{v}, u)$. Hence, Equation A.1.1 is sufficient to simulate the Reduce operation. \square

* $n^k = n \times (n - 1) \times \dots \times (n - k + 1)$.

†The sub-tuple does not have to be consecutive, but instead can be a any subset of the tuple that keeps the element order.

Theorem A.5. B -ary HO-GNNs are equally expressive as NLMs with maximum arity $B + 1$.

Proof. This is a direct conclusion by combining Lemma A.2 and Lemma A.4. \square

A.2 Expressiveness of hypergraph convolution and attention

There exist other variants of hypergraph neural networks. In particular, hypergraph convolution[2, 12, 35], attention[9] and message passing[17] focus on updating node features instead of tuple features through hyperedges. These approaches can be viewed as instances of hypergraph neural networks, and they have smaller time complexity because they do not model all high-arity tuples. However, they are less expressive than the standard hypergraph neural networks with equal max arity.

These approaches can be formulated to two steps at each iteration. At the first step, each hyperedge is updated by the features of nodes it connects.

$$h_{i,e} = \text{AGG}_{v \in e} f_{i-1,v} \tag{A.2.1}$$

At the second step, each node is updated by the features of hyperedges connecting it.

$$f_{i,v} = \text{AGG}_{v \in e} h_{i,e} \tag{A.2.2}$$

where $f_{i,v}$ is the feature of node v at iteration i , and $h_{i,v}$ is the aggregated message passing through hyperedge e at iteration $i + 1$.

It is not hard to see that Equation A.2.1 can be realized by B iterations of NLM layers with Expand operations where B is the max arity of hyperedges. This can be done by expanding each node feature to every high arity features that contain the node, and aggregate them at the tuple corresponding to each hyperedge. Then, Equation A.2.2 can also be realized by B iterations of NLM layers with Reduce

operations, as the tuple feature will finally be reduced to a single node contained in the tuple.

This approach has lower complexity compared to the GNNs we study applied on hyperedges, because it only requires communication between nodes and hyperedges connecting to them, which takes $O(|V| \cdot |E|)$ time at each iteration. Compared to them, NLMs takes $O(|V|^B)$ time because NLMs keep features of every tuple with max arity B , and allow communication from tuples to tuples instead of between tuples and single nodes. An example is provided below that this approach can not solve while NLMs can.

Consider a graph with 6 nodes and 6 edges forming two triangles $(1, 2, 3)$ and $(4, 5, 6)$. Because of the symmetry, the representation of each node should be identical throughout hypergraph message passing rounds. Hence, it is impossible for these models to conclude that $(1, 2, 3)$ is a triangle but $(4, 2, 3)$ is not, based only on the node representations, because they are identical. In contrast, NLMs with max arity 3 can solve them (as standard triangle detection problem in Table 3.1).

A.3 The Time and Space Complexity of NLMs

Handling high-arity features and using deeper models usually increase the computational cost in terms of time and space. As an instance that use the architecture of RelNN, NLMs with depth D and max arity B takes $O(Dn^B)$ time when applying to graphs with size n . This is because both Expand and Reduce operation have linear time complexity with respect to the input size (which is $O(n^B)$ at each iteration). If we need to record the computational history (which is typically the case when training the network using back propagation), the space complexity is the same as the time complexity.

GNNs applied to $(B - 1)$ -ary hyperedges and depth D are equally expressive as RelNNs with depth $O(D)$ and max arity B . Though up to $(B - 1)$ -ary features are kept in their architecture, the broadcast message passing scheme scale up the complexity by a factor of $O(n)$, so they also have time and space complexity $O(Dn^B)$.

Here the length of feature tensors W is treated as a constant.

Appendix B

Arity and Depth Hierarchy: Proofs and Analysis

B.1 Proof of Theorem 3.1: Arity Hierarchy.

[26] have connected high-dimensional GNNs with high-dimensional WL tests. Specifically, they showed that the B -ary HO-GNNs are equally expressive as B -dimensional WL test on graph isomorphism test problem. In Theorem A.5 we proved that B -ary HO-GNNs are equivalent to NLM of maximum arity $B + 1$ in terms of expressiveness. Hence, NLM of maximum arity $B + 1$ can distinguish if two non-isomorphic graphs if and only if B -dimensional WL test can distinguish them.

However, Cai et al. [7] provided an construction that can generate a pair of non-isomorphic graphs for every B , which can not be distinguished by $(B - 1)$ -dimensional WL test but can be distinguished by B -dimensional WL test. Let G_B^1 and G_B^2 be such a pair of graph.

Since NLM of maximum arity $B + 1$ is equally expressive as B -ary HO-GNNs, there must be such a NLM that classify G_B^1 and G_B^2 into different label. However, such NLM can not be realized by any NLM of maximum arity B because they are proven to have identical outputs on G_B^1 and G_B^2 .

In the other direction, NLMs of maximum arity $B + 1$ can directly realize NLMs of maximum arity B , which completes the proof.

B.2 Upper Depth Bound for Unbounded-Precision NLM.

The idea for proving an upper bound on depth is to connect NLMs to WL-test, and use the $O(n^B)$ upper bound on number of iterations for B -dimensional test [19], and FOC formula is the key connection.

For any fixed n , B -dimensional WL test divide all graphs of size n , $\mathcal{G}_{=n}$, into a set of equivalence classes $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_m\}$, where two graphs belong to the same class if they can not be distinguished by the WL test. We have shown that NLMs of maximum arity $(B + 1)$ must have the same input for all graphs in the same equivalence class. Thus, any NLM of maximum arity $B + 1$ can be view as a labeling over $\mathcal{C}_1, \dots, \mathcal{C}_m$.

Stated by Cai et al. [7], B -dimensional WL test are as powerful as FOC_{B+1} in differentiating graphs graphs. Combined with the $O(n^B)$ upper bound of WL test iterations, for each \mathcal{C}_i , there must be an FOC_{B+1} formula of quantifier depth $O(n^B)$ that exactly recognize \mathcal{C}_i over $\mathcal{G}_{=n}$.

Finally, with unbounded precision, for any $f(n)$, NLM of maximum arity $B + 1$ and depth $f(n)$ can compute all FOC_{B+1} formulas with quantifier depth $f(n)$. Note that there are finite number of such formula because the superscript of counting quantifiers is bounded by n .

For any graph in some class \mathcal{C}_i , the class can be determined by evaluating these FOC formulas, and then the label is determined. Therefore, any NLM of maximum arity $B + 1$ can be realized by a NLM of maximum arity $B + 1$ and depth $O(n^B)$.

B.3 Graph Problems

We list a number of examples for graph classification and regression tasks, and we provide the definitions and the current best known NLMs for learning these tasks from data. For some of the problems, we will also show why they can not be solved by a simpler problems, or indicate them as open problems.

Node Color Majority. Each node is assigned a color $c \in \mathcal{C}$ where \mathcal{C} is a finite set of all colors. The model needs to predict which color the most nodes have.

Using a single layer with sum aggregation, the model can count the number of nodes of color c for each $c \in \mathcal{C}$ on its global representation.

Count Red Nodes. Each node is assigned a color of red or blue. The model needs to count the number of red nodes.

Similarly, using a single layer with sum aggregation, the model can count the number of red nodes on its global representation.

3-Link Detection. Given an unweighted, undirected graph, the model needs to detect whether there is a triple of nodes (a, b, c) such that $a \neq c$ and (a, b) and (b, c) are edges.

This is equivalent to check whether there exists a node with degree at least 2. We can use a Reduction operation with sum aggregation to compute the degree for each node, and then use a Reduction operation with max aggregation to check whether the maximum degree of nodes is greater than or equal to 2.

Note that this can not be done with 1 layer, because the edge information is necessary for the problem, and they require at least 2 layers to be passed to the global representation.

4-Link Detection. Given an unweighted undirected graph, the model needs to detect whether there is a 4-tuple of nodes (a, b, c, d) such that $a \neq c, b \neq d$ and $(a, b), (b, c), (c, d)$ are edges (note that a triangle is also a 4-link).

This problem is equivalent to check whether there is an edge between two nodes with degrees ≥ 2 . We can first reduce the edge information to compute the degree for each node, and then expand it back to 2-dimensional representations, so we can check for each edge if the degrees of its ends are ≥ 2 . Then the results are reduced to the global representation with existential quantifier (realized by max aggregation) in 2 layers.

Triangle Detection. Given a unweighted undirected graph, the model is asked to determine whether there is a triangle in the graph i.e. a tuple (a, b, c) so that $(a, b), (b, c), (c, a)$ are all edges.

This problem can be solved by NLM [4,3]: we first expand the edge to 3-dimensional representations, and determine for each 3-tuple if they form a triangle. The results of 3-tuples require 3 layers to be passed to the global representation.

We can prove that Triangle Detection indeed requires breadth at least 3. Let k -regular graphs be graphs where each node has degree k . Consider two k -regular graphs both with n nodes, so that exactly one of them contains a triangle*. However, NLMs of breadth 2 has been proven not to be stronger than WL test on distinguish graphs, and thus can not distinguish these two graphs (WL test can not distinguish any two k -regular graphs with equal size).

4-Clique Detection and Counting. Given an undirected graph, check existence of, or count the number of tuples (a, b, c, d) so that there are edges between every pair of nodes in the tuple.

This problem can be easily solved by a NLM with breadth 4 that first expand the edge information to the 4-dimensional representations, and for each tuple determine whether its is a 4-clique. Then the information of all 4-tuples are reduced 4 times to the global representation (sum aggregation can be used for counting those).

Though we did not find explicit counter-example construction on detecting 4-cliques with NLMs of breadth 3, we suggest that this problem can not be solved with NLMs with 3 or lower breadth.

Connectivity. The connectivity problems are defined on unweighted undirected graphs. S-T connectivity problems provides two nodes S and T (labeled with specific colors), and the model needs to predict if they are connected by some edges. All pair connectivity problem require the model to answer for every pair of nodes. Connectivity- k problems have an additional requirement that the distance between the pair of nodes can not exceed k .

S-T connectivity- k can be solved by a NLM of breadth 2 with k iterations. Assume S is colored with color c , at every iteration, every node with color c will spread the color to its neighbors. Then, after k iterations, it is sufficient to check whether T has

*Such construction is common. One example is $k = 2, n = 6$, and the graph may consist of two separated triangles or one hexagon

the color c .

With NLMs of breadth 3, we can use $O(\log k)$ matrix multiplications to solve connectivity- k between every pair of nodes. Since the matrix multiplication can naturally be realized by NLMs of breadth 3 with two layers. All-pair connectivity problems can all be solved with $O(\log k)$ layers.

Theorem B.1 (S-T connectivity- k with NLM). S-T connectivity- k can not be solved by a NLM of maximum arity within $o(k)$ iterations.

Proof. We construct two graphs each has $2k$ nodes $u_1, \dots, u_k, v_1, \dots, v_k$. In both graph, there are edges (u_i, u_{i+1}) and (v_i, v_{i+1}) for $1 \leq i \leq k-1$ i.e. there are two links of length k . We then set $S = u_1, T = u_n$ and $S = u_1, T = v_n$ the the two graphs.

We will analysis GNNs as NLMs are proved to be equivalent to them by scaling the depth by a constant factor. Now consider the node refinement process where each node x is refined by the multiset of labels of x 's neighbors and the multiset of labels of x 's non-neighbors.

Let $C_j^{(i)}(x)$ be the label of x in graph j after i iterations, at the beginning, WLOG, we have

$$C_1^{(0)}(u_1) = 1, C_1^{(0)}(u_n) = 2, C_1^{(0)}(u_1) = 1, C_1^{(0)}(v_n) = 2$$

and all other nodes are labeled as 0.

Then we can prove by induction: after $i \leq \frac{k}{2} - 1$ iterations, for $1 \leq t \leq i + 1$ we have

$$C_1^{(u_t)} = C_2^{(i)}(u_t), C_1^{(v_t)} = C_2^{(i)}(v_t)$$

$$C_1^{(u_{k-t+1})} = C_2^{(i)}(v_{k-t+1}), C_1^{(v_{k-t+1})} = C_2^{(i)}(u_{k-t+1})$$

and for $i + 2 \leq t \leq k - i - 1$ we have

$$C_1^{(u_t)} = C_2^{(i)}(u_t), C_1^{(v_t)} = C_2^{(i)}(v_t)$$

This is true because before $\frac{k}{2}$ iterations are run, the multiset of all node labels

are identical for the two graphs (say $S^{(i)}$). Hence each node x is actually refined by its neighbors and $S^{(i)}$ where $S^{(i)}$ is the same for all nodes. Hence, before running $\frac{k}{2}$ iterations when the message between S and T finally meets in the first graph, GNN can not distinguish the two graphs, and thus can not solve the connectivity with distance $k - 1$. \square

Max Degree. The max degree problem gives a graph and ask the model to output the maximum degree of its nodes.

Like we mentioned in 3-link detection, one layer for computing the degree for each node, and another layer for taking the max operation over nodes should be sufficient.

Max Flow. The Max Flow problem gives a directional graph with capacities on edges, and indicate two nodes S and T . The models is then asked to compute the amount of max-flow from S to T .

Notice that the Breadth First Search (BFS) component in Dinic’s algorithm[10] can be implemented on NLMs as they does not require node identities (all new-visited nodes can augment to their non-visited neighbors in parallel). Since the BFS runs for $O(n)$ iteration, and the Dinic’s algorithm runs BFS $O(n^2)$ times, the max-flow can be solved by NLMs with in $O(n^3)$ iterations.

Distance. Given a graph with weighted edges, compute the length of the shortest between specified node pair (S-T Distance) or all node pairs (All-pair Distance).

Similar to Connectivity problems, but Distance problems now additionally record the minimum distance from S (for S-T) or between every node pairs (for All-pair), which can be updated using min operator (using Min-plus matrix multiplication for All-pair case).

Bibliography

- [1] Waiss Azizian and Marc Lelarge. Expressive power of invariant and equivariant graph neural networks. In *ICLR*, 2021. 14, 39
- [2] Song Bai, Feihu Zhang, and Philip HS Torr. Hypergraph convolution and hypergraph attention. *Pattern Recognition*, 110:107637, 2021. 47
- [3] Pablo Barceló, Egor V Kostylev, Mikael Monet, Jorge Pérez, Juan Reutter, and Juan Pablo Silva. The logical expressiveness of graph neural networks. In *ICLR*, 2020. 24, 25, 39
- [4] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. 13
- [5] Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Lio, Guido F Montufar, and Michael Bronstein. Weisfeiler and lehman go cellular: Cw networks. In *NeurIPS*, 2021. 14
- [6] Davide Buffelli, Pietro Liò, and Fabio Vandin. Szeshtreg: a regularization method for improving size-generalization in graph neural networks. In *NeurIPS*, 2022. 40
- [7] 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. 25, 51, 52
- [8] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. In *ICLR*, 2019. 25, 39
- [9] Kaize Ding, Jianling Wang, Jundong Li, Dingcheng Li, and Huan Liu. Be more with less: Hypergraph attention networks for inductive text classification. In *EMNLP*, 2020. 47
- [10] Efim A Dinic. Algorithm for solution of a problem of maximum flow in networks with power estimation. In *Soviet Math. Doklady*, volume 11, pages 1277–1280, 1970. 56

- [11] Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. Neural logic machines. In *ICLR*, 2019. 9, 13, 17, 18, 19, 21, 34, 39, 40, 41
- [12] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *AAAI*, 2019. 47
- [13] Martin Fürer. Weisfeiler-lehman refinement requires at least a linear number of iterations. In *ICALP*, 2001. 26
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017. 13
- [15] E Mark Gold. Language identification in the limit. *Inf. Control.*, 10(5):447–474, 1967. 30
- [16] Juris Hartmanis and Richard E Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965. 25
- [17] Jing Huang and Jie Yang. Unignn: a unified framework for graph and hypergraph neural networks. In *IJCAI*, 2021. 47
- [18] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990. 27, 39
- [19] Sandra Kiefer and Pascal Schweitzer. Upper bounds on the quantifier depth for graph differentiation in first order logic. In *LICS*, 2016. 52
- [20] Janne H Korhonen and Jukka Suomela. Towards a complexity theory for the congested clique. In *SPAA*, 2018. 26
- [21] AA Leman and B Weisfeiler. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya*, 2(9):12–16, 1968. 24
- [22] Yuan Li, Alexander Razborov, and Benjamin Rossman. On the ac^0 complexity of subgraph isomorphism. *SIAM J. Comput.*, 46(3):936–971, 2017. 39
- [23] Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 2017. 13
- [24] Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020. 39
- [25] Christian Merkwirth and Thomas Lengauer. Automatic generation of complementary descriptors with molecular graph networks. *Journal of chemical information and modeling*, 45(5):1159–1168, 2005. 13

- [26] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *AAAI*, 2019. 16, 17, 21, 24, 51
- [27] Shreyas Pai and Sriram V Pemmaraju. Connectivity lower bounds in broadcast congested clique. In *PODC*, 2019. 27
- [28] Benjamin Rossman. *Average-case complexity of detecting cliques*. PhD thesis, Massachusetts Institute of Technology, 2010. 39
- [29] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008. 13
- [30] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *ESWC*, 2018. 13
- [31] Petar Veličković, Rex Ying, Matilde Padovano, Raia Hadsell, and Charles Blundell. Neural execution of graph algorithms. In *ICLR*, 2020. 13
- [32] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *ICLR*, 2019. 29, 39
- [33] Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S Du, Ken-ichi Kawarabayashi, and Stefanie Jegelka. What can neural networks reason about? In *ICLR*, 2020. 31, 39, 40, 42
- [34] Keyulu Xu, Mozhi Zhang, Jingling Li, Simon S Du, Ken-ichi Kawarabayashi, and Stefanie Jegelka. How neural networks extrapolate: From feedforward to graph neural networks. In *ICLR*, 2021. 31, 39, 40, 42
- [35] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. Hypergen: A new method of training graph convolutional networks on hypergraphs. In *NeurIPS*, 2019. 47