

Towards Fair Graph Neural Networks via Graph Counterfactual

Zhimeng Guo The Pennsylvania State University zzg5107@psu.edu Jialiang Li New Jersey Institute of Technology jl2356@njit.edu Teng Xiao The Pennsylvania State University tengxiao@psu.edu

Yao Ma Rensselaer Polytechnic Institute may13@rpi.edu Suhang Wang* The Pennsylvania State University szw494@psu.edu

ABSTRACT

Graph neural networks have shown great ability in representation (GNNs) learning on graphs, facilitating various tasks. Despite their great performance in modeling graphs, recent works show that GNNs tend to inherit and amplify the bias from training data, causing concerns of the adoption of GNNs in high-stake scenarios. Hence, many efforts have been taken for fairness-aware GNNs. However, most existing fair GNNs learn fair node representations by adopting statistical fairness notions, which may fail to alleviate bias in the presence of statistical anomalies. Motivated by causal theory, there are several attempts utilizing graph counterfactual fairness to mitigate root causes of unfairness. However, these methods suffer from non-realistic counterfactuals obtained by perturbation or generation. In this paper, we take a causal view on fair graph learning problem. Guided by the casual analysis, we propose a novel framework CAF, which can select counterfactuals from training data to avoid non-realistic counterfactuals and adopt selected counterfactuals to learn fair node representations for node classification task. Extensive experiments on synthetic and realworld datasets show the effectiveness of CAF. Our code is available at https://github.com/TimeLovercc/CAF-GNN.

CCS CONCEPTS

• Computing methodologies \rightarrow Machine learning.

KEYWORDS

Graph neural networks; Counterfactual fairness; Causal learning

ACM Reference Format:

Zhimeng Guo, Jialiang Li, Teng Xiao, Yao Ma, and Suhang Wang. 2023. Towards Fair Graph Neural Networks via Graph Counterfactual. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23), October 21–25, 2023, Birmingham, United Kingdom.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3583780.3615092

*Corresponding Author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0124-5/23/10...\$15.00 https://doi.org/10.1145/3583780.3615092

1 INTRODUCTION

Graphs are pervasive in real-world, such as knowledge graphs [22], social networks [20] and biological networks [23]. Recently, graph neural networks (GNNs) [18, 35] have shown great ability in modeling graph-structural data. Generally, GNNs adopt the message passing mechanism, which updates a node's representation by iteratively aggregating its neighbors' representations. The resulting representation preserves both node attributes and local graph structure information, facilitating various downstream tasks such as node classification [11, 18] and link prediction [44]. Despite their great performance, recent studies [5, 19, 25] show that GNNs tend to inherit bias from training data, which may result in biased predictions towards sensitive attributes, such as age, gender and race. In addition, the message passing mechanism of GNNs and graph structure could magnify the bias [5]. For example, in social networks, nodes of the same race are more likely to connect to each other. The message passing of GNNs would make the representation of linked nodes similar, resulting in a high correlation of node representation with race, hence the biased prediction. The biased prediction has raised concerns from ethical and societal perspectives, which severely limits the adoption of GNNs in high-stake decision-making systems [38], such as job applicants ranking [28] and criminal prediction [15].

Hence, many efforts have been taken for fair GNNs [5, 19, 25, 29]. However, most existing methods are based on statistical fairness notions, which aim to make statistically fair predictions for different sub-groups or individuals [24]. Several works have pointed out such fairness notions fail to detect discrimination in the presence of statistical anomalies [21, 27]. Therefore, there has been a recent shift toward counterfactual fairness in graph modeling [1, 25]. This approach aims to eradicate the root causes of bias by mapping the causal relationships among variables. The identified causal structure allows for the adjustment of sensitive data to generate counterfactuals, ensuring that the prediction remains unaltered by the sensitive information through the utilization of these counterfactuals. For example, NIFTY [1] perturbs sensitive attributes to obtain counterfactuals and maximizes the similarity between original representations and perturbed representations to make representations invariant to sensitive attributes. GEAR [25] adopts GraphVAE [17] to generate counterfactuals and minimizes the discrepancy between original representations and counterfactual representations to get rid of the influence of sensitive attributes. Despite their superior performance, existing graph counterfactual fairness works need to flip sensitive attributes or generate counterfactuals with GraphVAE, which can easily result in non-realistic counterfactuals. Such non-realistic

counterfactuals may disrupt the underlying latent semantic structure, thereby potentially undermining the model's performance. This is because simply flipping sensitive attributes cannot model the influence on other features or graph structure causally caused by sensitive attributes [1], and the generative approach lacks supervision of real counterfactuals and could be over-complicated [25].

Motivated by the discussion above, in this paper, we investigate whether one can obtain counterfactuals within the training data. For example, if a female applicant was rejected by a college, we aim to find another male applicant who has a similar background as the counterfactual applicant. Thus, we can get realistic counterfactuals and avoid the ill-supervised generation process. To achieve our goal, we are faced with several challenges: (i) Graph data is quite complex, thus it is infeasible to directly find counterfactuals in the original data space. Besides, some guidance or rules are needed to find the counterfactuals. (ii) To achieve graph counterfactual fairness, learned representation should be invariant to sensitive attributes and information causally influenced by sensitive attributes. It is critical to design proper supervision to help models get rid of sensitive information. To tackle the aforementioned challenges, we propose a casual view of the graph, label and sensitive attribute. The causal interpretation guides us to find counterfactuals and learn disentangled representations, where the disentangled content representations are informative to the labels and invariant to the sensitive attributes. Guided by the causal analysis, we propose a novel framework, Counterfactual Augmented Fair GNN (CAF), to simultaneously learn fair node representations for graph counterfactual fairness and keep the performance on node classification tasks. Specifically, based on the causal interpretation, we derive several constraints to enforce the learned representations being invariant across different sensitive attributes. To obtain proper counterfactuals to guide representation learning, we utilize labels and sensitive attributes as guidance to filter out potential counterfactuals in representation space. Our main contributions are:

- We provide a causal formulation of the fair graph learning process and fair node representation learning task.
- We propose a novel framework CAF to learn node representations for graph counterfactual fairness. Specifically, we find counterfactuals in representation space and design novel constraints to learn the content representations.
- We conduct extensive experiments on real-world datasets and synthetic dataset to show the effectiveness of our model.

2 RELATED WORKS

Graph Neural Networks. Graph neural networks (GNNs) have dominated various tasks on graph-structured data, such as node classification [3, 18, 34, 36, 37], graph classification [33] and link prediction [44, 45]. Existing GNNs can be categorized into spatial-based GNNs and spectral-based GNNs. Spatial-based GNNs leverage the graph structure directly, focusing on the relationships between nodes and their immediate neighbors to inform feature learning. On the other hand, spectral-based GNNs operate in the spectral domain defined by the graph Laplacian and its eigenvectors, making them better suited to capture global properties of the graph. The superior performance of GNNs has greatly extended their application scenarios [10]. For example, banks may leverage GNNs

to process transaction networks to detect the abnormal behavior of users [6]. The applications in critical decision-making systems place higher requirements for GNNs, such as being fair and interpretable [42]. Despite their extensive utility and efficacy, recent studies [4, 5, 19] show that GNNs can harbor implicit biases on different groups, which can lead to skewed or unfair outcomes. This bias issue is particularly critical when GNNs are deployed in high-stake scenarios, making it necessary to ensure fairness in the modeling process [26]. Thus, mitigating bias and promoting fairness in GNNs are active and necessary research areas [6]. The source of bias in Graph Neural Networks (GNNs) primarily originates from two areas. First, it comes from the inherent bias in the input data, which may contain unequal representation or prejudiced information about nodes or connections in the graph. Second, the bias can stem from the algorithmic design of the GNN itself, which may unintentionally emphasize certain features or connections over others during the learning process. Therefore, there is a trend for the research community to design fairer GNN models to deal with graph-based tasks [5, 6, 14, 24].

Fairness in GNNs. Fairness is a widely-existed issue of machine learning systems [26, 28, 30, 46]. Researchers evaluate the fairness of models with many kinds of fairness notions, including group fairness [12, 43], individual fairness [8] and counterfactual fairness [21]. The metrics can also be used to measure the fairness performance of Graph Neural Networks [1, 25]. The commonly used fairness notions in GNNs are statistical parity [43] and equal opportunity [12]. FairGNN [5] utilizes adversarial training to establish fairness in graph-based models, refining its representation through an adversary tasked with predicting sensitive attributes. EDITS [7], on the other hand, is a pre-processing technique that focuses on ensuring fairness in graph learning. It aims to eliminate sensitive information from the graph data by correcting any inherent biases present within the input network. However, these methods and their metrics are developed based on correlation [27], which has been found to be unable to deal with statistical anomalies, such as Simpson's paradox [32]. Based on the causal theory, counterfactual fairness can model the causal relationships and gets rid of the correlation-induced abnormal behavior [21, 27]. There is an increasing interest to apply counterfactual fairness on graphs to design fairer GNNs [1, 25]. NIFTY [1] perturbs sensitive attributes for each node to obtain counterfactuals and omits the causal relationships among variables. GEAR [25] uses GraphVAE [17] to generate the graph structure and node features causally caused by the sensitive attributes. For more details about counterfactual learning on graphs, please refer to the survey [9].

Our paper is inherently different from existing work: (i) Unlike existing works that might generate unrealistic counterfactuals, our work avoids the generation process and selects counterfactuals with sensitive attributes and labels as guidance; and (ii) We propose a causal view to understand the source of bias. Based on the causal interpretation, we also design several constraints to help our model learn the fair node representations.

3 PRELIMINARIES

In this section, we start by introducing the necessary notation and defining the problem at hand. Following this, we employ the Structural Causal Model to frame the issue, which will then motivate our solution - the disentangled fair representation learning method.

3.1 Notations and Problem Definition

Throughout the paper, we use italicized uppercase letters to represent random variables (e.g., S, E) and use italicized lowercase letters to denote the specific value of scalars (e.g., s, y_i). Non-italicized bold lowercase and uppercase letters are used to denote specific values of vectors (e.g., \mathbf{x}_i) and matrices (e.g., \mathbf{X}), respectively.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ denote an attributed graph, where $\mathcal{V} = \{v_1, ..., v_N\}$ is the set of N nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges, $\mathbf{X} \in \mathbb{R}^{N \times D}$ is the node attribute matrix. The i-th row of \mathbf{X} , i.e., \mathbf{x}_i is the feature vector of node v_i . $\mathbf{A} \in \{0,1\}^{N \times N}$ is the adjacency matrix of the graph \mathcal{G} , where $\mathbf{A}_{ij} = 1$ if nodes v_i and v_j are connected; otherwise $\mathbf{A}_{ij} = 0$. We use $\mathbf{s} \in \{0,1\}^{N \times 1}$ to denote the sensitive attributes, where s_i is the sensitive attribute of v_i . Following [25], we only consider binary sensitive attributes and leave the extension of multi-category sensitive attributes as future work. We use $\mathbf{y} \in \{1, ..., c\}^{N \times 1}$ to denote the ground-truth node labels, where y_i is the label of v_i . In this paper, we assume that both target labels and sensitive attributes are binary variables for convenience.

For the semi-supervised node classification task, only part of nodes $\mathcal{V}_L \in \mathcal{V}$ are labeled for training and the remaining nodes $\mathcal{V}_U = \mathcal{V} \backslash \mathcal{V}_L$ are unlabeled. The goal is to train a classifier f to predict the labels of unlabeled nodes, which has satisfied node classification performance and fairness performance simultaneously. Given X, A and Y_L , the goal of semi-supervised node classification is to learn a mapping function f to predict the labels of unlabeled nodes, i.e., $f: (A, X) \to \mathcal{Y}_U$, where \mathcal{Y}_U the set of predicted labels of unlabeled nodes \mathcal{V}_U .

3.2 The Desiderata for Fair Graph Learning

GNNs have shown remarkable capabilities in the realm of semisupervised node classification. However, they are not immune to bias issues, primarily stemming from imbalanced or prejudiced input data, and potentially from the structural design of the GNNs themselves, which may inadvertently prioritize certain features or connections. Therefore, substantial efforts have been directed towards developing fairness-aware methodologies within GNNs. The majority of these methods strive to ensure correlation-based fairness notions, such as demographic parity or equality of opportunity. However, these correlation-based fairness notions can be inherently flawed, particularly in the presence of statistical anomalies, which calls for more nuanced and robust approaches to achieve fairness in GNNs. Recent advance [27] shows that causal-based fairness notions can help resolve this issue. Thus, to help design a fair GNN classifier, we take a deep causal look under the observed graph. Without loss of generality, in this work, we focus on the node classification task and construct a Structural Causal Model [31] in Figure 1. It presents the causal relationships among five variables: sensitive attribute S, ground-truth label Y, environment feature E, content feature C and ego-graph G for each node. Each link denotes a deterministic causal relationship between two variables. We list the following explanations for the SCM:

 S → E. The variable E denotes latent environment features that are determined by the sensitive attribute S. For example, people

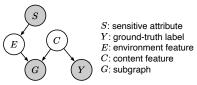


Figure 1: Structural Causal Model for model prediction. We use white color to denote latent variables and use gray color to denote the observed variables.

of different genders will have different heights or other physical characteristics, where *S* is the sensitive attribute of genders and *E* is physical characteristics that are causally determined by the sensitive attribute. This relationship will lead to bias in latent feature space, which we will explain shortly.

- C → Y. The variable C denotes the content feature that determines ground-truth label Y. Taking the credit scoring as an example, ideally, we assign credit scores using personal information not related to the sensitive attribute, i.e., we use content feature C instead of E to assign credit score Y.
- *E* → *G* ← *C*. The ego-graph *G* is determined by the content feature *C* and the environment feature *E*, which are two disjoint parts. *E* and *C* are latent features and *G* is the observed ego-graph. Considering a one-hop ego-graph, it contains the social connections of center node and the observed feature of center node. The causal relationship indicates environment feature *E* and content feature *C* can determine one's social connections and personal features (node attributes).

The SCM paves us a way to understand the source of bias and how to design a fair GNN classifier. Next, we give details about source of bias and disentangled learning. Our objective is to approximate the content feature C with a content representation denoted as \hat{C} , and similarly, approximate the environment feature E with an environment representation denoted as \hat{E} . To streamline our discussion, we will slightly abuse notation by also employing the symbols C and E to signify the corresponding content and environment representations throughout the remainder of the paper.

3.2.1 Source of Bias. From the causal graph, we can observe that the sensitive variable S and the label variable Y are independent with each other, i.e., the only path from S to Y, $S \to E \to G \leftarrow C \leftarrow Y$ is blocked by the collider G. However, it is worthy noting that S and Y are **dependent conditioned on** G, i.e.,

$$P(Y, S|G) \neq P(Y|G)P(S|G). \tag{1}$$

The conditional dependency of Y and S on G is one major reason that leads to biased prediction. If we directly learn a GNN model that aims to predict Y based on G, as Y and S are dependent given G, the learned label Y will have correlation with S, resulting in the biased prediction on sensitive attribute S.

Alternatively, we can understand the bias by treating existing GNNs as composed of a feature extractor g and a classifier c. The feature extractor g takes the subgraph centered at a node as input and learns node representation as $\mathbf{z} = g(G)$. Then the classifier c uses the representation \mathbf{z} to predict the label as $\hat{y} = c(\mathbf{z})$. As G is dependent on E and C, the learned representation \mathbf{z} is likely to contain mixed information of both E and C. Hence, the predicted label \hat{y} is also likely to have correlation with S.

3.2.2 Disentangled Fair Representation Learning. From the above analysis motivates, we can observe that in order to have fair prediction, we need to learn disentangled representation E and C to block the path from S to Y conditioned on G, and only use the content information C to predict Y, i.e., P(Y|C). As C determines Y, it contains all the label information to predict Y. Meanwhile, observing E and C can block the conditional path from S to Y, i.e., P(Y,S|E,C,G) = P(Y|C,E,G)P(S|C,E,G). Note that observing C blocks the path from E to Y and the path from E to E. This gives us

$$P(Y, S|E, C, G) = P(Y|C)P(S|E).$$
(2)

The above equation shows that observing E and C would make Y and S independent and P(Y|C) is unbiased.

Hence, if we can learn disentangled latent representation E and C, we would be able to use C for fair classification. However, the main challenge is we do not have ground-truth E and C to help us train a model that can learn disentangled representation. With a slight abuse of notation, we also use C to denote the learned content representation and use E to denote the learned environment representation. Fortunately, we can use the SCM to derive several properties of the optimal representation, which would be used to help learn the latent representation of C and E:

- Invariance: C ⊥ E. This property can be understood in two perspectives. That is, the content representations should be independent to the sensitive attributes and the environment representation induced by the sensitive attribute. Meanwhile, the environment representations should be independent to the labels and the content representation which is informative to the labels.
- Sufficiency: (C, E) → G. The combined representation can used to reconstruct the observed graph.
- *Informativeness:* $C \rightarrow Y$. The content representations should have the capacity to give accurate predictions of labels Y.

4 METHODOLOGY

The causal view suggests us to learn disentangled representation c and e for node v, with c capturing the content information that is useful for label prediction and irrelevant to sensitive attributes, and e capturing the environment information depends on sensitive attribute only. With the disentanglement, c can be used to give fair predictions. However, how to effectively disentangle c and e remains a question given that we do not have ground-truth of disentangled representation. Intuitively, for a node v with sensitive attribute s, its content representation c should remain the same when the sensitive attribute is flipped to 1 - s while its environment representation e should change correspondingly. Hence, if we know the counterfactual of node v, we will be able to utilize the counterfactual to help learn disentangled representation for fair classification; while the counterfactual is not observed. To address the challenges, we propose a novel framework CAF as shown in Figure 2 (a), which is composed of: (i) a GNN encoder that takes ego-graph \mathcal{G} of node v to learn disentangled representation \mathbf{c} and e; (ii) the counterfactual augmentation module, which aims to discover counterfactual for each factual observation and utilize the

counterfactual to help learn disentangled representation; (iii) a fair classifier which takes **c** as input for fair classification. Next, we give the details of each component.

4.1 Disentangled Representation Learning

For each node v_i , the content representation \mathbf{c}_i should capture the important node attribute and neighborhood information for predicting the label while the environment representation \mathbf{e}_i should capture all important information relevant to sensitive attribute. As GNNs have shown great ability in modeling graph structured data, we adopt GNNs to learn \mathbf{c}_i and \mathbf{e}_i . Instead of adopting two GNNs to learn \mathbf{c}_i and \mathbf{e}_i separately, to reduce the number of parameters, we adopt one GNN to learn \mathbf{c}_i and \mathbf{e}_i . We empirically found that using two GNNs and one GNN have similar performance due to constraints we designed to disentangle \mathbf{c}_i and \mathbf{e}_i , which will be introduced later. Specifically, the GNN f_θ parameterized by θ takes \mathcal{G} as input and learns representation as:

$$[C, E] = H = f_{\theta}(A, X), \tag{3}$$

where $\mathbf{H} \in \mathbb{R}^{N \times d}$ is the learned representation matrix with the i-th row, i.e., \mathbf{h}_i , as the representation of node v_i . We treat the first d_c columns as the content representation matrix \mathbf{C} and use the next d_e columns as the environment representation matrix \mathbf{E} . Note that $d = d_c + d_e$. In our implementation, we set $d_c = d_e$. $\mathbf{C} \in \mathbb{R}^{N \times d_c}$ is the content representation matrix with the i-th row, i.e., \mathbf{c}_i , as the content representation of node v_i . Similarly, $\mathbf{E} \in \mathbb{R}^{N \times d_e}$ is the environment representation matrix with the i-the row, i.e., \mathbf{e}_i as the environment representation of node v_i . f_θ is flexible to be various GNNs such as GCN [18] and GraphSAGE [11].

To make sure \mathbf{c}_i captures the content information for fair label prediction, and \mathbf{e}_i and \mathbf{c}_i are disentangled, based on the causal analysis in Section 3, we add following constraints:

Informativeness Constraint. First, the content representation \mathbf{c}_i should be informative to the downstream tasks, i.e., $C \to Y$. Hence, for node v_i , we should be able to get accurate label prediction from \mathbf{c}_i . Thus, we introduce a classifier f_ϕ with model parameter ϕ . It takes \mathbf{c}_i as input and predicts the class distribution of v_i as:

$$\hat{\mathbf{y}}_i = f_\phi(\mathbf{c}_i). \tag{4}$$

The loss function for training the classifier is given as:

$$\mathcal{L}_{\text{pred}} = \frac{1}{|\mathcal{V}_L|} \sum_{v_i \in \mathcal{V}_L} \ell(\hat{\mathbf{y}}_i, \mathbf{y}_i), \tag{5}$$

where y_i is the one-hot encoding of ground-truth label of v_i . $\ell(\hat{y}_i, y_i)$ denotes the cross entropy between \hat{y}_i and y_i .

Sufficiency Constraint. As shown in our causal view, the representation $(\mathbf{c}_i \text{ and } \mathbf{e}_i)$ should be sufficient to reconstruct the observed factual graph \mathcal{G}_i . In disentangled representation learning research, the reconstruction supervision is usually adopted to guide the learning process [13, 40]. However, existing graph counterfactual fairness approaches [1, 25] fail to provide supervision to preserve graph information in the representations. Thus, they put their models at a risk of being stuck in trivial solutions to merely get spurious information in the representations, which contradicts the SCM and is not sufficient to reconstruct the observed graph \mathcal{G}_i . In our model, we formalize the sufficiency constraint as a reconstruction of the graph structure. Specifically, for a pair of nodes (v_i, v_i) , we predict

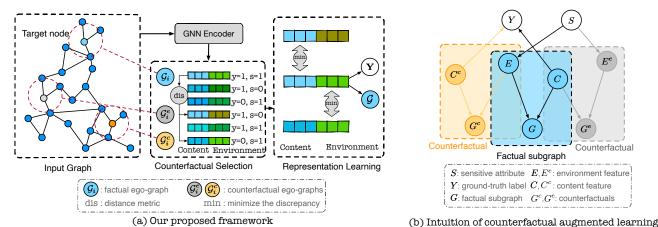


Figure 2: An illustration of (a) our proposed framework; (b) intuition of counterfactual augmented learning.

the link existence probability as $p_{ij} = \sigma(\mathbf{h}_i \mathbf{h}_i^T)$, where $\mathbf{h}_i = [\mathbf{c}_i, \mathbf{e}_i]$ is the node representation of node v_i . The sufficiency constraint is

$$\mathcal{L}_{\text{suf}} = \frac{1}{|\mathcal{E}| + |\mathcal{E}^-|} \sum_{(v_i, v_j) \in \mathcal{E} \cup \mathcal{E}^-} -e_{ij} \log p_{ij} - (1 - e_{ij}) \log p_{ij},$$
(6)

where \mathcal{E}^- is the set of sampled negative edges. $e_{ij} = 1$ if node v_i and v_i are connected; otherwise $e_{i,i} = 0$.

Orthogonal Constraint. The above model can help to learn c_i that captures graph information for label prediction, however, it doesn't guarantee that c_i doesn't contain sensitive attribute information. To make sure that c_i and e_i are disentangled, i.e., c_i doesn't contain any environment information relevant to sensitive attribute, we further impose the orthogonal constraint, i.e., $\mathbf{c}_i^T \mathbf{e}_i = 0$.

4.2 Counterfactual Augmented Learning

As we do not have ground-truth of c_i and e_i , we used several constraints to implicitly supervise the learning of c_i and e_i . To fully learn disentangled c_i and e_i , we propose to learn better e_i and c_i that follows the counterfactual constraints. As shown in Figure 2 (b), generally, for a node v_i with observe the factual sensitive attribute s_i and label y_i , its content representation c_i should remain similar when the sensitive attribute is flipped to $1 - s_i$ but its environment representation e_i should change correspondingly, which forms the counterfactual subgraph \mathcal{G}_{i}^{e} . Similarly, when flip label y_{i} but keep the sensitive attribute s_i unchanged, then v_i 's environment representation e_i remain the same, while its content representation should change accordingly, leading to the counterfactual subgraph \mathcal{G}_{i}^{c} . Thus, if we know \mathcal{G}_{i}^{e} and \mathcal{G}_{i}^{c} , we would be able to use these counterfactual graphs together with factual graph G_i to guide the learning of c_i and e_i . However, in real-world, we can only observe factual graphs. To solve this challenge, we propose to find potential candidate counterfactuals with the observed factual graphs.

The sensitive attribute and label are used to find counterfactuals in our model. Considering the fair credit scoring problem, when someone was assigned a low score, straightforward thinking is to know the results of people who have a similar background to her but of a different gender. For example, Sarah, a female, got a low credit score. Then she may ask, what if I were a male, what will my credit score be? This thinking inspires us to directly find counterfactuals from the observed node samples instead of performing perturbation or generating [1, 25]. The advantages of selecting counterfactuals from the observed node samples are twofold: (1) It avoids making assumptions about the graph generation process with sensitive attributes. (2) It does not need additional supervision signal.

Another problem comes: selecting counterfactuals from the original data space is also challenging due to the complexity of graph distance calculation. To get counterfactual \mathcal{G}_i^e , we need to find some nodes which have different sensitive attribute and the same label. Similarly, we find some nodes with the same sensitive attribute and different labels as counterfactual \mathcal{G}_{i}^{c} . The task can be formalized as:

$$\mathcal{G}_{i}^{c} = \arg\min_{\mathcal{G}_{i} \in \mathbb{G}} \{ m(\mathcal{G}_{i}, \mathcal{G}_{j}) \mid y_{i} \neq y_{j}, s_{i} = s_{j} \}, \tag{7}$$

$$G_i^c = \underset{\mathcal{G}_j \in \mathbb{G}}{\operatorname{arg \, min}} \{ m(\mathcal{G}_i, \mathcal{G}_j) \mid y_i \neq y_j, s_i = s_j \},$$

$$G_i^e = \underset{\mathcal{G}_j \in \mathbb{G}}{\operatorname{arg \, min}} \{ m(\mathcal{G}_i, \mathcal{G}_j) \mid y_i = y_j, s_i \neq s_j \},$$

$$(8)$$

where $\mathbb{G} = \{G_i | v_i \in \mathcal{V}\}\$ and $m(\cdot, \cdot)$ is a metric of measuring the distance between a pair of subgraphs. Nevertheless, the problem of computing the distance of pairs of graphs is inefficient and infeasible due to the complex graph structure and large search space [45]. As we already have node representations $\mathbf{h}_i = [\mathbf{c}_i, \mathbf{e}_i]$ that capture the graph structure and node attribute information, we propose to measure the distance in the latent space, which can greatly reduce the computation burden. Then the counterfactual graph searching problem in Eq. (7) and Eq. (8) is converted to the problem below:

$$\mathbf{h}_{i}^{c} = \arg\min_{\mathbf{h}_{i} \in \mathbb{H}} \{ \|\mathbf{h}_{i} - \mathbf{h}_{j}\|_{2}^{2} | y_{i} \neq y_{j}, s_{i} = s_{j} \}, \tag{9}$$

$$\mathbf{h}_{i}^{c} = \arg\min_{\mathbf{h}_{j} \in \mathbb{H}} \{ \|\mathbf{h}_{i} - \mathbf{h}_{j}\|_{2}^{2} | y_{i} \neq y_{j}, s_{i} = s_{j} \},$$

$$\mathbf{h}_{i}^{e} = \arg\min_{\mathbf{h}_{j} \in \mathbb{H}} \{ \|\mathbf{h}_{i} - \mathbf{h}_{j}\|_{2}^{2} | y_{i} = y_{j}, s_{i} \neq s_{j} \},$$

$$\mathbf{h}_{j} \in \mathbb{H}$$
(10)

where $\mathbb{H} = \{\mathbf{h}_i | v_i \in \mathcal{V}\}$ and we use L2 distance to find counterfactuals. A problem is that we only have limited labels in the training set. So we first pre-train the backbone model. With pretrained model, we can obtain the prediction for unlabeled nodes as pseudo-labels. The pseudo-labels work as the guidance of the counterfactual searching problem. Note that for each factual input we can also get multiple counterfactuals by selecting a set of counterfactuals in Eq. (9) and Eq. (10) instead of one. Thus, the counterfactual \mathcal{G}_{i}^{c} can be naturally extended to a set of K counterfactuals

 $\{\mathcal{G}_i^{c_k}|k=1,...,K\}$ and \mathcal{G}_i^e can be extended to $\{\mathcal{G}_i^{e_k}|k=1,...,K\}$. We fix K to 10 in our implementation.

We can utilize the counterfactuals to supervise the disentanglement of \mathbf{c}_i and \mathbf{e}_i . Specifically, as shown in Figure 2 (b), counterfactual $\mathcal{G}_i^{e_k}$ shares the same content information with factual graph \mathcal{G}_i and has different environment information. Without supervision, the factual content representation $\mathbf{c}_i^{e_k}$ may contain both the counterfactual content representation tinformation. When we minimize the discrepancy of the learned representations with $\mathrm{dis}(\mathbf{c}_i, \mathbf{c}_i^{e_k})$, f_θ will tend to merely keep the content information and squeeze the sensitive information out of learned representations. In a similar manner, we can use $\mathrm{dis}(\mathbf{e}_i, \mathbf{e}_i^{e_k})$ to make the environment representation \mathbf{e}_i be invariant to the content information stored in \mathbf{c}_i . Also, we put the orthogonal constraint here to encourage \mathbf{c}_i and \mathbf{e}_i to store different information in representation space. The invariance constraint is given as:

$$\mathcal{L}_{inv} = \frac{1}{|\mathcal{V}| \cdot K} \sum_{v_i \in \mathcal{V}} \sum_{k=1}^{K} \left[\operatorname{dis}(\mathbf{c}_i, \mathbf{c}_i^{e_k}) + \operatorname{dis}(\mathbf{e}_i, \mathbf{e}_i^{c_k}) + \gamma K \cdot |\operatorname{cos}(\mathbf{c}_i, \mathbf{e}_i)| \right],$$
(11)

where $\operatorname{dis}(\cdot, \cdot)$ is a distance metric, such as the cosine distance and L2 distance in our implementation. $|\cos(\cdot, \cdot)|$ is the absolute value of cosine similarity and we optimize this term to approximate $\mathbf{c}_i^T \mathbf{e}_i = 0$. γ is the hyper-parameter to control the orthogonal constraint.

4.3 Final Objective Function of CAF

Putting the disentangled representation learning module and the counterfactual selection module together, the final objective function of the proposed CAF framework is:

$$\min_{\theta,\phi} \mathcal{L} = \mathcal{L}_{\text{pred}} + \alpha \mathcal{L}_{\text{inv}} + \beta \mathcal{L}_{\text{suf}}, \qquad (12)$$

where θ and ϕ are parameters for the GNN encoder and the prediction head, respectively. α and β are hyper-parameters controlling the invariance constraint and the sufficiency constraint.

4.4 Training Algorithm

The whole process of CAF is summarized in Algorithm 1. Our method relies on the counterfactuals in the representation space to guide the disentanglement. However, the randomly initialized representation at the first several epochs may degrade the performance of our model. Therefore, we first pre-train a plain node representation learning model Y = $g_{\Theta,\Phi}(\mathbf{A},\mathbf{X})$ only with $\mathcal{L}_{\mathrm{pred}}$. Then we use the optimized parameters $\Theta^*, \Phi^* = \min_{\Theta, \Phi} \mathcal{L}_{pred}$ to initialize the parameters θ and ϕ of our model and use the aforementioned framework to get the desired disentangled representations. We do not necessarily update the counterfactuals for each epoch. We update the counterfactuals once for t epochs and t = 10 in our implementation. As shown in Algorithm 1, we first pre-train $g_{\Theta,\Phi}$ and use the optimized parameter to initialize f_{θ} and ϕ from line 1 to line 2. Then we iteratively optimize f_{θ} and ϕ from line 3 to line 10. In each iteration, we first perform forward propagation to get node representations in line 4. And then for each t epoch we update the selected counterfactuals once from line 5 to line 7. Afterwards, we compute the overall objective and perform backpropagation to

Algorithm 1: Training Algorithm of CAF.

```
Input: g = (\mathcal{V}, \mathcal{E}, \mathbf{X}), \mathbf{Y}_L, t, \alpha, \beta, \gamma, K \text{ and } num\_epoch
Output: f_{\theta} and f_{\phi}
  1: Pre-train g_{\Theta,\Phi} based on \mathcal{L}_{\mathrm{pred}} with Eq. (5)
  2: Use the optimized \Theta^* and \Phi^* to initialize f_{\theta} and f_{\phi}
  3: for epoch in range(num_epoch) do
          Compute representations H and predicted labels \hat{\mathbf{Y}} with f_{\theta}
          and f_{\phi} by Eq. (3) and Eq. (4)
          if num\_epoch \% t = 0 then
  5:
              Obtain two sets of counterfactuals \{\mathcal{G}_i^{c_k} \mid k=1,\ldots,K\} and \{\mathcal{G}_i^{e_k} \mid k=1,\ldots,K\} with Eq. (9) and (10).
  6:
  7:
          Compute the overall objective \mathcal{L} with Eq. (12)
          Update \theta and \phi according to the objective \mathcal{L}
 10: end for
 11: return f_{	heta} and f_{\phi}
```

optimize the parameters θ and ϕ from line 8 to line 9. After training, we obtain the desired fair model f_{θ} and f_{ϕ} in line 11.

5 EXPERIMENTS

In this section, we conduct experiments to evaluate the effectiveness of the proposed method and compare it with state-of-the-art fair GNNs. Specifically, we aim to answer the following questions:

- (RQ 1) How effective is the proposed CAF for fair node classification task on both synthetic datasets and real-world datasets?
- (RQ 2) Can the proposed CAF find appropriate counterfactuals?
- (RQ 3) How do the proposed modules work? How can each regularization term affect the model performance?

5.1 Experiment Settings

5.1.1 Real-World Datasets. We conduct experiments on three widely used real-world datasets, namely German Credit [2], Credit Defaulter [41], Bail [16]. The statistics of the datasets can be found in Table 2. The details of the datasets are as follows:

- German Credit [2]: the nodes in the dataset are clients and two
 nodes are connected if they have high similarity of the credit
 accounts. The task is to classify the credit risk level as high or
 low with the sensitive attribute "gender".
- Credit Defaulter [41]: the nodes in the dataset are used to represent the credit card users and the edges are formed based on the similarity of the payments information. The task is to classify the default payment method with sensitive attribute "age".
- Bail [16]: these datasets contain defendants released on bail during 1990-2009 as nodes. The edges between two nodes are connected based on the similarity of past criminal records and demographics. The task is to classify whether defendants are on bail or not with the sensitive attribute "race".

5.1.2 Synthetic Dataset. Real-world datasets do not offer ground-truth counterfactuals, prompting us to construct a synthetic dataset based on the Structural Causal Model (SCM) as depicted in Figure 1. The primary advantage of a synthetic dataset is that it provides us with ground-truth counterfactuals for each node, which

Dataset	Metrics	GCN	GraphSAGE	GIN	FairGNN	EDITS	NIFTY	GEAR	CAF
	AUC (↑)	74.00±1.51	74.54±0.86	72.69±1.02	65.85±9.49	69.76±5.46	72.05±2.15	65.80±3.00	71.87±1.33
	F1 (†)	80.05±1.20	81.15 ± 0.97	$82.62{\scriptstyle\pm1.55}$	82.29±0.32	$81.04{\scriptstyle\pm1.09}$	79.20±1.19	$78.04 \scriptstyle{\pm 2.07}$	82.16±0.22
German	$\Delta_{SP}(\downarrow)$	41.94±5.52	23.79 ± 6.70	$14.85{\scriptstyle\pm4.64}$	7.65±8.07	8.42±7.35	7.74±7.80	8.60±3.47	6.60±1.66
	$\Delta_{EO}(\downarrow)$	31.11±4.40	15.13 ± 5.74	8.26±6.72	4.18±4.86	5.69 ± 2.16	5.17±2.38	6.34 ± 2.31	1.58±1.14
	AUC (†)	88.50±1.80	90.50±2.10	$77.30{\scriptstyle \pm 6.90}$	88.20±3.50	89.07±2.26	92.04±0.89	89.60±1.60	91.39±0.34
	F1 (↑)	78.20±2.30	80.40±3.20	$65.60{\scriptstyle \pm 8.40}$	78.40±2.10	77.83 ± 3.79	77.81±6.03	80.00±3.10	83.09±0.98
Bail	$\Delta_{SP}(\downarrow)$	7.50±1.40	8.60±3.90	$6.50{\scriptstyle \pm 3.40}$	7.40±2.60	3.74 ± 3.54	5.74±0.38	$5.80{\scriptstyle\pm1.70}$	2.29±1.06
	$\Delta_{EO}(\downarrow)$	2.30±1.90	3.90 ± 2.20	4.10±2.30	4.60±1.30	4.46±3.50	4.07±1.28	1.90±2.30	1.17±0.52
	AUC (†)	68.40±1.90	75.60±1.10	$70.60{\scriptstyle \pm 1.00}$	68.00±2.10	75.04 ± 0.12	72.89±0.44	74.00±0.80	73.42±1.89
	F1 (†)	79.40±2.70	$82.10{\scriptstyle \pm 0.80}$	$80.50{\scriptstyle\pm1.60}$	78.00±4.20	$82.41{\scriptstyle \pm 0.52}$	82.60±1.25	$\underline{83.50 \pm 0.80}$	83.63±0.89
Credit	$\Delta_{SP}(\downarrow)$	10.80±3.10	10.90 ± 3.00	13.00±3.70	18.70±3.60	$11.34{\scriptstyle\pm6.36}$	10.65±1.65	$\underline{10.40{\scriptstyle\pm1.30}}$	8.63±2.13
	$\Delta_{EO}(\downarrow)$	8.70±3.50	9.40±3.30	$12.10{\scriptstyle\pm4.20}$	17.50±3.50	9.38±5.39	8.10±1.91	$8.60{\scriptstyle \pm 1.80}$	6.85±1.55
Avg. (Rank)		5.58	4.42	5.75	5.92	4.67	3.83	4.08	1.75

Table 1: Node classification performance and group fairness performance on real-world datasets.

Table 2: Real-world dataset statistics.

Dataset	German Credit	Bail	Credit Defaulter
# Nodes	1,000	18,876	30,000
# Edges	22,242	321,308	1,436,858
# Attributes	27	18	13
Sens.	Gender	Race	Age
Label	Credit status	Bail decision	Future default

enables us to assess the quality of the obtained counterfactuals. In our approach, we consider settings with binary sensitive attributes and binary labels. A graph with 2000 nodes is sampled in our implementation. To generate the desired counterfactuals, we maintain the same sampled value of noise variables and use consistent causal relationships for each node. The sensitive attributes and labels are sampled from two different Bernoulli distributions, with $s_i \sim \mathcal{B}(p)$ and $y_i \sim \mathcal{B}(q)$, respectively. This results in generating vectors $\mathbf{s}_i = [(s_i)_{\times N}]$ and $\mathbf{y}_i = [(y_i)_{\times N}]$. Next, environment and content features, e_i and c_i , are sampled from normal distributions $\mathbf{e}_i \sim \mathcal{N}(\mathbf{s}_i, \mathbf{I})$ and $\mathbf{c}_i \sim \mathcal{N}(\mathbf{y}_i, \mathbf{I})$, respectively. These features are combined to form the overall latent feature $z_i = [c_i, e_i]$. The observed feature for each node v_i , denoted as \mathbf{x}_i , is computed as $\mathbf{x}_i = \mathbf{W}\mathbf{z}_i + \mathbf{b}_i$, where $\mathbf{W}_{ij} \sim \mathcal{N}(1, 1)$, and $\mathbf{W} \in \mathbb{R}^{d_2 \times 2d_1}$, with $\mathbf{b}_i \sim \mathcal{N}(0, \mathbf{I}) \in \mathbb{R}^{d_2}$. The adjacency matrix **A** is defined such that $A_{ij} = 1 \text{ if } \sigma(\cos(\mathbf{z}_i, \mathbf{z}_j) + \epsilon_{ij}) \geqslant \alpha \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ and } i \neq j, \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1), \text{ with } \epsilon_{ij} \sim \mathcal{N}(0, 1)$ $A_{ij} = 0$ otherwise. Here, $\sigma(\cdot)$ denotes the Sigmoid function, and the threshold α controls the edge number. We have the freedom to set sensitive attribute probability p, label probability q, latent feature dimension $2d_1$, observed feature dimension d_2 , node number N, and threshold α to control the biased graph generation process. Note that in the SCM we have $C \to Y$ instead of $Y \to C$, thus a better way is to first generate content features and then assign labels to the features. Intuitively, we argue that when using an optimal classifier to deal with content features with different means will assign the same label in our generation process. Therefore, to simplify the generation process, we use $C \rightarrow Y$ in our dataset design.

The synthetic dataset comes with notable advantages. Firstly, it gives us access to exact counterfactuals. After generating the initial graph, we keep all noise variables and unrelated variables unchanged, then adjust the sensitive attribute s_i or label y_i to calculate the precise counterfactual through the same graph generation

procedure. Secondly, the synthetic dataset enables adjustable bias levels, providing us control over the extent of bias in our models. As a result, we can undertake a comprehensive and detailed evaluation of our model's fairness and prediction quality.

5.1.3 Baselines. To evaluate the effectiveness of CAF, we include representative and state-of-the-art methods, which can be categorized into three categories: (1) plain node classification methods: GCN [18], GraphSAGE [11] and GIN [39]; (2) fair node classification methods: FairGNN [5], EDITS [7]; (3) graph counterfactual fairness methods: NIFTY [1] and GEAR [25]. Unless otherwise specified, we use GraphSAGE as the model backbone except for baseline GCN and GIN. We use SAGE to denote GraphSAGE. The detailed descriptions about the datasets are as follows:

- GCN [18]: GCN is a popular spectral GNN, which adopts a localized first-order approximation of spectral graph convolutions.
- GraphSAGE [11]: GraphSAGE is a method for inductive learning that leverages node feature information to generate unsupervised embeddings for nodes in large graphs, even if they were not included in the initial training.
- GIN [39]: Graph Isomorphism Network (GIN) is a graph-based neural network that can capture different topological structures by injecting the node's identity into its aggregation function.
- FairGNN [5]: FairGNN uses adversarial training to achieve fairness on graphs. It trains the learned representation via an adversary which is optimized to predict the sensitive attribute.
- EDITS [7]: EDITS is a pre-processing method for fair graph learning. It aims to debias the input network to remove the sensitive information in the graph data.
- NIFTY [1]: It simply performs a flipping on the sensitive attributes to get counterfactual data. It regularizes the model to be invariant to both factual and counterfactual data samples.
- GEAR [25]: GEAR is a method for counterfactual fairness on graphs. It utilizes a variational auto-encoder to synthesize counterfactual samples to achieve counterfactual fairness for graphs.

5.1.4 Evaluation Metrics. We evaluate the model performance from three perspectives: classification performance, group fairness and counterfactual fairness. (i) For classification performance, we use AUC and the F1 score to measure node classification performance. (ii) For fairness, following [5], we adopt two commonly used group

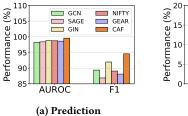
fairness metrics, i.e., statistical parity (SP) Δ_{SP} and equal opportunity (EO) Δ_{EO} , which are computed as $\Delta_{SP} = |P(\hat{y}_u = 1 \mid s = 0) - P(\hat{y}_u = 1 \mid s = 1)|$ and $\Delta_{EO} = |P(\hat{y}_u = 1 \mid y_u = 1, s = 0) - P(\hat{y}_u = 1 \mid y_u = 1, s = 1)|$. The smaller Δ_{SP} and Δ_{EO} are, the fairer the model is. (iii) For counterfactual fairness, as we have the ground-truth counterfactuals on the synthetic dataset, Following [25], we use the counterfactual fairness metric δ_{CF} , i.e., $\delta_{CF} = |P\left((\hat{y}_i)_{S \leftarrow s'} \mid \mathbf{X}, \mathbf{A}\right) - P\left((\hat{y}_i)_{S \leftarrow s'} \mid \mathbf{X}, \mathbf{A}\right)|$, where $s, s' \in \{0, 1\}^N$ are the sensitive attributes and s' = 1 - s. $(\hat{y}_i)_{S \leftarrow s'}$ is the computed ground-truth counterfactual label with the same data generation process as shown in Figure 1. We use subscript $s \leftarrow s'$ to denote counterfactual computation [31], i.e., keeping the same data generation process and values of random noise variable. Counterfactual fairness of the graph is only measured on synthetic dataset.

5.1.5 Setup. For German Credit, Credit Defaulter and Bail , we follow train/valid/test split in [1]. For the constructed synthetic dataset, we use a 50/25/25 split for training/validation/testing data. We randomly initialize the parameters. For each combination of the hyper-parameters configuration, we run the experiments with 10 random seeds and grid search for the best configuration.

5.2 Performance Comparison

To answer RQ1, we conduct experiments on real-world datasets and synthetic dataset with comparison to baselines.

- *5.2.1 Performance on Real-World Datasets.* Table 1 shows the average performance with standard deviation of ten runs on real-world datasets. The best results are highlighted in **bold** and the runner-up results are underlined. From Table 1, we observe:
- CAF can improve the group fairness performance. Across three datasets, Table 1 shows CAF can make fairer predictions than other baseline methods. CAF beats all the baselines with respect to the group fairness metrics.
- There exists a trade-off between group fairness and prediction performance. Plain node classification methods, such as GCN, GraphSAGE and GIN, tend to have better prediction performance and worse group fairness performance. Fair node classification methods, including FairGNN, EDITS, NIFTY, GEAR and CAF, tend to suffer from a prediction performance drop and the group fairness performance is better.
- CAF achieves best performance on the prediction-fairness tradeoff. We use the average rank of two prediction metrics and two group fairness metrics to know the performance of the tradeoff. Our model ranks 1.75 and the runner-up model ranks 3.83. Our model outperforms the state-of-the-art node representation learning methods, which shows the effectiveness of our model.
- Graph counterfactual fairness methods, such as NIFTY, GEAR and CAF, achieved better performance than other baselines. Correlationbased counterfactual notions can capture the causal relationships and help to boost the group fairness performance.
- 5.2.2 Performance on Synthetic Dataset. Figure 3 reports the performance on the synthetic dataset. On the synthetic dataset, we have the desired ground-truth counterfactuals, which can be used to measure the performance of graph counterfactual fairness. We compare our model with plain node classification models and counterfactual fairness models. The observations are as follows:



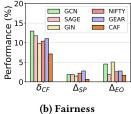
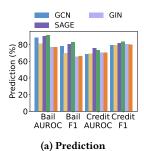


Figure 3: Node classification performance and group fairness performance on synthetic datasets.



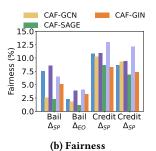


Figure 4: Comparison of the prediction performance and fairness performance of different backbones.

- CAF beats all the models with respect to the prediction, group fairness and counterfactual fairness metrics. We argue that in our assumed biased generation process, our model can effectively find invariant, sufficient and informative representations to make accurate and fair predictions.
- Other graph counterfactual fairness-based methods, including NIFTY and GEAR, cannot consistently outperform other methods. These methods design their model without considering meaningful causal relationships. NIFTY simply perturbs the sensitive attribute and omits the further influence on features and graph structure. GEAR adopts an GraphVAE to model the causal relationships, which may fail to generate meaningful counterfactuals.

5.3 Flexibility of CAF for Various Backbones

To show the flexibility of CAF in improving the fairness of various backbones while maintaining high classification accuracy, other than GraphSAGE, we also plug our model in GCN and GIN. Figure 4 shows the classification performance and fairness performance on Bail and Credit. From Figure 4, we observe that compared with the backbones, CAF can significantly improve the fairness with no or marginal decrease in classification performance. For example, on Bail dataset, the prediction performance with GIN backbone drops by 0.54% on AUROC but the Δ_{SP} drops by 1.37% and the Δ_{EO} drops by 0.86%, which is an improvement on fairness performance. This demonstrates the flexibility of CAF in benefiting various backbones.

5.4 Quality of Counterfactuals

To answer RQ2, we compare the counterfactuals obtained by CAF with ground-truth counterfactuals to investigate whether we can obtain the desired counterfactuals. We conduct experiments on the synthetic dataset which has ground-truth counterfactuals. We first use CAF to obtain counterfactuals. To measure the discrepancy of

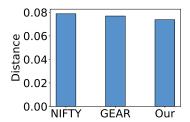


Figure 5: Discrepancy between learned counterfactual representation and ground-truth counterfactual representation.

Table 3: Ablation study.

Models					
Midueis	AUROC ↑	F1 ↑	$\delta_{\mathrm{CF}}\downarrow$	$\Delta_{\mathrm{DP}}\downarrow$	$\Delta_{\mathrm{EO}}\downarrow$
SAGE	98.40±0.31	86.89±3.59	11.76±2.92	$1.87 \scriptstyle{\pm 0.82}$	1.85±1.41
CAF	99.57 ± 0.06	$94.58 \scriptstyle{\pm 0.58}$	7.12 ± 0.41	$0.65{\scriptstyle \pm 0.41}$	1.66 ± 0.57
CAF-NA	99.33±0.27	94.72±2.15	$10.42{\scriptstyle\pm1.98}$	1.63 ± 0.41	1.83 ± 0.74
CAF-NB	98.47 ± 0.28	88.64±4.31	8.31±3.22	0.91 ± 0.28	1.63±0.51
CAF-NP	98.36±0.39	84.72±3.07	13.47 ± 2.62	1.37 ± 0.55	1.96±1.78
CAF-NS	98.81±0.73	91.98±3.25	9.72 ± 2.05	$1.35{\scriptstyle \pm 0.18}$	$1.73{\scriptstyle \pm 0.55}$

the obtained counterfactuals with respect to the feature and structure in the ego graph, we compare the learned counterfactual representations and the ground-truth counterfactual representations. We compare our model with two graph counterfactual fairness baselines, i.e., NIFTY [1] and GEAR [25]. NIFTY simply flips the sensitive attributes to get their counterfactuals. GEAR uses a Graph-VAE to generate the counterfactuals based on self-perturbation and neighbor-perturbation. Figure 5 shows the average result for all the nodes on the synthetic dataset. We show that CAF can find better counterfactuals than other graph counterfactual fairness models, i.e., smaller discrepancy to ground-truth counterfactuals. The result also shows there is still space for existing methods to improve the performance of getting appropriate counterfactuals.

5.5 Ablation Study

In our model, the pre-trained model can provide pseudo-labels for the nodes in the unlabeled set. Thus, we can select counterfactuals from the entire dataset. The model trained from scratch, without any pre-training, is denoted as CAF-NP. Without pseudo-labels, we can only select counterfactuals from the training set, which is denoted as the variant CAF-NS. We evaluate the performance on synthetic dataset. The results are reported in Table 3. We find the model CAF-NS performs worse than the CAF but better than CAF-NP. The result shows the pseudo-labels can also boost the performance of our model. Usually, the training set is small and the model may not obtain desired counterfactuals from the limited data points. Although pseudo-labels may contain some noisy information, they can also help improve the our model performance.

We further delve into how the constraints impact performance. When merely setting $\alpha=0$ or $\beta=0$, we denote the model as CAF-NA and CAF-NB, respectively. The models CAF-NA and CAF-NB outperform SAGE, yet fall short when compared to CAF. This indicates that both the sufficiency and invariance constraints collectively contribute to the superior performance of our model.

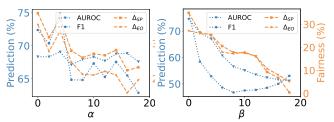


Figure 6: Hyper-parameter study on German dataset.

5.6 Hyper-Parameter Sensitivity Analysis

There are two important hyperparameters in CAF, i.e., α and β . α controls the contribution of the invariance regularization \mathcal{L}_{inv} and β controls the contribution of the sufficiency regularization. To understand the impact of α and β on CAF, we fix β as 5 and vary α as $\{0,1,\ldots,18\}$. Similarly, we fix α as 1 and vary β as $\{0,1,\ldots,18\}$. We report the result on German dataset in Figure 6. From Figure 6, we have the following observations: there exists a trade-off between prediction performance and fairness performance. The trend is that when we increase the α and β , we will get worse prediction performance and better fairness performance. We argue that without these regularizations, the model may rely on sensitive information. When we decrease the regularization budget, we can disentangle content representations. Thus, the prediction performance will get worse and the fairness performance will be better.

6 CONCLUSION AND FUTURE WORK

In this paper, we study the problem of learning fair node representations with GNNs. We first formalize the biased graph generation process with an SCM. Motivated by causal theory, we propose a novel framework CAF to learn fair node presentations which meet graph counterfactual fairness criteria and can achieve good prediction-fairness performance. Specifically, we align the model design with the data generation process and convert the problem to learn content representations. We derive several properties of the optimal content representation from the causal graph, i.e., invariance, sufficiency and informativeness. To get appropriate supervision for the invariance regularization, we design a counterfactual selection module. Extensive experiments demonstrate that CAF can achieve state-of-the-art performance on synthetic dataset and real-world datasets with respect to the prediction-fairness trade-off.

There are several interesting directions worth exploring. First, in this paper, we mainly focus on binary classification and binary sensitive attribute. We will extend the work to multi-class classification and multi-category sensitive attributes. Second, in this paper, we focus on static graphs while there are many different kinds of graphs in real-world. Thus, we aim to extend our model to more complex graph learning settings, such as dynamic graphs, multi-value sensitive attributes and labels.

ACKNOWLEDGMENTS

This material is based upon work supported by, or in part by the National Science Foundation (NSF) under grants number IIS-1909702, IIS-2153326, and IIS-2212145, Army Research Office (ARO) under grant number W911NF-21-1-0198, Department of Homeland Security (DHS) CINA under grant number E205949D, and Cisco Faculty Research Award.

REFERENCES

- Chirag Agarwal, Himabindu Lakkaraju, and Marinka Zitnik. 2021. Towards a unified framework for fair and stable graph representation learning. In *Uncertainty* in Artificial Intelligence. PMLR, 2114–2124.
- [2] Arthur Asuncion and David Newman. 2007. UCI machine learning repository. (2007).
- [3] Zhengyu Chen, Teng Xiao, and Kun Kuang. 2022. Ba-gnn: On learning bias-aware graph neural network. In 2022 IEEE 38th International Conference on Data Engineering (ICDE). IEEE, 3012–3024.
- [4] Enyan Dai, Limeng Cui, Zhengyang Wang, Xianfeng Tang, Yinghan Wang, Monica Cheng, Bing Yin, and Suhang Wang. 2023. A unified framework of graph information bottleneck for robustness and membership privacy. In KDD.
- [5] Enyan Dai and Suhang Wang. 2021. Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information. In WSDM. 680–688.
- [6] Enyan Dai, Tianxiang Zhao, Huaisheng Zhu, Jun Xu, Zhimeng Guo, Hui Liu, Jiliang Tang, and Suhang Wang. 2022. A Comprehensive Survey on Trustworthy Graph Neural Networks: Privacy, Robustness, Fairness, and Explainability. ArXiv abs/2204.08570 (2022).
- [7] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. 2022. EDITS: Modeling and Mitigating Data Bias for Graph Neural Networks.. In *The Web Conference* (WWW). ACM, 1259–1269.
- [8] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In Proceedings of the 3rd innovations in theoretical computer science conference. 214–226.
- [9] Zhimeng Guo, Teng Xiao, Charu Aggarwal, Hui Liu, and Suhang Wang. 2023.Counterfactual Learning on Graphs: A Survey. arXiv abs/2304.01391 (2023).
- [10] William L. Hamilton. 2020. Graph Representation Learning. Vol. 14. Morgan & Samp; Claypool Publishers. 1–159 pages.
- [11] William L. Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs.. In Conference on Neural Information Processing Systems (NeurIPS). 1024–1034.
- [12] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In NeurIPS. 3315–3323.
- [13] Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework.. In International Conference on Learning Representations (ICLR).
- [14] Zhimeng Jiang, Xiaotian Han, Chao Fan, Zirui Liu, Na Zou, Ali Mostafavi, and Xia Hu. 2022. Fmp: Toward fair graph message passing against topology bias. ArXiv abs/2202.04187 (2022).
- [15] Guangyin Jin, Qi Wang, Cunchao Zhu, Yanghe Feng, Jincai Huang, and Jiangping Zhou. 2020. Addressing crime situation forecasting task with temporal graph convolutional neural network approach. In 2020 12th International Conference on Measuring Technology and Mechatronics Automation (ICMTMA). IEEE, 474–478.
- [16] Kareem L Jordan and Tina L Freiburger. 2015. The effect of race/ethnicity on sentencing: Examining sentence type, jail length, and prison length. *Journal of Ethnicity in Criminal Justice* 13, 3 (2015), 179–196.
- [17] Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. ArXiv abs/1611.07308 (2016).
- [18] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks.. In *ICLR*.
- [19] Öykü Deniz Kose and Yanning Shen. 2021. Fairness-aware node representation learning. ArXiv abs/2106.05391 (2021).
- [20] Sanjay Kumar, Abhishek Mallik, Anavi Khetarpal, and BS Panda. 2022. Influence maximization in social networks using graph embedding and graph neural network. *Information Sciences* 607 (2022), 1617–1636.
- [21] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. Advances in neural information processing systems 30 (2017).
- [22] Rui Li, Jianan Zhao, Chaozhuo Li, Di He, Yiqi Wang, Yuming Liu, Hao Sun, Senzhang Wang, Weiwei Deng, Yanming Shen, Xing Xie, and Qi Zhang. 2022. HousE: Knowledge Graph Embedding with Householder Parameterization.. In International Conference on Machine Learning (ICML). 13209–13224.
- [23] Yi Liu, Limei Wang, Meng Liu, Yuchao Lin, Xuan Zhang, Bora Oztekin, and Shui-wang Ji. 2022. Spherical message passing for 3d molecular graphs. In *International Conference on Learning Representations (ICLR)*.
- [24] Donald Loveland, Jiayi Pan, Aaresh Farrokh Bhathena, and Yiyang Lu. 2022. FairEdit: Preserving Fairness in Graph Neural Networks through Greedy Graph

- Editing. ArXiv abs/2201.03681 (2022).
- [25] Jing Ma, Ruocheng Guo, Mengting Wan, Longqi Yang, Aidong Zhang, and Jundong Li. 2022. Learning fair node representations with graph counterfactual fairness. In Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining. 695–703.
- [26] David Madras, Elliot Creager, Toniann Pitassi, and Richard S. Zemel. 2018. Learning Adversarially Fair and Transferable Representations.. In ICML. 3381–3390.
- [27] Karima Makhlouf, Sami Zhioua, and Catuscia Palamidessi. 2020. Survey on
- causal-based machine learning fairness notions. ArXiv abs/2010.09553 (2020).
 [28] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. Comput. Surveys 54, 6 (2021), 1–35.
- [29] Tahleen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. 2019. Fairwalk: Towards Fair Graph Embedding. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence Organization, 3289–3295.
- [30] Harjit Singh Sekhon, Sanjit Kumar Roy, and James Devlin. 2016. Perceptions of fairness in financial services: an analysis of distribution channels. *International Journal of Bank Marketing* 34, 2 (2016), 171–190.
- [31] Ramalingam Shanmugam. 2018. Elements of causal inference: foundations and learning algorithms. Journal of Statistical Computation and Simulation 88, 16 (2018), 3248–3248.
- [32] Edward H Simpson. 1951. The interpretation of interaction in contingency tables. Journal of the Royal Statistical Society: Series B (Methodological) 13, 2 (1951), 238–241
- [33] Yongduo Sui, Xiang Wang, Jiancan Wu, Min Lin, Xiangnan He, and Tat-Seng Chua. 2022. Causal attention for interpretable and generalizable graph classification. In KDD. 1696–1705.
- [34] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks.. In International Conference on Learning Representations (ICLR).
- [35] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. IEEE transactions on neural networks and learning systems 32, 1 (2020), 4–24.
- [36] Teng Xiao, Zhengyu Chen, Zhimeng Guo, Zeyang Zhuang, and Suhang Wang. 2022. Decoupled self-supervised learning for graphs. Advances in Neural Information Processing Systems 35 (2022), 620–634.
- [37] Teng Xiao, Zhengyu Chen, Donglin Wang, and Suhang Wang. 2021. Learning How to Propagate Messages in Graph Neural Networks. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. ACM, 1894–1903.
- [38] Teng Xiao and Donglin Wang. 2021. A General Offline Reinforcement Learning Framework for Interactive Recommendation. In AAAI Conference on Artificial Intelligence (AAAI), Vol. 35. 4512–4520.
- [39] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In International Conference on Learning Representations (ICLR).
- [40] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. 2021. Causalvae: Disentangled representation learning via neural structural causal models. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 9593–9602.
- [41] I-Cheng Yeh and Che-hui Lien. 2009. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. Expert systems with applications 36, 2 (2009), 2473–2480.
- [42] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. 2022. Explainability in graph neural networks: A taxonomic survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022).
- [43] Richard S. Zemel, Yu Wu, Kevin Swersky, Toniann Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations.. In *International Conference on Machine Learning (ICML)*. 325–333.
- [44] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. NeurIPS 31 (2018).
- [45] Tong Zhao, Gang Liu, Daheng Wang, Wenhao Yu, and Meng Jiang. 2022. Learning from Counterfactual Links for Link Prediction.. In International Conference on Machine Learning (ICML), 26911–26926.
- [46] Huaisheng Zhu, Guoji Fu, Zhimeng Guo, Zhiwei Zhang, Teng Xiao, and Suhang Wang. 2023. Fairness-aware Message Passing for Graph Neural Networks. arXiv preprint arXiv:2306.11132 (2023).