# Non-Local Graph Neural Networks

Meng Liu\*, Zhengyang Wang\*, and Shuiwang Ji, Senior Member, IEEE

Abstract—Modern graph neural networks (GNNs) learn node embeddings through multilayer local aggregation and achieve great success in applications on assortative graphs. However, tasks on disassortative graphs usually require non-local aggregation. In addition, we find that local aggregation is even harmful for some disassortative graphs. In this work, we propose a simple yet effective non-local aggregation framework with an efficient attention-guided sorting for GNNs. Based on it, we develop various non-local GNNs. We perform thorough experiments to analyze disassortative graph datasets and evaluate our non-local GNNs. Experimental results demonstrate that our non-local GNNs significantly outperform previous state-of-the-art methods on seven benchmark datasets of disassortative graphs, in terms of both model performance and efficiency.

Index Terms—Graph neural networks, non-local aggregation, attention mechanism, disassortative graphs.

### 1 Introduction

RAPH neural networks (GNNs) process graphs and map each node to an embedding vector [1], [2], [3], [4], [5], [6], [7]. These node embeddings can be directly used for node-level applications, such as node classification [8], link prediction [9]. In addition, they can be used to learn the graph representation vector with graph pooling [10], [11], [12], [13], [14], [15], [16], in order to fit graph-level tasks [17]. Many variants of GNNs have been proposed, such as ChebNets [18], GCNs [8], GraphSAGE [19], GATs [20], LGCN [21] and GINs [22]. Their advantages have been shown on various graph datasets and tasks [23]. However, these GNNs share a multilayer local aggregation framework, which is similar to convolutional neural networks (CNNs) [24] on grid-like data such as images and texts.

The importance of non-local aggregation has been recently demonstrated in many applications in the field of computer vision [25], [26] and natural language processing [27]. In particular, the attention mechanism has been widely explored to achieve non-local aggregation and capture long-range dependencies from distant locations. Basically, the attention mechanism measures the similarity between every pair of locations and enables information to be communicated among distant but similar locations. In terms of graphs, non-local aggregation is also crucial for disassortative graphs in which nodes with the same label are distant from each other, while previous studies of GNNs focus on assortative graph datasets (Section 2.2). In addition, we find that local aggregation is even harmful for some disassortative graphs (Section 4.3). The recently proposed Geom-GCN [28] explores to capture long-range dependencies in disassortative graphs. It contains an attention-like step that computes the Euclidean distance between every pair of nodes. However, this step is computationally prohibitive

for large-scale graphs, as the computational complexity is quadratic in the number of nodes. In addition, Geom-GCN employs pre-trained node embeddings [29], [30], [31] that are not task-specific, limiting the effectiveness and flexibility.

In this work, we propose a simple yet effective non-local aggregation framework for GNNs. At the heart of the framework lies an efficient attention-guided sorting, which enables non-local aggregation through classic local aggregation operators in general deep learning. The proposed framework can be flexibly used to augment common GNNs with low computational costs. Based on the framework, we build various efficient non-local GNNs. In addition, we perform detailed analysis on existing disassortative graph datasets, and apply different non-local GNNs accordingly. Experimental results show that our non-local GNNs significantly outperform previous methods on node classification tasks on seven benchmark datasets of disassortative graphs.

#### 2 BACKGROUND AND RELATED WORK

#### 2.1 Graph Neural Networks

We focus on learning the embedding vector for each node through graph neural networks (GNNs). Most existing GNNs follow a local aggregation framework. In general, each layer of GNNs scans every node in the graph and aggregates local information from directly connected nodes, i.e., the 1-hop neighbors. Specifically, a common layer of GNNs performs a two-step processing similar to the depthwise separable convolution [32]: spatial aggregation and feature transformation. The first step updates each node embedding using embedding vectors of spatially neighboring nodes. For example, GCNs [8] and GATs [20] compute a weighted sum of node embeddings within the 1hop neighborhood, where weights come from the degree of nodes and the interaction between nodes, respectively. GraphSAGE [19] applies the max pooling, while GINs [22] simply sums the node embeddings. The feature transformation step is similar to the  $1 \times 1$  convolution, where each node embedding vector is mapped into a new feature space through a shared linear transformation [8], [19], [20] or multilayer perceptron (MLP) [22]. Different from these

Meng Liu and Shuiwang Ji are with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843, USA. (E-mail: sji@tamu.edu)

Zhengyang Wang was with the Department of Computer Science and Engineering, Texas A&M University, and is currently with Amazon.com Services I.I.C.

Meng Liu and Zhengyang Wang contribute equally to this work, and their names are in alphabetical order.

studies, LGCN [21] explores to directly apply the regular convolution through top-k ranking.

Nevertheless, each layer of these GNNs only aggregates local information within the 1-hop neighborhood. While stacking multiple layers can theoretically enable communication between nodes across the multi-hop neighborhood, the receptive field of a multiple-layer GCN usually includes more noise than useful information. In addition, deep GNNs usually suffer from the over-smoothing problem [33], [34], [35] and the over-squashing issue [36].

#### 2.2 Assortative and Disassortative Graphs

There are many kinds of graphs in the literature, such as citation networks [37], community networks [35], co-occurrence networks [38], and webpage linking networks [39]. We focus on graph datasets for the node classification tasks. In particular, we categorize graph datasets into assortative and disassortative ones [31], [40] according to the node homophily in terms of labels, *i.e.*, how likely nodes with the same label are near each other in the graph.

Assortative graphs refer to those with a high node homophily. Common assortative graph datasets are citation networks and community networks. On the other hand, graphs in disassortative graph datasets contain more nodes that have the same label but are distant from each other. Example disassortative graph datasets are co-occurrence networks and webpage linking networks.

As shown in [31], distant nodes with the same label in disassortative graphs could be structurally similar. Thus, they are informative to each other. Most existing GNNs perform local aggregation only and achieve good performance on assortative graphs [8], [19], [20], [21]. However, they may fail on disassortative graphs, where informative nodes in the same class tend to be out of the local multi-hop neighborhood and non-local aggregation is needed. Thus, in this work, we explore the non-local GNNs.

#### 2.3 Attention Mechanism

The attention mechanism [27] has been widely used in GNNs [20], [41], [42] as well as other deep learning models [25], [26], [43]. A typical attention mechanism takes three groups of vectors as inputs, namely the query vector q, key vectors  $(k_1, k_2, \ldots, k_n)$ , value vectors  $(v_1, v_2, \ldots, v_n)$ . Note that key and value vectors have a one-to-one correspondence and can be the same sometimes. The attention mechanism computes the output vector o as

$$o = \sum_i a_i v_i,$$
 
$$a_i = \text{Attend}(q, k_i) \in \mathbb{R}, \ i = 1, 2, \dots, n,$$
 (1)

where the  $Attend(\cdot)$  function could be any function that outputs a scalar attention score  $a_i$  from the interaction between q and  $k_i$ , such as dot product [41] or even a neural network [20]. The definition of the three groups of input vectors depends on the models and applications.

Notably, existing GNNs usually use the attention mechanism for local aggregation [20], [41]. Specifically, when aggregating information for node v, the query vector is the embedding vector of v while the key and value vectors come

from node embeddings of v's directly connected nodes. Note that the attention mechanism can be easily extended for non-local aggregation [25], [26], by letting the key and value vectors correspond to all the nodes in the graph when aggregating information for each node. However, it is computationally prohibitive given large-scale graphs, as iterating it for each node in a graph of n nodes requires  $O(n^2)$  time. In this work, we propose a novel non-local aggregation method that only requires  $O(n\log n)$  time.

#### 3 THE PROPOSED METHOD

## 3.1 Non-Local Aggregation with Attention-Guided Sorting

We consider a graph  $\mathcal{G}=(V,E)$ , where V is the set of nodes and E is the set of edges. Each edge  $e\in E$  connects two nodes so that  $E\subseteq V\times V$ . Each node  $v\in V$  has a node feature vector  $x_v\in\mathbb{R}^d$ . The k-hop neighborhood of v refers to the set of nodes  $\mathcal{N}_k(v)$  that can reach v within k edges. For example, the set of v's directly connected nodes is its 1-hop neighborhood  $\mathcal{N}_1(v)$ .

As illustrated in Figure 1, our proposed non-local aggregation framework is composed of three steps, namely local embedding, attention-guided sorting, and non-local aggregation. In the following, we describe them one by one.

**Local Embedding:** Our proposed framework is built upon a local embedding step that extracts local node embeddings from the node feature vectors. The local embedding step can be as simple as

$$z_v = \text{MLP}(x_v) \in \mathbb{R}^f, \ \forall v \in V.$$
 (2)

The MLP(·) function is a multilayer perceptron (MLP), and f is the dimension of the local node embedding  $z_v$ . Note that the MLP(·) function is shared across all the nodes in the graph. Applying MLP only takes the node itself into consideration without aggregating information from the neighborhood. This property is very important on some disassortative graphs, as shown in Section 4.3.

On the other hand, graph neural networks (GNNs) can be used as the local embedding step as well, so that our proposed framework can be easily employed to augment existing GNNs. As introduced in Section 2.1, modern GNNs perform multilayer local aggregation. Typically, for each node, one layer of a GNN aggregates information from its 1-hop neighborhood. Stacking L such local aggregation layers allows each node to access information that is L hops away. To be specific, for each node  $v \in V$ , the  $\ell$ -th layer of a L-layer GNN ( $\ell = 1, 2, \ldots, L$ ) can be described as

$$z_v^{(\ell)} = \mathbf{T}^{(\ell)} \left( \mathbf{A}^{(\ell)} \left( \{ z_u^{(\ell-1)} : u \in \mathcal{N}_1(v) \cup v \} \right) \right) \in \mathbb{R}^f, \quad (3)$$

where  $z_v^{(0)}=x_v$ , and  $z_v=z_v^{(L)}$  is the local node embedding.  $A^{(\ell)}(\cdot)$  and  $T^{(\ell)}(\cdot)$  functions represent the spatial aggregation and feature transformation step introduced in Section 2.1, respectively. With the above framework, GNNs can capture the node feature information from nodes within a local neighborhood as well as the structural information.

When either MLP or GNNs is used as the local embedding step, the local node embedding  $z_v$  only contains local information of a node v. However,  $z_v$  can be used to guide non-local aggregation, as distant but informative nodes are

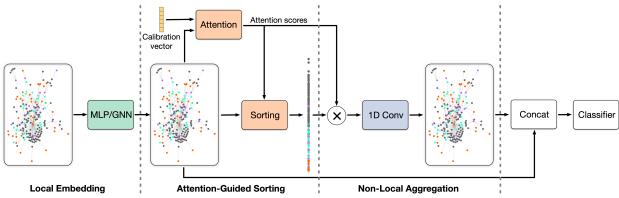


Fig. 1. An illustration of the proposed non-local aggregation framework. It consists of three steps, including local embedding, attention-guided sorting, and non-local aggregation. Details are described in Section 3.1.

likely to have similar node features and local structures. Based on this intuition, we propose the attention-guided sorting to enable the non-local aggregation.

**Attention-Guided Sorting:** The basic idea of the attention-guided sorting is to learn an ordering of nodes, where distant but informative nodes are put near each other. Specifically, given the local node embedding  $z_v$  obtained through the local embedding step, we compute one set of attention scores by

$$a_v = \text{Attend}(c, z_v) \in \mathbb{R}, \ \forall v \in V,$$
 (4)

where c is a calibration vector that is randomly initialized and jointly learned during training [43]. In this attention operator, c serves as the query vector and  $z_v$  are the key vectors. In addition, we also treat  $z_v$  as the value vectors. However, unlike the attention mechanism introduced in Section 2.3, we use the attention scores to sort the value vectors instead of computing a weighted sum to aggregating them. Note that originally there is no ordering among nodes in a graph. To be specific, as  $a_v$  and  $z_v$  have one-to-one correspondence through Eq. (4), sorting the attention scores in non-decreasing order into  $(a_1, a_2, \ldots, a_n)$  provides an ordering among nodes, where n = |V| is the number of nodes in the graph. The resulting sequence of local node embeddings can be denoted as  $(z_1, z_2, \ldots, z_n)$ .

The attention process in Eq. (4) can be also understood as a projection of local node embeddings onto a 1-dimensional space. The projection depends on the concrete ATTEND(·) function and the calibration vector c. As indicated by its name, the calibration vector c is used to calibrate the 1-dimensional space, in order to push distant but informative nodes close to each other in this space. This goal is fulfilled through the following non-local aggregation step and the training of the calibration vector c, as demonstrated below.

**Non-Local Aggregation:** We point out that, with the attention-guided sorting, the non-local aggregation can be achieved by convolution, the most common local aggregation operator in deep learning. Specifically, given the sorted sequence of local node embeddings  $(z_1, z_2, \ldots, z_n)$ , we compute

$$(\hat{z}_1, \hat{z}_2, \dots, \hat{z}_n) = \text{Conv}(z_1, z_2, \dots, z_n),$$
 (5)

where the  $CONV(\cdot)$  function represents a 1D convolution with appropriate padding. Note that  $CONV(\cdot)$  can be re-

placed by a 1D convolutional neural network as long as the number of input and output vectors remains the same.

To see how the  $CONV(\cdot)$  function performs non-local aggregation with the attention-guided sorting, we take an example where the  $CONV(\cdot)$  function is a 1D convolution of kernel size 2s + 1. In this case,  $\hat{z}_i$  is computed from  $(z_{i+s},\ldots,z_{i-s})$ , corresponding to the receptive field of the  $CONV(\cdot)$  function. As a result, if the attention-guided sorting leads to  $(z_{i+s}, \ldots, z_{i-s})$  containing nodes that are distant but informative to  $z_i$ , the output  $\hat{z}_i$  aggregates nonlocal information. Another view is that we can consider the attention-guided sorting as re-connects nodes in the graph, where  $(z_{i+s}, \ldots, z_{i-s})$  can be treated as the 1-hop neighborhood of  $z_i$ . After the CONV(·) function,  $\hat{z}_i$  and  $z_i$ are concatenated as the input to a classifier to predict the label of the corresponding node, where both non-local and local dependencies can be captured. In order to enable the end-to-end training of the calibration vector c, we modify Eq. (5) into

$$(\hat{z}_1, \hat{z}_2, \dots, \hat{z}_n) = \text{Conv}(a_1 z_1, a_2 z_2, \dots, a_n z_n),$$
 (6)

where we multiply the attention score with the corresponding local node embedding. As a result, the calibration vector  $\boldsymbol{c}$  receives back-propagated gradients through the attention scores during training.

The remaining question is how to make sure that the attention-guided sorting pushes distant but informative nodes together. We can understand this intuitively by comparing our attention-guided sorting with the attention mechanism that uses learnable query [43]. In such attention mechanism, attention scores obtained by attending learnable query vector to key vectors are used to compute the weighted sum of value vectors. In this case, the learnable query vector can be optimized towards its goal, i.e., assigning appropriate weights to value vectors, since the attention scores serve as weights directly. In our attention-guided sorting, the calibration vector is trained towards its goal, *i.e.*, deriving a good sorting, in an indirect manner. Specifically, the attention scores can be treated as unnormalized weights. At the beginning of training, the calibration vector yields random weights and a correspondingly random sorting. Thus the near nodes might not be informative. Notably, nodes in the same receptive field of the subsequent 1D convolutional neural network tends to have similar weights as all the nodes are already sorted according to the weights.

Hence, if the near nodes are not informative, the gradient back-propagated from the classification loss tends to tune the calibration vector to produce better weights and a correspondingly useful sorting. Hence, the calibration vector can be trained to obtain a good sorting effectively.

Also, note that the requirement of non-local aggregation depends on the concrete graphs. In fact, our proposed framework grants GNNs the ability of non-local aggregation but lets the end-to-end training process determine whether to use non-local information. The back-propagation from the supervised loss will tune the calibration vector c effectively and encourage  $\hat{z}_i$  to capture useful information that is not encoded by  $z_i$ . In the case of disassortative graphs,  $\hat{z}_i$ usually needs to aggregate information from distant but informative nodes. Hence, the calibration vector c tends to arrange the attention-guided sorting to put distant but informative nodes together, as demonstrated experimentally in Section 4.5. On the other hand, nodes within the local neighborhood are usually much more informative than distant nodes in assortative graphs. In this situation,  $\hat{z}_i$  may simply perform local aggregation that is similar to GNNs.

## 3.2 Time Complexity Analysis

We perform theoretical analysis of the time complexity of our proposed framework. As discussed in Section 2.3, using the attention mechanism [25], [26], [27] to achieve non-local aggregation requires  $O(n^2)$  time for a graph of n nodes. Essentially, the  $O(n^2)$  time complexity is due to the fact that the ATTEND(·) function needs to be computed between every pair of nodes. In particular, the recently proposed Geom-GCN [28] contains a similar non-local aggregation step. For each  $v \in V$ , Geom-GCN finds the set of nodes from which the Euclidean distance to v is less than a pre-defined number, where the Euclidean distance between every pair of nodes needs to be computed. As the computation of the the Euclidean distance between two nodes can be understood as the ATTEND(·) function, Geom-GCN has at least  $O(n^2)$  time complexity.

In contrast, our proposed non-local aggregation framework requires only  $O(n\log n)$  time. To see this, note that the ATTEND(·) function in Eq. (4) only needs to be computed once, instead of iterating it for each node. As a result, computing the attention scores only takes O(n) time. Therefore, the time complexity of sorting, *i.e.*  $O(n\log n)$ , dominates the total time complexity of our proposed framework. In Section 4.6, we compare the real running time on different datasets among common GNNs, Geom-GCN, and our non-local GNNs as introduced in the next section.

#### 3.3 Efficient Non-Local Graph Neural Networks

We apply our non-local aggregation framework to build efficient non-local GNNs. Recall that our proposed framework starts with the local embedding step, followed by the attention-guided sorting and the non-local aggregation step.

In particular, the local embedding step can be implemented by either MLP or common GNNs, such as GCNs [8] or GATs [20]. MLP extracts the local node embedding only from the node feature vector and excludes the information from nodes within the local neighborhood. This property can be helpful on some disassortative graphs, where nodes

within the local neighborhood provide more noises than useful information. On other disassortative graphs, informative nodes locate in both local neighborhood and distant locations. In this case, GNNs are more suitable as the local embedding step. Depending on the disassortative graphs in hand, we build different non-local GNNs with either MLP or GNNs as the local embedding step. In Section 4.3, we show that these two categories of disassortative graphs can be distinguished through simple experiments, where we apply different non-local GNNs accordingly. Specifically, the number of layers is set to 2 for both MLP and GNNs.

In terms of the attention-guided sorting, we need to specify the ATTEND( $\cdot$ ) function in Eq. (4). In order to make it as efficient as possible, we choose it as

$$a_v = \text{Attend}(c, z_v) = c^T z_v \in \mathbb{R}, \ \forall v \in V,$$
 (7)

where c is part of the training parameters.

With the attention-guided sorting, we can implement the non-local aggregation step through convolution, as explained in Section 3.1. Specifically,  $\mathsf{CONV}(\cdot)$  function is set as a 2-layer convolutional neural network composed of two 1D convolutions. The kernel size is treated as a hyperparameter. The activation function is ReLU [44].

Finally, we use a linear classifier that takes the concatenation of  $\hat{z}_i$  and  $z_i$  as inputs and makes prediction for the corresponding node. Depending on the local embedding step, we build three efficient non-local GNNs, namely non-local MLP (NLMLP), non-local GCN (NLGCN), and non-local GAT (NLGAT). The models can be end-to-end trained with the classification loss.

#### 4 EXPERIMENTS

#### 4.1 Datasets

We perform experiments on seven disassortative graph datasets [28], [38], [39], [45] (snap-patents, Chameleon, Squirrel, Actor, Cornell, Texas, Wisconsin) and four assortative graph datasets [37], [46] (Cora, Citeseer, Pubmed, ogbn-arxiv). These datasets are commonly used to evaluate GNNs on node classification tasks [8], [20], [21], [28]. We provide detailed descriptions of disassortative graph datasets below.

snap-patents is an extremely large patent graph [47] where nodes are patents and edges denote citation relationships. Node features are extracted from patent metadata. Each node is labeled by the year when the corresponding patent war granted.

Chameleon and Squirrel are Wikipedia networks [39] where nodes represent web pages from Wikipedia and edges indicate mutual links between pages. Node feature vectors are bag-of-word representations of informative nouns in the corresponding pages. Each node is labeled with one of five classes according to the number of average monthly traffic of the web page.

Actor is an actor co-occurrence network, where nodes denote actors and edges indicate co-occurrence on the same web page from Wikipedia. It is extracted from the film-director-actor-writer network in [38]. Node feature vectors are bag-of-word representations of keywords in the actors' Wikipedia pages. Each node is labeled with one of five classes according to the topic of the actor's Wikipedia page.

TABLE 1 Statistics of the eleven datasets used in our experiments. The definition of  $H(\mathcal{G})$  is provided in Section 4.1.  $H(\mathcal{G})$  can be used to distinguish assortative and disassortative graph datasets.

	Assortative										
Datasets	Cora	Citeseer	Pubmed	ogbn-arxiv	snap-patents	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin
H(G)	0.83	0.71	0.79	0.64	0.22	0.25	0.22	0.24	0.11	0.06	0.16
Splits	60/20/20	60/20/20	60/20/20	54/18/28	50/25/25	60/20/20	60/20/20	60/20/20	60/20/20	60/20/20	60/20/20
#Ñodes	2708	3327	19717	169343	2923922	2277	5201	7600	183	183	251
#Edges	5429	4732	44338	1166243	13975788	36101	217073	33544	295	309	499
#Features	1433	3703	500	128	269	2325	2089	931	1703	1703	1703
#Classes	7	6	3	40	5	5	5	5	5	5	5

Cornell, Texas, and Wisconsin come from the WebKB dataset collected by Carnegie Mellon University. Nodes represent web pages and edges denote hyperlinks between them. Node feature vectors are bag-of-word representations of the corresponding web pages. Each node is labeled as student, project, course, staff, or faculty.

In order to distinguish assortative and disassortative graph datasets, [28] proposes a metric to measure the homophily of a graph  $\mathcal{G}$ , defined as

$$H(\mathcal{G}) = \frac{1}{|V|} \sum_{v \in V} \frac{|\{u : u \in \mathcal{N}_1(v) \text{ and } l(u) = l(v)\}|}{|\mathcal{N}_1(v)|}, \quad (8)$$

where  $|\{u:u\in\mathcal{N}_1(v)\text{ and }l(u)=l(v)\}|$  denotes the number of v's directly connected nodes who have the same label as v and  $|\mathcal{N}_1(v)|$  represents the number of v's directly connected nodes. Intuitively, a large  $H(\mathcal{G})$  indicates an assortative graph, and vice versa. The  $H(\mathcal{G})$  and other statistics for all datasets are summarized in Table 1.

#### 4.2 Baselines

We compare our proposed non-local MLP (NLMLP), nonlocal GCN (NLGCN), and non-local GAT (NLGAT) with various baselines: (1) MLP is the simplest deep learning model. It makes prediction solely based on the node feature vectors, without aggregating any local or non-local information. (2) GCN [8] and GAT [20] are the most common GNNs. As introduced in Section 2.1, they only perform local aggregation. (3) Geom-GCN [28] is a recently proposed GNN that can capture long-range dependencies. Geom-GCN requires the use of different node embedding methods, such as Isomap [29], Poincare [30], and struc2vec [31]. We simply report the best results from [28] for Geom-GCN and the following two variants without specifying the node embedding method. (4) Geom-GCN-g [28] is a variant of Geom-GCN that performs local aggregation only. It is similar to common GNNs. (5) Geom-GCN-s [28] is a variant of Geom-GCN that does not force local aggregation. The designed functionality is similar to our NLMLP. (6) In addition, we also consider several recently developed methods on disassortative graphs as baselines. They are H<sub>2</sub>GCN [48], SimP-GCN [49], FAGCN [50], and CPGCN [51]. Since H<sub>2</sub>GCN and CPGCN has multiple variants, we report its best result on

We implement MLP, GCN, GAT, and our methods using Pytorch [52] and Pytorch Geometric [53]. As has been discussed<sup>1</sup>, in fair settings, the results of GCN and GAT differ from those in [28].

1. https://openreview.net/forum?id=S1e2agrFvS&noteId=8tGKV1oSzCr

The experiments on *snap-patents* are repeatedly conducted by 5 times, following [45], and experiments on other datasets are repeatedly performed 10 times. The average test accuracy over runs are reported unless stated specifically. Testing is performed when validation accuracy achieves maximum on each run. Apart from the details specified in Section 3.3, we tune the following hyperparameters on validation set for our proposed models: (1) the number of hidden unit  $\in \{16, 48, 96, 128, 256\}$ , (2) the kernel size of convolution  $\in \{3, 5, 7\}$ , (3) dropout rate  $\in \{0, 0.5, 0.8\}$ , (4) weight decay  $\in \{0, 5e-4, 5e-5, 5e-6\}$ , and (5) learning rate  $\in \{0.001, 0.01, 0.05\}$ .

#### 4.3 Analysis of Disassortative Graph Datasets

As discussed in Section 3.3, the disassortative graph datasets can be divided into two categories. Nodes within the local neighborhood provide more noises than useful information in disassortative graphs belonging to the first category. Therefore, local aggregation should be avoided in models on such disassortative graphs. As for the second category, informative nodes locate in both local neighborhood and distant locations. Intuitively, a graph with lower  $H(\mathcal{G})$  is more likely to be in the first category. However, it is not an accurate way to determine the two categories.

Knowing the exact category of a disassortative graph is crucial, as we need to apply non-local GNNs accordingly. As analyzed above, the key difference lies in whether the local aggregation is useful. Hence, we can distinguish two categories of disassortative graphs by comparing the performance between MLP and common GNNs (GCN, GAT) on each of the seven disassortative graph datasets. To be rigorous, we compare the average of the best validation accuracy over runs in Table 2.

We can see that *Actor*, *Cornell*, *Texas*, and *Wisconsin* fall into the first category, while *snap-patents*, *Chameleon*, and *Squirrel* belong to the second category. We add the performance on assortative graph datasets for reference, where the local aggregation is effective so that GNNs tend to outperform MLP.

#### 4.4 Comparisons with Baselines

According to the insights from Section 4.3, we apply different non-local GNNs according to the category of disassortative graph datasets, and make comparisons with corresponding baselines. In our experiments, we focus on comparing the model performance on disassortative graph datasets, in order to demonstrate the effectiveness of our non-local aggregation framework. The performances on assortative graph datasets are also included in Table 5 for

TABLE 2
Comparisons between MLP and common GNNs in terms of the best validation accuracy. These analytical experiments are used to determine the two categories of disassortative graph datasets, as introduced in Section 4.3.

	Assortative				Disassortative						
Datasets	Cora	Citeseer	Pubmed	ogbn-arxiv	snap-patents	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin
MLP	$77.2{\pm}1.6$	$74.4{\scriptstyle\pm1.4}$	$87.4 \pm 0.5$	$56.6 {\pm} 0.1$	$31.2 \pm 0.1$	$47.8{\scriptstyle\pm1.7}$	$33.3 {\pm} 0.8$	<b>35.9</b> ±0.9	<b>83.1</b> ±3.3	<b>81.5</b> ±3.9	<b>86.7</b> ±2.4
GCN	$88.8 \pm 1.0$	<b>78.1</b> ±1.3	$88.7 \pm 0.5$	<b>71.5</b> $\pm$ 0.1	$46.2 \pm 0.6$	<b>67.7</b> $\pm$ 1.3	$54.2 \pm 1.1$	$32.2 {\pm} 0.8$	$58.5 {\pm} 2.7$	$68.1 \pm 1.6$	$62.3{\pm}3.5$
GAT	<b>89.8</b> $\pm 1.4$	$77.6{\pm}1.3$	$87.7{\pm}0.4$	$66.5{\pm}2.5$	$45.6{\pm}0.5$	$66.4{\scriptstyle\pm1.3}$	$51.8{\pm}2.6$	$30.6{\pm}0.6$	$61.9{\pm}1.6$	$63.2{\pm}4.7$	$64.3{\pm}3.5$

TABLE 3

Comparisons between our NLMLP and strong baselines on the four disassortative graph datasets belonging to the first category as defined in Section 4.3.

Datasets	Actor	Cornell	Texas	Wisconsin
MLP	$35.1{\pm}0.8$	$81.6{\pm}6.3$	$81.3{\pm}7.1$	$84.9{\pm}5.3$
Geom-GCN Geom-GCN-s	31.6 34.6	60.8 75.4	67.6 73.5	64.1 80.4
H <sub>2</sub> GCN SimP-GCN	$35.9{\pm}1.0 \\ 36.2$	82.2±6.0 84.1	84.9±6.8 81.6	$86.7{\pm}4.7$ 85.5
NLMLP (ours)	<b>37.9</b> ±1.3	<b>84.9</b> ±5.7	<b>85.4</b> ±3.8	<b>87.3</b> ±4.3

TABLE 4
Comparisons between our NLGCN, NLGAT and strong baselines on the three disassortative graph datasets belonging to the second category as defined in Section 4.3.

Datasets	snap-patents	Chameleon	Squirrel
GCN GAT	$^{45.7\pm 0.0}_{45.4\pm 0.4}$	$67.6{\pm}2.4\\65.0{\pm}3.7$	$54.9{\pm}1.9 \\ 51.3{\pm}2.5$
Geom-GCN Geom-GCN-g	- -	60.9 68.0	38.1 46.0
H <sub>2</sub> GCN FAGCN	- -	$59.4{\pm}2.0$ $61.7$	$37.9{\pm}2.0\\39.7$
NLGCN (ours) NLGAT (ours)	$egin{array}{c} {\bf 50.8} {\pm 0.4} \ {48.3} {\pm 0.2} \end{array}$	$70.1 \pm 2.9$ $65.7 \pm 1.4$	$59.0 \pm 1.2$ $56.8 \pm 2.5$

reference, indicating that the proposed framework will not hurt the performance when non-local aggregation is not strongly desired.

Specifically, we employ NLMLP on Actor, Cornell, Texas, and Wisconsin. The corresponding baselines are MLP, Geom-GCN, and Geom-GCN-s, as Table 2 has shown that GCN and GAT perform much worse than MLP on these datasets. And Geom-GCN-g is similar to GCN and has worse performance than Geom-GCN-s, which is shown in Table 5. We also include H<sub>2</sub>GCN and SimP-GCN as baselines since they also conduct experiments on these datasets. The comparison results are reported in Table 3. We find that MLP consistently outperforms Geom-GCN-s by large margins. In particular, although Geom-GCN-s does not explicitly perform local aggregation, it is still outperformed by MLP. A possible explanation is that Geom-GCN-s uses pre-trained node embeddings, which aggregates information from the local neighborhood implicitly. In contrast, our NLMLP is built upon MLP with the proposed non-local aggregation framework, which excludes the local noises and collects useful information from non-local informative nodes. The NLMLP sets the new state-of-the-art performance on these disassortative graph datasets.

On snap-patents, Chameleon, and Squirrel that belong to

the second category of disassortative graph datasets, we apply NLGCN and NLGAT accordingly. The baselines are GCN, GAT, Geom-GCN, Geom-GCN-g, H<sub>2</sub>GCN, FAGCN, and CPGCN. On these datasets, these baselines that explicitly perform local aggregation show advantages over MLP and Geom-GCN-s, as shown in Table 2 and 5. As shown in Table 4, our proposed NLGCN achieves the best performance on all three datasets. Note that CPGCN uses different splits as other baselines. For comparing with CPGCN, we additionally run our NLGCN using the same splits as CPGCN. On Chameleon and Squirrel, our NLGCN achieves better test accuracies of 41.3 and 60.0 respectively, compared to 37.0 and 56.9 of CPGCN. In addition, it is worth noting that our NLGCN and NLGAT are built upon GCN and GAT, respectively. They show improvements over their counterparts, which indicates that the advantages of our proposed non-local aggregation framework are general for common GNNs.

We summarize the results on all datasets in Table 5 for reference.

#### 4.5 Analysis of the Attention-Guided Sorting

We analyze the results of the attention-guided sorting in our proposed framework, in order to show that our non-local GNNs indeed perform non-local aggregation.

Suppose the attention-guided sorting leads to the sorted sequence  $(z_1, z_2, \ldots, z_n)$ , which goes through a convolution or CNN into  $(\hat{z}_1, \hat{z}_2, \ldots, \hat{z}_n)$ . As discussed in Section 3.1, we can consider the sequence  $(z_1, z_2, \ldots, z_n)$  as a re-connected graph  $\hat{\mathcal{G}}$ , where we treat nodes within the receptive field of  $\hat{z}_i$  as directly connected to  $z_i$ , i.e.  $z_i$ 's 1-hop neighborhood. The information within this new 1-hop neighborhood will be aggregated. If our non-local GNNs indeed perform non-local aggregation, the homophily of the re-connected graph should be larger than the original graph. Therefore, we compute  $H(\hat{\mathcal{G}})$  for several datasets to verify this statement. Following Section 4.4, we apply NLMLP on *Actor*, *Cornell*, *Texas*, and *Wisconsin* and NLGCN on *Chameleon* and *Squirrel*.

Figure 2 compares  $H(\mathcal{G})$  with  $H(\mathcal{G})$ . We can observe that  $H(\hat{\mathcal{G}})$  is much larger than  $H(\mathcal{G})$ , indicating that distant but informative nodes are near each other in the re-connected graph  $\hat{\mathcal{G}}$ . We also provide the visualizations of the sorted sequence for *Cornell* and *Texas*. We can see that nodes with the same label tend to be clustered together. These facts indicate that our non-local GNNs indeed perform non-local aggregation with the attention-guided sorting.

#### 4.6 Efficiency Comparisons

As analyzed in Section 3.2, our proposed non-local aggregation framework is more efficient than previous methods

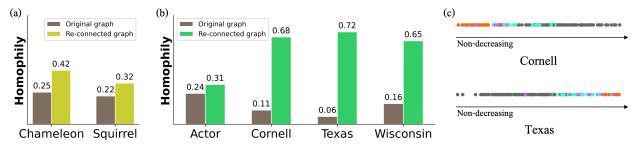


Fig. 2. (a) Comparisons of the homophily between the original graph and the re-connected graph given by our NLGCN on *Chameleon* and *Squirrel*. (b) Comparisons of the homophily between the original graph and the re-connected graph given by our NLMLP on *Actor*, *Cornell*, *Texas*, and *Wisconsin*. (c) Visualization of sorted node sequence after the attention-guided sorting for *Cornell* and *Texas*. The colors denote node labels. Details are explained in Section 4.5.

TABLE 5
Comparisons between our NLMLP, NLGCN, NLGAT and baselines on all the 11 datasets. †We can only use smaller hidden dimensions for GAT and NLGAT on *ogbn-arxiv* and *snap-patents* due to the huge memory requirement. ‡Uses a GPU with 48 GB of memory.

	Assortative			Disassortative						****	
Datasets	Cora	Citeseer	Pubmed	ogbn-arxiv	snap-patents	Chameleon	Squirrel	Actor	Cornell	Texas	Wisconsin
MLP	$76.5{\scriptstyle\pm1.3}$	$73.6{\scriptstyle\pm1.9}$	$87.5{\pm}0.4$	$54.1{\scriptstyle\pm0.1}$	$31.3 \pm 0.1$	$48.5{\pm}3.0$	$31.5{\pm}1.4$	$35.1{\pm}0.8$	$81.6{\pm}6.3$	$81.3{\pm}7.1$	$84.9 {\pm} 5.3$
GCN	$88.2 \pm 1.2$	$75.7{\pm}1.3$	$88.4 \pm 0.6$	$70.4 \pm 0.2$	$45.7 \pm 0.0$	$67.6 \pm 2.4$	$54.9 \pm 1.9$	$30.3 \pm 1.6$	$54.2 {\pm} 7.3$	$61.1 \pm 7.0$	$59.6{\pm}4.5$
GAT	$88.4{\pm}1.4$	$76.1{\scriptstyle\pm1.0}$	$87.0{\pm}0.3$	$65.0{\pm}2.9^{\dagger}$	$45.4 {\pm} 0.4^{\dagger}$	$65.0{\pm}3.7$	$51.3{\pm}2.5$	$29.4{\pm}1.2$	$56.3{\pm}4.3$	$57.9{\pm}6.1$	$57.8 \pm 4.3$
Geom-GCN	85.3	78.0	90.1	-	-	60.9	38.1	31.6	60.8	67.6	64.1
Geom-GCN-g	87.0	80.6	90.7	-	-	68.0	46.0	32.0	67.0	73.1	69.4
Geom-GCN-s	73.3	72.2	87.0	-	-	61.6	38.0	34.6	75.4	73.5	80.4
NLMLP (ours)	$76.9 \pm 1.8$	$73.4 \pm 1.9$	$88.2 {\pm} 0.5$	$54.0 {\pm} 0.3$	$32.1\pm0.1^{\ddagger}$	$50.7{\pm}2.2$	$33.7{\pm}1.5$	<b>37.9</b> ±1.3	<b>84.9</b> ±5.7	<b>85.4</b> ±3.8	<b>87.3</b> ±4.3
NLGCN (ours)	$88.1 \pm 1.0$	$75.2{\pm}1.4$	$89.0{\pm}0.5$	<b>70.6</b> $\pm$ 0.3	<b>50.8</b> ±0.4 <sup>‡</sup>	$70.1 \pm 2.9$	<b>59.0</b> $\pm$ 1.2	$31.6{\scriptstyle\pm1.0}$	$57.6{\pm}5.5$	$65.5{\pm}6.6$	$60.2 \pm 5.3$
NLGAT (ours)	$\textbf{88.5} {\pm} 1.8$	$76.2{\pm}1.6$	$88.2{\pm}0.3$	$66.7{\pm}2.4^{\dagger}$	$48.3\pm0.2^{\dagger\ddagger}$	$65.7{\pm}1.4$	$56.8{\pm}2.5$	$29.5{\pm}1.3$	$54.7{\pm}7.6$	$62.6{\pm}7.1$	$56.9{\pm}7.3$

TABLE 6
Comparisons in terms of real running time (*milliseconds*). <sup>†</sup>We have memory issue if we use the same dimension for GAT as GCN and NLGCN on *snap-patents*.

	snap-patents	Chameleon	Squirrel
GCN	2193.0 (1.0×)	22.2 (1.0×)	$14.3 (1.0 \times)$
GAT	$OOM^{\dagger}$	$33.2 (1.5 \times)$	$163.3 (11.4 \times)$
Geom-GCN	-	$3615.0 \ (163.1 \times)$	$10430.0 \ (727.3 \times)$
NLGCN (ours)	2437.4 (1.1×)	26.3 (1.2×)	39.6 (2.8×)

based on the original attention mechanism, such as Geom-GCN [28]. Concretely, our method requires only  $O(n \log n)$  computation time in contrast to  $O(n^2)$ . In this section, we compare the real running time to verify our analysis. Specifically, we compare NLGCN with Geom-GCN as well as GCN and GAT. For Geom-GCN, we use the code provided in [28]. Each model is trained for 500 epochs on each dataset and the average training time per epoch is reported.

The results are shown in Table 6. Although our NLGCN is built upon GCN, it is just slightly slower than GCN and faster than GAT, showing the efficiency of our non-local aggregation framework. On the other hand, Geom-GCN is significantly slower due to the fact that it has  $O(n^2)$  time complexity.

#### 5 CONCLUSION

In this work, we propose a simple yet effective non-local aggregation framework for GNNs. The core of the framework is an efficient attention-guided sorting, which enables non-local aggregation through convolution. The proposed framework can be easily used to build non-local GNNs with low computational costs. We perform thorough experiments on node classification tasks to evaluate our pro-

posed method. In particular, we experimentally analyze existing disassortative graph datasets and apply different non-local GNNs accordingly. The results show that our non-local GNNs significantly outperform previous state-of-theart methods on all benchmark datasets of disassortative graphs, in terms of both accuracy and speed.

#### **ACKNOWLEDGMENTS**

This work was supported in part by National Science Foundation grants IIS-1908198 and DBI-1922969.

#### REFERENCES

- [1] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017
- [2] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proceedings of the 34th international conference on machine learning*, 2017, pp. 1263–1272.
- [3] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner et al., "Relational inductive biases, deep learning, and graph networks," arXiv preprint arXiv:1806.01261, 2018.
- [4] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," arXiv preprint arXiv:1812.08434, 2018.
- [5] S. Zhang, H. Tong, J. Xu, and R. Maciejewski, "Graph convolutional networks: Algorithms, applications and open challenges," in *International Conference on Computational Social Networks*. Springer, 2018, pp. 79–91.
- [6] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," arXiv preprint arXiv:1901.00596, 2019.
- [7] Z. Zhang, P. Cui, and W. Zhu, "Deep learning on graphs: A survey," IEEE Transactions on Knowledge and Data Engineering, 2020.

- [8] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations*, 2017.
- [9] M. Zhang and Y. Chen, "Link prediction based on graph neural networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 5165–5175.
- [10] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, and J. Leskovec, "Hierarchical graph representation learning with differentiable pooling," in *Advances in neural information processing systems*, 2018, pp. 4800–4810.
- [11] M. Zhang, Z. Cui, M. Neumann, and Y. Chen, "An end-to-end deep learning architecture for graph classification," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [12] H. Gao and S. Ji, "Graph U-Nets," in International Conference on Machine Learning, 2019, pp. 2083–2092.
- [13] Y. Ma, S. Wang, C. C. Aggarwal, and J. Tang, "Graph convolutional networks with eigenpooling," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 723–731.
- [14] J. Lee, I. Lee, and J. Kang, "Self-attention graph pooling," in International Conference on Machine Learning, 2019, pp. 3734–3743.
- [15] H. Yuan and S. Ji, "Structpool: Structured graph pooling via conditional random fields," in *International Conference on Learning Representations*, 2020.
- [16] F. M. Bianchi, D. Grattarola, and C. Alippi, "Spectral clustering with graph neural networks for graph pooling," in *International Conference on Machine Learning*. PMLR, 2020, pp. 874–883.
- [17] P. Yanardag and S. Vishwanathan, "Deep graph kernels," in Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, 2015, pp. 1365–1374.
- [18] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing Systems*, 2016, pp. 3844–3852.
- [19] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in Advances in Neural Information Processing Systems, 2017, pp. 1024–1034.
- [20] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," in *International Conference* on Learning Representation, 2018.
- [21] H. Gao, Z. Wang, and S. Ji, "Large-scale learnable graph convolutional networks," in Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018, pp. 1416–1424.
- [22] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *International Conference on Learning Representations*, 2019.
- [23] F. Errica, M. Podda, D. Bacciu, and A. Micheli, "A fair comparison of graph neural networks for graph classification," in *International Conference on Learning Representations*, 2020.
- [24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [25] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proceedings of the IEEE conference on computer vision* and pattern recognition, 2018, pp. 7794–7803.
- [26] Z. Wang, N. Zou, D. Shen, and S. Ji, "Non-local U-Nets for biomedical image segmentation," in *Thirty-Fourth AAAI Conference* on Artificial Intelligence, 2020.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [28] H. Pei, B. Wei, K. C.-C. Chang, Y. Lei, and B. Yang, "Geom-GCN: Geometric graph convolutional networks," in *International Conference on Learning Representations*, 2020.
- [29] J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," science, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [30] M. Nickel and D. Kiela, "Poincaré embeddings for learning hierarchical representations," in Advances in Neural Information Processing Systems, 2017, pp. 6338–6347.
- [31] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo, "struc2vec: Learning node representations from structural identity," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2017, pp. 385–394.

- [32] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1251–1258.
- [33] K. Xu, C. Li, Y. Tian, T. Sonobe, K.-i. Kawarabayashi, and S. Jegelka, "Representation learning on graphs with jumping knowledge networks," in *International Conference on Machine Learning*, 2018, pp. 5449–5458.
- [34] Q. Li, Z. Han, and X.-M. Wu, "Deeper insights into graph convolutional networks for semi-supervised learning," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [35] D. Chen, Y. Lin, W. Li, P. Li, J. Zhou, and X. Sun, "Measuring and relieving the over-smoothing problem for graph neural networks from the topological view," in *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [36] U. Alon and E. Yahav, "On the bottleneck of graph neural networks and its practical implications," arXiv preprint arXiv:2006.05205, 2020.
- [37] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad, "Collective classification in network data," AI magazine, vol. 29, no. 3, pp. 93–93, 2008.
- [38] J. Tang, J. Sun, C. Wang, and Z. Yang, "Social influence analysis in large-scale networks," in *Proceedings of the 15th ACM SIGKDD* international conference on Knowledge discovery and data mining, 2009, pp. 807–816.
- [39] B. Rozemberczki, C. Allen, and R. Sarkar, "Multi-scale attributed node embedding," arXiv preprint arXiv:1909.13021, 2019.
  [40] M. E. Newman, "Assortative mixing in networks," Physical review
- [40] M. E. Newman, "Assortative mixing in networks," Physical review letters, vol. 89, no. 20, p. 208701, 2002.
- [41] H. Gao and S. Ji, "Graph representation learning via hard and channel-wise attention networks," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 741–749.
- [42] B. Knyazev, G. W. Taylor, and M. Amer, "Understanding attention and generalization in graph neural networks," in Advances in Neural Information Processing Systems, 2019, pp. 4204–4214.
- [43] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy, "Hierarchical attention networks for document classification," in Proceedings of the 2016 Conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies, 2016, pp. 1480–1489.
- [44] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.
  [45] D. Lim, X. Li, F. Hohne, and S.-N. Lim, "New bench-
- [45] D. Lim, X. Li, F. Hohne, and S.-N. Lim, "New benchmarks for learning on non-homophilous graphs," arXiv preprint arXiv:2104.01404, 2021.
- [46] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec, "Open graph benchmark: Datasets for machine learning on graphs," Neural Information Processing Systems (NeurIPS), 2020.
- [47] J. Leskovec and A. Krevl, "Snap datasets: Stanford large network dataset collection," 2014.
- [48] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, "Beyond homophily in graph neural networks: Current limitations and effective designs," *Advances in Neural Information Processing* Systems, vol. 33, 2020.
- [49] W. Jin, T. Derr, Y. Wang, Y. Ma, Z. Liu, and J. Tang, "Node similarity preserving graph convolutional networks," in *Proceedings of the* 14th ACM International Conference on Web Search and Data Mining, 2021, pp. 148–156.
- [50] D. Bo, X. Wang, C. Shi, and H. Shen, "Beyond low-frequency information in graph convolutional networks," in *Proceedings of* the AAAI Conference on Artificial Intelligence, vol. 35, no. 5, 2021, pp. 3950–3957.
- [51] J. Zhu, R. A. Rossi, A. Rao, T. Mai, N. Lipka, N. K. Ahmed, and D. Koutra, "Graph neural networks with heterophily," in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, no. 12, 2021, pp. 11 168–11 176.
- [52] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [53] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in ICLR Workshop on Representation Learning on Graphs and Manifolds, 2019.