

LSGNN: Towards General Graph Neural Network in Node Classification by Local Similarity

Yuhan Chen¹, Yihong Luo^{2,3}, Jing Tang^{2,3*},
Liang Yang⁴, Siya Qiu^{2,3}, Chuan Wang^{5*} and Xiaochun Cao⁶

¹School of Computer Science and Engineering, Guangzhou Campus of Sun Yat-sen University

²The Hong Kong University of Science and Technology (Guangzhou)

³The Hong Kong University of Science and Technology

⁴School of Artificial Intelligence, Hebei University of Technology

⁵Institute of Information Engineering, Chinese Academy of Sciences

⁶School of Cyber Science and Technology, Shenzhen Campus of Sun Yat-sen University
draym@qq.com, yluocg@connect.ust.hk, jingtang@ust.hk, yangliang@vip.qq.com,
sqial@connect.ust.hk, wangchuan@iie.ac.cn, caoxiaochun@mail.sysu.edu.cn

Abstract

Heterophily has been considered as an issue that hurts the performance of Graph Neural Networks (GNNs). To address this issue, some existing work uses a graph-level weighted fusion of the information of multi-hop neighbors to include more nodes with homophily. However, the heterophily might differ among nodes, which requires to consider the local topology. Motivated by it, we propose to use the local similarity (LocalSim) to learn node-level weighted fusion, which can also serve as a plug-and-play module. For better fusion, we propose a novel and efficient Initial Residual Difference Connection (IRDC) to extract more informative multi-hop information. Moreover, we provide theoretical analysis on the effectiveness of LocalSim representing node homophily on synthetic graphs. Extensive evaluations over real benchmark datasets show that our proposed method, namely Local Similarity Graph Neural Network (LSGNN), can offer comparable or superior state-of-the-art performance on both homophilic and heterophilic graphs. Meanwhile, the plug-and-play model can significantly boost the performance of existing GNNs.

1 Introduction

Graph Neural Network (GNN) has received significant interest in recent years due to its powerful ability in various real-world applications based on graph-structured data, i.e., node classification [Kipf and Welling, 2016], graph classification [Errica *et al.*, 2020], and link prediction [Zhang and Chen, 2018]. Combining convolutional network and graph signal processing, numerous Graph Convolutional Network (GCN) architectures [Scarselli *et al.*, 2009; Defferrard *et al.*, 2016; Hamilton *et al.*, 2017; Velickovic *et al.*, 2017;

Kipf and Welling, 2016] have been proposed and show superior performance in the above application domains. Recent work [Battaglia *et al.*, 2018] believes that the success of GCN and its variants is built on the homophily assumption: connected nodes tend to have the same class label [Hamilton, 2020]. This assumption provides proper inductive bias, raising the general message-passing mechanism: aggregating neighbour information to update ego-node feature—a special form of low-pass filter [Bo *et al.*, 2021].

However, the homophily assumption does not always hold [McPherson *et al.*, 2001; Jiang *et al.*, 2013]. In such cases, empirical evidence shows that GCN may even be worse than simple Multi-Layer Perceptrons (MLPs) that use merely node features as input [Chien *et al.*, 2021] (heterophily issue). A potential explanation is that the low-pass filter is believed to hurt the performance of GCN on heterophilic graph. Recently, the heterophily issue has received the community’s interest, and various methods [Pei *et al.*, 2020; Chien *et al.*, 2021; He *et al.*, 2021; Li *et al.*, 2022] have been proposed to address the issue. Many existing works [Chien *et al.*, 2021; He *et al.*, 2021; Li *et al.*, 2022] propose to fuse intermediate representations, i.e., outputs of different layers, for heterophilic graphs learning. However, they only consider graph-level heterophily. Moreover, intermediate representations extracted by traditional propagation methods [Kipf and Welling, 2016] may drop some important information, which limits the effectiveness of information fusion among neighbors.

We show that in real-world datasets the heterophily among nodes can be significantly different (Section 3.3), which suggests that only considering graph-level heterophily is insufficient. Although Liu *et al.* consider node-level heterophily, they simply map each node’s representation to its node-level weight, do not take local topology into consideration, yielding suboptimal results. Motivated by the deficiency of existing techniques, we propose to use the local similarity to indicate node homophily to conduct better node-level weighted fusion adaptively. Moreover, we empirically and theoretically show its capabilities on synthetic graphs (Section 4). Fur-

*Corresponding authors: Jing Tang and Chuan Wang.

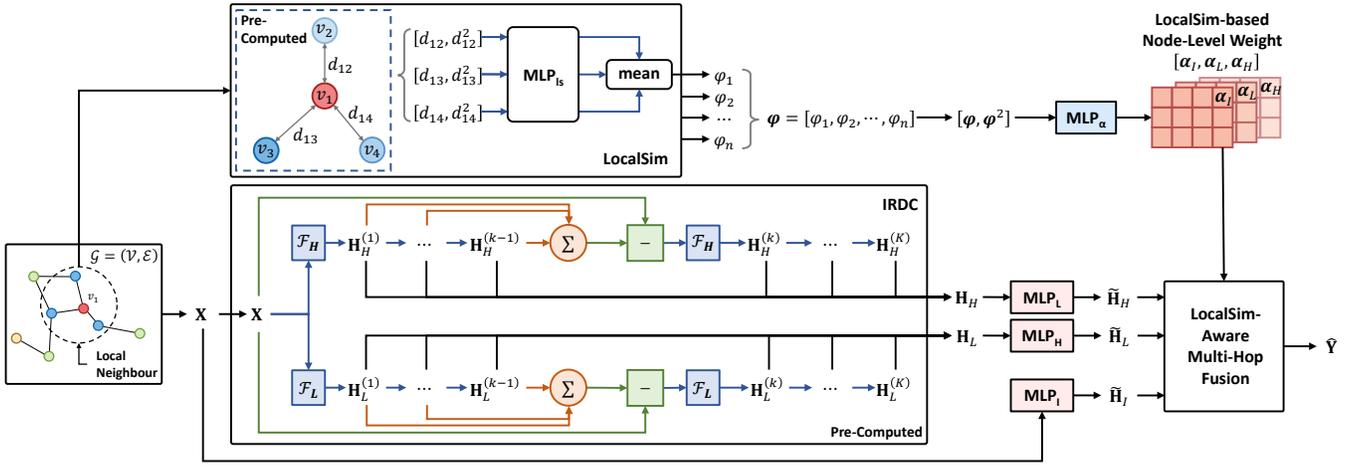


Figure 1: An illustration of LSGNN framework. ‘Pre-Computed’ means the part can be pre-computed without training.

thermore, to obtain more informative intermediate representations for better fusion, we propose a novel Initial Residual Difference Connection (IRDC), which can leverage all the input information. Meanwhile, the IRDC propagation preserves the key property of SGC [Wu *et al.*, 2019] on removing the nonlinear transformations between propagation, which is time- and GPU-consuming in GCNs. This property makes the feature propagation of LSGNN efficient and deep-learning free. Based on the above novel designs, LSGNN offers high performance and high efficiency in evaluations (Section 6).

To summarize, we make the following contributions: **1)** We study the local node homophily of real-world datasets and suggest using LocalSim as an indicator of node homophily. Moreover, we empirically and theoretically show the effectiveness of using LocalSim to indicate node homophily on synthetic graphs. **2)** We propose LSGNN, which contains a LocalSim-aware multi-hop fusion to guide the learning of node-level weight for intermediate representations and an IRDC to extract more informative intermediate representations for better fusion. **3)** LocalSim-based node-level weight is a plug-and-play module and can significantly boost the performance of state-of-the-art models, such as H₂GCN and GPRGNN. **4)** We conduct extensive experiments to demonstrate the superiority of our method, LSGNN, which can offer comparable or superior performance against 13 other methods over homophilic and heterophilic graphs.

2 Preliminaries

2.1 Notations

We denote an undirected graph without self-loops as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_i\}_{i=1}^n$ is the node set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set. Let \mathcal{N}_i denotes the neighbours of node v_i . Denote by $\mathbf{A} \in \mathbb{R}^{n \times n}$ the adjacency matrix, and by \mathbf{D} the diagonal matrix standing for the degree matrix such that $\mathbf{D}_{ii} = \sum_{j=1}^n \mathbf{A}_{ij}$. Let $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{D}}$ be the corresponding matrices with self-loops, i.e., $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ and $\tilde{\mathbf{D}} = \mathbf{D} + \mathbf{I}$, where \mathbf{I} is the identity matrix. Denote by $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n \in \mathbb{R}^{n \times d}$ the initial node feature matrix, where d is the number of dimen-

sions, and by $\mathbf{H}^{(k)} = \{\mathbf{h}_i^{(k)}\}_{i=1}^n$ the representation matrix in the k -th layer. We use $\mathbf{Y} = \{\mathbf{y}_i\}_{i=1}^n \in \mathbb{R}^{n \times C}$ to denote the ground-truth node label matrix, where C is the number of classes and \mathbf{y}_i is the one-hot encoding of node v_i 's label.

2.2 Simple Graph Convolution (SGC)

SGC [Wu *et al.*, 2019] claims that the nonlinearity between propagation in GCN (i.e., the ReLU activation function) is not critical, whereas the majority of benefit arises from the propagation itself. Therefore, SGC removes the nonlinear transformation between propagation so that the class prediction $\hat{\mathbf{Y}}$ of a K -layer SGC becomes

$$\hat{\mathbf{Y}} = \text{softmax}(\mathbf{S}^K \mathbf{X} \mathbf{W}), \quad (1)$$

where \mathbf{S} is a GNN filter, e.g., $\mathbf{S} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$ is used in the vanilla SGC, and \mathbf{W} is the learned weight matrix. Removing the nonlinearities, such a simplified linear model yields orders of magnitude speedup over GCN while achieving comparable empirical performance.

2.3 Heterophily Issue

Existing GNN models are usually based on the homophily assumption that connected nodes are more likely to belong to the same class. However, many real-world graphs are heterophilic. For example, the matching couples in dating networks are usually heterosexual. Empirical evidence shows that when the homophily assumption is broken, GCN may be even worse than simple Multi-Layer Perceptrons (MLPs) that use merely node features as input [Chien *et al.*, 2021]. Therefore, it is essential to develop a general GNN model for both homophilic and heterophilic graphs.

Pei *et al.* [2020] introduce a simple index to measure node homophily/heterophily. That is, the homophily $\mathcal{H}(v_i)$ of node v_i is the ratio of the number of v_i 's neighbours who have the same labels as v_i to the number of v_i 's neighbours, i.e.,

$$\mathcal{H}(v_i) = \frac{|\{v_j \mid v_j \in \mathcal{N}_i, \mathbf{y}_j = \mathbf{y}_i\}|}{|\mathcal{N}_i|}, \quad (2)$$

and the homophily in a graph is the average of all nodes, i.e.,

$$\mathcal{H}(\mathcal{G}) = \frac{1}{n} \sum_{v_i \in \mathcal{V}} \mathcal{H}(v_i). \quad (3)$$

Similarly, the heterophily of node v_i is measured by $1 - \mathcal{H}(v_i)$, and the heterophily in a graph is $1 - \mathcal{H}(\mathcal{G})$.

3 Methodology

3.1 Overview

In this section, we present the **Local Similarity Graph Neural Network (LSGNN)**, consisting of a novel propagation method named **Initial Residual Difference Connection (IRDC)** for better fusion (**Section 3.2**), and a **LocalSim-aware multi-hop fusion (Section 3.4)** and thus cope with both homophilic and heterophilic graphs. We also demonstrate the motivation of using **LocalSim** as an effective indicator of homophily (**Section 3.3**). Figure 1 shows the framework of LSGNN.

3.2 Initial Residual Difference Connection (IRDC)

Simply aggregating representation from the previous layer (as applied by most GNNs) might lose important information and consequently incur over-smoothing issues. To obtain more informative intermediate representations, i.e., outputs of different layers, we propose a novel propagation method, namely **Initial Residual Difference Connection (IRDC)**.

General Form. Inspired by SGC, we remove nonlinearity between layers and formulate IRDC as follows:

$$\begin{aligned} \mathbf{H}^{(1)} &= \text{IRDC}^{(1)}(\mathbf{S}, \mathbf{X}) = \mathbf{S}\mathbf{X}, \\ \mathbf{H}^{(k)} &= \text{IRDC}^{(k)}(\mathbf{S}, \mathbf{X}) = \mathbf{S} \left((1 - \gamma)\mathbf{X} - \gamma \sum_{\ell=1}^{k-1} \mathbf{H}^{(\ell)} \right), \end{aligned} \quad (4)$$

for $k = 2, \dots, K$. \mathbf{S} is a GNN filter, and $\gamma \in [0, 1]$ is a hyperparameter. The term $\sum_{\ell=1}^{k-1} \mathbf{H}^{(\ell)}$ in Equation (4) can be seen as the total information extracted by the previous $k - 1$ layers. In this way, IRDC feeds the information that has never been processed from the original input into the next layer, thus can fully exploit all the information from the initial node features. A normalization operation is conducted following IRDC to handle the potential value scale issue. See Appendix B for comparison between different residual connections.

Parameterization in Practice. Many GCNs use adjacency matrix with self-loops as the low-pass GNN filter to extract similar information from neighbours. However, the self-loops added in GCNs may not always be helpful [Zhu *et al.*, 2020b]. Instead, we use enhanced filters [Bo *et al.*, 2021] of both low-pass and high-pass ($\mathcal{F}_L + \mathcal{F}_H = \mathbf{I}$) by adding weighted identity matrix \mathbf{I} (i.e., weighted self-loops),

$$\begin{aligned} \mathcal{F}_L &= \beta \mathbf{I} + \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}, \\ \mathcal{F}_H &= \mathbf{I} - \mathcal{F}_L = (1 - \beta) \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}, \end{aligned} \quad (5)$$

where $\beta \in [0, 1]$ is a hyperparameter. Intuitively, combining low-pass and high-pass filters together can learn better representations for both homophilic and heterophilic graphs. In a

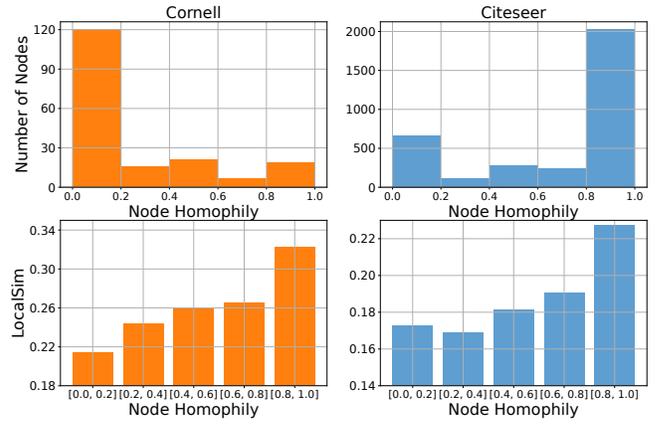


Figure 2: Positive relation between node homophily and LocalSim on Cornell (left column) and Citeseer (right column) datasets. The upper figures show the histogram of node homophily, while the lower figures show the result for node homophily vs. LocalSim.

K -layer LSGNN, the output $\mathbf{H}_L^{(k)}, \mathbf{H}_H^{(k)} \in \mathbb{R}^{n \times d}$ of the k^{th} low-pass and high-pass graph filter layer are

$$\mathbf{H}_L^{(k)} = \text{IRDC}^{(k)}(\mathcal{F}_L, \mathbf{X}) \text{ and } \mathbf{H}_H^{(k)} = \text{IRDC}^{(k)}(\mathcal{F}_H, \mathbf{X}),$$

where $k = 1, 2, \dots, K$. Next, $\mathbf{H}_L^{(k)}$ and $\mathbf{H}_H^{(k)}$ are transformed into z dimensions by learnable $\mathbf{W}_L^{(k)}, \mathbf{W}_H^{(k)} \in \mathbb{R}^{d \times z}$, i.e.,

$$\tilde{\mathbf{H}}_L^{(k)} = \sigma \left(\mathbf{H}_L^{(k)} \mathbf{W}_L^{(k)} \right) \text{ and } \tilde{\mathbf{H}}_H^{(k)} = \sigma \left(\mathbf{H}_H^{(k)} \mathbf{W}_H^{(k)} \right),$$

where σ is the ReLU function. To leverage the the initial node features, we also obtain a transformation $\tilde{\mathbf{H}}_I$ of \mathbf{X} such that

$$\tilde{\mathbf{H}}_I = \sigma(\mathbf{X} \mathbf{W}_I),$$

where $\mathbf{W}_I \in \mathbb{R}^{d \times z}$ is a learned weight matrix. In what follows, $\tilde{\mathbf{H}}_I, \tilde{\mathbf{H}}_L^{(k)}, \tilde{\mathbf{H}}_H^{(k)} \in \mathbb{R}^{n \times z}$ ($k = 1, 2, \dots, K$) are delivered to the LocalSim-aware multi-hop fusion (Section 3.4) to generate the final node representations.

3.3 LocalSim: An Indicator of Homophily

Based on the hypothesis that nodes with similar features are likely to belong to the same class, we start out by defining a simple version of Local Similarity (LocalSim) ϕ_i of node v_i to indicate its homophily, i.e.,

$$\phi_i = \frac{1}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} d_{ij} = \frac{1}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} \text{sim}(\mathbf{x}_i, \mathbf{x}_j), \quad (6)$$

where $\text{sim}(\cdot, \cdot) : (\mathbb{R}^d, \mathbb{R}^d) \mapsto \mathbb{R}$ is a similarity measure, e.g.,

$$\text{sim}(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} \frac{\mathbf{x}_i^\top \mathbf{x}_j}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|}, & \text{Cosine similarity,} \\ -\|\mathbf{x}_i - \mathbf{x}_j\|_2, & \text{Euclidean similarity.} \end{cases}$$

Figure 2 shows the relation between node homophily and LocalSim (using Cosine similarity) on Cornell (left column) and Citeseer (right column) datasets. From the upper figures, we observe that the homophily/heterophily among nodes are different. In particular, Cornell is a heterophilic graph, while Citeseer is a homophilic graph. Moreover, the lower figures show a clear positive correlation between node homophily and LocalSim, which indicates that LocalSim can be used to represent node homophily appropriately.

3.4 LocalSim-Aware Multi-Hop Fusion

In Section 3.3, we show that LocalSim, leveraging the local topology information, can identify the homophily of nodes. We apply LocalSim to learn the adaptive fusion of intermediate representations for each node.

The naive LocalSim in Equation (6) simply averages the similarity between the ego-node and its neighbors. In what follows, we refine the definition of LocalSim by incorporating nonlinearity, i.e., d_{ij}^2 . Specifically, the refined LocalSim φ_i of node i is given by

$$\varphi_i = \frac{1}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} \text{MLP}_{\text{ls}}([d_{ij}, d_{ij}^2]), \quad (7)$$

where $\text{MLP}_{\text{ls}} \rightarrow \mathbb{R}^2 \mapsto \mathbb{R}$ is a 2-layer perceptron. It is trivial to see that incorporating d_{ij}^2 via MLP_{ls} can encode the variance of series $\{d_{ij}\}_{v_j \in \mathcal{N}_i}$, since $\text{Var}[d_{ij}] = \mathbb{E}[d_{ij}^2] - \mathbb{E}^2[d_{ij}]$ with respect to $v_j \in \mathcal{N}_i$. As a consequence, φ_i can better characterize the real ‘‘local similarity’’ than ϕ_i . To illustrate, we compare φ_i and ϕ_i through a simple example. Suppose that node v_1 has two neighbours v_2 and v_3 such that $d_{12} = 0$ and $d_{13} = 1$, while node v_4 has two neighbours v_5 and v_6 such that $d_{45} = d_{46} = 0.5$. Then, $\phi_1 = \phi_4 = 0.5$, which cannot characterize the distinct neighbour distributions. On the other hand, by introducing nonlinearity, we can easily distinguish φ_1 and φ_4 . Finally, we have the LocalSim vector for all nodes $\varphi = [\varphi_1, \varphi_2, \dots, \varphi_n] \in \mathbb{R}^n$.

LocalSim can guide the intermediate representations fusion, i.e., using φ as local topology information when learning weights for fusion. In addition, we also introduce nonlinearity by using φ and φ^2 together to generate weights, i.e.,

$$[\alpha_I, \alpha_L, \alpha_H] = \text{MLP}_\alpha([\varphi, \varphi^2]), \quad (8)$$

where $\alpha_I, \alpha_L, \alpha_H \in \mathbb{R}^{n \times K}$ are the weights for node representation from each channel and $\text{MLP}_\alpha: \mathbb{R}^2 \mapsto \mathbb{R}^{3K}$ is a 2-layer perceptron. Then, the node representations for the k -th layer is computed as follows:

$$\mathbf{Z}^{(k)} = \alpha_I^{(k)} \odot \tilde{\mathbf{H}}_I + \alpha_L^{(k)} \odot \tilde{\mathbf{H}}_L^{(k)} + \alpha_H^{(k)} \odot \tilde{\mathbf{H}}_H^{(k)}, \quad (9)$$

where $\mathbf{Z}^{(k)} \in \mathbb{R}^{n \times z}$ and $\alpha_I^{(k)}, \alpha_L^{(k)}, \alpha_H^{(k)} \in \mathbb{R}^n$ for $k = 1, 2, \dots, K$, and $\alpha_I^{(k)} \odot \tilde{\mathbf{H}}_I$ is the i -th entry of $\alpha_I^{(k)}$ times the i -th row of $\tilde{\mathbf{H}}_I$ for $i = 1, 2, \dots, n$ (so do $\alpha_L^{(k)} \odot \tilde{\mathbf{H}}_L^{(k)}$ and $\alpha_H^{(k)} \odot \tilde{\mathbf{H}}_H^{(k)}$). Then, we obtain the final representation as $\tilde{\mathbf{H}}_I \parallel \mathbf{Z}^{(1)} \parallel \mathbf{Z}^{(2)} \parallel \dots \parallel \mathbf{Z}^{(K)}$, where \parallel is the concatenation function. Finally, the class probability prediction matrix $\hat{\mathbf{Y}} \in \mathbb{R}^{n \times C}$ is computed by

$$\hat{\mathbf{Y}} = \left[\tilde{\mathbf{H}}_I \parallel \mathbf{Z}^{(1)} \parallel \mathbf{Z}^{(2)} \parallel \dots \parallel \mathbf{Z}^{(K)} \right] \mathbf{W}_{\text{out}}, \quad (10)$$

where $\mathbf{W}_{\text{out}} \in \mathbb{R}^{(K+1)z \times C}$ is a learned matrix.

4 Case Study of Toy Example

In this section, we conduct a case study using a simple synthetic graph to demonstrate the superiority of our LocalSim-based node-level weight.

4.1 Empirical Investigation

Inspired by the Featured Stochastic Block Model (FSBM) [Chanpuriya and Musco, 2022], in the following, we introduce FSBM with mixture of heterophily that can generate a simple synthetic graph with different node heterophily.

Definition 1 (FSBM with Mixture of Heterophily). *A Stochastic Block Model (SBM) graph \mathcal{G} with mixture of heterophily has n nodes partitioned into r communities C_1, C_2, \dots, C_r and consisting of t subgraphs $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_t$ where there is no edge between any two subgraphs, with intra-community and inter-community edge probabilities p_k and q_k in subgraph \mathcal{G}_k . Let $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_r \in \{0, 1\}^n$ be indicator vectors for membership in each community, i.e., the j -th entry of \mathbf{c}_i is 1 if the j -th node is in C_i and 0 otherwise. An FSBM is such a graph model \mathcal{G} , plus a feature vector $\mathbf{x} = \mathbf{f} + \boldsymbol{\eta} \in \mathbb{R}^n$, where $\boldsymbol{\eta} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is zero-centered, isotropic Gaussian noise and $\mathbf{f} = \sum_i \mu_i \mathbf{c}_i$ for some $\mu_1, \mu_2, \dots, \mu_r \in \mathbb{R}$, which are the expected feature values of each community.*

According to Definition 1 and Equation (2), nodes in the same subgraph have identical homophily in expectation, while nodes from different subgraphs have distinct homophily. We consider FSBMs with $n = 1000$, 2 equally-sized communities C_1 and C_2 , 2 equally-sized subgraphs \mathcal{G}_1 and \mathcal{G}_2 , feature means $\mu_1 = 1, \mu_2 = -1$, and noise variance $\sigma = 1$. We generate different graphs by varying $\lambda_1 = \frac{p_1}{p_1 + q_1}$ and $\lambda_2 = \frac{p_2}{p_2 + q_2}$, where $\lambda_1, \lambda_2 \in [0, 1]$ and high λ_τ means high homophily, with the expected degree of all nodes to 10 (which means $\frac{1}{4}(p_\tau + q_\tau)n = 10$ for $\tau = 1, 2$).

We employ graph-level weight and LocalSim-based node-level weight to classify the nodes. For LocalSim-based node-level weight, we use the proposed method in Section 3. For graph-level weight, we replace α in Equation (8) with learnable graph-level value. For the similarity measure, we use $\text{sim}(x, y) = -(x - y)^2$, where x and y are scalar values. Finally, considering the simplicity of the synthetic networks, we use the simple LocalSim in Equation (6), and only use adjacency matrix with self-loop $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ and set the layers of IRDC to one.

As shown in Figure 3, our LocalSim-based node-level weight consistently outperforms graph-level weight. In particular, when node homophily is significantly different among subgraphs, the graph-level weight performs poorly, while our approach still performs very well which can recognize the node homophily well and adaptively provide suitable filters.

4.2 Theoretical Analysis

In addition to the empirical investigation, we theoretically confirm the capabilities of LocalSim node-level weight. For simplicity, we assume the number of inter-community and intra-community edges for all the nodes is the same as their expectation and without self-loop. Suppose that the optimal graph-level weight can provide a global filter that achieves the best average accuracy among all subgraphs (not the best in every subgraph), while the optimal node-level weight can provide the optimal filter for each subgraph which can achieve

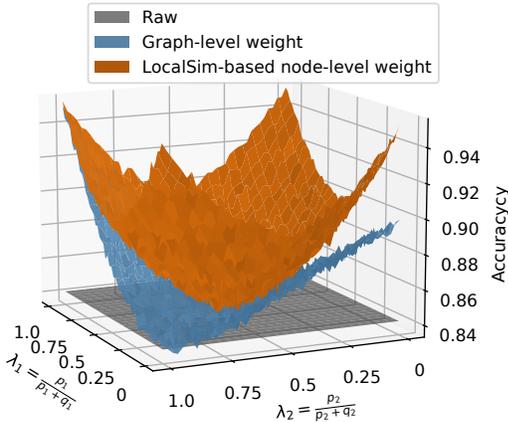


Figure 3: Accuracy (higher is better) on the synthetic graphs using graph-level weight and LocalSim-based node-level weight. Node homophily is measured by λ_i . ‘Raw’ shows the result directly using raw feature.

the best accuracy in every subgraph. However, to achieve the optimal node-level weight, we need to estimate the node homophily λ , and LocalSim is used for the aim. Next, we investigate the effectiveness of using LocalSim to indicate λ .

Theorem 1 (Effectiveness of LocalSim Indicating λ). *Consider FSBM with 2 equally-sized communities and 2 equally-sized subgraphs, and assume that the number of inter-community and intra-community edges for all the nodes is the same as their expectation and without self-loop. Consider that LocalSim is defined as $\phi(v_i) = -\frac{1}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} (x_i - x_j)^2$. For any node v_i that belongs to \mathcal{G}_τ for $\tau = 1, 2$, we have $\mathbb{E}[\phi(v_i)] = -2\sigma^2 - (1 - \lambda_\tau)(\mu_1 - \mu_2)^2$, where $\lambda_\tau = \frac{p_\tau}{p_\tau + q_\tau}$. Moreover, for any two nodes $v_i \in \mathcal{G}_1$ and $v_j \in \mathcal{G}_2$, the expectation of L1-norm distance between $\phi(v_i)$ and $\phi(v_j)$ is no less than to $|\lambda_1 - \lambda_2|(\mu_1 - \mu_2)^2$.*

We defer the proof to Appendix A. Theorem 1 states that $\mathbb{E}[\phi(v_i)]$ linearly correlates to λ_τ when $v_i \in \mathcal{G}_\tau$, which suggests that LocalSim can represent the homophily λ_τ of node v_i unbiasedly. We can also get that the expectation of L1-norm distance positively correlates to $|\lambda_1 - \lambda_2|$. This indicates that our estimation is likely to perform better when $|\lambda_1 - \lambda_2|$ goes up. When the gap between λ_1 and λ_2 shrinks, the two subgraphs become similar to each other, so does the optimal filter for each community. These results confirm the effectiveness of LocalSim indicating node homophily, which yields high performance in our empirical investigation.

5 Additional Related Work

In this section, we discuss some related work that attempts to address the heterophily issue. Mixhop [Abu-El-Haija *et al.*, 2019] proposes to extract features from multi-hop neighborhoods to get more information. CPGNN [Zhu *et al.*, 2020a] modifies feature propagation based on node classes in GCN to accommodate heterophily. Geom-GCN [Pei *et al.*, 2020] proposes to use the graph structure defined by geometric relationship to perform aggregation process to ad-

dress heterophily. FAGCN [Bo *et al.*, 2021] proposes to use the attention mechanism to learn the edge-level aggregation weights of low-pass and high-pass filters to adaptively address the heterophily issue, while we use LocalSim to parameterize the node-level weight which achieves less computation complexity and better performance. H₂GCN [Zhu *et al.*, 2020b] and GPRGNN [Chien *et al.*, 2021] propose fusing intermediate representation at graph-level based on their analysis of graphs. ACM-GCN [Luan *et al.*, 2022] divides feature propagation into three channels: low-pass, high-pass, and identity, and then adaptively mix the three channels during propagation but does not fuse different intermediate representations. Compared with H₂GCN, GPRGNN, and ACM-GCN, our work only needs propagation once which means the propagation can be pre-computed, while their works are trained via back-propagation through repeated feature propagation which yields huge computation costs. ASGC [Chandhuriya and Musco, 2022] directly uses the least square to filter the features obtained by SGC [Chen *et al.*, 2020] for classification, while our work involves multiple intermediate information, which is fused by node-level weight. These designs yield better performance based on pre-propagated features. Different from all these methods, our work uses LocalSim to parameterize the node-level weight for better adaptive fusion efficiently, and performs IRDC as an efficient propagation method for better informative representation, thus boosting the performance while keeping high efficiency. Moreover, we show that our designed LocalSim-based node-level weight can be used to improve the performance of H₂GCN and GPRGNN, which confirms the superiority of our method.

6 Experiment

6.1 Datasets and Setups

We evaluate LSGNN on nine small real-world benchmark datasets including both homophilic and heterophilic graphs. For homophilic graphs, we adopt three widely used citation networks, Cora, Citeseer, and Pubmed [Sen *et al.*, 2008; Namata *et al.*, 2012]. For heterophilic graphs, Chameleon and Squirrel are page-page Wikipedia networks [Rozemberczki *et al.*, 2021], Actor is a actor network [Pei *et al.*, 2020], and Cornell, Texas and Wisconsin are web pages networks [Pei *et al.*, 2020]. Statistics of these datasets can be seen in Appendix D.

For all datasets, we use the feature matrix, class labels, and 10 random splits (48%/32%/20%) provided by Pei *et al.* [2020]. Meanwhile, we use Optuna to tune the hyperparameter 200 times in experiments. For LSGNN, we set the number of layer $K = 5$.

There are three types of baseline methods, including non-graph models (MLP), traditional GNNs (GCN, GAT [Velickovic *et al.*, 2017], etc.), and GNNs tolerating heterophily (Geom-GCN, GPRGNN, etc.).

6.2 Evaluations on Real Benchmark Datasets

As shown in Table 1, LSGNN outperforms most of the baseline methods. On homophilic graphs, LSGNN achieves the SOTA performance on the Cora and Pubmed datasets and

	Cora	Citeseer	Pubmed	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin	Avg
² MLP	74.75	72.41	86.65	46.36	29.68	35.76	81.08	81.89	85.29	65.99
² GCN	87.28	76.68	87.38	59.82	36.89	30.26	57.03	59.46	59.80	61.62
² GAT	82.68	75.46	84.68	54.69	30.62	26.28	58.92	58.38	55.29	58.56
² MixHop	87.61	76.26	85.31	60.50	43.80	32.22	73.51	77.84	75.88	68.10
³ GCNII	<u>88.37</u>	77.33	<u>90.15</u>	63.86	38.47	37.44	77.86	77.57	80.39	70.16
¹ Geom-GCN	85.27	77.99	90.05	60.90	38.14	31.63	60.81	67.57	64.12	64.05
¹ H ₂ GCN	87.81	77.07	89.59	59.39	37.90	35.86	82.16	84.86	86.67	71.26
¹ WRGNN	88.20	76.81	88.52	65.24	48.85	36.53	81.62	83.62	86.98	72.93
³ GPRGNN	87.95	77.13	87.54	46.58	31.61	34.63	80.27	78.38	82.94	67.45
³ LINKX	84.64	73.19	87.86	68.42	<u>61.81</u>	36.10	77.84	74.60	75.49	71.11
³ ACM-GCN	87.91	77.32	90.00	66.93	54.40	36.28	85.14	<u>87.84</u>	<u>88.43</u>	74.92
³ GGCN	87.95	77.14	89.15	71.14	55.17	<u>37.54</u>	85.68	84.86	86.86	75.05
¹ GloGNN	88.33	<u>77.41</u>	89.62	<u>71.21</u>	57.88	37.70	<u>85.95</u>	84.32	88.04	<u>75.61</u>
LSGNN	88.49	76.71	90.23	79.04	72.81	36.18	88.92	90.27	90.20	79.21

Table 1: Results on real-world benchmark datasets: Mean accuracy (%). **Boldface** and underline numbers mark the top-1 and top-2 results, respectively. ¹ denotes results from their respective papers, ² from the H₂GCN paper, and ³ from the GloGNN paper.

	$\mathcal{H}(\mathcal{G})$	GCN	SGC	GCNII	LINKX	GloGNN	LSGNN
arXiv-year	0.22	46.02	32.83	47.21	56.00	54.79	56.42
ogbn-arXiv	0.66	71.74	69.39	72.74	54.45*	49.22*	72.64
Avg	NA	58.88	51.11	59.98	55.23	52.01	64.53

Table 2: Average accuracy (%) over 5 runs on large datasets. The field marked * denotes our implementation.

is lower than the best result by only 1.28% on the Citeseer dataset. On heterophilic graphs except Actor, LSGNN significantly outperforms all the other baseline models by 1.77%–11.00%. Specifically, the accuracy of LSGNN is 7.83% and 2.97% higher than GloGNN, the existing SOTA method on heterophilic graphs, on Chameleon and Cornell, is 11.00% higher than LINKX [Lim *et al.*, 2021] on Squirrel, and is 2.43% and 1.77% higher than ACM-GCN on Texas and Wisconsin. In particular, LSGNN achieves the best average accuracy, higher than the second-best by 3.60%. This demonstrates that LSGNN can achieve comparable or superior SOTA performance on both homophilic and heterophilic graphs.

6.3 Evaluations on Large Benchmark Datasets

We also conduct experiments on two large datasets: ogbn-arXiv (homophilic) and arXiv-year (heterophilic), both with 170k nodes and 1M edges but different labels, following settings in [Hu *et al.*, 2020] and [Lim *et al.*, 2021], see Appendix D for setting details. Table 2 shows that our method has competitive performance on both homophilic and heterophilic larger graphs.

6.4 Training Cost Comparison

LSGNN only needs to propagate once during training, and the networks for node-level weight are small. Thus the efficiency of LSGNN can be close to SGC’s. As shown in Figure 4(a), on the average training time over nine datasets, our proposed method is only slower than SGC, a simple and effective model, and is nearly 2× faster than GCN and GAT. This confirms the high training efficiency of LSGNN. See Appendix C for more setting details and detailed comparisons of training costs on each dataset.

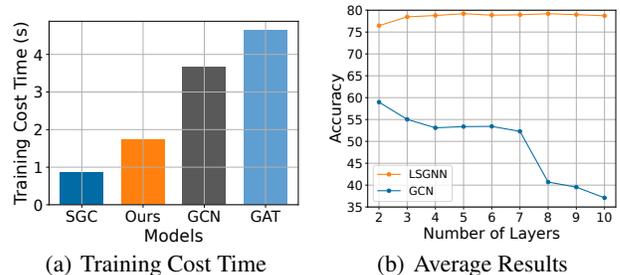


Figure 4: Comparison of the average Training Cost Time (s) on nine real-world datasets (Left). Average results with various model depth on nine real-world datasets (Right).

6.5 Alleviate Over-Smoothing Issue

To validate whether LSGNN can alleviate the over-smoothing issue, we compare the performance between vanilla GCN and LSGNN under different layers of propagation (Figure 4(b)), see Appendix E for full results. It can be seen that GCN achieves the best performance at 2 layers, and its performance decreases rapidly as layers increase. In contrast, the results of LSGNN keep stable and high. The reasons are two-fold: (i) our IRDC can extract better informative representations, and (ii) our LocalSim-based fusion can adaptively remain and drop distinguishable and indistinguishable representations, respectively. Through these two designs, LSGNN can perform well as layers increase, which indicates the capability of LSGNN to prevent the over-smoothing issue.

6.6 LocalSim-based Node-level Weight as a Plug-and-Play Module

LocalSim-based node-level weight (Section 3.4) can be regarded as a plug-and-play module, which can be added to existing GNNs that also fuse intermediate representations without involving extra hyperparameters, such as H₂GCN, GPRGNN, and DAGNN.

H₂GCN and GPRGNN use learned graph-level weight to fuse the intermediate representations, which is only replaced

	Cora	Citeseer	Pubmed	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin	Avg
H ₂ GCN	87.81	77.07	89.59	59.39	37.90	35.86	82.16	84.86	86.67	71.26
w/ LocalSim	88.42	76.77	89.91	64.91	41.73	36.29	85.68	85.68	88.82	73.13
GPRGNN	87.95	77.13	87.54	46.58	31.61	34.63	80.27	78.38	82.94	67.45
w/ LocalSim	88.83	76.60	90.26	65.39	51.96	35.59	85.14	87.84	87.65	74.36
DAGNN	82.63	75.47	88.50	58.03	39.59	30.20	61.89	58.92	55.10	61.15
w/ LocalSim	84.12	75.29	89.97	61.10	44.91	34.13	79.73	77.03	85.29	70.17

Table 3: LocalSim-based node-level weight as a plug-and-play module added to other models. **w/ LocalSim** denotes the LocalSim-based node-level weight is used, replacing the graph-level weight used in H₂GCN and GPRGNN, and replacing the random node-level weight (learned without any local topology information) in DAGNN. The best results are in **boldface**.

	Components			Dataset										
	Rand	LocalSim	w-self-loop	IRDC	Cora	Citeseer	Pubmed	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin	Avg
Baseline					86.33	74.53	89.99	67.28	47.73	35.27	86.47	88.92	88.04	73.84
	✓				87.93	75.32	90.08	67.72	48.99	35.65	78.38	88.51	88.43	73.45
		✓	✓		87.62	76.60	90.20	69.96	49.95	35.80	87.57	88.92	89.80	75.16
		✓	✓		87.65	76.89	89.98	75.95	65.50	35.80	87.84	88.11	89.80	77.50
LSGNN		✓	✓	✓	88.49	<u>76.71</u>	90.23	79.04	72.81	36.18	88.92	90.27	90.20	79.21

Table 4: Ablation study on 9 real-world benchmark datasets. Cells with ✓ mean that the corresponding components are applied to the baseline model. **Boldface** and underlined mark the top-1 and top-2 results. ‘Rand’ denotes random node-level weight (i.e., learned without any local topology information). ‘LocalSim’ means that the added node-level weight is learned under the guidance of LocalSim. ‘w-self-loop’ denotes weighted self-loops. ‘IRDC’ stands for the propagation via IRDC.

by LocalSim-based node-level weight here. Note that we only add LocalSim-based node-level weight into H₂GCN and GPRGNN, but not the enhanced filters or the high-pass filters. As shown in Table 3, the performance of both H₂GCN and GPRGNN increases significantly, indicating that taking each node’s local topology into consideration when fusing intermediate representations is more reasonable and effective than merely considering graph-level structure.

DAGNN also fuse intermediate representations by node-level weight, which however is simply mapped from each node’s representation and is learned without involving any local topology information, leading to a suboptimal result. We use LocalSim-based node-level weight in DAGNN similarly. The dramatic improving in Table 3 suggests that using LocalSim as local topology information can ensure the weight being learned more effectively. Although DAGNN is not designed for heterophilic graphs, it can also perform well on heterophilic graphs after using LocalSim-based node-level weight, and is even superior to H₂GCN on Squirrel.

6.7 Ablation Study

In what follows, we investigate the effectiveness of LocalSim-based node-level weight, weighted self-loops, and IRDC. Specifically, in the baseline model, we use graph-level weight learned without any local topology information to fuse the intermediate representations, add self-loops when conducting graph filtering, and apply the propagation method used in GCNs. Then, the above components are added one by one to the baseline model for performance comparison. As shown in Table 4, all the components are effective and can bring great benefits to the model in terms of performance improvement.

Note that learning node-level weight without any local topology information (‘Rand’) might even bring negative effect on the model (lower than baseline by 0.39% on average).

Once LocalSim is used as local topology information, the model can easily learn an optimal node-level weight (75.16% on average, higher than the baseline by 1.32%), which indicates the superiority of LocalSim-Aware Multi-Hop Fusion.

When replacing the fixed self-loops with weighted self-loops, the performance increases significantly (by 2.34% to 77.50% on average), especially on heterophilic graphs, such as Chameleon and Squirrel. This is because self-loops might not always be helpful on heterophilic graphs.

IRDC also brings a significant improvement (by an increase of 1.71% to 79.21% on average). This suggests that such a propagation method can extract more informative representation for better performance.

7 Conclusion

In this paper, LSGNN is proposed for better performance on both homophilic and heterophilic graphs. Many GNNs use graph-level weight to fuse intermediate representations, which does not fully consider the local structure of different nodes. Some GNNs use node-level weight learned without involving topology information, yielding suboptimal results. Our empirical study and theoretical analysis on synthetic graphs demonstrate the importance of node-level weight considering the local topology information of nodes. The proposed LocalSim-aware multi-hop fusion uses local similarity as guidance to generate a more appropriate node-level weight for intermediate representations fusion, and it can also be used as a plug-and-play module to improve the performance of existing GNNs. For better fusion, IRDC is proposed to extract more informative intermediate representations boosting the performance. Evaluations over real-world benchmark datasets show the superiority of LSGNN in handling both homophilic and heterophilic graphs and the effectiveness of all the proposed components.

Acknowledgements

This work is partially supported by the National Key R&D Program of China under Grant No. 2022YFB2703303, by the National Natural Science Foundation of China (NSFC) under Grant No. U22B2060, Grant No. 61972442, Grant No. 62102413, and Grant No. U2001202, by Guangzhou Municipal Science and Technology Bureau under Grant No. 2023A03J0667, by HKUST(GZ) under a Startup Grant, by the S&T Program of Hebei under Grant No. 20350802D, by the Natural Science Foundation of Hebei Province of China under Grant No. F2020202040, and by the Natural Science Foundation of Tianjin of China under Grant No. 20JCY-BJC00650.

Contribution Statement

Yuhan Chen and Yihong Luo are equally contributing to this work, as both of them largely contributed to the proposed method, code implementation, experiments, and the writing of the manuscript. Jing Tang significantly revised and polished the manuscript, improving the overall quality of the writing, and offered guidance and mentorship to Yihong throughout the research process. Liang Yang provided valuable feedback and suggestions for manuscript revision. Siya Qiu carried out a thorough grammar check. Both Chuan Wang and Xiaochun Cao offered guidance and mentorship to Yuhan throughout the research process.

References

- [Abu-El-Haija *et al.*, 2019] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *International Conference on Machine Learning*, pages 21–29. PMLR, 2019.
- [Battaglia *et al.*, 2018] 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.
- [Bo *et al.*, 2021] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. *arXiv preprint arXiv:2101.00797*, 2021.
- [Chanpuriya and Musco, 2022] Sudhanshu Chanpuriya and Cameron N Musco. Simplified graph convolution with heterophily. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [Chen *et al.*, 2020] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, pages 1725–1735. PMLR, 2020.
- [Chien *et al.*, 2021] Eli Chien, Jianhao Peng, Pan Li, and Olga Milenkovic. Adaptive universal generalized pagerank graph neural network. In *International Conference on Learning Representations*, 2021.
- [Defferrard *et al.*, 2016] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in Neural Information Processing Systems*, 29, 2016.
- [Errica *et al.*, 2020] Federico Errica, Marco Podda, Davide Bacciu, and Alessio Micheli. A fair comparison of graph neural networks for graph classification. In *International Conference on Learning Representations*. OpenReview.net, 2020.
- [Hamilton *et al.*, 2017] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. *arXiv*, abs/1706.02216, 2017.
- [Hamilton, 2020] William L Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- [He *et al.*, 2021] Mingguo He, Zhewei Wei, Hongteng Xu, et al. Bernnet: Learning arbitrary graph spectral filters via bernstein approximation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Hu *et al.*, 2020] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33:22118–22133, 2020.
- [Jiang *et al.*, 2013] Yuexin Jiang, Daniel I Bolnick, and Mark Kirkpatrick. Assortative mating in animals. *The American Naturalist*, 181(6):E125–E138, 2013.
- [Kipf and Welling, 2016] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv*, abs/1609.02907, 2016.
- [Li *et al.*, 2022] Xiang Li, Renyu Zhu, Yao Cheng, Caihua Shan, Siqiang Luo, Dongsheng Li, and Weining Qian. Finding global homophily in graph neural networks when meeting heterophily. In *International Conference on Machine Learning*, pages 13242–13256. PMLR, 2022.
- [Lim *et al.*, 2021] Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and Ser Nam Lim. Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods. *Advances in Neural Information Processing Systems*, 34:20887–20902, 2021.
- [Liu *et al.*, 2020] Meng Liu, Hongyang Gao, and Shuiwang Ji. Towards deeper graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 338–348, 2020.
- [Luan *et al.*, 2022] Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. Revisiting heterophily for

- graph neural networks. *arXiv preprint arXiv:2210.07606*, 2022.
- [McPherson *et al.*, 2001] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [Namata *et al.*, 2012] Galileo Namata, Ben London, Lise Getoor, Bert Huang, and U Edu. Query-driven active surveying for collective classification. In *10th international workshop on mining and learning with graphs*, volume 8, page 1, 2012.
- [Pei *et al.*, 2020] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020.
- [Rozemberczki *et al.*, 2021] Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-Scale Attributed Node Embedding. *Journal of Complex Networks*, 9(2), 2021.
- [Scarselli *et al.*, 2009] 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, 2009.
- [Sen *et al.*, 2008] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *Ai Magazine*, 2008.
- [Velickovic *et al.*, 2017] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv*, abs/1710.10903, 2017.
- [Wu *et al.*, 2019] Felix Wu, Tianyi Zhang, Amauri Holanda de Souza Jr, Christopher Fifty, Tao Yu, and Kilian Q Weinberger. Simplifying graph convolutional networks. *arXiv preprint arXiv:1902.07153*, 2019.
- [Zhang and Chen, 2018] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- [Zhu *et al.*, 2020a] Jiong Zhu, Ryan A Rossi, Anup Rao, Tung Mai, Nedim Lipka, Nesreen K Ahmed, and Danai Koutra. Graph neural networks with heterophily. *arXiv preprint arXiv:2009.13566*, 2020.
- [Zhu *et al.*, 2020b] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33, 2020.