

Normalize Then Propagate: Efficient Homophilous Regularization for Few-shot Semi-Supervised Node Classification

Baoming Zhang¹, Mingcai Chen², Jianqing Song¹, Shuangjie Li¹, Jie Zhang¹, Chongjun Wang^{1*}

¹ State Key Laboratory of Novel Software Technology, Nanjing University, China

² Nanjing University of Posts and Telecommunications, China

zhangbm@smail.nju.edu.cn, chenmc@njupt.edu.cn, {sjq, shuangjieli}@smail.nju.edu.cn, guyuexiao95@gmail.com, chjwang@nju.edu.cn

Abstract

Graph Neural Networks (GNNs) have demonstrated remarkable ability in semi-supervised node classification. However, most existing GNNs rely heavily on a large amount of labeled data for training, which is labor-intensive and requires extensive domain knowledge. In this paper, we first analyze the restrictions of GNNs generalization from the perspective of supervision signals in the context of few-shot semi-supervised node classification. To address these challenges, we propose a novel algorithm named NormProp, which utilizes the homophily assumption of unlabeled nodes to generate additional supervision signals, thereby enhancing the generalization against label scarcity. The key idea is to efficiently capture both the class information and the consistency of aggregation during message passing, via decoupling the direction and Euclidean norm of node representations. Moreover, we conduct a theoretical analysis to determine the upper bound of Euclidean norm, and then propose homophilous regularization to constraint the consistency of unlabeled nodes. Extensive experiments demonstrate that NormProp achieve state-of-the-art performance under low-label rate scenarios with low computational complexity.

1 Introduction

Graphs are powerful data structures that allow us to easily express various relationships (i.e., edges) between objects (e.g., nodes). Graph Neural Networks (GNNs) (Defferrard, Bresson, and Vandergheynst 2016; Kipf and Welling 2017; Veličković et al. 2018) are a family of machine learning models specifically designed to model non-Euclidean graph data. Modern GNNs have emerged as the forefront approach for acquiring representations at various levels within graphs, encompassing node-level (Kipf and Welling 2017; Klicpera, Bojchevski, and Günnemann 2019), edge-level (Zhang and Chen 2018), and graph-level tasks (Ying et al. 2018; Han et al. 2022). A typical GNN model usually follows the message passing framework (Gilmer et al. 2017), which mainly contains two operators: feature transformation and feature propagation. In semi-supervised node classification settings (Zhu and Ghahramani 2002), only a subset of nodes have labels, and the parameters of the GNN are trained using the

classification loss function that relies on the representations of labeled nodes (Kipf and Welling 2017).

However, in numerous real-world scenarios of semi-supervised learning, obtaining a substantial amount of accurate labels for training can be challenging due to the labor-intensive data annotation (Li et al. 2019b). Furthermore, in the problem of semi-supervised node classification with limited labeled nodes (Ding et al. 2022), the scarcity of labeled nodes per class not only exacerbates that GNNs are prone to overfitting, but also magnifies additional challenge. Recent studies (Ma et al. 2022; Hu et al. 2020; Tang et al. 2020) have shown that GNNs exhibit certain biases originating from the perspective of graph topology. Ma, Deng, and Mei (2021) study the subgroup generalization of GNNs and find that the shortest path distance to labeled nodes can also affect the GNNs’ performance. Specifically, Han et al. (2024) uncover the label position bias, which indicates that *the node closer to the labeled nodes tends to perform better*. Considering the citation networks Cora and Citeseer (Planetoid split) (Yang, Cohen, and Salakhudinov 2016), the label signal is only relevant to 61.4% and 32.8% of the nodes, respectively, in the case of two-layer GNN. When only a few nodes are labeled, it exacerbates label position bias, as a larger number of nodes lack adequate label information during message passing.

To mitigate the problem of overfitting due to limited labels, a straightforward idea is to enlarge the receptive field of aggregation by incorporating additional representations from unlabeled high-order nodes into the calculation of the classification loss. Unfortunately, most of the existing GNNs have been designed with shallow architectures primarily due to the *over-smoothing issue* (Zhou et al. 2021). With the increasing of GNN layers, the Dirichlet energy converge to zero (Cai and Wang 2020). As a result, node representations become increasingly similar, blurring the distinction between different classes. Moreover, high-order neighbor aggregation further increases the complexity of message passing and may fail to acquire information from remote nodes on another connected component.

An example on a two-layer GNN (Kipf and Welling 2017) is shown in Figure 1. The receptive field of classification loss solely focuses on the labeled node and its two-hop neighbors (loss-relevant nodes). We refer to nodes that do not participate in loss calculation as *remote nodes*. Previous works

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

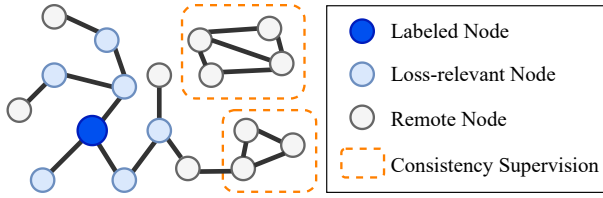


Figure 1: The source of the loss signal for a two-layer GNN. Suppose that there is only one labeled node in the graph, and that the graph consists of two connected components.

on limited labels in GNN mainly utilize remote nodes from three perspectives: high-order neighbors aggregation (Ding et al. 2022), pseudo labeling (Sun, Lin, and Zhu 2020) and graph structure learning (Xie, Tam, and Lau 2023). However, whether using pseudo-label iterative training or establishing connections between labeled nodes and remote nodes, there is a higher computational cost. In order to acquire more training signals for GNN with low complexity, we directly generate supervision signals from remote nodes instead of using message-passing framework to aggregate more information from unlabeled nodes. As shown in Figure 1, additional supervision signals are derived from the neighbor consistency of remote nodes, based on the classic homophily assumption (Nt and Maehara 2019).

To address the above challenges, we propose **NormProp**, a novel GNN model that enhances the model’s generalization ability by incorporating the consistency information from commonly ignored nodes. With the aim of efficiently obtaining a consistency metric synchronized with message passing, our key idea is to *decouple the direction and norm of node representation vectors*. The direction and Euclidean norm of node representation vector, extracted by NormProp, respectively denote the class information and neighbor consistency. Moreover, the core operation is “**Normalize then Propagate**”, which maps node features to the unit hypersphere and subsequently applies a low-pass filter (Wu et al. 2019). The classification loss relies on computing the cosine similarities between the labeled node representations and their corresponding class prototypes. Additionally, We conduct a theoretical analysis to determine the lower and upper bounds of Euclidean norm for representation vectors after propagation. Through the above theory, we propose **Homophilous Regularization** based on homophily assumption, which constrains the neighbor consistency of remote nodes. Training supervision involves both the classification loss incurred from labeled nodes and the homophilous regularization applied to the unlabeled nodes. Extensive experiments on node classification tasks with limited labeled nodes demonstrate that NormProp not only exhibits excellent performance but also requires low computational complexity.

In summary, we list our contributions as follows:

- We analyze the problem of few-shot semi-supervised node classification from the perspective of supervision signals, and consequently provide an efficient idea to enhance the generalization by leveraging the homophily of remote nodes.

- We propose a novel GNN, NormProp, where the direction and Euclidean norm of node representations learned by our model respectively indicate class information and consistency of aggregation. We conduct theoretical analysis on the upper bound of Euclidean norm, and propose homophilous regularization to constrain the neighbor consistency of unlabeled nodes.
- Extensive experimental results indicate the effectiveness and efficiency of NormProp, providing compelling validation for the advantages of NormProp over state-of-the-art methods.

2 Preliminaries and Related Works

2.1 Notations

We consider transductive semi-supervised unweighted and undirected graphs node classification. Given a graph $G = (V, E)$ with $n = |V|$ nodes and $m = |E|$ edges, where $V = \{v_1, v_2, \dots, v_n\}$ is the set of nodes and $E \subseteq V \times V$ is the set of edges between nodes in V . Let $X \in \mathbb{R}^{n \times d}$ denote the node feature matrix, with x_i presenting a d -dimensional feature vector of node v_i . The topology information of the entire graph is described by the adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, where $\mathbf{A}_{ij} = 1$ if $\langle v_i, v_j \rangle \in E$, otherwise $\mathbf{A}_{ij} = 0$. Let $\mathcal{N}(i)$ denote the neighboring nodes set of node v_i . The diagonal degree matrix is denoted as $\mathbf{D} \in \mathbb{R}^{n \times n}$, where $D_{ii} = |\mathcal{N}(i)|$ represents the degree of node v_i . Moreover, the adjacency matrix and degree matrix of graph with self-loops are respectively denoted as $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_n$ and $\tilde{\mathbf{D}} = \mathbf{D} + \mathbf{I}_n$. For few-shot semi-supervised node classification, only a few training nodes have labels.

2.2 Decoupled GCN

GCN (Kipf and Welling 2017) stacks layers of first-order spectral filters, with each message-passing layer followed by a learnable linear transform $\mathbf{W}^{(k)}$ and a nonlinear activation function $\sigma(\cdot)$. The k -th layer of GCN is defined as follow:

$$\mathcal{H}^{(k+1)} = \sigma \left(\hat{\mathbf{A}} \mathcal{H}^{(k)} \mathbf{W}^{(k)} \right), \quad (1)$$

where $\hat{\mathbf{A}} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$ is the symmetrically normalized adjacency matrix, $\mathcal{H}^{(k)}$ denotes node representations of the k -th GCN layer. Recent decoupled GCN separate the transformation and propagation in each layer, such as APPNP (Klicpera, Bojchevski, and Günnemann 2019) and GPR-GNN (Chien et al. 2020). Formally, the decoupled GCN can be written as:

$$\hat{Y} = \text{softmax} \left(\bar{\mathbf{A}} f_{\theta} (X) \right), \quad (2)$$

where $f_{\theta}(\cdot)$ is a feature encoder, $\bar{\mathbf{A}} = \sum \alpha_k \hat{\mathbf{A}}^k$ is the propagation strategy, which involves the k -order neighbors. Our theoretical analysis is based on the decoupled GCN form.

2.3 Related Works

As mentioned in the introduction, with only a few labeled nodes, the expressiveness of GNNs to utilize information in graph data is limited. M3S (Sun, Lin, and Zhu 2020) propose a multi-stage training framework, adding confident data with

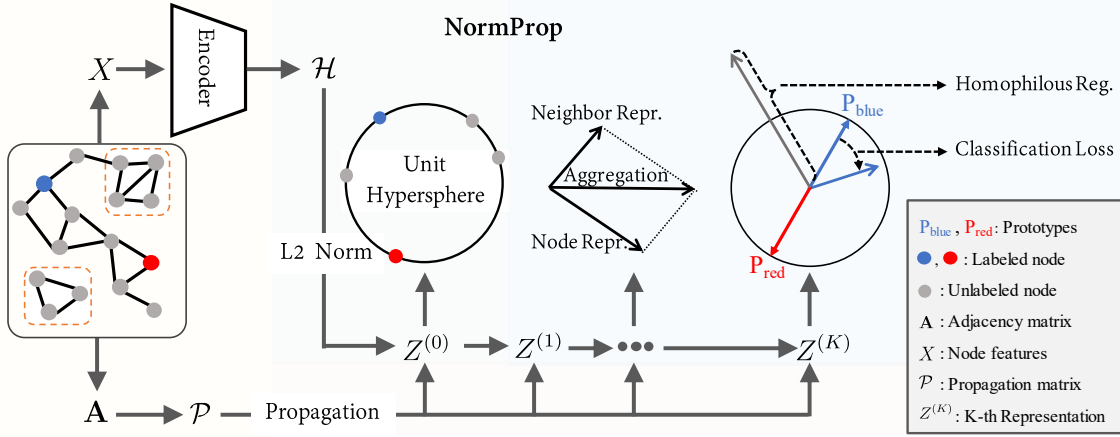


Figure 2: An illustration of the proposed NormProp (Normalize then Propagate).

virtual labels to the labeled set to enlarge the training set. Meta-PN (Ding et al. 2022) infers high-quality pseudo labels on unlabeled nodes with a meta-learned label propagation strategy, augmenting rare labeled nodes and enabling broad receptive fields during training. Currently, there are some techniques available to modify structure and enlarge receptive field, thereby amplifying the influence of label signals. To counter the over-smoothing issue, decoupled GCNs have gained popularity due to their simplicity. APPNP (Klicpera, Bojchevski, and Günnemann 2019) combines GNN with the concept of Personalized PageRank (Page et al. 1999), effectively decoupling the feature transformation and propagation with restart probability. DAGNN (Liu, Gao, and Ji 2020) aggregate all the previous embeddings at the final layers. The most time-consuming part of many GNNs is propagation (Chen, Ma, and Xiao 2018), and expanding the scope of propagation will aggravate the problem. Graph Structure Learning (GSL) (Zhu et al. 2021) can additionally facilitate the propagation of label signals. AM-GCN (Wang et al. 2020) propose an adaptive multi-channel GCN that combines the original structures with the k NN graph generated based on the similarity of node features. Violin (Xie, Tam, and Lau 2023) creates virtual connections between labeled and unlabeled nodes. However, many GSL methods require adjusting graph structure by modeling nodes similarity, which entails quadratic complexity.

3 NormProp: Normalize then Propagate

We utilize the decoupled GCN framework, which separates representation transformation and propagation (Eq.2). The above-mentioned framework eliminates the linear transformation and non-linear activation in each graph convolution layer, enabling us delve deeper into the node representation changes caused by propagation, especially consistency of the nodes encountered during propagation.

Expanding on the aforementioned decoupled GCN framework, we map node features to embeddings in the hyperspherical space, and present a new approach called NormProp. In our approach, we propagate L_2 hypersphere embed-

dings and generate supervision signals from both directions (class information) and Euclidean norm (consistency of aggregation), as illustrated in Fig.2. Unlike many few-shot node classification methods that rely on pseudo-labels (Ding et al. 2022) and multi-stage training (Sun, Lin, and Zhu 2020), the proposed NormProp is end-to-end. Moreover, the proposed framework consists of the following parts:

(I) **Hyperspherical Node Feature Encoder.** We map the vanilla node feature onto a unit hypersphere. Precisely, we apply a multilayer perceptron $f_\phi: \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ operates only on nodes' feature, and then generate $\mathcal{H} = [\mathcal{H}_1^T, \dots, \mathcal{H}_n^T] \in \mathbb{R}^{n \times d'}$. Subsequently, we project network outputs \mathcal{H} to the hypersphere through L_2 normalization, resulting the normalized node representations $Z^{(0)} \in \mathbb{R}^{n \times d'}$.

(II) **Propagation.** Since we apply the decoupled GCN framework in order to better study the impact of propagation, the propagation corresponds to a fixed low-pass filter (Nt and Maehara 2019) applied to each node hyperspherical representation dimension, akin to the approach employed in SGC (Wu et al. 2019). The ultimate representation of node v_i is essentially the weighted sum of its K -order neighborhoods' hyperspherical representation, e.g., $Z_i^{(K)} = \sum_{j \in \mathcal{N}^K(i)} \alpha_{ij} Z_j^{(0)}$, where $\mathcal{N}^K(i)$ denotes the K -order neighbors of node v_i .

In summary, the final representation of the nodes extracted through our NormProp can be succinctly expressed using the following mathematical formula:

$$\mathcal{H}_i = f_\phi(X_i) \quad \triangleright \text{Encode node features} \quad (3)$$

$$Z_i^{(0)} = \frac{\mathcal{H}_i}{\|\mathcal{H}_i\|_2} \quad \triangleright L_2 \text{ Normalization} \quad (4)$$

$$Z^{(K)} = \mathcal{P}^K Z^{(0)} \quad \triangleright \text{Propagation} \quad (5)$$

where $X_i \in \mathbb{R}^d$ denotes the vanilla feature of node v_i , $\mathcal{H}_i \in \mathbb{R}^{d'}$ denotes the non-normalized representation of v_i obtained by MLP $f_\phi(\cdot)$ with parameter set ϕ , $Z^{(0)} \in \mathbb{R}^{n \times d'}$ denote the hyperspherical representations after L_2 normalization, $Z^{(K)} \in \mathbb{R}^{n \times d'}$ denotes the final node representations achieved by aggregating K -order neighbors, and

the corresponding propagation weight matrix is denote as $\mathcal{P}^K \in \mathbb{R}^{n \times n}$. As shown in Fig.2, the representation vector $Z^{(K)}$ can be decouple into direction and Euclidean norm, corresponding to class information and consistency of aggregation, respectively. Similar to Eq.21 in Appendix A.1, classification is performed based on the final node representation $Z^{(K)}$ by finding the nearest class prototype in $\mathbb{P} = [P_1, \dots, P_C] \in \mathbb{R}^{C \times d'}$, using *cosine similarity* as the metric. Formally, the inference of node v_i can be defined as:

$$c^* = \arg \max_{c \in C} \left(\cos \theta_{(Z_i^{(K)}, P_c)} \right). \quad (6)$$

Additionally, we implement the predefined hyperspherical prototypes following Mettes et al. (Mettes, Van der Pol, and Snoek 2019). The solution of the hyperspherical prototypes is achieved through efficient data-independent optimization, as briefly described in Appendix A.1.

4 Homophilous Regularization: Consistent Loss based on Euclidean Norm

In this section, we establish the upper bound of Euclidean Norm in our NormProp, which is equivalent to the random walk on the graph. In the following, in order to further enhance the consistency of certain subgraphs, we propose the homophilous regularization based on homophily assumption. Finally, we take a different angle and analyze that homophilous regularization fundamentally symbolizes the agreement between original graph and augmented graph across propagation.

4.1 Upper Bound of Euclidean Norm

In the previously proposed framework, the directions of node representations $Z^{(K)}$ effectively describe their class information using the metric of cosine similarity. In other ways, we perform a comprehensive analysis to further explore the relationship between the Euclidean Norm of $Z^{(K)}$ and the consistency of node representations encountered during message passing.

Definition 1 (Normed Vector Space) *Let E be a vector space over a field K , where K is either the field \mathbb{R} of reals, or the field \mathbb{C} of complex numbers. A norm on E is a function $\|\cdot\| : E \rightarrow \mathbb{R}_+$, assigning a nonnegative real number $\|u\|$ to any vector $u \in E$, and satisfying the following conditions for all $x, y, z \in E$:*

$$\|x\| \geq 0 \quad \triangleright \text{Positivity} \quad (7)$$

$$\|\lambda x\| = |\lambda| \|x\| \quad \triangleright \text{Scaling} \quad (8)$$

$$\|x + y\| \leq \|x\| + \|y\| \quad \triangleright \text{Triangle inequality} \quad (9)$$

A vector space E together with a norm $\|\cdot\|$ is called a normed vector space.

In this paper, we consider Euclidean Norm on the real field, specifically the l_2 norm of node representations during propagation. The outcome of the propagation can be understood as a weighted sum of the representations from neighboring nodes. According to the properties of Definition 1,

We can derive an upper bound on the Euclidean norm of the representation propagation:

$$\left\| \sum_{j \in \mathcal{N}(i)} \alpha_{ij} \cdot Z_j^{(0)} \right\|_2 \leq \sum_{j \in \mathcal{N}(i)} \alpha_{ij} \|Z_j^{(0)}\|_2, \quad (10)$$

where $\sum_{j \in \mathcal{N}(i)} \alpha_{ij} Z_j^{(0)}$ denotes the propagation result of node v_i ; $\alpha_{ij} \geq 0$ denotes the aggregation coefficient. As shown in Eq.4, node representations $Z^{(0)}$ are distributed over the surface of hypersphere, so $\|Z_j^{(0)}\|_2 = 1$. Finally, the l_2 norm upper bound of Z^K can be formulated as:

$$\|Z_i^{(K)}\|_2 \leq [\mathcal{P}^K \mathbf{1}_{n \times 1}]_i, \quad (11)$$

where $\mathbf{1}_{n \times 1}$ denotes an all-ones matrix of size $n \times 1$; $\mathcal{P}^K \in \mathbb{R}^{n \times n}$ denotes the weight matrix of aggregation.

It is important to highlight that in Eq.11, the requirement for equality is that all the representation vectors involved in the aggregation are parallel. Furthermore, the Euclidean norm of the final node representation is positively correlated with consistency of aggregation. Additionally, the upper bound $\mathcal{P}^K \mathbf{1}_{n \times 1}$ in Eq.11 can be interpreted as the result of a random walk (Lovász 1993) on the graph. Because of Eq.4, all nodes have equal norm (unnormalized probability), and the norm of node representations after propagation (Eq.5) reflects the importance of nodes in the random walk.

4.2 Homophilous Regularization

Building upon the foundation of the bound, we will introduce the consistent metric based on Euclidean norm, and explore additional supervisory signals from unlabeled nodes.

Definition 2 (Consistent Metric) *For node v_i , the lower and upper bounds of its Euclidean norm in the propagation of NormProp could be relaxed as:*

$$0 \leq \|Z_i^{(K)}\|_2 \leq [\mathcal{P}^K \mathbf{1}_{1 \times n}]_i. \quad (12)$$

Moreover, because of the non-negativity property of the norm and its positive correlation with consistency of the subgraph, we normalize it using the upper bound. For node v_i , we define consistent metric $\zeta \in [0, 1]$ as follow:

$$\zeta = \frac{\|Z_i^{(K)}\|_2}{[\mathcal{P}^K \mathbf{1}_{1 \times n}]_i}. \quad (13)$$

Many theoretical foundations of GNNs are the homophily assumption (Ciotti et al. 2016), which states that nodes tend to connect with others who are ‘‘similar’’ to them. However, previous works (Chen et al. 2021; Hou et al. 2020) have pointed out that the nodes close to the class boundary have the risk of information conflict. Therefore, we consider that the node v_i exhibits a high level confidence, e.g., cosine similarity between $Z_i^{(K)}$ and any prototype $P \in \mathbb{P}$ is greater than threshold τ , indicating a high level consistency during message passing. Let $\Omega_\tau = \{v_j : \cos(\theta_{P, Z_j^{(K)}}) \geq \tau\}$. Formally, the homophilous regularization \mathcal{L}_h is defined based on consistent metric ζ :

$$\mathcal{L}_h = 1 - \frac{1}{|\Omega_\tau|} \sum_{i \in \Omega_\tau} \frac{\|Z_i^{(K)}\|_2}{[\mathcal{P}^K \mathbf{1}_{1 \times n}]_i} \quad (14)$$

Homophilous regularization aims to enhance the consistency of certain specific subgraphs. The term "Homophilous Regularization" is used to describe \mathcal{L}_h in Eq.14 because its mathematical form is similar to L_2 regularization $\|\cdot\|_2$.

4.3 Reinterpretation of Homophilous Regularization

In Eq.14, we view the homophilous regularization as the ratio of Euclidean norm. Furthermore, we will provide a more intuitive explanation from a different perspective. \mathcal{L}_h can be expressed over the entire graph as a matrix formulation:

$$\mathcal{L}_h = 1 - \frac{1}{\|\mathcal{M}\|_1} \text{tr} \left(\frac{\left(Z^{(K)\top} Z^{(K)} \mathcal{M} \right)^{-\frac{1}{2}}}{\mathcal{P}^K \mathbf{1}_{n \times n}} \right), \quad (15)$$

where $\text{tr}(\cdot)$ denotes trace of a matrix; $\mathbf{1}_{n \times n}$ denotes an all-ones matrix of size $n \times n$; $\mathcal{M} = \text{diag}(M_1, \dots, M_n) \in \mathbb{R}^{n \times n}$ denotes the diagonal masking matrix, $M_i = 1$ if $v_i \in \Omega_\tau$, or $M_i = 0$ otherwise. The portion of Eq.15 that is related to the node representations can be transformed into the following form.

$$Z^{(K)\top} Z^{(K)} \mathcal{M} = (\mathcal{P}^K Z^{(0)})^\top (\mathcal{P}^K Z^{(0)}) \mathcal{M} \quad (16)$$

$$= (\mathcal{P}^K Z^{(0)})^\top \mathcal{P}^K (Z^{(0)} \mathcal{M}) \quad (17)$$

$$= (\mathcal{P}^K Z^{(0)})^\top (\mathcal{P}^K \tilde{Z}^{(0)}), \quad (18)$$

where $\tilde{Z}^{(0)} = Z^{(0)} \mathcal{M}$ denotes feature masking operation (You et al. 2020), e.g., sets the norm of nodes with low confidence to zero. On one hand, $\mathcal{P}^K Z^{(0)}$ denotes the propagation using all nodes, while on the other hand, $\mathcal{P}^K \tilde{Z}^{(0)}$ represents the propagation using only the features from high-confident nodes. Essentially, \mathcal{L}_h serves as the contrastive loss, ensuring the consistency of the inner product between the augmented graph and the original graph.

5 Model Implementation and Training

5.1 Implementation of NormProp

Following the graph convolution design of GCN (Kipf and Welling 2017) and SGC (Wu et al. 2019), we employ the graph convolution matrix with the renormalization trick as aggregation operator, e.g., $\mathcal{P} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$. Due to the removal of linear transformation and activation function in message passing in our framework, the final form of propagation is equivalent to the fixed K -order low-pass filter in SGC. During the computation of node representations (Eq.3 ~ Eq.5) in the implementation of NormProp, the propagation rule (Eq.5) can be written as:

$$Z^{(K)} = \left(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \right)^K Z^{(0)}, \quad (19)$$

where $Z^{(K)}$ denotes the final node representations. Because of the over-smoothing problem (Cai and Wang 2020), we utilize a shallow aggregation scheme to extract the information of subgraphs in Eq.19. Moreover, we implement node feature encoder $f_\phi(\cdot)$ using a two-layer MLP.

5.2 Model Training

The classification loss \mathcal{L}_c in the supervision signal is achieved through cosine similarity between $Z_i^{(K)}$ of the labeled node v_i and its corresponding class prototype P_{y_i} . Training and inference is achieved through cosine similarities between examples and prototypes.

Ultimately, we combine classification loss \mathcal{L}_c and homophilous regularization \mathcal{L}_h into a unified training loss \mathcal{L} :

$$\mathcal{L} = \underbrace{\frac{1}{|\Omega_{\text{train}}|} \sum_{i \in \Omega_{\text{train}}} \left(1 - \cos \theta_{(Z_i^{(K)}, P_{y_i})} \right)}_{\text{Classification Loss from Labeled Nodes}} + \lambda \underbrace{\left(1 - \frac{1}{|\Omega_\tau|} \sum_{j \in \Omega_\tau} \frac{\|Z_j^{(K)}\|_2}{[\mathcal{P}^K \mathbf{1}_{1 \times n}]_j} \right)}_{\text{Homophilous Regularization from Unlabeled Nodes}}, \quad (20)$$

where λ is a mixing coefficient. The classification loss \mathcal{L}_c is calculated by measuring the cosine similarity between the final representation $Z_i^{(K)}$ of the training nodes with labels (e.g., $v_i \in \Omega_{\text{train}}$) and its corresponding prototype P_{y_i} . To tackle the challenge of limited availability of high-confidence nodes during the early stages of training, we implement a warm-up strategy. This strategy entails training the model using only classification loss \mathcal{L}_c during the initial phase. After reaching a certain number of epochs, we add the homophilous regularization and utilize \mathcal{L} in Eq.20 to train the model.

6 Experiments

We first give the experimental setup, then compare our NormProp with the state-of-the-art methods on a wide variety of semi-supervised node classification datasets in the limited labeled settings. We analysis the impact of hyperparameters by conducting an ablation study of NormProp. Afterwards, we provide complexity analysis and empirical study on Ogbn-arxiv. The code of NormProp is available at <https://github.com/Pallaksch/NormProp>.

6.1 Experiment Setup

Datasets. We utilize a total of 6 real-world datasets in our experiments. For the classical semi-supervised node classification, we adhere to the standard fixed split (Yang, Cohen, and Salakhudinov 2016) of previous work (Planetoid split, such as GCN, Violin (Xie, Tam, and Lau 2023) for Cora, Citeseer and Pubmed. Additionally, we evaluate NormProp on few-shot semi-supervised node classification using the Cora-ML, Citeseer, MS-CS and Ogbn-arxiv with limited label, following the Meta-PN settings. More details of the datasets can be found in the Appendix A.2.

Baselines. We compare NormProp with various state-of-the-art baselines: (I) Classical methods: GCN (Kipf and Welling 2017), GIN (Xu et al. 2019), GraphSAGE (Hamilton, Ying, and Leskovec 2017), SGC (Wu et al. 2019) and GAT (Veličković et al. 2018). (II) Graph structure learning methods (Zhu et al. 2021): GRAND (Feng et al. 2020), GAM (Stretcu et al. 2019) and Violin (Xie, Tam, and Lau

Method	Cora-ML		Citeseer		MS-CS		Ogbn-arxiv
	3-shot	5-shot	3-shot	5-shot	3-shot	5-shot	2.5% label ratio
MLP	41.07 ± 0.76	51.12 ± 0.61	43.34 ± 0.56	44.90 ± 0.60	70.33 ± 0.37	79.41 ± 0.31	37.68 ± 0.34
GCN	48.02 ± 0.89	67.32 ± 1.02	53.60 ± 0.86	62.60 ± 0.58	69.24 ± 0.94	84.43 ± 0.89	61.36 ± 0.51
SGC	49.60 ± 0.55	67.24 ± 0.86	57.37 ± 0.98	61.55 ± 0.53	72.11 ± 0.76	87.51 ± 0.27	59.38 ± 0.06
LP	62.07 ± 0.71	68.01 ± 0.62	54.07 ± 0.59	55.73 ± 1.19	57.96 ± 0.69	62.98 ± 0.61	54.30 ± 0.00
GLP	65.57 ± 0.26	71.26 ± 0.31	65.76 ± 0.49	71.36 ± 0.18	86.10 ± 0.21	86.94 ± 0.23	N/A
IGCN	66.60 ± 0.29	72.50 ± 0.20	67.47 ± 0.29	72.92 ± 0.10	85.83 ± 0.06	87.01 ± 0.05	N/A
APPNP	72.39 ± 0.98	78.32 ± 0.58	67.55 ± 0.77	71.08 ± 0.61	86.65 ± 0.42	90.13 ± 0.86	61.60 ± 0.28
DAGNN	71.86 ± 0.75	77.20 ± 0.69	66.62 ± 0.27	70.55 ± 0.12	86.32 ± 0.57	90.30 ± 0.66	63.01 ± 0.39
C&S	68.93 ± 0.86	73.37 ± 0.24	63.02 ± 0.72	64.72 ± 0.53	85.86 ± 0.45	87.99 ± 0.24	56.69 ± 0.12
GPR-GNN	70.98 ± 0.84	75.18 ± 0.52	64.32 ± 0.81	65.28 ± 0.52	86.12 ± 0.37	90.29 ± 0.38	58.5 ± 0.8 [†]
M3S	64.66 ± 0.31	69.64 ± 0.18	65.12 ± 0.20	68.18 ± 0.18	84.96 ± 0.18	86.83 ± 0.29	N/A
Meta-PN	74.94 ± 0.25	79.88 ± 0.15	70.48 ± 0.34	74.14 ± 0.50	88.99 ± 0.29	91.31 ± 0.22	62.66 ± 0.14
NormProp	<u>73.39 ± 0.75</u>	81.07 ± 0.49	<u>69.67 ± 0.62</u>	<u>72.92 ± 0.86</u>	89.50 ± 0.62	91.71 ± 0.33	64.87 ± 0.23

Table 1: Test accuracy (%) on few-shot semi-supervised node classification. (‘†’: results are collected from published papers, ‘N/A’: results are not reported.)

2023). (III) Deep GNNs: APPNP (Klicpera, Bojchevski, and Günnemann 2019), JKNet (Xu et al. 2018), DAGNN (Liu, Gao, and Ji 2020), GCNII (Chen et al. 2020), C&S (Huang et al. 2020) and GPR-GNN (Chien et al. 2020). (IV) Label propagation based methods: LP(Zhu and Ghahramani 2002), GLP and IGCN(Li et al. 2019a). (V) Pseudo label based methods for few-shot node classification: M3S (Sun, Lin, and Zhu 2020) and Meta-PN (Ding et al. 2022). The implementation details can be found in the Appendix A.4.

Parameter Setup. We implemented NormProp using Pytorch and Pytorch Geometric. The node feature encoder $f_\phi(\cdot)$ of our proposed methods is implemented using a MLP with two layers, and the hidden size in $\{32, 64, 256, 512\}$. We tune the following hyper-parameters: $K \in \{2, 3\}$; $\lambda \in [0, 2]$; weight decay $\in \{1e-2, 5e-3, 1e-3, 5e-4, 1e-4, 0\}$; similarity threshold $\tau \in [0, 1.0]$. We use the Adam

Method	Cora	Citeseer	Pubmed
MLP	58.51 ± 0.80	55.64 ± 0.46	72.71 ± 0.61
GIN	78.83 ± 1.45	66.87 ± 0.96	77.83 ± 0.42
SGC	80.70 ± 0.55	71.94 ± 0.07	78.82 ± 0.04
GraphSAGE	81.73 ± 0.58	69.86 ± 0.62	77.20 ± 0.40
GCN	82.52 ± 0.60	71.02 ± 0.83	79.16 ± 0.35
GAT	82.76 ± 0.88	71.87 ± 0.53	77.74 ± 0.34
JK-Net	80.35 ± 0.58	67.29 ± 1.02	78.36 ± 0.31
APPNP	83.13 ± 0.58	71.39 ± 0.68	80.30 ± 0.17
GCNII	84.17 ± 0.40	72.46 ± 0.74	79.85 ± 0.34
GRAND	84.50 ± 0.30	74.20 ± 0.30	80.00 ± 0.30
GAM	84.80 ± 0.06	72.46 ± 0.44	81.00 ± 0.09
M3S	79.57 ± 0.59	69.04 ± 0.84	77.49 ± 0.50
Violin	84.49 ± 0.66	74.26 ± 0.40	81.23 ± 0.42
Meta-PN	82.07 ± 1.34	71.84 ± 1.56	79.37 ± 0.37
NormProp	85.46 ± 0.51	74.33 ± 0.57	80.72 ± 1.09

Table 2: Test accuracy (%) on standard semi-supervised node classification.

(Kingma and Ba 2014) optimizer with a learning rate of 0.01. More details can be found in Appendix A.4.

6.2 Experiment Results

Standard semi-supervised node classification. Following the typical setup of GNN experiments, we present the average and standard deviation obtained after 10 runs, as summarized in Table 2. We can observe that NormProp significantly outperform the SOTA methods in the Cora and Citeseer, even using a fixed graph convolution without adjusting graph structure. Moreover, NormProp demonstrates a significant improvement of about 3% over both GCN and SGC when applied to Cora and Citeseer, even if they utilize the same low-pass filter (e.g., $\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$ in Eq.19) and undergo the same steps of propagation. Our NormProp, in contrast to deep GNNs that enlarge the receptive field to enhance the influence of label signals, acquires additional supervisory signals via homophilous regularization solely for optimizing the model’s parameters.

Few-shot semi-supervised node classification. Following the settings in Meta-PN (Ding et al. 2022), we report mean accuracy and 95% confidence interval after 50 runs in Table 1. We can observe that NormProp is competitive with other SOTA methods. In the evaluation on the large scale real-world graph ogbn-arxiv, NormProp outperforms all others under the 2.5% label ratio. Meta-PN achieves the best performance on some datasets but comes with a high complexity due to pseudo-label generation and long-distance propagation (as shown in Table 3). On the ogbn-arxiv, our NormProp’s training time is just 17% of Meta-PN’s, with inference time at 7%, further analysis is provided in Section 6.4.

6.3 Ablation Study

To validate the effectiveness of \mathcal{L}_h , we compared the bias of NormProp for all nodes during training when $\lambda = 0$ and $\lambda = 1.0$. We define the global bias of node representations $Z^{(k)}$ generated by NormProp: $\frac{1}{|V|} \sum_{i \in V} \left(1 - \cos \theta_{(Z_i^{(k)}, P_{y_i})}\right)$.

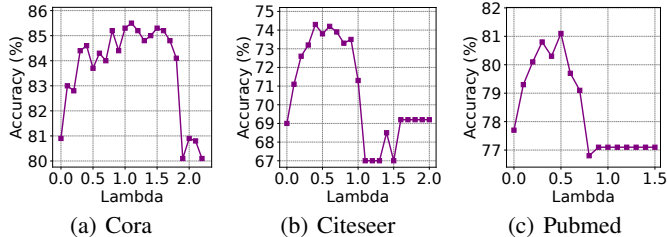


Figure 3: Impact of hyper-parameter λ .

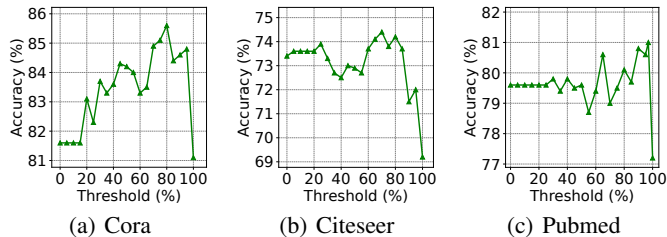


Figure 4: Impact of hyper-parameter threshold τ .

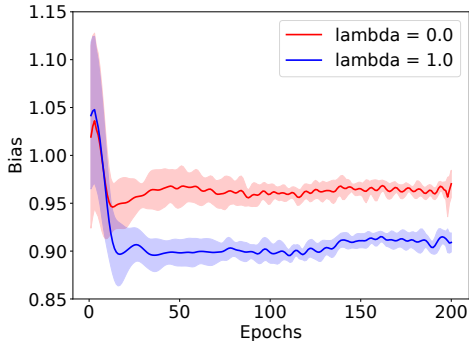


Figure 5: Global bias of NormProp on Cora.

Global bias is actually the calculation of classification loss \mathcal{L}_c for all nodes. The curve representing bias during training process is depicted in Figure 5. We make the following observation: After the warm-up phase (> 10 epochs), homophilous regularization can effectively reduce the global bias of node representations.

To understand the hyperparameter impacts on NormProp, we conduct experiments with different values of mixing coefficient λ and similarity threshold τ . We present the hyperparameter study in Figure 3 and Figure 4 for Planetoid datasets (Cora, Citeseer, Pubmed). When the λ is too large, the model tends to favor trivial solutions, resulting in smoother representations of the nodes and hampers the separability between different classes. In Figure 4, when the τ is small, there will be more nodes with poor consistency affecting the optimization of model. The value of τ involves the selection of unlabeled nodes involved in homophilous regularization, which is a typical quantity-quality trade-off.

6.4 Complexity Analysis

Considering that NormProp is an end-to-end method, we analysis the time complexity of one training epoch based on Eq.3 ~ Eq.5 and Eq.20. We assume that the sizes of different layers in the MLP are not significantly different. First, node representation $Z^{(0)}$ need to be prepared according Eq.3 and Eq.4. Let $|V|$ denotes the numbers of nodes, $|E|$ denotes the numbers of edges, F' denotes the dim of hidden layer, C denotes the class number. The complexity of Eq.3 and Eq.4 is $\mathcal{O}(F^2|V|)$. Second, there are K steps of propagation in

Method	Training time	Inference time
GCN	89.97 ± 0.18	0.0598 ± 0.0002
APPNP	107.21 ± 0.05	0.0725 ± 0.0003
DAGNN	141.05 ± 0.02	0.0942 ± 0.0003
Meta-PN	479.26 ± 17.74	0.3395 ± 0.0001
NormProp	79.45 ± 0.04	0.0239 ± 0.0002

Table 3: Training / Inference time in seconds for different methods on ogbn-arxiv.

Eq.5, where the time cost is $\mathcal{O}(KF|E|)$. Third, the complexity of loss calculation Eq.20 is $\mathcal{O}(FC|V|)$. Based on the above three components, the complexity of NormProp is $\mathcal{O}(F^2|V| + KF|E| + FC|V|)$. The time complexity of our method exhibits a linear relationship with both the number of nodes $|V|$ and the number of edges $|E|$. We provide efficiency empirical study with some SOTA methods in Table 3. NormProp demonstrates significantly greater computational efficiency compared to the other methods, both in training time and inference time. More details can be found in Appendix A.4 and A.5.

7 Conclusion

In this paper, we begin by analyzing the limitations of GNN generalization in the context of the semi-supervised node classification problem with limited label, focusing on the perspective of supervision signals. In order to effectively utilize the information from unlabeled nodes to improve the generalization ability of the model, we propose a novel GNN named NormProp. By employing “Normalize then Propagate” schema, the direction and Euclidean norm of node representations extracted by NormProp respectively indicate class information and consistency of aggregation. Crucially, we provide a theoretical analysis on the upper bound of Euclidean norm. Based on the bound of Euclidean norm and the homophily assumption, we propose homophilous regularization to improve the homophily of certain specific subgraphs. According to our comprehensive experiments, our NormProp achieves state-of-the-art performance. The experiments show that homophilous regularization can effectively and efficiently utilize the information from unlabeled nodes, thereby enhancing generalization ability.

Acknowledgments

This paper is supported by the National Natural Science Foundation of China (Grant No. 62192783, 62376117), the National Social Science Fund of China (Grant No. 23BJL035), the Science and Technology Major Project of Nanjing (comprehensive category) (Grant No. 202309007), and the Collaborative Innovation Center of Novel Software Technology and Industrialization at Nanjing University.

References

- Cai, C.; and Wang, Y. 2020. A Note on Over-Smoothing for Graph Neural Networks. *ArXiv*, abs/2006.13318.
- Chen, D.; Lin, Y.; Zhao, G.; Ren, X.; Li, P.; Zhou, J.; and Sun, X. 2021. Topology-Imbalance Learning for Semi-Supervised Node Classification. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 29885–29897.
- Chen, J.; Ma, T.; and Xiao, C. 2018. Fastgcn: fast learning with graph convolutional networks via importance sampling. *arXiv preprint arXiv:1801.10247*.
- Chen, M.; Wei, Z.; Huang, Z.; Ding, B.; and Li, Y. 2020. Simple and Deep Graph Convolutional Networks. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 1725–1735.
- Chien, E.; Peng, J.; Li, P.; and Milenkovic, O. 2020. Adaptive universal generalized pagerank graph neural network. *arXiv preprint arXiv:2006.07988*.
- Ciotti, V.; Bonaventura, M.; Nicosia, V.; Panzarasa, P.; and Latora, V. 2016. Homophily and missing links in citation networks. *EPJ Data Science*, 5: 1–14.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 3837–3845.
- Ding, K.; Wang, J.; Caverlee, J.; and Liu, H. 2022. Meta propagation networks for graph few-shot semi-supervised learning. In *Proceedings of the AAAI conference on artificial intelligence*.
- Feng, W.; Zhang, J.; Dong, Y.; Han, Y.; Luan, H.; Xu, Q.; Yang, Q.; Kharlamov, E.; and Tang, J. 2020. Graph random neural networks for semi-supervised learning on graphs. *Advances in neural information processing systems*, 33: 22092–22103.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning*.
- Hamilton, W. L.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, 1024–1034.
- Han, H.; Liu, X.; Shi, F.; Torkamani, M.; Aggarwal, C.; and Tang, J. 2024. Towards Label Position Bias in Graph Neural Networks. *Advances in Neural Information Processing Systems*, 36.
- Han, X.; Jiang, Z.; Liu, N.; and Hu, X. 2022. G-mixup: Graph data augmentation for graph classification. In *International Conference on Machine Learning*, 8230–8248. PMLR.
- Hou, Y.; Zhang, J.; Cheng, J.; Ma, K.; Ma, R. T. B.; Chen, H.; and Yang, M.-C. 2020. Measuring and Improving the Use of Graph Information in Graph Neural Networks. In *International Conference on Learning Representations*.
- Hu, S.; Xiong, Z.; Qu, M.; Yuan, X.; Côté, M.-A.; Liu, Z.; and Tang, J. 2020. Graph policy network for transferable active learning on graphs. *Advances in Neural Information Processing Systems*, 33: 10174–10185.
- Huang, Q.; He, H.; Singh, A.; Lim, S.-N.; and Benson, A. R. 2020. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
- Klicpera, J.; Bojchevski, A.; and Günnemann, S. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *International Conference on Learning Representations*.
- Li, Q.; Wu, X.-M.; Liu, H.; Zhang, X.; and Guan, Z. 2019a. Label efficient semi-supervised learning via graph filtering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9582–9591.
- Li, X.; Sun, Q.; Liu, Y.; Zhou, Q.; Zheng, S.; Chua, T.-S.; and Schiele, B. 2019b. Learning to Self-Train for Semi-Supervised Few-Shot Classification. In *Advances in Neural Information Processing Systems*.
- Liu, M.; Gao, H.; and Ji, S. 2020. Towards Deeper Graph Neural Networks. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, 338–348.
- Lovász, L. 1993. Random walks on graphs. *Combinatorics, Paul erdos is eighty*, 2(1-46): 4.
- Ma, J.; Deng, J.; and Mei, Q. 2021. Subgroup generalization and fairness of graph neural networks. *Advances in Neural Information Processing Systems*, 34: 1048–1061.
- Ma, J.; Ma, Z.; Chai, J.; and Mei, Q. 2022. Partition-based active learning for graph neural networks. *arXiv preprint arXiv:2201.09391*.
- Mettes, P.; Van der Pol, E.; and Snoek, C. 2019. Hyperspherical prototype networks. *Advances in neural information processing systems*, 32.
- Nt, H.; and Maehara, T. 2019. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*.
- Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1999. The PageRank Citation Ranking : Bringing Order to the Web. In *The Web Conference*.

- Stretcu, O.; Viswanathan, K.; Movshovitz-Attias, D.; Platanios, E.; Ravi, S.; and Tomkins, A. 2019. Graph agreement models for semi-supervised learning. *Advances in Neural Information Processing Systems*, 32.
- Sun, K.; Lin, Z.; and Zhu, Z. 2020. Multi-stage self-supervised learning for graph convolutional networks on graphs with few labeled nodes. In *Proceedings of the AAAI conference on artificial intelligence*.
- Tang, X.; Yao, H.; Sun, Y.; Wang, Y.; Tang, J.; Aggarwal, C.; Mitra, P.; and Wang, S. 2020. Investigating and mitigating degree-related biases in graph convolutional networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 1435–1444.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Wang, X.; Zhu, M.; Bo, D.; Cui, P.; Shi, C.; and Pei, J. 2020. Am-gcn: Adaptive multi-channel graph convolutional networks. In *Proceedings of the 26th ACM SIGKDD International conference on knowledge discovery & data mining*, 1243–1253.
- Wu, F.; Jr., A. H. S.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. Q. 2019. Simplifying Graph Convolutional Networks. In *International Conference on Machine Learning*, volume 97, 6861–6871.
- Xie, S.; Tam, D. S. H.; and Lau, W. C. 2023. Violin: virtual overbridge linking for enhancing semi-supervised learning on graphs with limited labels. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 4451–4459.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- Xu, K.; Li, C.; Tian, Y.; Sonobe, T.; Kawarabayashi, K.; and Jegelka, S. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80, 5449–5458.
- Yang, Z.; Cohen, W.; and Salakhudinov, R. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, 40–48. PMLR.
- Ying, Z.; You, J.; Morris, C.; Ren, X.; Hamilton, W.; and Leskovec, J. 2018. Hierarchical graph representation learning with differentiable pooling. *Advances in neural information processing systems*, 31.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33: 5812–5823.
- Zhang, M.; and Chen, Y. 2018. Link prediction based on graph neural networks. *Advances in neural information processing systems*, 31.
- Zhou, K.; Huang, X.; Zha, D.; Chen, R.; Li, L.; Choi, S.-H.; and Hu, X. 2021. Dirichlet energy constrained learning for deep graph neural networks. *Advances in Neural Information Processing Systems*, 34: 21834–21846.
- Zhu, X.; and Ghahramani, Z. 2002. Learning from labeled and unlabeled data with label propagation. *ProQuest number: information to all users*.
- Zhu, Y.; Xu, W.; Zhang, J.; Du, Y.; Zhang, J.; Liu, Q.; Yang, C.; and Wu, S. 2021. A survey on graph structure learning: Progress and opportunities. *arXiv preprint arXiv:2103.03036*.