

FairGP: A Scalable and Fair Graph Transformer Using Graph Partitioning

Renqiang Luo¹, Huafei Huang², Ivan Lee², Chengpei Xu³, Jianzhong Qi⁴, Feng Xia^{5*}

¹Dalian University of Technology

²University of South Australia

³The University of New South Wales

⁴The University of Melbourne

⁵RMIT University

{lrnqiang, hhuafei}@outlook.com, Ivan.Lee@unisa.edu.au,
chengpei.xu@unsw.edu.au, jianzhong.qi@unimelb.edu.au, f.xia@ieee.org

Abstract

Recent studies have highlighted significant fairness issues in Graph Transformer (GT) models, particularly against subgroups defined by sensitive features. Additionally, GTs are computationally intensive and memory-demanding, limiting their application to large-scale graphs. Our experiments demonstrate that graph partitioning can enhance the fairness of GT models while reducing computational complexity. To understand this improvement, we conducted a theoretical investigation into the root causes of fairness issues in GT models. We found that the sensitive features of higher-order nodes disproportionately influence lower-order nodes, resulting in sensitive feature bias. We propose Fairness-aware scalable GT based on Graph Partitioning (FairGP), which partitions the graph to minimize the negative impact of higher-order nodes. By optimizing attention mechanisms, FairGP mitigates the bias introduced by global attention, thereby enhancing fairness. Extensive empirical evaluations on six real-world datasets validate the superior performance of FairGP in achieving fairness compared to state-of-the-art methods.

Introduction

Despite the reported capabilities of Graph Transformer (GT) models in graph representation learning, recent studies have exposed fairness concerns associated with these models (Luo et al. 2024c). In critical applications such as medical diagnoses (Pan et al. 2024), loan approval (Zhang et al. 2024b), and e-commerce (Agarwal et al. 2024), features such as gender, race, age, and region are legally protected to prevent discrimination and are thus considered sensitive features (Wang et al. 2024). Unfortunately, GTs generate biased predictions that discriminate against specific subgroups characterized by these sensitive features.

In addition, GTs are computationally demanding and memory-intensive when applied on large-scale graphs due to the complexity of global attention mechanisms (Dao et al. 2022; Chen et al. 2023). Consequently, training deep GTs on large-scale graphs can be resource-intensive, necessitating sophisticated techniques for partitioning nodes into smaller mini-batches to alleviate the computational overhead (Wu

et al. 2023b). Partitioning strategies impact the performance of GTs and reduce the time and space complexity (Blondel et al. 2008; Traag, Waltman, and Van Eck 2019). Additionally, these strategies may also be subject to fairness issues.

Will GT with graph partitioning exhibit the same fairness issues as a vanilla GT? Various efforts have been devoted to developing fairness-aware algorithms, aiming to control the degree to which a model depends on sensitive features, measured by independence criteria such as statistical parity (measured by Δ_{SP}) and equality opportunity (measured by Δ_{EO}) (Zhang et al. 2024a). We calculate the fairness metrics for both vanilla GT and GT with graph partitioning separately, as shown in Figure 1. The experimental settings are detailed in Appendix A.1, along with results from three additional graph partitioning strategies (Luo et al. 2024a). The results show that graph partitioning improves fairness in GT models, indicating its potential to address fairness issues within GTs. This observation raises a natural and fundamental question: *how does graph partitioning promote fairness?* To answer the question, we conducted a theoretical investigation into the underlying mechanisms.

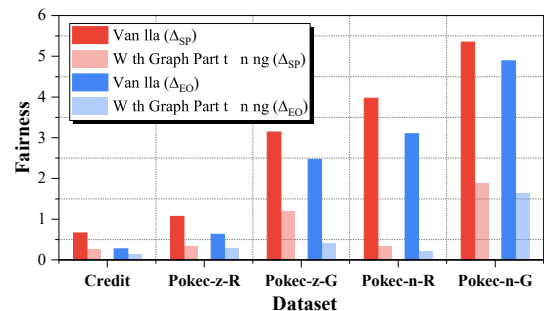


Figure 1: Graph partitioning increases GT fairness.

Intuitively, the global attention mechanism in GTs tends to favour higher-order nodes, causing the sensitive features of these nodes to disproportionately influence lower-order nodes, resulting in sensitive feature bias (Shehzad et al. 2024). Our theoretical analysis indicates that global attention among nodes with different sensitive features significantly impacts GT fairness. Minimizing the influence of

*Corresponding Author

higher-order nodes on nodes with varying sensitive features can effectively improve GT fairness. In addition, we categorize the attention into inter-cluster and intra-cluster attention. To enhance fairness, we optimize both inter-cluster and intra-cluster attention, thereby reducing interactions between nodes with different sensitive features. By adapting the graph partitioning, we enable a large number of lower-order nodes to avoid the undue influence from higher-order nodes. This effectively addresses the fairness issue, ensuring that the sensitive features of lower-order nodes are less affected by those of higher-order.

Building on this insight, we propose **Fairness-aware scalable GT based on Graph Partitioning (FairGP)**. FairGP reduces the influence range of higher-order nodes in global attention by partitioning the graph and focusing attention within respective subgraphs. Additionally, FairGP optimizes attention to enhance the similarity between the distribution of the original sensitive features and the distribution of their embeddings. This optimization minimizes the negative influence between nodes with different sensitive features, thereby enhancing GT fairness. The main contributions of this work are summarized as follows:

- We show that GT fairness is influenced by the attention of higher-order nodes and the attention between nodes with different sensitive features. To the best of our knowledge, this work is the first to uncover the underlying cause of fairness issues in GTs.
- We propose a novel fairness-aware scalable GT, FairGP, which partitions the graph to mitigate the influence of higher-order node attention and optimizes inter-cluster attention to mitigate bias in nodes with different sensitive features.
- To verify the effectiveness of the proposed FairGP, we conduct extensive empirical experiments comparing with state-of-the-art models. The results show that FairGP consistently improves fairness by at least 40% across both statistical parity and equality opportunity over six real-world datasets.

Related Work

Fairness-aware Graph Neural Network (GNNs)

Fairness in GNNs has gained substantial attention, particularly in efforts to identify and mitigate biases associated with specific sensitive features (Zhang et al. 2024c). Various fairness-aware GNN studies aim to preserve the independence of sensitive features through pre-processing and in-processing techniques (Dong et al. 2023; Luo et al. 2024d).

Graphair (Ling et al. 2023) is a recent pre-processing algorithm. It automatically identifies fairness-aware augmentations from input graph data, aiming to circumvent sensitive features while preserving the integrity of other valuable information. FairAC (Guo, Chu, and Li 2023) proposes a fairness feature completion approach. It completes missing information by adversarial learning and achieves fair node embeddings for graphs lacking certain features. Unlike GNN methods, GTs that depend on the attention score matrix face challenges in directly identifying sensitive features in the input, making it difficult to apply pre-processing

techniques effectively. In-processing methods, on the other hand, focus on message-passing or convolution propagation. FairSIN (Yang et al. 2024) introduces a neutralization-based paradigm, which incorporates additional fairness-enhancing features into node features or representations before message-passing. FMP (Jiang et al. 2024) directly addresses the use of sensitive features during forward propagation for node classification tasks using cross-entropy loss, eliminating the need for data pre-processing. However, the global attention mechanism in GTs, which facilitates direct node-to-node interactions, differs from the message-passing or convolution used in GNNs (Yin and Zhong 2023). In-processing techniques may not be directly applicable to GTs.

Scalable Graph Transformer

Recent GT studies are shifting their focus to handling the large-scale graph (Xing et al. 2024). DIFFormer (Wu et al. 2023a) introduces a scalable GT framework guided by data-invariant functions, enabling efficient processing of large-scale graphs. Similarly, SGFormer (Wu et al. 2023b) simplifies the Transformer architecture to boost computational efficiency on large-scale graphs. Polynormer (Deng, Yue, and Zhang 2024) advances the field by developing a polynomial expressive GT that can capture features in complex graph structures more effectively.

Despite these advancements, existing GT methods often lack explicit consideration of algorithmic fairness. When applied to real-world datasets, they exhibit significant fairness issues. FairGT (Luo et al. 2024c) addresses the fairness problem within GT based on a sensitive feature complete graph. It faces challenges in deployment on large graphs due to its reliance on global attention in two whole graphs (original and sensitive feature complete graph). Scalable techniques such as mini-batch and graph partitioning are challenging to deploy in FairGT while ensuring guaranteed performance, because the sensitive feature complete graph relies on all edges. Addressing the fairness issues of GTs on large-scale graphs remains a pressing and current challenge.

Preliminaries

Notations

Unless otherwise specified, we denote sets with copperplate uppercase letters (e.g., \mathcal{A}), matrices with bold uppercase letters (e.g., \mathbf{A}), and vectors with bold lowercase letters (e.g., \mathbf{x}). We denote the number of elements in the set \mathcal{A} as $|\mathcal{A}|$. We denote a graph as $\mathcal{G} = \{\mathcal{V}, \mathbf{A}, \mathbf{H}\}$, where \mathcal{V} is the set of nodes in the graph and $|\mathcal{V}| = n$, $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the adjacency matrix, and $\mathbf{H} \in \mathbb{R}^{n \times d}$ is the node feature matrix. We follow the convention of NumPy in Python for matrix and vector indexing: $\mathbf{A}[i, j]$ represents the entry of matrix \mathbf{A} at the i -th row and the j -th column; $\mathbf{A}[i, :]$ and $\mathbf{A}[:, j]$ represent the i -th row and the j -th column of matrix \mathbf{A} , respectively. $\mathbf{H}[:, s]$ denotes a sensitive feature of nodes, and $\mathbf{H}[i, s]$ denotes a sensitive feature of node v_i .

Fairness Evaluation Metrics

We present two definitions of fairness for a binary ground truth label $y \in \{0, 1\}$ and a binary sensitive feature $\mathbf{H}[i, s] \in$

$\{0, 1\}$. We use $\hat{y} \in \{0, 1\}$ to represent the predicted label.

Definition 1. Statistical Parity (i.e., Demographic Parity, Independence) (Dwork et al. 2012). Statistical parity requires the predictions to be independent of the sensitive features. The prediction can be formally written as:

$$\mathbb{P}(\hat{y} = 1 | \mathbf{H}[i, s] = 0) = \mathbb{P}(\hat{y} = 1 | \mathbf{H}[i, s] = 1). \quad (1)$$

When both the predicted labels and the sensitive features are binary, the extent of statistical parity can be quantified by Δ_{SP} , defined as follows:

$$\Delta_{\text{SP}} = |\mathbb{P}(\hat{y} = 1 | \mathbf{H}[i, s] = 0) - \mathbb{P}(\hat{y} = 1 | \mathbf{H}[i, s] = 1)|. \quad (2)$$

The Δ_{SP} measures the acceptance rate difference between two sensitive subgroups with $\mathbf{H}[i, s] = 0$ and $\mathbf{H}[i, s] = 1$.

Definition 2. Equal Opportunity (Hardt, Price, and Srebro 2016). Equal opportunity requires that the probability of a positive instance leading to a positive outcome be the same across all subgroups. For individuals with positive ground truth labels, positive predictions should not depend on sensitive features. This can be formulated as follows:

$$\begin{aligned} &\mathbb{P}(\hat{y} = 1 | y = 1, \mathbf{H}[i, s] = 0) \\ &= \mathbb{P}(\hat{y} = 1 | y = 1, \mathbf{H}[i, s] = 1). \end{aligned} \quad (3)$$

Fairness-aware algorithms prevent the allocation of unfavourable predictions to individuals solely based on their sensitive subgroup affiliation. In particular, Δ_{EO} quantifies the extent of deviation in predictions from equal opportunity. To quantitatively assess equal opportunity:

$$\begin{aligned} \Delta_{\text{EO}} &= |\mathbb{P}(\hat{y} = 1 | y = 1, \mathbf{H}[i, s] = 0) \\ &\quad - \mathbb{P}(\hat{y} = 1 | y = 1, \mathbf{H}[i, s] = 1)|. \end{aligned} \quad (4)$$

Both Δ_{SP} and Δ_{EO} are evaluated on the test set.

Transformer

The Transformer architecture consists of a composition of Transformer layers. Each Transformer layer has two parts: a self-attention module and a position-wise feed-forward network (FFN). Let $\mathbf{X} = [\mathbf{h}_1^\top, \dots, \mathbf{h}_i^\top]^\top \in \mathbb{R}^{i \times d_h}$ be the input of the self-attention module where d_h is the hidden dimension and $\mathbf{h}_j \in \mathbb{R}^{1 \times d_h}$ is the hidden representation at position j . The input \mathbf{X} is projected by three matrices $\mathbf{W}_Q \in \mathbb{R}^{d_h \times d_K}$, $\mathbf{W}_K \in \mathbb{R}^{d_h \times d_K}$, and $\mathbf{W}_V \in \mathbb{R}^{d_h \times d_V}$ to the corresponding representations \mathbf{Q} , \mathbf{K} , \mathbf{V} . The self-attention is then calculated as:

$$\text{Attn}(\mathbf{X}) = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_K}} \right) \mathbf{V}, \quad (5)$$

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \mathbf{K} = \mathbf{X}\mathbf{W}_K, \mathbf{V} = \mathbf{X}\mathbf{W}_V.$$

To simplify, we consider single-head self-attention and assume $d_K = d_V = d_h$. The extension to multi-head attention is straightforward, and we omit the bias terms for simplicity. The attention score matrix is $\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}$:

$$\hat{\mathbf{A}} = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_K}} \right). \quad (6)$$

$\hat{\mathbf{A}}[u, v]$ represents the attention score between nodes u and v . Due to the Softmax function, we have $\sum_{i=1}^n \hat{\mathbf{A}}[i, :] = 1$ and $\sum_{i=1}^n \hat{\mathbf{A}}[:, i] = 1$.

The Fairness Issues in GT

In this section, we analyze the underlying cause of fairness issues in GTs. Firstly, we present our empirical observations that the global attention mechanism negatively impacts fairness in GT. Because, global attention tends to focus on the proportions of sensitive features among higher-order nodes, while fairness-aware algorithms often protect the proportions of sensitive features among all nodes (most of the nodes are lower-order nodes). Methods that can truncate the attention score matrix, such as graph partitioning, can effectively mitigate this negative impact. Then, we present the theoretical findings on GT fairness and explain how graph partitioning enhances fairness. A theoretical foundation is provided, enabling us to propose a GT fairness algorithm (FairGP) based on graph partitioning.

Datasets	All Nodes		Higher-order		Prediction ($\hat{y} = 1$)	
	$s = 1$	$s = 0$	$s = 1$	$s = 0$	$s = 1$	$s = 0$
Credit	10.17	1	2.25	1	2.94	1
Pokec-z-R	1.82	1	1	1.29	1	2.63
Pokec-z-G	1	1	1	1.25	1	3.10
Pokec-n-R	2.12	1	2.51	1	3.50	1
Pokec-n-G	1.26	1	1	1.37	1	2.69
AMine-L	1.18	1	1	1.36	1	6.81

Table 1: Distribution of different sensitive features. We denote $\mathbf{H}[i, s] = 1$ as $s = 1$, and $\mathbf{H}[i, s] = 0$ as $s = 0$.

Empirical Observations

We empirically show how the global attention mechanism in GT tends to favour higher-order nodes. The experimental settings are detailed in Appendix A.2. We focus on the proportion of different sensitive features across different nodes, analyzing how these distributions are affected by GT training. This analysis includes the proportion of different sensitive features among all nodes, higher-order nodes, and prediction results. We normalize the minority subgroups as 1. Thus, if $|\mathbf{H}[i, s] = 1| > |\mathbf{H}[i, s] = 0|$, for $i \in \{1, 2, \dots, n\}$:

$$\text{All Nodes}(\mathbf{H}[i, s] = 1) = \frac{|\mathbf{H}[i, s] = 1|}{|\mathbf{H}[i, s] = 0|}, \quad (7)$$

$$\text{All Nodes}(\mathbf{H}[i, s] = 0) = 1,$$

and, if $|\mathbf{H}[i, s] = 0| > |\mathbf{H}[i, s] = 1|$:

$$\text{All Nodes}(\mathbf{H}[i, s] = 1) = 1,$$

$$\text{All Nodes}(\mathbf{H}[i, s] = 0) = \frac{|\mathbf{H}[i, s] = 0|}{|\mathbf{H}[i, s] = 1|}. \quad (8)$$

The definition is similar in higher-order nodes and prediction results, and the details are shown in Appendix B.

In Table 1, higher proportions of sensitive features will be highlighted in red ($\mathbf{H}[i, s] = 1$) and blue ($\mathbf{H}[i, s] = 0$) for clarity. When the overall distribution of sensitive features among all nodes differs from that of higher-order nodes, GTs tend to learn to predict based on the distribution of higher-order nodes. It may even lead to the opposite prediction. The majority population in both distributions is the same,

the proportion of sensitive features in the prediction result will align more closely with the higher-order nodes. Consequently, the proportions of sensitive features among higher-order nodes can affect the GT fairness.

Theoretical Findings

We denote the similarity between the distribution of original sensitive features and the distribution of sensitive features when they are mapped to node embeddings as sensitive feature similarity. Node embeddings are computed by the global attention mechanism. We select the Euclidean Norm to measure similarity. A higher similarity indicates greater independence of sensitive attributes during training, aligning with algorithmic fairness. This analysis helps us understand how the attention mechanism impacts the preservation or alteration of sensitive feature distributions during the embedding process. By comparing these distributions, we can assess whether a GT model introduces or mitigates bias, thereby uncovering the underlying causes of fairness issues.

Theorem 1 *The sensitive feature similarity is bounded by the attention scores between nodes with different sensitive features.*

$$\|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 \leq \sum_{u, v \in \mathcal{V}} \sum_{\mathbf{H}[u, s] \neq \mathbf{H}[v, s]} \hat{\mathbf{A}}[u, v] \quad (9)$$

Proof. Based on the Triangle Inequality:

$$\begin{aligned} & \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 \\ &= \sqrt{\sum_{u \in \mathcal{V}} (\mathbf{H}[u, s] - \sum_{v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] \mathbf{H}[v, s])^2} \\ &\leq \sum_{u \in \mathcal{V}} \sqrt{(\mathbf{H}[u, s] - \sum_{v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] \mathbf{H}[v, s])^2} \\ &= \sum_{u \in \mathcal{V}} \|\mathbf{H}[u, s] - \sum_{v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] \mathbf{H}[v, s]\|_2 \\ &= \sum_{u \in \mathcal{V}} \left\| \sum_{v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right\|_2 \\ &\leq \sum_{u, v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] \|\mathbf{H}[u, s] - \mathbf{H}[v, s]\|_2 \\ &= \sum_{u, v \in \mathcal{V}} \sum_{\mathbf{H}[u, s] \neq \mathbf{H}[v, s]} \hat{\mathbf{A}}[u, v]. \end{aligned}$$

□

Since global attention tends to concentrate on higher-order nodes (Xing et al. 2024), algorithmic fairness is predominantly influenced by these higher-order nodes with different sensitive features.

Next, we consider GT with partitioning. We denote $\hat{\mathbf{A}}'$ as the GT attention score matrix with graph partitioning. We assume even partitions, and the number of partitions (i.e. clusters) is c . Let α_{pq} represent the attention score between clusters p and q . Then the approximate attention score between node u and v (in different clusters) with graph partitioning

can be expressed as $\hat{\mathbf{A}}'[u, v] = \frac{\alpha_{pq}}{n/c}$ (Xing et al. 2024). Thus,

$$\hat{\mathbf{A}}'[u, v] = \begin{cases} \hat{\mathbf{A}}[u, v], & u \in \mathcal{V}_p, v \in \mathcal{V}_p \\ \frac{\alpha_{pq}}{n/c}, & u \in \mathcal{V}_p, v \in \mathcal{V}_q \end{cases} \quad (10)$$

Lemma 1 The sensitive feature similarity are lower than \sqrt{n} , whether using graph partitioning.

$$\begin{aligned} \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 &\leq \sqrt{n}, \\ \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 &\leq \sqrt{n}. \end{aligned} \quad (11)$$

The proof of **Lemma 1** is shown in Appendix C.

Theorem 2 *The change in sensitive feature similarity in GT with graph partitioning is constrained by the attention mechanism between different clusters.*

$$\begin{aligned} & \left| \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 - \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 \right| \\ &\geq \frac{1}{2\sqrt{n}} \left| \sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_q} \hat{\mathbf{A}}[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s])^2 \right) \right. \\ &\quad \left. - \sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_q} \hat{\mathbf{A}}'[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s])^2 \right) \right|, \end{aligned} \quad (12)$$

where $p \neq q$.

Proof. Based on **Lemma 1**,

$$\|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 + \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 \leq 2\sqrt{n}.$$

In addition,

$$\|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 \geq 0, \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 \geq 0.$$

Thus,

$$\begin{aligned} & \left| \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 - \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 \right| \\ &= \frac{\left| \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2^2 - \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2^2 \right|}{\|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 + \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2} \\ &\geq \frac{1}{2\sqrt{n}} \left| \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2^2 - \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2^2 \right| \\ &= \frac{1}{2\sqrt{n}} \left| \left(\sum_{u \in \mathcal{V}} \left(\sum_{v \in \mathcal{V}} \hat{\mathbf{A}}[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right)^2 \right) \right. \\ &\quad \left. - \sum_{u \in \mathcal{V}} \left(\sum_{v \in \mathcal{V}} \hat{\mathbf{A}}'[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right)^2 \right|. \end{aligned}$$

Because,

$$\begin{aligned} & \sum_{u \in \mathcal{V}_p} \sum_{v \notin \mathcal{V}_p} \hat{\mathbf{A}}[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s]) = 0, \\ & \sum_{u \in \mathcal{V}_p} \sum_{v \notin \mathcal{V}_p} \hat{\mathbf{A}}'[u, v] (\mathbf{H}[u, s] - \mathbf{H}[v, s]) = 0. \end{aligned}$$

In addition, based on equation (10), if $u \in \mathcal{V}_p, v \in \mathcal{V}_p$:

$$\hat{\mathbf{A}}[u, v] = \hat{\mathbf{A}}'[u, v].$$

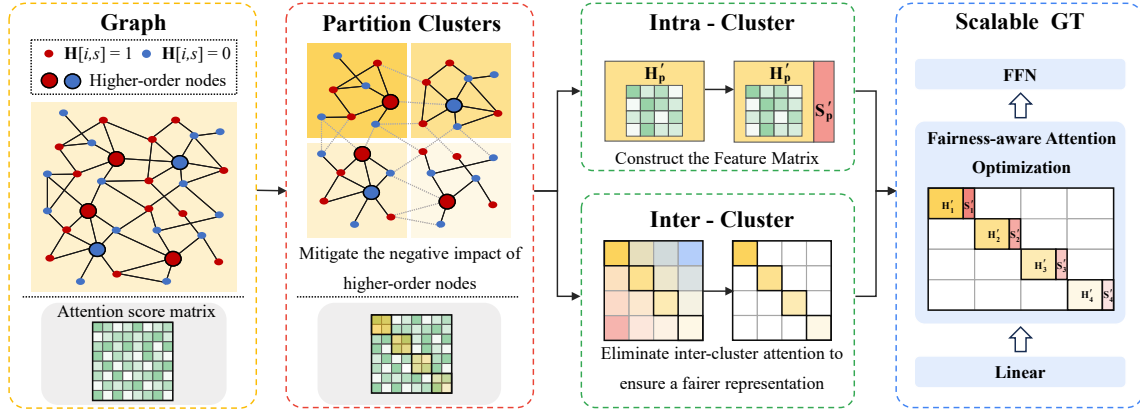


Figure 2: The illustration of FairGP.

Thus,

$$\begin{aligned}
& \left| \|\mathbf{H}[:, s] - \hat{\mathbf{A}}\mathbf{H}[:, s]\|_2 - \|\mathbf{H}[:, s] - \hat{\mathbf{A}}'\mathbf{H}[:, s]\|_2 \right| \\
& \geq \frac{1}{2\sqrt{n}} \left| \left(\sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_p} \hat{\mathbf{A}}[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right) \right. \right. \\
& + \sum_{v \notin \mathcal{V}_p} \hat{\mathbf{A}}[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s])^2 \\
& - \left. \left(\sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_p} \hat{\mathbf{A}}'[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right) \right. \right. \\
& + \left. \left. \sum_{v \notin \mathcal{V}_p} \hat{\mathbf{A}}'[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s])^2 \right) \right| \\
& = \frac{1}{2\sqrt{n}} \left| \sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_q} \hat{\mathbf{A}}[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right)^2 \right. \\
& - \left. \sum_{u \in \mathcal{V}_p} \left(\sum_{v \in \mathcal{V}_q} \hat{\mathbf{A}}'[u, v](\mathbf{H}[u, s] - \mathbf{H}[v, s]) \right)^2 \right|.
\end{aligned}$$

□

The sensitive feature similarity bound is maximized when $\sum_{u \in \mathcal{V}_p} \sum_{v \in \mathcal{V}_q} \hat{\mathbf{A}}'[u, v] = 0$ ($p \neq q$). Thus, setting inter-cluster attention scores to zero enhances GT fairness.

FairGP

We are now ready to present our model FairGP to enhance GT fairness with graph partitioning. As shown in Figure 2, we first encode the structural topology by constructing the node feature matrix. Next, we partition the graph into different clusters, which helps reduce attention computation time complexity and enhance fairness. Finally, we mitigate inter-cluster attention to address fairness issues arising from the higher-order node bias.

Constructing the Feature Matrix

For the node feature matrix, we select the eigenvectors corresponding to the t largest eigenvalues to construct the structure matrix $\mathbf{S} \in \mathbb{R}^{n \times t}$. Then, we combine the original feature matrix \mathbf{H} with the structure matrix \mathbf{S} to preserve both node features and structural information:

$$\mathbf{H}' = \mathbf{H} \parallel \mathbf{S}, \quad (13)$$

where \parallel represents the concatenation operator, and $\mathbf{H}' \in \mathbb{R}^{n \times (d+t)}$ denotes the fused feature matrix. FairGT (Luo et al. 2024c) shows that the largest eigenvalues are closely tied to structural information, which tends to be fairer. Thus, the feature matrix contributes to a fairer encoding.

Graph Partitioning

According to the empirical observations and **Theorem 1**, graph partitioning improves GT fairness by reducing the biased influence of higher-order nodes. FairGP partitions the graph into c non-overlapping clusters using METIS (Karypis and Kumar 1998). Let $\mathcal{G}_p = \{\mathcal{V}_p, \mathcal{E}_p\}$ represents a subgraph of \mathcal{G} , satisfying $\bigcup_{p=1}^c \mathcal{G}_p = \mathcal{G}$ and $\bigcap_{p=1}^c \mathcal{G}_p = \emptyset$.

Fairness-aware Attention Optimization

GTs utilize global attention mechanism to capture information between node pairs, causing the higher-order node bias. To obtain the updated hidden representations $\mathbf{I} \in \mathbb{R}^{n \times n}$.

$$\mathbf{I} = \text{FFN} \left(\text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{n}} \right) \mathbf{V} \right), \quad (14)$$

$$\mathbf{Q} = \mathbf{H}'\mathbf{W}^\mathbf{Q}, \mathbf{K} = \mathbf{H}'\mathbf{W}^\mathbf{K}, \mathbf{V} = \mathbf{H}'\mathbf{W}^\mathbf{V},$$

where $\mathbf{W}^\mathbf{Q}$, $\mathbf{W}^\mathbf{K}$, and $\mathbf{W}^\mathbf{V} \in \mathbb{R}^{n \times n}$. The $\mathbf{W}^\mathbf{Q}$, $\mathbf{W}^\mathbf{K}$, and $\mathbf{W}^\mathbf{V}$ are trainable weights of the linear layers in the Transformer, and FFN represents a Feed-Forward Neural Network.

FairGP separates the attention into inter-cluster and intra-cluster. For the intra-cluster attention $\hat{\mathbf{A}}'[u, v]$ where $u, v \in \mathcal{V}_p$, the attention is the same as the whole attention score:

$$\hat{\mathbf{A}}'[u, v] = \hat{\mathbf{A}}[u, v]. \quad (15)$$

Methods	Pokec-z-R				Pokec-z-G			
	ACC(%) ↑	AUC(%) ↑	$\Delta_{SP}(\%)$ ↓	$\Delta_{EO}(\%)$ ↓	ACC(%) ↑	AUC(%) ↑	$\Delta_{SP}(\%)$ ↓	$\Delta_{EO}(\%)$ ↓
GCN	66.24 ± 2.12	72.78 ± 0.37	7.32 ± 1.48	7.60 ± 1.87	65.09 ± 0.92	71.29 ± 0.73	6.68 ± 2.57	2.31 ± 0.61
GAT	66.01 ± 0.69	70.76 ± 1.30	4.71 ± 1.05	2.72 ± 0.85	66.49 ± 0.23	70.14 ± 0.19	11.21 ± 1.56	7.02 ± 1.17
APPNP	65.24 ± 1.26	70.91 ± 1.46	4.52 ± 1.02	1.78 ± 0.34	64.91 ± 0.60	71.08 ± 2.18	6.17 ± 2.40	3.44 ± 1.90
FairGNN	64.04 ± 0.90	71.75 ± 1.65	4.95 ± 0.21	4.29 ± 0.20	66.21 ± 2.21	70.12 ± 0.39	2.39 ± 0.48	2.62 ± 0.32
FairSIN	67.76 ± 0.71	72.49 ± 1.54	1.49 ± 0.74	<u>0.59 ± 0.50</u>	65.91 ± 1.04	70.12 ± 0.39	2.00 ± 2.49	2.62 ± 1.54
FMP	64.96 ± 0.50	72.42 ± 0.40	0.81 ± 0.40	1.73 ± 1.03	65.13 ± 1.12	70.21 ± 1.82	2.39 ± 0.84	2.26 ± 0.23
FUGNN	<u>68.38 ± 0.43</u>	72.53 ± 1.13	<u>0.53 ± 0.27</u>	1.32 ± 0.95	66.64 ± 1.45	71.93 ± 1.14	3.83 ± 0.81	2.79 ± 0.43
DIFFormer	66.42 ± 1.18	<u>73.27 ± 1.15</u>	1.63 ± 1.13	1.94 ± 0.46	66.17 ± 1.16	71.39 ± 2.26	2.50 ± 0.04	<u>1.64 ± 0.40</u>
SGFormer	65.43 ± 0.11	72.41 ± 0.85	2.32 ± 0.06	1.76 ± 0.79	66.29 ± 0.70	<u>71.50 ± 0.48</u>	<u>1.91 ± 1.02</u>	2.33 ± 0.82
Polynormer	65.22 ± 1.35	71.89 ± 0.44	2.34 ± 0.06	2.40 ± 0.42	65.82 ± 0.58	69.94 ± 0.53	5.49 ± 0.41	2.59 ± 1.03
CoBFormer	66.72 ± 1.04	73.00 ± 0.44	2.45 ± 0.49	1.89 ± 0.39	67.23 ± 0.39	70.57 ± 0.18	2.40 ± 0.07	2.38 ± 1.32
FairGT	OOM	OOM	OOM	OOM	OOM	OOM	OOM	OOM
FairGP	68.50 ± 0.43	73.83 ± 0.25	0.33 ± 0.12	0.27 ± 0.16	<u>66.82 ± 0.03</u>	71.00 ± 0.47	1.18 ± 0.24	0.40 ± 0.24

Methods	Credit				AMiner-L			
	ACC(%) ↑	AUC(%) ↑	$\Delta_{SP}(\%)$ ↓	$\Delta_{EO}(\%)$ ↓	ACC(%) ↑	AUC(%) ↑	$\Delta_{SP}(\%)$ ↓	$\Delta_{EO}(\%)$ ↓
GCN	73.87 ± 2.48	<u>61.45 ± 2.34</u>	12.86 ± 1.84	10.63 ± 0.24	92.61 ± 0.74	90.87 ± 1.19	7.21 ± 0.17	4.62 ± 1.14
GAT	68.29 ± 1.84	60.86 ± 1.43	9.74 ± 2.04	6.42 ± 0.26	90.31 ± 0.41	90.46 ± 1.91	6.57 ± 0.20	3.40 ± 1.17
APPNP	74.19 ± 0.79	61.06 ± 1.08	13.30 ± 0.94	9.47 ± 1.86	<u>92.75 ± 1.29</u>	90.49 ± 0.67	7.00 ± 0.07	4.00 ± 0.64
FairGNN	68.29 ± 2.25	61.73 ± 1.39	9.74 ± 0.28	8.83 ± 0.46	91.51 ± 1.86	90.16 ± 0.29	5.65 ± 0.34	3.78 ± 0.32
FairSIN	<u>78.09 ± 0.27</u>	60.36 ± 0.79	1.85 ± 0.42	1.23 ± 0.28	OOM	OOM	OOM	OOM
FMP	74.51 ± 1.41	59.94 ± 0.06	2.31 ± 0.44	1.27 ± 0.01	92.95 ± 0.11	91.11 ± 1.11	6.81 ± 0.29	<u>3.23 ± 0.78</u>
FUGNN	77.02 ± 0.07	59.26 ± 1.24	<u>0.62 ± 0.48</u>	<u>0.18 ± 0.09</u>	90.85 ± 1.70	90.56 ± 0.32	<u>4.57 ± 0.60</u>	3.45 ± 0.04
DIFFormer	77.41 ± 0.85	60.94 ± 1.77	1.94 ± 0.72	2.47 ± 0.15	90.43 ± 0.86	90.42 ± 0.15	5.81 ± 0.79	3.60 ± 0.83
SGFormer	75.85 ± 0.69	60.95 ± 0.87	2.27 ± 0.85	2.32 ± 0.27	90.81 ± 0.66	90.02 ± 0.66	5.48 ± 1.38	3.27 ± 0.27
Polynormer	76.13 ± 0.83	59.95 ± 2.61	1.81 ± 0.14	1.67 ± 0.10	91.75 ± 1.27	<u>91.69 ± 0.13</u>	6.87 ± 1.15	3.98 ± 0.60
CoBFormer	76.40 ± 0.21	60.84 ± 2.41	2.92 ± 0.96	2.39 ± 0.60	90.00 ± 1.73	89.11 ± 0.14	5.00 ± 1.08	4.02 ± 0.25
FairGT	77.85 ± 0.64	61.38 ± 0.28	1.89 ± 0.59	2.05 ± 1.02	OOM	OOM	OOM	OOM
FairGP	78.13 ± 0.08	60.05 ± 2.98	0.25 ± 0.13	0.14 ± 0.03	91.35 ± 1.28	<u>91.12 ± 1.19</u>	4.44 ± 0.13	2.97 ± 0.37

Table 2: Comparison on utility (ACC and AUC) and fairness (Δ_{SP} and Δ_{EO}) in percentage (%). ↑ denotes the larger, the better; ↓ denotes the opposite. The best results are red and bold, runner-ups are blue and underlined. OOM means out-of-memory.

According to **Theorem 2**, the attention between different clusters is zero, which helps improve fairness. To achieve this goal, we optimize the attention mechanism accordingly:

$$\hat{\mathbf{A}}'[u, v] = 0, \quad u \in \mathcal{V}_p \quad \text{and} \quad v \in \mathcal{V}_q. \quad (16)$$

This can nullify inter-cluster attention, enhancing GT fairness. By focusing on intra-cluster attention and eliminating inter-cluster attention, FairGP ensures a fairer representation of the graph’s structural information.

Complexity of FairGP

Dividing the graph into smaller clusters can reduce the computational burden while preserving essential structural information. For each cluster, the time and space complexity is $O(\left(\frac{n}{c}\right)^2)$, making FairGP’s overall complexity $O(\frac{1}{c^2}n^2)$. This makes our approach scalable and efficient, ensuring that FairGP can handle large-scale graphs effectively.

Experiments

Datasets and Implementation Details

In our experiment, the task is node classification, tested on four datasets: **Credit**, **Pokec-z**, **Pokec-n**, and **AMiner-L**. Specifically, the datasets are labelled as **Pokec-z-R** and **Pokec-n-R** when living region is the sensitive feature, and

Pokec-z-G and **Pokec-n-G** when gender is the sensitive feature. More details are shown in Appendix D.

Baselines

We use three types of baseline: **GNNs** (GCN (Kipf and Welling 2017), GAT (Velickovic et al. 2018), APPNP (Klicpera, Bojchevski, and Gunnemann 2019)), **fairness-aware GNNs** (FairGNN (Dai and Wang 2023), FairSIN (Yang et al. 2024), FMP (Jiang et al. 2024), FUGNN (Luo et al. 2024b)), and **GTs** (DIFFormer (Wu et al. 2023a), SGFormer (Wu et al. 2023b), Polynormer (Deng, Yue, and Zhang 2024), CoBFormer (Xing et al. 2024), FairGT (Luo et al. 2024c)).

Multi-Class Sensitive Features

There are multi-class sensitive features in dataset **AMiner-L**. The previous fairness evaluation metrics are formulated for binary scenarios, but they can be easily extended into multi-class sensitive feature scenarios following previous work (Jiang et al. 2023). The main idea is to guarantee fairness across all pairs of sensitive subgroups. Quantification strategies can be applied by leveraging either the variance across different sensitive subgroups. Specifically, the fairness evaluation metrics in multi-class sensitive features are

defined as follows:

$$\begin{aligned}\Delta_{SP} &= \text{Var}_{i=1}^m (|\mathbb{P}(\hat{y} = 1|s = i)|), \\ \Delta_{EO} &= \text{Var}_{i=1}^m (|\mathbb{P}(\hat{y} = 1|y = 1, s = i)|),\end{aligned}\quad (17)$$

where class number denotes m .

Comparison Results

Table 2 offers a detailed comparison of the fairness evaluation metrics for our proposed FairGP method against various baseline models across four real-world datasets. Further comparison results for other datasets are included in Appendix E. Notably, **Pokec-z-R** and **Pokec-z-G** are the same datasets examined under two different sensitive features. FairGP effectively addresses the multi-sensitive feature issue, as confirmed by the correlation results. We present metrics such as overall Accuracy, AUC, Δ_{SP} , and Δ_{EO} . The table clearly indicates that FairGP consistently achieves superior fairness, validating the efficacy of our approach in fairness-aware node classification. FairGP not only improves fairness compared to traditional GTs but also surpasses existing fairness-aware GNN methods in fairness metrics. This highlights its effectiveness in addressing fairness-related issues in scalable GTs, demonstrating the practicality of FairGP for real-world applications.

Ablation Study

FairGP contains three key components: feature matrix, graph partitioning and attention optimization. To assess their contributions, we conducted an ablation study, removing each component independently to evaluate its impact on prediction fairness.

Dataset	Metric	FairGP	w/o FM	w/o GP	w/o AO
Pokec-z-R	ACC \uparrow	68.50 ± 2.01	68.16 ± 1.06	66.45 ± 1.98	68.20 ± 2.35
	AUC \uparrow	73.83 ± 0.25	73.16 ± 1.23	71.97 ± 2.24	73.11 ± 1.16
	Δ_{SP} \downarrow	0.33 ± 0.12	1.07 ± 0.31	2.14 ± 0.69	2.44 ± 1.36
	Δ_{EO} \downarrow	0.27 ± 0.16	0.65 ± 0.39	5.44 ± 0.92	1.19 ± 0.61
Pokec-z-G	ACC \uparrow	66.82 ± 0.03	66.21 ± 0.74	64.83 ± 1.90	66.52 ± 0.24
	AUC \uparrow	71.00 ± 0.47	70.57 ± 0.41	71.61 ± 3.16	70.26 ± 0.80
	Δ_{SP} \downarrow	1.18 ± 0.40	3.14 ± 0.31	5.84 ± 0.69	2.37 ± 0.59
	Δ_{EO} \downarrow	0.40 ± 0.24	2.47 ± 0.90	4.99 ± 1.04	1.85 ± 0.39
AMiner-L	ACC \uparrow	91.35 ± 1.28	91.03 ± 0.33	89.47 ± 0.70	90.47 ± 0.76
	AUC \uparrow	91.12 ± 1.19	91.21 ± 0.29	93.76 ± 0.33	90.60 ± 0.80
	Δ_{SP} \downarrow	4.44 ± 0.13	5.05 ± 0.21	5.95 ± 0.85	5.06 ± 0.12
	Δ_{EO} \downarrow	2.97 ± 0.37	3.18 ± 0.58	3.32 ± 0.60	4.29 ± 1.13

Table 3: Ablation study of FairGP.

The FairGP without graph partitioning is denoted as **w/o GP**. Since GT without graph partitioning encounters out-of-memory issues in large-scale graphs, we employ mini-batch processing, a typical scalable technique, to keep it running.

Even when compared to results using the mini-batch technique, FairGP shows a significant improvement in fairness, with enhancements of at least 25% across all datasets, and in some instances, exceeding 90%. These results validate our findings that graph partitioning improves GT fairness.

The FairGP with a feature matrix that excludes structural features is denoted as **w/o FM**, and without attention optimization is denoted as **w/o AO**. Notably, when comparing **w/o FM**, and **w/o AO**, FairGP consistently outperforms them in terms of accuracy and fairness. The results shown in Table 3 reinforce the necessity of graph partitioning and attention optimization within FairGP. More ablation experiments for other datasets are shown in Appendix F.

Training Cost Comparison

FairGP, partitioning a graph into clusters, exhibits efficiency compared to GTs. To validate the efficiency of FairGP, we compare its training time with GT baselines. For a fair comparison, we standardize key parameters across all methods, setting the number of hidden dimensions to 128, the number of layers to 1, and the number of heads to 1. In addition, all models are evaluated with a runtime of 100 epochs, and the units are in seconds. The runtimes of different sensitive features in the same dataset are identical. Therefore, we combine the results of **Pokec-z-R** and **Pokec-z-G** into **Pokec-z**. The same approach is applied to **Pokec-n**. The runtimes of FairGP and GT baselines are shown in Table 4, and the runtimes of computing **S** are shown in Appendix G.

	DIFFormer	SGFormer	Polynormer	CoBFormer	FairGP
Pokec-z	138.73	80.60	18.34	15.66	6.88
Pokec-n	124.43	79.39	15.69	11.38	6.14
Aminer-L	830.44	708.34	70.03	66.84	32.77

Table 4: Runtimes (s).

Conclusion

We showed how graph partitioning addresses GT fairness issues on large-scale graphs. We proposed FairGP, a novel algorithm to enhance scalable GT fairness. FairGP is built on graph partitioning, while graph structure information is encoded in the node features to enhance the sensitive feature similarity and the attention computation is adapted to decrease the negative impact of global attention. We also present an empirical study of the efficacy of the proposed approach. In diverse real-world datasets, FairGP earns the best Δ_{SP} and Δ_{EO} comparing state-of-the-art methods. Moreover, in the vast majority of cases, FairGP not only improved fairness but also achieved leading performance, further demonstrating its effectiveness and superiority in large-scale graph applications. In future work, we will further expand FairGP to address situations with limited sensitive features. Since FairGP relies on access to the entire graph for partitioning, it is less suited for applications where privacy preservation or significant misinformation is a concern. In the future, we aim to extend FairGP to be applicable in attribute-constrained and more complex scenarios.

References

- Agarwal, N.; Sirohi, A. K.; Kumar, S.; and Jayadeva. 2024. No Prejudice! Fair Federated Graph Neural Networks for Personalized Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10): P10008.
- Chen, J.; Gao, K.; Li, G.; and He, K. 2023. NAGphormer: A Tokenized Graph Transformer for Node Classification in Large Graphs. In *International Conference on Learning Representations*.
- Dai, E.; and Wang, S. 2023. Learning Fair Graph Neural Networks with Limited and Private Sensitive Attribute Information. *IEEE Transactions on Knowledge and Data Engineering*, 35(7): 7103–7117.
- Dao, T.; Fu, D.; Ermon, S.; Rudra, A.; and Re, C. 2022. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness. In *Advances in Neural Information Processing Systems*.
- Deng, C.; Yue, Z.; and Zhang, Z. 2024. Polynormer: Polynomial-Expressive Graph Transformer in Linear Time. In *International Conference on Learning Representations*.
- Dong, Y.; Ma, J.; Wang, S.; Chen, C.; and Li, J. 2023. Fairness in Graph Mining: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 35(01): 10583–10602.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. S. 2012. Fairness through Awareness. In *Proceedings of Innovations in Theoretical Computer Science*.
- Guo, D.; Chu, Z.; and Li, S. 2023. Fair Attribute Completion on Graph with Missing Attributes. In *International Conference on Learning Representations*.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*.
- Jiang, Z.; Han, X.; Fan, C.; Liu, Z.; Mostafavi, A.; and Hu, X. 2023. Fairness and Explainability: Bridging the Gap towards Fair Model Explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Jiang, Z.; Han, X.; Fan, C.; Liu, Z.; Mostafavi, A.; and Hu, X. 2024. Chasing Fairness in Graphs: A GNN Architecture Perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Karypis, G.; and Kumar, V. 1998. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. *SIAM Journal on Scientific Computing*, 20(1): 359–392.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
- Klicpera, J.; Bojchevski, A.; and Gunnemann, S. 2019. Predict then Propagate: Graph Neural Networks Meet Personalized PageRank. In *International Conference on Learning Representations*.
- Ling, H.; Jiang, Z.; Luo, Y.; Ji, S.; and Zou, N. 2023. Learning Fair Graph Representations via Automated Data Augmentations. In *International Conference on Learning Representations*.
- Luo, R.; Huang, H.; Lee, I.; Xu, C.; Qi, J.; and Feng, X. 2024a. FairGP: A Scalable and Fair Graph Transformer Using Graph Partitioning. *arXiv preprint arXiv:2412.10669*.
- Luo, R.; Huang, H.; Yu, S.; Han, Z.; He, E.; Zhang, X.; and Xia, F. 2024b. FUGNN: Harmonizing Fairness and Utility Graph Neural Networks. In *Proceedings of the 30th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Luo, R.; Huang, H.; Yu, S.; Zhang, X.; and Xia, F. 2024c. FairGT: A Fairness-aware Graph Transformer. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence*.
- Luo, R.; Tang, T.; Xia, F.; Liu, J.; Xu, C.; Zhang, L. Y.; Xiang, W.; and Zhang, C. 2024d. Algorithmic Fairness: A Tolerance Perspective. *arXiv preprint arXiv:2405.09543*.
- Pan, C.; Xu, J.; Yu, Y.; Yang, Z.; Wu, Q.; Wang, C.; Chen, L.; and Yang, Y. 2024. Towards Fair Graph Federated Learning via Incentive Mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Shehzad, A.; Xia, F.; Abid, S.; Peng, C.; Yu, S.; Zhang, D.; and Verspoor, K. 2024. Graph Transformers: A Survey. *arXiv preprint arXiv:2407.09777*.
- Traag, V. A.; Waltman, L.; and Van Eck, N. J. 2019. From Louvain to Leiden: Guaranteeing Well-connected Communities. *Scientific reports*, 9(1): 1–12.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Wang, Z.; Xie, Y.; Li, Z.; Jia, X.; Jiang, Z.; Jia, A.; and Xu, S. 2024. SimFair: Physics-guided Fairness-aware Learning with Simulation Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wu, Q.; Yang, C.; Zhao, W.; He, Y.; Wipf, D.; and Yan, J. 2023a. DIFFormer: Scalable Graph Transformers Induced by Energy Constrained Diffusion. In *International Conference on Learning Representations*.
- Wu, Q.; Zhao, W.; Yang, C.; Zhang, H.; Nie, F.; Jiang, H.; Bian, Y.; and Yan, J. 2023b. SGFormer: Simplifying and Empowering Transformers for Large-Graph Representations. In *Advances in Neural Information Processing Systems*.
- Xing, Y.; Wang, X.; Li, Y.; Huang, H.; and Shi, C. 2024. Less is More: On the Over-Globalizing Problem in Graph Transformers. In *International Conference on Machine Learning*.
- Yang, C.; Liu, J.; Yan, Y.; and Shi, C. 2024. FairSIN: Achieving Fairness in Graph Neural Networks through Sensitive Information Neutralization. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Yin, S.; and Zhong, G. 2023. LGI-GT: Graph Transformers with Local and Global Operators Interleaving. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence*.

Zhang, B.; Dong, Y.; Chen, C.; Zhu, Y.; Luo, M.; and Li, J. 2024a. Adversarial Attacks on Fairness of Graph Neural Networks. In *International Conference on Learning Representations*.

Zhang, G.; Cheng, D.; Yuan, G.; and Zhang, S. 2024b. Learning Fair Representations via Rebalancing Graph Structure. *Information Processing and Management*, 61: 103570.

Zhang, Z.; Zhang, M.; Yu, Y.; Yang, C.; Liu, J.; and Shi, C. 2024c. Endowing Pre-trained Graph Models with Provable Fairness. In *Proceedings of the Web Conference*.