

Marginal Nodes Matter: Towards Structure Fairness in Graphs

Xiaotian Han¹ Kaixiong Zhou² Ting-Hsiang Wang¹ Jundong Li³ Fei Wang⁴ Na Zou¹

¹Texas A&M University, ²Rice University, ³University of Virginia, ⁴Weill Cornell Medicine
{han, thwang1231, nzou1}@tamu.edu, jundong@virginia.edu, few2001@med.cornell.edu

ABSTRACT

In social network, a person located at the periphery region (marginal node) is likely to be treated unfairly when compared with the persons at the center. While existing fairness works on graphs mainly focus on protecting sensitive attributes (e.g., age and gender), the fairness incurred by the graph structure should also be given attention. On the other hand, the information aggregation mechanism of graph neural networks amplifies such structure unfairness, as marginal nodes are often far away from other nodes. In this paper, we focus on novel fairness incurred by the graph structure on graph neural networks, named *structure fairness*. Specifically, we first analyzed multiple graphs and observed that marginal nodes in graphs have a worse performance of downstream tasks than others in graph neural networks. Motivated by the observation, we propose **Structural Fair Graph Neural Network (SFairGNN)**, which combines neighborhood expansion based structure debiasing with hop-aware attentive information aggregation to achieve structure fairness. Our experiments show SFairGNN can significantly improve structure fairness while maintaining overall performance in the downstream tasks.

1. INTRODUCTION

Recent years have witnessed a surge of research interests in graph machine learning for different applications, such as social networks [26, 55], protein-protein interaction networks [3], knowledge graphs [20, 45, 73], and Spatio-Temporal forecasting [14]. Among them, graph neural networks (GNNs) [63, 64, 72, 74], as a family of deep learning based algorithms, have become an essential paradigm for graph analysis due to their superior modeling performance. Despite their success, it has been shown that these algorithms often suffer from fairness issues for decision-making [16, 18, 23, 24, 42, 48, 61, 70]. To mitigate this issue, a vast majority of existing works on graph fairness [11, 24, 42, 48, 54] focus on obtaining fair node representations that are invariant to the change of the protected attributes such as age, gender, and ethnicity.

Complementing existing efforts on learning fair representations regarding protected attributes [11, 24, 42, 48, 54], we investigate a novel problem on the fairness issue of GNNs from the *structure* perspective. In particular, graph structure fairness should be paid more attention to because of the following key reasons: 1) Marginal nodes in the graph should not be automatically depreciated by graph learning

algorithms. For instance, in the Twitter network, a tweet posted by a user with few followers may receive limited attention, although the tweet could be insightful and valuable for the whole community. As another example, in a job market network such as the LinkedIn network, people should not be neglected just because they are at the periphery of the social network, a situation commonly faced by individuals from underprivileged communities. 2) As one of the most popular methods in graph analysis, GNNs exacerbate structure unfairness because their built-in information aggregation mechanism exploits the egocentric structure of nodes. In other words, each node only receives information from nearby nodes in information aggregation.

Despite its importance and necessity, understanding and achieving structure fairness on GNNs is non-trivial due to the following two challenges. First, a proper measurement of structure fairness is missing, although a quantitative measurement is essential to detect and compensate for *structurally* underprivileged nodes. Second, given that GNNs use an information aggregation mechanism to update node representation from neighbors, but the information aggregation mechanism implicitly causes unfair decision-making, it remains a challenging problem to design an effective graph neural network framework to mitigate the unfairness problem of information aggregation.

To overcome the challenges mentioned above, we first validate that existing graph neural networks suffer from structure fairness problems and then propose to address them accordingly. We first conduct an experiment to validate that structure fairness is highly correlated to the performance of downstream tasks in existing GNNs. Since the information aggregation mechanism of GNNs essentially enriches the node information by aggregating information from their neighbors, centrality [10, 43] is an intuitive indicator of graph structure, which contains the information between nodes and their neighbors. To this end, we adopt two centrality measurements (closeness centrality [36] and eigenvector centrality [9]) to identify the marginal nodes.

To address structure fairness in graph neural networks, we propose a simple-yet-effective **Structural Fair Graph Neural Network (SFairGNN)**, which mitigates the structure unfairness issue of GNNs without sacrificing their performance on downstream tasks. Specifically, SFairGNN has two essential components: 1) neighborhood expansion and 2) hop-aware attentive information aggregation. Neighborhood expansion compensates marginal nodes by introducing new neighbors to them during the training phase. On the other hand, hop-aware attentive information aggregation optimizes how

the new information should be properly and effectively integrated into each node based on downstream tasks. Our investigation and proposed model will help researchers and society understand how to identify and mitigate structure unfairness in graph-based decision-making. In summary, the **contributions** of this work are as follows:

- We conducted an analysis of the performance of downstream tasks on marginal nodes and found that they were significantly disadvantaged, leading to unfair decision-making. Our investigation led us to the initial observation that the graph structure was responsible for the unfairness suffered by these nodes.
- To address the issue of unfairness caused by the graph structure, we proposed a new graph neural network model called SFairGNN. Our proposed model is specifically designed to mitigate the impact of structural unfairness in graph neural networks.
- To evaluate the performance of SFairGNN, we conducted experiments on real-world datasets. Our results showed that SFairGNN outperformed state-of-the-art GNN models in terms of fairness, without sacrificing performance. Our experiments demonstrate the effectiveness of SFairGNN in reducing unfairness and improving overall performance in graph-based machine learning models.

The rest of the paper is organized as follows. In [Section 3](#), we provide an overview of the necessary background for SFairGNN, which includes an introduction to graph neural networks and metrics for measuring the structure of graphs. [Section 3](#) delves into the issue of fairness in graph structures. Next, in [Section 4](#), we present our proposed fairness-aware graph neural network, SFairGNN. We discuss the experimental findings and analysis in [Section 5](#). Finally, in [Sections 6](#) and [7](#), we discuss the related work and provide concluding remarks and suggestions for future research.

2. PRELIMINARIES

In this section, we introduce the essential preliminaries for our proposed methods, including graph neural networks, as well as two centrality metrics of a graph: closeness centrality and eigenvector centrality.

2.1 Graph Neural Networks

Most graph neural networks (GNNs) [[30,39,59,62](#)] adopt a message-passing mechanism to update node representations over the graph iteratively. Without loss of generality, we present a general framework of GNNs, which can be divided into two steps: aggregation and transformation. Generally, at the k -th layer in GNNs, the hidden representation $\mathbf{x}_i^{(k)}$ of node n_i is updated as follows:

$$\begin{aligned} \mathbf{h}_i^{(k)} &= \text{AGGREGATE}(\{a_{ij}^{(k)} \mathbf{W}^{(k)} \mathbf{x}_j^{(k-1)} : j \in \mathcal{N}_i\}), \\ \mathbf{x}_i^{(k)} &= \text{ACT}(\text{COMBINE}(\mathbf{x}_i^{(k-1)}, \mathbf{h}_i^{(k)})), \end{aligned}$$

where $\mathbf{W}^{(k)}$ denotes the trainable weight, and $a_{ij}^{(k)}$ denotes the edge weight between nodes i and j , and \mathcal{N}_i denotes the set of neighbors adjacent to node i . $\text{AGGREGATE}(\cdot)$ is an aggregation function, which is applied to aggregate the node representations of all neighbors. Finally, $\text{COMBINE}(\cdot)$ and $\text{ACT}(\cdot)$ are applied to combine neighbor information and activate node representations, respectively.

2.2 Indicators of Structure

To account for the impact of graph structure toward fair decision-making in GNNs, we first need to identify the marginal nodes in the graph. To reflect node sociological origin [[36](#)] and the graph structure at the node level, we use two important centralities (i.e., closeness centrality and eigenvector centrality) to quantify the structure of each node. Closeness centrality [[36](#)] of a node measures its average nearness (inverse distance) to all other nodes in the graph. Eigenvector centrality [[9](#)] of a node measures the influence it has on all other nodes in the graph.

Closeness Centrality Closeness centrality [[5,49](#)] of a node is the reciprocal of the sum of the shortest path distances from the node to all other nodes on the graph. The nodes with a high closeness centrality score have the shortest distances to all other nodes, which easily access other nodes and influence other nodes. For instance, in a social network, a person with a lower mean distance from others might find that their opinions reach others more quickly than the opinions of someone with a higher mean distance. Formally, the mathematical expression of closeness centrality of node n_i is:

$$c(n_i) = \frac{N-1}{\sum_{n_j} d(n_i, n_j)}, \quad (1)$$

where $d(n_i, n_j)$ is the shortest-path distance between nodes i and j , N is the number of nodes in the graph.

Eigenvector Centrality Eigenvector centrality is a measure of the influence of each node in a graph [[9](#)]. A high eigenvector centrality score for a node indicates that it is connected to many other nodes with high scores. The eigenvector centrality score for node i is calculated as the i -th element of the vector \mathbf{x} , which is defined by the equation $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$, where \mathbf{A} is the adjacency matrix of the graph and λ is the corresponding eigenvalue. According to the Perron-Frobenius theorem [[47](#)], if λ is the largest eigenvalue of \mathbf{A} , then there exists a unique solution \mathbf{x} , all of whose entries are positive. [[29](#)] Throughout this paper, we use the indicator variable s_i to represent the centrality metric of node i . Specifically, s_i can take either closeness centrality or eigenvector centrality, depending on the context.

3. PROBLEM FORMULATION

In the following subsections, we introduce and analyze the structure fairness in graph neural networks and propose how to measure the structure fairness.

3.1 Preliminary Experiments

To explore the structure unfairness in GNNs, we perform node classification task with representative GNNs, i.e., graph convolutional networks (GCN) [[35](#)] and graph attention networks (GAT) [[59](#)]. Specially, we cluster nodes in the graph into multiple bins based on their indicators of structure (i.e., closeness centrality and eigenvector centrality). Then we calculate the average accuracy of each bin and compare the accuracy difference between bins.

Experimental Results. We present the experimental results on CoraFull and PubMed dataset with closeness and eigenvector centrality in [Figure 1](#). The X-axis represents different levels of closeness/eigenvector centrality, while Y-axis represents the prediction accuracy.

Results Analysis. The lower classification accuracy of marginal nodes shows that they are treated unfairly by GCN




Figure 1: The distribution of node classification accuracy with respect to closeness/eigenvector centrality on CoraFull and PubMed datasets. The X-axis indicates closeness/eigenvector centrality and Y-axis indicates classification accuracy. We can see that the performance of GCN and GAT is biased towards nodes with higher closeness/eigenvector centrality. In contrast, SFairGNN has a more balanced classification performance distribution and implies a lower correlation between downstream task performance and closeness/eigenvector centrality, implying a lower correlation between downstream task performance and centrality scores. \rightarrow on the top-left figure clearly indicates that GCN and GAT are biased towards nodes with higher closeness/eigenvector centrality, while our method does not exhibit this bias.

and GAT. Therefore, it is clear that structure unfairness in GCN and GAT can lead to discriminative results for marginal individuals, especially when life-changing decisions are at stake. Such unfairness can lead to discriminatory results for marginalized individuals, particularly in cases where significant decisions are at stake. As shown in the figures on the top row, GCN and GAT tend to achieve higher classification accuracy for nodes with higher closeness centrality and lower classification accuracy for nodes with lower closeness centrality. The experimental results presented in this section indicate that SFairGNN is capable of achieving structural fairness for downstream tasks.

Similarly, the experiment results on another indicator of structure (i.e., eigenvector centrality) have a similar trend. From the figure, GCN and GAT bias towards nodes with higher eigenvector centrality. In contrast, SFairGNN has comparable average node classification accuracy and has more even performance distribution and a lower correlation between performance and eigenvector centrality. The figures in the second row in Figure 1 show that eigenvector centrality is indeed highly correlated to classification accuracy with GCN and GAT models. The observation here is quite similar to the experiment using closeness centrality as an indicator of structure.

3.2 Why GNNs are Structurally Unfair?

Here, we analyze the root cause of the unfair decision-making of GNNs. As we have discussed, the neighborhood aggregation mechanism of GNNs treats nodes with different structures unfairly, especially for the marginal nodes. To further illustrate the unfair neighbor aggregation mechanism, we use Figure 2a as a toy example. There are three groups of nodes with different egocentric structures, which are marked in yellow, blue, and green. The yellow indicates the central nodes, while the green indicates the marginal nodes. And the blue nodes situate between them. Since the green nodes and the blue nodes obtain less information than the yellow nodes, the GNNs will treat these three groups of nodes differently. Through the analysis of the toy example, we conclude that the aggregation mechanism of GNNs can lead to structure unfairness and put the marginal nodes into a disadvantaged

situation since it is inherently difficult for marginal nodes to receive the global information of the graph. Although the aggregation mechanism in GNNs improves the performance of the graph analysis, it amplifies the structure fairness.

3.3 Measurement of Structure Fairness

Since centralities and classification accuracy are both continuous variables, we propose the following two metrics to measure structure fairness. One is the correlation between the indicator of graph structure and the probability of correct classification. Another one is the variation of the average classification accuracy for nodes in different bins within different egocentric structure ranges.

Pearson Correlation Coefficient (PCC) is a statistical metric that measures the linear correlation between two variables [6]. In our scenario, we use PCC to assess the linear correlation between closeness centrality and the probability of correct classification. A PCC value in the range of $[-1, 0]/[0, 1]$ indicates that two variables are negatively/positively correlated, respectively. A low PCC value suggests that the algorithm is preserving structural fairness.

Standard Deviation (STD) quantifies the variation of average classification accuracy of nodes in different bins within different ranges of centrality values. Specifically, we cluster nodes into different bins based on their centrality values. We first calculate the mean classification accuracy for each bin and then calculate their STD. A small STD value means that a model has similar classification accuracy for nodes in different bins, which implies that nodes are treated fairly. From Figure 1, the STD of the GCN, GAT, and SFairGNN are 0.2480, 0.2935, 0.1583 for the CoraFull dataset, respectively. As our proposed SFairGNN has the lowest STD, it beats GNN and GAT in terms of structure fairness based on the closeness/eigenvector centrality.

4. THE PROPOSED METHOD

In this section, we introduce **Structural Fair Graph Neural Network (SFairGNN)**, whose fundamental intuition is to allow marginal nodes to better obtain the information from

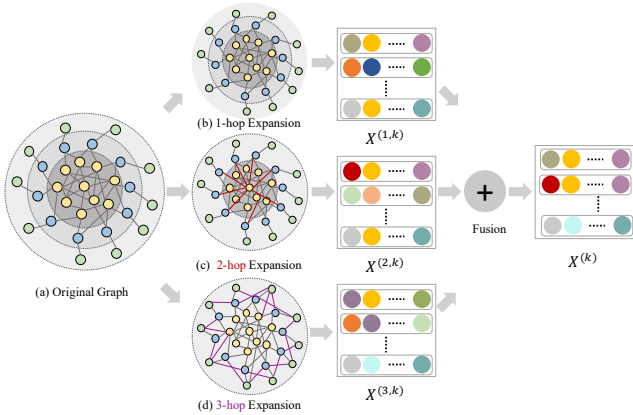


Figure 2: An illustration of our proposed method, specifically an example of the 2-layer SFairGNN. After the neighborhood expansion and hop-wise information aggregation, we obtain the embedding matrices for these three hops: $\{\mathbf{X}^{(1,k)}, \mathbf{X}^{(2,k)}, \mathbf{X}^{(3,k)}\}$. They are combined with a fusion function to output the final embedding $\mathbf{X}^{(k)}$ at k -th graph convolution layer.




Figure 3: The distribution of closeness centrality for different-hop neighbors. The X-axis represents nodes colored yellow, blue, and green, while the Y-axis denotes closeness centrality. In (a), we have the original graph, while (b) and (c) show the graphs after 2-hop and 3-hop neighborhood expansions, respectively. We observe that the distribution of closeness centrality becomes more stable as the hop number h increases.

the center of the graph. The framework of SFairGNN shown in Figure 2 is composed of two essential components: 1) the neighborhood expansion component that creates new neighbors for marginal nodes in the graph and brings them closer to the center of the graph; 2) the hop-aware attentive information aggregation component that leverages attention mechanism to fuse representations of neighbors at different hops towards a fair embedding for all nodes. In the following sections, we will introduce the details of neighborhood expansion and hop-aware attentive information aggregation. To illustrate the proposed SFairGNN, we use closeness centrality as an example indicator of graph structure.

4.1 Neighborhood Expansion Based Structure Debiasing

Before delving into the neighborhood expansion based structure debiasing, we first investigate the distribution of closeness centrality over a synthetic graph to better understand

the unfairness of the traditional information aggregation of GNNs. We plot the distribution of closeness centrality in Figure 3a based on the graph shown in Figure 2a. The synthetic graph in Figure 2a contains central, middle, and marginal nodes colored yellow, blue, and green. It is observed from Figure 3a that the closeness centrality varies significantly among the three colors, where the marginal nodes are much lower than the central nodes.

Motivated by the above observation, we propose the neighborhood expansion method to debias the structure unfairness by strengthening the marginal nodes' connections to the center of the graph. Specifically, a heuristic threshold is adopted to determine whether a node is marginal or not. We name this threshold as margin line and denote it as $line$. The margin line serves as an important hyperparameter of the proposed method. The nodes with closeness centrality $s_i \leq line$ are regarded as marginal nodes; otherwise, they are central nodes. For each node n_i , we expand its original neighbor set \mathcal{N}_i to a hyperset $\{\mathcal{N}_i^{(1)}, \mathcal{N}_i^{(2)}, \dots, \mathcal{N}_i^{(h)}\}$ to indicate neighbors within h hops. The 1-hop neighbor set $\mathcal{N}_i^{(1)}$ is given by \mathcal{N}_i , which is generated based on the original adjacency matrix \mathbf{A} . We define a debiased adjacency matrix $\tilde{\mathbf{A}}$ with each element given as:

$$\tilde{\mathbf{A}}_{i,j} = \begin{cases} \mathbf{A}_{i,j}, & \text{if } s_i \leq line \text{ or } s_j \leq line, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where s_i is the indicator of structure for node n_i . In the debiased matrix $\tilde{\mathbf{A}}$, we only strengthen the connections of marginal nodes. Then we use the h -order power of debiased matrix $\tilde{\mathbf{A}}^h$ to generate neighbor set $\mathcal{N}_i^{(h)}$. We obtain a series of neighbor sets of node n_i from different hops $\{\mathcal{N}_i^{(1)}, \mathcal{N}_i^{(2)}, \dots, \mathcal{N}_i^{(h)}\}$, which will be used for hop-aware attentive information aggregation.

Figure 2b, Figure 2c, and Figure 2d represent the graphs obtained by neighborhood expansions with $h = 1, 2$ and 3, respectively. Note that the graph with 1-hop expansion is the same as the original graph. To verify the effectiveness of neighborhood expansion, we illustrate the closeness centrality distributions of graphs with different hop expansions in Figure 3. As we can see, increasing hop numbers for neighborhood expansion shrinks the closeness centrality gap between different node groups. The empirical study shows that neighborhood expansion can effectively alleviate the structure unfairness in terms of closeness centrality.

4.2 Hop-Aware Attentive Aggregation

Since the neighbors of different hops contain different amounts of information, we propose hop-aware attentive information aggregation to provide customized information aggregation from neighbors of different hops. The overview is illustrated in Figure 2. Specifically, at the k -th graph convolution layer, let $\mathbf{X}^{(k-1)}$ and $\mathbf{X}^{(k)}$ denote the input and output embedding matrices of all the nodes, respectively. The signal processing at each layer includes two steps: 1) hop-wise information aggregation and 2) hop fusion. With the series of neighbor sets $\{\mathcal{N}_i^{(1)}, \mathcal{N}_i^{(2)}, \dots, \mathcal{N}_i^{(h)}\}$ obtained by the neighborhood expansion, we can compute the hop-wise embedding matrices $\{\mathbf{X}^{(1,k)}, \mathbf{X}^{(2,k)}, \dots, \mathbf{X}^{(h,k)}\}$ at the k -th layer. Then we fuse these embedding matrices to obtain the final output node representation $\mathbf{X}^{(k)}$. We present the details of the two steps in the following:

Hop-Wise Information Aggregation Given neighbor sets $\mathcal{N}_i^{(h)}$ of all the nodes, we aim to obtain embedding $\mathbf{X}^{(h,k)}$ at the k -th layer for h -th hop neighbors. Considering the different importance of expanded neighbors, we leverage the attention mechanism [59] to learn the attention coefficients of different hop neighbors automatically. Specifically, given neighbor j for node i , attention coefficients is calculated by:

$$\alpha_{ij}^{(h,k)} = \frac{\exp\left(f\left(\mathbf{W}^{(h,k)} \cdot \left[\mathbf{x}_i^{(k-1)} \parallel \mathbf{x}_j^{(k-1)}\right]\right)\right)}{\sum_{m \in \mathcal{N}_i^{(h)}} \exp\left(f\left(\mathbf{W}^{(h,k)} \cdot \left[\mathbf{x}_i^{(k-1)} \parallel \mathbf{x}_m^{(k-1)}\right]\right)\right)}, \quad (3)$$

where \parallel is the concatenation operation. $f(\cdot)$ is the activation function. $\mathbf{W}^{(h,k)}$ is the trainable weight matrix for the h -th hop neighbors at the k -th layer. $\mathbf{x}_i^{(k-1)}$ is the representation of node i , which is the i -th row of embedding matrix $\mathbf{X}^{(k-1)}$. With the attention coefficient $\alpha_{ij}^{(h,k)}$, the embedding of node n_i at the k -th layer within h -th hop neighbors is given by:

$$\mathbf{x}_i^{(h,k)} = \text{AGG}\left(\{\alpha_{ij}^{(h,k)} \mathbf{W}^{(h,k)} \mathbf{x}_j^{(k-1)} : j \in \mathcal{N}_i^{(h)}\}\right). \quad (4)$$

Then we have $\mathbf{X}^{(h,k)} = [\mathbf{x}_1^{(h,k)}, \mathbf{x}_2^{(h,k)}, \dots, \mathbf{x}_N^{(h,k)}]^T$, where N is the total number of nodes in the graph.

Hop Fusion After we obtain a series of hop-wise embedding matrices $\{\mathbf{X}^{(1,k)}, \mathbf{X}^{(2,k)}, \dots, \mathbf{X}^{(h,k)}\}$, we adopt fusion function \mathcal{F} to fuse them into final output embedding:

$$\mathbf{X}^{(k)} = \mathcal{F}\left(\mathbf{X}^{(1,k)}, \mathbf{X}^{(2,k)}, \dots, \mathbf{X}^{(h,k)}\right). \quad (5)$$

Fusion methods \mathcal{F} could be implemented by different functions, such as average or maximum functions. We perform experiments to compare the effects of different fusion methods [50] (i.e., SEQ, AVG, and MAX) in Section 5.3. Note that in the used datasets, the input node features usually play a deterministic role in deciding the nodes' labels. These informative node features will prohibit us from focusing on the influence of structure on the following classification task. Therefore, instead of taking the node features as input, we randomly initialize $\mathbf{X}^{(0)}$ as input node features and set it trainable during the training phase.

4.3 Discussion

Model Analysis. Here we analyze why the two key components of SFairGNN promote structure fairness: 1) neighborhood expansion strengthens the connection of marginal nodes in the graph, which positions them closer to the center of the graph and thus improves their accessibility to information from other nodes. 2) hop-aware attentive information aggregation enables SFairGNN to learn customized weights for information from neighbors of different hops away, which adaptively considers the varied strength of influence from neighboring nodes.

Avoid Constraining Loss over Fairness SFairGNN promotes the structure fairness of marginal nodes by neighbor expansion instead of explicitly optimizing fairness in loss function for two reasons. First, SFairGNN intends to improve fairness not for the measure's own sake, but for a "healthier" graph structure. This minimizes the trade-off between the overall performance and fairness for SFairGNN, as lowering the performance of "privileged" nodes is the shortest path to achieve fairness for direct optimization. From Table 2, we verify SFairGNN does not sacrifice the model performance

for fairness. Second, compared to direct optimization, neighbor expansion is drastically more efficient. This is because computing centrality scores for every training iteration is prohibitively expensive.

Comparison to Related Works. Our proposed method focuses on structure fairness, while the previous works [11, 17, 48, 57] mainly focus on learning fair node representations regarding protected attributes. For the two methods which resemble SFairGNN the most, we discuss the difference between them and our proposed method below. Fairwalk [48] proposes a fairness-aware embedding method based on the random walk technique for attributed networks. FairGNN [17] uses adversarial debiasing and covariance constraint to regularize the GNN to obtain fair node representations and predictions. Both of them focus on the categorical sensitive attributes, while our proposed method does not require any sensitive features and aims to achieve structure-level debias.

Comparison to Cold-Start Our problem is different from the cold-start problem to a large extent, which does not aim to improve fairness across users but improve the recommendation performance for new users. In addition, while cold-start in recommender systems gradually warms up, unfairness in graphs, especially in the social context, can worsen over time and thus an algorithmic solution is needed.

5. EXPERIMENTS

In this section, we empirically evaluate the effectiveness of SFairGNN in node classification tasks by answering the following questions:

- **Q1:** To what extent can SFairGNN outperform the current GNNs in simultaneously achieving high classification accuracy and maintaining structure fairness?
- **Q3:** What is the impact of using different fusion methods on the performance of SFairGNN?
- **Q2:** How do the hyperparameters of SFairGNN, such as the hop number and margin line, affect its ability to mitigate bias?

5.1 Experimental Setting

In this section, we outline our experimental setting for evaluating the effectiveness of SFairGNN in mitigating bias and achieving structural fairness. Specifically, we provide a comprehensive overview of the benchmark datasets, the baseline methods, and the implementation details both our proposed method and the baseline models.

Baselines In the experiments, we consider the widely used graph neural networks including graph convolutional networks (GCN) [35], graph attention networks (GAT) [59], and GraphSAGE [30] as baseline methods. The details of baseline methods are listed as follows:

- **GCN** [35] is the pioneer work of graph neural networks. GCN is a scalable approach for semi-supervised learning on graph-structured data that is based on an efficient variant of convolutional neural networks.
- **GAT** [59] incorporates the attention mechanism into the propagation step. This method computes the hidden states of each node by attending to its neighbors following a self-attention strategy.
- **GraphSAGE** [30] efficiently generates node representations for previously unseen data. Instead of training individual representations for each node, GraphSAGE learns

Table 1: The statistics of the datasets.

Dataset	#Nodes	#Edges	#Feature	Density	#Class
Cora	2708	10556	1433	0.0014	7
CiteSeer	3327	9104	3703	0.0008	6
PubMed	19717	88648	500	0.0002	3
CoraFull	19793	130622	8710	0.0003	70
Photo	7650	238162	745	0.0042	8
Physics	34493	495924	8415	0.0004	5

a function that generates representations by sampling and aggregating features from a node’s local neighborhood.

Datasets Following the previous works, we use six benchmark datasets to study the graph structure fairness on GNN models. They are Cora, CiteSeer, PubMed, CoraFull, Physics, and Photo [65]. The detailed statistics of these datasets are presented in Table 1.

- **Cora, CiteSeer, and PubMed** are citation networks [65], where nodes and edges denote papers and citations, respectively. Node features are bag-of-words for papers and node labels indicate the fields of papers.
- **CoraFull** [8] is a well-known citation network that contains labels based on the paper topic. CoraFull contains a network extracted from the entire citation network of Cora, while the Cora dataset is its subset.
- **Physics** [51] is co-authorship graph based on the Microsoft Academic Graph from the KDD Cup 2016 challenge. Nodes are authors and edges indicate co-authorship. Node features represent paper keywords for each author’s papers, and class labels indicate the most active fields of study for each author.
- **Photo** [51] is part of the Amazon co-purchase graph, where nodes represent goods, and edges indicate the frequency of two goods being bought together. The node features are bag-of-words encoded product reviews and class labels are given by the product category.

Implementation Details We implement both SFairGNN and the baselines with PyTorch [1] and PyTorch Geometric [27]. We initialize model weights with xavier [28] and adopt the Adam optimizer to minimize the cross-entropy loss. We set the initial learning rate to 0.005. We adopt the randomly initialized and trainable embedding vector for each node over a graph instead of node features provided by the dataset. We use 90/10 as the training/test dataset split.

5.2 Main Results on Structure Fairness

To address research question **Q1**, we conduct experiments to compare the proposed SFairGNN with baseline methods and evaluate their fairness performance using the proposed metrics: PCC (Pearson correlation coefficient) and STD (standard deviation). In this section, we utilize closeness centrality and eigenvector centrality as indicators of the graph structure in the experiments. Next, we present an analysis of the experimental results for both closeness centrality and eigenvector centrality.

Structure Fairness with Closeness Centrality Here, we show how the proposed SFairGNN tackles the fairness issue. As shown in Table 2, we can see that SFairGNN achieves the best performance in terms of PCC among all the baseline methods. The PCC values of SFairGNN are all close to zero, which indicates that the closeness centrality and clas-

sification accuracy are nearly independent. In other words, SFairGNN successfully tackles the unfairness issue triggered by the biased node egocentric structure. Furthermore, we remark that SFairGNN has a much smaller STD value than baseline methods on most benchmark datasets. This empirical result indicates that SFairGNN dramatically improves the fairness of classification decisions for nodes with different closeness centrality scores. Note that the traditional GNNs use an unfair neighbor aggregation strategy to update node representations, causing marginal nodes to receive less information. On the contrary, SFairGNN expands the neighbors of marginal nodes to debias the unfair node egocentric structures. Such neighborhood expansion strategy enforces structure fairness over the graph. From Table 2, we observe that the overall classification accuracy of SFairGNN is comparable to the baseline methods. SFairGNN does not decrease the overall performance even though helping the marginal nodes may harm the nodes at the center of the graph. On the other hand, SFairGNN can improve structure fairness since the SFairGNN performs better on both the STD and PCC metrics. Thus our proposed model can alleviate the unfairness issue while maintaining the overall model performance for downstream tasks.

Datasets	Metrics	GCN	GAT	SAGE	SFairGNN	↑
Cora	ACC	76.00	66.49	74.85	77.65	--
	STD	13.15	12.36	12.63	10.53	19.9%
	PCC	10.47	26.50	10.96	0.76	92.7%
CiteSeer	ACC	54.03	47.55	54.45	53.67	--
	STD	15.35	15.67	16.23	13.07	14.8%
	PCC	14.50	27.30	14.24	12.92	10.9%
PubMed	ACC	77.44	66.66	77.82	78.85	--
	STD	8.12	12.48	07.05	5.44	33.0%
	PCC	13.48	3.27	14.68	5.98	55.6%
CoraFull	ACC	49.23	40.53	51.86	50.24	--
	STD	9.65	10.95	09.37	7.17	25.7%
	PCC	24.80	29.35	24.20	15.83	36.2%
Physics	ACC	90.64	78.79	90.61	91.12	--
	STD	11.33	17.02	8.79	7.65	32.5%
	PCC	37.09	50.59	29.95	28.56	23.0%
Photo	ACC	90.66	89.15	90.81	88.89	--
	STD	8.71	10.16	8.91	8.87	-1.8%
	PCC	20.03	32.82	15.50	10.19	48.9%

sification accuracy are nearly independent. In other words, SFairGNN successfully tackles the unfairness issue triggered by the biased node egocentric structure. Furthermore, we remark that SFairGNN has a much smaller STD value than baseline methods on most benchmark datasets. This empirical result indicates that SFairGNN dramatically improves the fairness of classification decisions for nodes with different closeness centrality scores. Note that the traditional GNNs use an unfair neighbor aggregation strategy to update node representations, causing marginal nodes to receive less information. On the contrary, SFairGNN expands the neighbors of marginal nodes to debias the unfair node egocentric structures. Such neighborhood expansion strategy enforces structure fairness over the graph. From Table 2, we observe that the overall classification accuracy of SFairGNN is comparable to the baseline methods. SFairGNN does not decrease the overall performance even though helping the marginal nodes may harm the nodes at the center of the graph. On the other hand, SFairGNN can improve structure fairness since the SFairGNN performs better on both the STD and PCC metrics. Thus our proposed model can alleviate the unfairness issue while maintaining the overall model performance for downstream tasks.

Structure Fairness with Eigenvector Centrality. As shown in Table 3, we can see that SFairGNN achieves the best performance on PCC and STD among the baseline methods. The PCC values of SFairGNN are very low, which indicates the eigenvector centrality and classification accuracy are nearly independent. This observation shows SFairGNN successfully tackles the unfairness issue in the graph. In addition, SFairGNN has a much smaller STD value than the baseline methods, which means SFairGNN dramatically improves the structure fairness for nodes with different eigenvector centrality values.

5.3 Effect of Fusion Methods

To answer the research question **Q2**, we conduct a series of experiments to explore the effect of fusion methods in

Table 3: Comparison between the proposed SFairGNN and baseline methods on eigenvector centrality. The best scores of STD and PCC are printed in bold. **Improv \uparrow** represents the fairness improvement, which is calculated based on the baseline GCN. We run each experiment 5 times and report the means. The hop number is set to 3.

Datasets	Metrics	GCN	GAT	SAGE	SFairGNN	\uparrow
Cora	ACC	76.41	64.97	75.84	76.00	--
	STD	9.97	12.52	9.96	11.05	-10.49%
	PCC	9.89	14.20	10.37	0.03	99.70%
CiteSeer	ACC	53.38	47.61	51.55	53.42	--
	STD	16.55	16.18	17.47	11.85	28.43%
	PCC	16.60	19.69	14.41	5.96	64.10%
PubMed	ACC	77.39	65.93	78.44	77.28	--
	STD	8.24	13.72	7.23	6.03	26.63%
	PCC	4.82	6.43	4.53	5.04	-4.98%
CoraFull	ACC	50.05	40.64	51.90	47.74	--
	STD	9.21	10.61	9.18	6.94	24.85%
	PCC	9.57	17.34	8.99	5.43	43.18%
Physics	ACC	91.14	76.07	90.63	91.74	--
	STD	11.33	18.20	8.23	9.10	19.71%
	PCC	17.98	23.34	12.50	12.41	44.25%
Photo	ACC	91.02	88.48	91.09	88.11	--
	STD	8.11	9.51	8.89	8.06	-0.56%
	PCC	8.54	4.79	5.16	1.64	80.79%




Figure 4: Performance comparison w.r.t. different fusion methods on CiteSeer and CoraFull datasets. We use SEQ, MAX, AVG fusion methods. The experiments show that the fusion method MAX obtains the best performance.

SFairGNN. We consider replacing the fusion method with several different strategies as follows:

- **SEQ** represents the approach where the model aggregates information of different hop neighbors sequentially. This means the model first aggregates information from first hop neighbors and then moves on to aggregate information from other hop neighbors.
- **MAX** represents the approach where the model selects the maximum value of each element from different hop embeddings for each node. This approach highlights the most important neighbor among different neighbors.
- **AVG** represents the approach where the model takes the average of embeddings of different hop neighbors for each node. This approach treats all the neighbors equally during attentive information aggregation.

One can see from Figure 4 that the fusion method MAX obtains the best performance among the three fusion methods because the MAX fusion method is capable of selecting the most important neighbor automatically to fuse the embeddings. Since the expanded neighbor may introduce noise to the graph structure, the MAX strategy can also reduce the




Figure 5: The effect of hop number on the CiteSeer and CoraFull datasets with the closeness centrality. The results show that 1) the values of PCC and STD generally decrease when the hop number is less than 4. 2) the performance drops when the hop number is too high.




Figure 6: The effect of margin line on the CiteSeer and CoraFull datasets. The results show that the value of STD generally decreases and then increases, where the optimal fairness is reached at the middle of the margin line range.

influence of such adverse effects. On the other hand, the SEQ strategy obtains the worst performance on the metric PCC, which may be caused by the fact that SEQ treats different hop neighbors in the same way.

5.4 Hyperparameter Sensitivity Analysis

To answer the research question **Q₃**, we conduct experiments to explore the sensitivity of SFairGNN to different hyperparameters, including the hop number h and the margin line $line$. The hop number h indicates the extent of the neighborhood expansion. The margin line $line$ determines whether the nodes are marginal nodes or not. For simplicity, we use closeness centrality as the indicator of structure in the experiments for hyperparameter sensitivity analysis.

Effect of Neighbor Hop. The hop number h controls the information reception field of marginal nodes, making it a vital hyperparameter. We conduct experiments on the PubMed and CoraFull datasets considering a series of hop numbers $h \in \{1, \dots, 5\}$. We consider closeness centrality for brevity. The results are presented in Figure 5. We observed that 1) the values of PCC and STD generally decrease when the hop number is less than or equal to 4. This is because the closeness centrality of marginal nodes is improved with the neighborhood expansion, reducing the variation of the nodes' closeness centrality over the graph.

2) when the hop number is too high, the performance drops because the original graph structure is damaged—being much different from the graph after neighborhood expansion.

Effect of Margin Line. The margin line *line* decides whether a node is marginal or not. To explore its effect, we conduct experiments on the Cora and CiteSeer datasets with a series of margin line value $line \in \{0.1, \dots, 0.9\}$ and report the results in Figure 6. One can see from Figure 6 that the value of STD generally decreases and then increases, where the optimal fairness is reached at the middle of the margin line range. This is because when the margin line is too small, the marginal nodes obtain insufficient compensation. On the other hand, the neighborhood expansion incorporates too much noise, which hurts the fairness performance when the margin line is too large. In fact, it should be noted that the margin line has significant sociological meanings, so it should be set seriously for specific fairness issues.

6. RELATED WORKS

This section summarizes two categories of related works, including graph neural networks and fairness in graphs.

6.1 Graph Neural Networks

Graph neural networks (GNNs) [31, 52, 56, 58, 64, 69, 71, 72, 74] have been the new state-of-the-art method to analyze graph-structured data, which are widely applied to social networks [32], and academic citation networks [35], knowledge graphs [21, 73], and to name a few. Starting with the success of graph convolutional network in the semi-supervised node classification task, a wide variety of GNN variants have enhanced and improved the node representation learning and downstream learning tasks [30, 59, 62]. Most of GNNs follow a message-passing strategy to learn node representations over a graph. GNNs apply a neighbor aggregator to update node representations iteratively via combining representations of neighbors and that of the node itself. GraphSAGE [30] learns representations by sampling and aggregating neighbor nodes, whereas GAT [59] uses the attention mechanism to aggregate representations from all neighbors. Besides, there is another branch of GNNs, spectral GCNs [12, 19, 35], which use spectral filters over graph laplacian to define the convolution operation.

6.2 Fairness in Graphs

Fairness [13, 44, 46] in machine learning algorithms attracts a lot of interest from both academic and industrial communities since life-changing decision-making needs to be fair in many sensitive environments, such as loan applications, health care [2], and hiring [4, 7]. One mainstream approach to achieve fairness is to use the in-processing approaches, which reduce the bias in the training data. In-processing approaches typically incorporate one or more fairness constraints into the model optimization process to maximize the performance of downstream tasks and improve fairness simultaneously [25, 67, 68]. Recently, a few works [11, 17, 22, 33, 34, 37, 38, 41, 48] focus on learning fair node representation regarding protected attributes (e.g., gender, age, race). [11] proposes an adversarial framework to enforce fairness constraints on graph learning, which can flexibly accommodate different combinations of fairness constraints during inference. [48] aims to study potential bias issues inherent within graph embedding and propose a fairness-aware embedding method named Fairwalk. FairGNN [17] uses adversarial debiasing and covariance constraint to regularize the GNN to yield fair predictions. There is also a line of works focusing on the fairness in graph structure [15, 40, 53, 60].

The position-aware graph neural network [66] also investigates the “position” of the node in graphs, which is related to our proposed method. A more comprehensive review of related works of fairness in graphs can be found in [23].

7. CONCLUSION

In this work, we propose an SFairGNN model that is resilient to structure unfairness in graphs. We show that marginal nodes in the graph have lower prediction performance in existing GNN models, which indicates the existing GNNs suffer from the structure fairness issue. To this end, SFairGNN is proposed to preserve fairness without sacrificing the prediction performance of GNNs. Specifically, SFairGNN relies on a neighborhood expansion strategy and a hop-aware attentive information aggregation strategy to optimally introduce new information to structurally underprivileged nodes, i.e., marginal nodes. Moreover, we conduct extensive experiments on several graph datasets and validate that our SFairGNN can achieve fairer node classification performance than the existing GNNs. In the future, we will consider both structure fairness and node attribute fairness for GNNs since the fairness may be caused by the interactions of structure and attributes. Another promising extension of this work is to study the interpretability of structure fairness.

8. ACKNOWLEDGEMENTS

We would like to thank all the anonymous reviewers for their valuable suggestions. Na Zou is supported by NSF IIS-1900990, IIS-1939716, IIS-2239257. Fei Wang is supported by NSF 1750326 and 2212175 for this research.

9. REFERENCES

- [1] P. Adam, G. Sam, C. Soumith, C. Gregory, Y. Edward, D. Zachary, L. Zeming, D. Alban, A. Luca, and L. Adam. Automatic differentiation in pytorch. In *NeurIPS*, 2017.
- [2] M. A. Ahmad, A. Patel, C. Eckert, V. Kumar, and A. Teredesai. Fairness in machine learning for health-care. In *KDD*, 2020.
- [3] A. Airola, S. Pyysalo, J. Björne, T. Pahikkala, F. Ginter, and T. Salakoski. All-paths graph kernel for protein-protein interaction extraction with evaluation of cross-corpus learning. *BMC bioinformatics*, 2008.
- [4] G. S. Alder and J. Gilbert. Achieving ethics and fairness in hiring: Going beyond the law. *Journal of Business Ethics*, 2006.
- [5] A. Bavelas. Communication patterns in task-oriented groups. *The journal of the acoustical society of America*, 22(6):725–730, 1950.
- [6] J. Benesty, J. Chen, Y. Huang, and I. Cohen. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 1–4. Springer, 2009.
- [7] M. Bogen and A. Rieke. Help wanted: An examination of hiring algorithms, equity, and bias, 2018.
- [8] A. Bojchevski and S. Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815*, 2017.

- [9] P. Bonacich. Some unique properties of eigenvector centrality. *Social networks*, 29(4):555–564, 2007.
- [10] S. P. Borgatti. Centrality and network flow. *Social networks*, 27(1):55–71, 2005.
- [11] A. Bose and W. Hamilton. Compositional fairness constraints for graph embeddings. In *ICML*, pages 715–724. PMLR, 2019.
- [12] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. In *ICLR*, 2014.
- [13] S. Caton and C. Haas. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053*, 2020.
- [14] X. Chen, M. Hou, T. Tang, A. Kaur, and F. Xia. Digital twin mobility profiling: A spatio-temporal graph learning approach. In *2021 IEEE 23rd Int Conf on High Performance Computing & Communications*, pages 1178–1187, 2021.
- [15] Z. Chen, T. Xiao, and K. Kuang. Ba-gnn: On learning bias-aware graph neural network. In *ICDE*, pages 3012–3024. IEEE, 2022.
- [16] M. Choudhary, C. Laclau, and C. Largeron. A survey on fairness for machine learning on graphs. *arXiv preprint arXiv:2205.05396*, 2022.
- [17] E. Dai and S. Wang. Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information. In *WSDM*, 2021.
- [18] E. Dai, T. Zhao, H. Zhu, J. Xu, Z. Guo, H. Liu, J. Tang, and S. Wang. A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. *arXiv preprint arXiv:2204.08570*, 2022.
- [19] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NeurIPS*, 2016.
- [20] J. Dong, Q. Zhang, X. Huang, K. Duan, Q. Tan, and Z. Jiang. Hierarchy-aware multi-hop question answering over knowledge graphs. In *International World Wide Web Conference*, pages 2519–2527, 2023.
- [21] J. Dong, Q. Zhang, X. Huang, Q. Tan, D. Zha, and Z. Zihao. Active ensemble learning for knowledge graph error detection. In *ACM International Conference on Web Search and Data Mining*, pages 877–885, 2023.
- [22] Y. Dong, N. Liu, B. Jalaian, and J. Li. Edits: Modeling and mitigating data bias for graph neural networks. In *Proceedings of the ACM Web Conference 2022*, pages 1259–1269, 2022.
- [23] Y. Dong, J. Ma, C. Chen, and J. Li. Fairness in graph mining: A survey. *arXiv preprint arXiv:2204.09888*, 2022.
- [24] X. Du, Y. Pei, W. Duivesteijn, and M. Pechenizkiy. Fairness in network representation by latent structural heterogeneity in observational data. In *AAAI*, 2020.
- [25] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- [26] W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin. Graph neural networks for social recommendation. In *The world wide web conference*, 2019.
- [27] M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [28] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 2010.
- [29] A. Hagberg, P. Swart, and D. S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- [30] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *NeurIPS*, 2017.
- [31] X. Han, Z. Jiang, N. Liu, Q. Song, J. Li, and X. Hu. Geometric graph representation learning via maximizing rate reduction. In *Proceedings of the ACM Web Conference 2022*, pages 1226–1237, 2022.
- [32] X. Huang, Q. Song, Y. Li, and X. Hu. Graph recurrent networks with attributed random walks. In *KDD*, 2019.
- [33] Z. Jiang, X. Han, C. Fan, Z. Liu, X. Huang, N. Zou, A. Mostafavi, and X. Hu. Topology matters in fair graph learning: a theoretical pilot study. 2022.
- [34] Z. Jiang, X. Han, C. Fan, Z. Liu, N. Zou, A. Mostafavi, and X. Hu. Fmp: Toward fair graph message passing against topology bias. *arXiv preprint arXiv:2202.04187*, 2022.
- [35] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [36] D. Koschützki, K. A. Lehmann, L. Peeters, S. Richter, D. Tenfelde-Podehl, and O. Zlotowski. Centrality indices. In *Network analysis*, pages 16–61. Springer, 2005.
- [37] O. D. Kose and Y. Shen. Fair node representation learning via adaptive data augmentation. *arXiv preprint arXiv:2201.08549*, 2022.
- [38] P. Li, Y. Wang, H. Zhao, P. Hong, and H. Liu. On dyadic fairness: Exploring and mitigating bias in graph connections. In *International Conference on Learning Representations*, 2021.
- [39] H. Ling, Z. Jiang, M. Liu, S. Ji, and N. Zou. Graph mixup with soft alignments. *arXiv preprint arXiv:2306.06788*, 2023.
- [40] Z. Liu, T.-K. Nguyen, and Y. Fang. On generalized degree fairness in graph neural networks. *arXiv preprint arXiv:2302.03881*, 2023.

- [41] J. Ma, J. Deng, and Q. Mei. Subgroup generalization and fairness of graph neural networks. *NeurIPS*, 34:1048–1061, 2021.
- [42] J. Ma, R. Guo, M. Wan, L. Yang, A. Zhang, and J. Li. Learning fair node representations with graph counterfactual fairness. In *WSDM*, pages 695–703, 2022.
- [43] P. V. Marsden. Egocentric and sociocentric measures of network centrality. *Social networks*, 2002.
- [44] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- [45] D. Q. Nguyen, V. Tong, D. Phung, and D. Q. Nguyen. Node co-occurrence based graph neural networks for knowledge graph link prediction. In *WSDM*, 2022.
- [46] D. Pessach and E. Shmueli. Algorithmic fairness. *arXiv preprint arXiv:2001.09784*, 2020.
- [47] S. U. Pillai, T. Suel, and S. Cha. The perron-frobenius theorem: some of its applications. *IEEE Signal Processing Magazine*, 22(2):62–75, 2005.
- [48] T. Rahman, B. Surma, M. Backes, and Y. Zhang. Fairwalk: towards fair graph embedding. In *IJCAI*, 2019.
- [49] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966.
- [50] D. Scherer, A. Müller, and S. Behnke. Evaluation of pooling operations in convolutional architectures for object recognition. In *International conference on artificial neural networks*, pages 92–101. Springer, 2010.
- [51] O. Shchur, M. Mumme, A. Bojchevski, and S. Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- [52] Y. Shi, Y. Dong, Q. Tan, J. Li, and N. Liu. Gigamae: Generalizable graph masked autoencoder via collaborative latent space reconstruction. In *ACM CIKM*, 2023.
- [53] H. Shomer, W. Jin, W. Wang, and J. Tang. Toward degree bias in embedding-based knowledge graph completion. *arXiv preprint arXiv:2302.05044*, 2023.
- [54] W. Song, Y. Dong, N. Liu, and J. Li. Guide: Group equality informed individual fairness in graph neural networks. In *SIGKDD*, pages 1625–1634, 2022.
- [55] Q. Tan, N. Liu, and X. Hu. Deep representation learning for social network analysis. *Frontiers in Big Data*, 2:2, 2019.
- [56] Q. Tan, X. Zhang, N. Liu, D. Zha, L. Li, R. Chen, S.-H. Choi, and X. Hu. Bring your own view: Graph neural networks for link prediction with personalized subgraph selection. In *ACM International Conference on Web Search and Data Mining*, pages 625–633, 2023.
- [57] X. Tang, H. Yao, Y. Sun, Y. Wang, J. Tang, C. Aggarwal, P. Mitra, and S. Wang. Investigating and mitigating degree-related biases in graph convolutional networks. In *CIKM*, 2020.
- [58] P. Veličković. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology*, 79:102538, 2023.
- [59] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [60] R. Wang, X. Wang, C. Shi, and L. Song. Uncovering the structural fairness in graph contrastive learning. *arXiv preprint arXiv:2210.03011*, 2022.
- [61] Y. Wang, Y. Zhao, Y. Dong, H. Chen, J. Li, and T. Derr. Improving fairness in graph neural networks via mitigating sensitive attribute leakage. In *ACM SIGKDD*, 2022.
- [62] F. Wu, T. Zhang, A. H. d. Souza Jr, C. Fifty, T. Yu, and K. Q. Weinberger. Simplifying graph convolutional networks. *arXiv preprint arXiv:1902.07153*, 2019.
- [63] L. Wu, P. Cui, J. Pei, L. Zhao, and L. Song. Graph neural networks. In *Graph Neural Networks: Foundations, Frontiers, and Applications*, pages 27–37. Springer, 2022.
- [64] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip. A comprehensive survey on graph neural networks. *IEEE TNNLS*, 2020.
- [65] Z. Yang, W. W. Cohen, and R. Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. *arXiv preprint arXiv:1603.08861*, 2016.
- [66] J. You, R. Ying, and J. Leskovec. Position-aware graph neural networks. In *International conference on machine learning*, pages 7134–7143. PMLR, 2019.
- [67] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Fairness constraints: Mechanisms for fair classification. *arXiv preprint arXiv:1507.05259*, 2015.
- [68] B. H. Zhang, B. Lemoine, and M. Mitchell. Mitigating unwanted biases with adversarial learning. In *AIES*, 2018.
- [69] C. Zhang, C. Huang, Y. Li, X. Zhang, Y. Ye, and C. Zhang. Look twice as much as you say: Scene graph contrastive learning for self-supervised image caption generation. In *CIKM*, 2022.
- [70] W. Zhang, J. C. Weiss, S. Zhou, and T. Walsh. Fairness amidst non-iid graph data: A literature review. *arXiv preprint arXiv:2202.07170*, 2022.
- [71] X. Zhang, Q. Tan, X. Huang, and B. Li. Graph contrastive learning with personalized augmentation. *arXiv preprint arXiv:2209.06560*, 2022.
- [72] Z. Zhang, P. Cui, and W. Zhu. Deep learning on graphs: A survey. *IEEE TKDE*, 2020.
- [73] Z. Zhang, F. Zhuang, H. Zhu, Z. Shi, H. Xiong, and Q. He. Relational graph neural network with hierarchical attention for knowledge graph completion. In *AAAI*, 2020.
- [74] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.