

FairMILE: Towards an Efficient Framework for Fair Graph Representation Learning

Yuntian He
The Ohio State University
Columbus, Ohio, USA
he.1773@osu.edu

Saket Gurukar*
Samsung Research America
Mountain View, California, USA
saket.gurukar@gmail.com

Srinivasan Parthasarathy
The Ohio State University
Columbus, Ohio, USA
srini@cse.ohio-state.edu

ABSTRACT

Graph representation learning models have demonstrated great capability in many real-world applications. Nevertheless, prior research indicates that these models can learn biased representations leading to discriminatory outcomes. A few works have been proposed to mitigate the bias in graph representations. However, most existing works require exceptional time and computing resources for training and fine-tuning. To this end, we study the problem of efficient fair graph representation learning and propose a novel framework FairMILE. FairMILE is a multi-level paradigm that can efficiently learn graph representations while enforcing fairness and preserving utility. It can work in conjunction with any unsupervised embedding approach and accommodate various fairness constraints. Extensive experiments across different downstream tasks demonstrate that FairMILE significantly outperforms state-of-the-art baselines in terms of running time while achieving a superior trade-off between fairness and utility.

CCS CONCEPTS

• Computing methodologies → Machine learning

KEYWORDS

Fairness, Machine Learning, Graph Representation Learning

1 INTRODUCTION

A critical task in graph learning is to learn the hidden representations of the graph, also known as *graph embedding*. The goal of graph embedding is to preserve both structure properties and node features in the graph. Such embeddings can be used to characterize individual users (e.g. Amazon and Netflix) and to promote new connections (e.g. LinkedIn). Various methods have been developed for this purpose [17, 19, 40, 42], including those based on graph neural networks (GNNs) [20, 29, 52]. Such models have been effective in many real-world applications, such as crime forecasting [24], fraud detection [50], and recommendation [14, 18].

Given the high-stake decision-making scenarios that such models are typically deployed in, it is critical to ensure that the decisions made by these models are fair. Prior studies [1, 10, 43] reveal that graph representation learning models may inherit the bias from the underpinning graph data. A common source of bias is node features which may contain historical bias in sensitive attributes or other correlated attributes [13]. Another cause of bias is the *homophily effect* - promoting links that may lead to increased segregation. Such bias can lead to a biased distribution in the embedding space [12] and cause unfair treatment towards particular sensitive

attributes such as gender and ethnicity [13]. There is a clear need to alleviate such bias, ideally without impacting the bottom line of model performance.

Recent efforts to address this problem seek to enhance fairness by adapting existing GNN models [1, 9, 10, 25, 34]. However, such adaptations often add to the models' complexity – on large-scale graphs, these models either cannot finish execution in a reasonable amount of time or often result in an out-of-memory error. These concerns are amplified by recent articles that suggest that the training time for many AI models is simply becoming unsustainable [36, 47] with respect to both compute and emission costs. To tackle this issue, a naive solution is to apply scalability improvement techniques such as the multi-level framework [21, 26, 35]. However, these solutions lack fairness considerations. Figure 1 evaluates these approaches in terms of efficiency, fairness, and utility. The results demonstrate that: (1) Prior fairness-aware models are time-consuming for fair graph representation learning; (2) Scalable approaches like MILE [35] cannot enhance the fairness in embeddings. These observations highlight the challenges of balancing efficiency, fairness, and utility in the problem of fair graph representation learning.

In addition to inefficiency, there are some other challenges with existing work. First, some works adapt existing unsupervised graph embedding approaches for fairness [27, 43], but it is challenging to accommodate all such models. Second, many fair representation learning methods only consider a single, binary sensitive attribute, while real-world graphs usually have multiple multi-class sensitive attributes - limiting their applicability.

To address the above-mentioned issues, we present **Fair MultI Level Embedding** framework (FairMILE). FairMILE is a general framework for fair and efficient graph representation learning. It adopts a multi-level framework used by recent scalable embedding methods [2, 6, 21, 22, 35]. However, unlike other multi-level frameworks, our framework incorporates fairness as a first-class citizen. Our framework is method agnostic in that it can accommodate any unsupervised graph embedding method treating it as a black box. Moreover, unlike a majority of fair graph representation learning models, FairMILE can learn fair embeddings with respect to multiple multi-class sensitive attributes simultaneously. To summarize, our main contributions are:

- **Novelty:** To the best of our knowledge, this is the first work that seeks to improve the efficiency issue present in fair graph representation learning. To that end, we develop a general-purpose framework called FairMILE.
- **Model-agnostic:** FairMILE can easily accommodate any unsupervised graph embedding methods and improve their fairness while preserving the utility.

*Work done when the author was a graduate student at OSU.

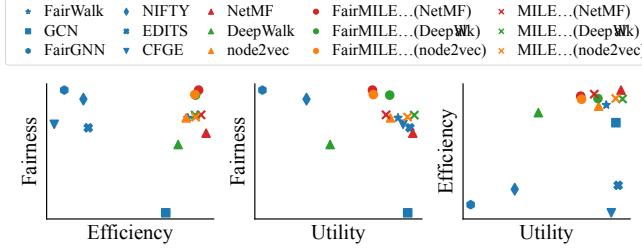


Figure 1: Evaluation of fairness-aware methods, graph embedding methods, MILE, and this work (FairMILE) in terms of efficiency, fairness, and utility. The top right corner of each plot corresponds to the best performance. This demonstrates that existing fair representation learning works are inefficient and scalable embedding methods are unable to enforce fairness, which highlights our novelty and contributions.

- **Efficiency and versatility:** Compared with existing approaches, FairMILE successfully improves the efficiency of fair graph representation learning. In addition, FairMILE can achieve fairness towards *multiple* and *non-binary* sensitive attributes, which most prior works fail to consider.
- **Evaluation:** We demonstrate both the efficacy and efficiency of FairMILE across both node classification and link prediction settings. Our results show that FairMILE can improve efficiency by up to two orders of magnitude and fairness by several factors while realizing comparable accuracy to competitive strawman.

2 PRELIMINARIES

2.1 Notations

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph, where \mathcal{V} is the set of nodes, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. Let \mathbf{A} be the graph adjacency matrix, where $A_{u,v}$ denotes the weight of edge (u, v) . $A_{u,v} = 0$ means u, v are not connected. $\delta(u)$ denotes the degree of node u . The graph also contains a set of sensitive attributes \mathcal{F} (e.g., gender and race). Each attribute may have a binary or multi-class value associated with a particular demographic group.

2.2 Fairness Metrics

In this paper, we focus on group fairness and use two metrics for evaluation. These metrics have been widely adopted in prior works [1, 10, 11, 13, 43]. Without loss of generality, we first introduce them in a binary prediction scenario with a binary sensitive attribute, and then we extend them to a general multi-class case.

Definition 1. Demographic Parity (also known as *statistical parity*) [13] requires that each demographic group should receive an *advantaged* outcome (i.e., $\hat{Y} > 0$) at the same rate, which is formulated as $P(\hat{Y} = 1|S = 0) = P(\hat{Y} = 1|S = 1)$, where \hat{Y} is the predicted label and $S \in \mathcal{F}$ is a binary sensitive attribute. To quantify how demographic parity is achieved, prior works [1, 10, 32] define $\Delta_{DP,\text{binary}}$ as

$$\Delta_{DP,\text{binary}} = |P(\hat{Y} = 1|S = 0) - P(\hat{Y} = 1|S = 1)|$$

To extend demographic parity for multi-class sensitive attributes, Rahman et al. [43] measure the variance of positive rates among all

groups. Here, we further extend it to a multi-class predicted label scenario by averaging the standard deviations (denoted as σ) across all advantaged classes. The formulation is given as

$$\Delta_{DP} = \frac{1}{|\mathcal{Y}^+|} \sum_{y \in \mathcal{Y}^+} \sigma \left(\{P(\hat{Y} = y|S = s) : \forall s\} \right)$$

where \mathcal{Y}^+ denotes a set of advantaged classes.

Definition 2. Equality of Opportunity requires that each demographic group has an identical probability to receive a specific advantaged outcome for its members with this advantaged ground-truth label. In binary classification tasks with respect to a binary sensitive attribute, existing works [1, 10, 11] compute the difference of true positive rates across two groups to measure the equality of opportunity, which is formulated as

$$\Delta_{EO,\text{binary}} = |P(\hat{Y} = 1|Y = 1, S = 0) - P(\hat{Y} = 1|Y = 1, S = 1)|$$

where Y is the ground truth label. Similarly, we define a new metric Δ_{EO} to extend equality of opportunity to a general scenario with multi-class labels and attributes:

$$\Delta_{EO} = \frac{1}{|\mathcal{Y}^+|} \sum_{y \in \mathcal{Y}^+} \sigma \left(\{P(\hat{Y} = y|Y = y, S = s) : \forall s\} \right)$$

Note that the previously adopted metrics for binary predicted labels and binary sensitive attributes are a special case of our proposed measures, i.e., $\Delta_{DP} = \frac{1}{2} \Delta_{DP,\text{binary}}$ and $\Delta_{EO} = \frac{1}{2} \Delta_{EO,\text{binary}}$.

2.3 Problem Statement

Problem 1. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the embedding dimensionality d , and a set of sensitive attributes \mathcal{F} , the problem of **Fair Graph Representation Learning** aims to learn a fair embedding model $f: \mathcal{V} \rightarrow \mathbb{R}^d$ with less inherent bias towards attributes in \mathcal{F} where the present bias is measured with Δ_{DP} and Δ_{EO} .

3 RELATED WORK

3.1 Fairness in Machine Learning

Since machine learning techniques are deployed to make decisions that have societal or ethical implications [7, 50, 51], serious concerns over their fairness are raised. There have been various definitions of fairness in machine learning. In this paper, we focus on the most popular definition *group fairness* [13], which requires that an algorithm should treat each demographic group equally. The groups are associated with a single or multiple *sensitive attributes*, such as gender and race. There are also other definitions of fairness including *individual fairness* [13] and *counterfactual fairness* [31].

Unfair outcomes are mostly caused by data bias and algorithmic bias [38]. There exists a wide range of biases in data. For example, features like home address can be associated with specific races and lead to unfair decisions indirectly [12]. The design of machine learning algorithms may also unintentionally amplify the bias in data. To address this concern, several fair machine learning algorithms have been proposed in recent years [3, 4, 16, 45, 53]. A comprehensive survey on fair machine learning is given in [38].

Table 1: Summary of fair graph representation learning methods.

Method	Method-agnostic	Multiple Sensitive Attributes	Non-binary Attributes
FairGNN [10]	×	×	×
NIFTY [1]	×	×	×
FairAdj [34]	×	×	✓
FairWalk [43]	×	×	✓
CFGE [5]	×	✓	✓
EDITS [11]	✓	×	✓
FairMILE (This work)	✓	✓	✓

3.2 Fair Graph Representation Learning

In the graph context, models trained to realize representations accounting for the connectivity and topology inherent to the network (e.g. homophily bias) can lead to biased representations. Downstream tasks that operate on such representations can lead to unfair recommendations [49], and even biased and unjust outcomes [37]. However, the incorporation of fairness with graph-based learning is challenging because of the non-i.i.d nature of the data and the homophily effect of graph data [10, 12, 23].

Recently, several methods have been proposed to learn fair graph representations. FairGNN [10] leverages adversarial learning to train fair GNNs for node classification. NIFTY [1] adds a fairness loss to the GNN objective as regularization. FairAdj [34] accommodates the VGAE [28] model for fair link prediction. For task-agnostic embedding, FairWalk [43] learns fair embeddings by adapting an embedding algorithm node2vec [17]. Specifically, it modifies the random walk process and adjusts the probability of selecting nodes in each sensitive group for fairness. To consider multiple sensitive attributes, CFGE [5] employs a set of adversaries with the encoder for compositional fairness constraints. Unlike these approaches, EDITS [11] is a pre-processing solution that reduces the bias in graph structure and node attributes, then trains vanilla GNNs on the debiased graph. However, since most methods are GNN-based, they require exceptional time for training and fine-tuning.

In addition to inefficiency, there are three other major drawbacks of existing works. First, some of them are not *method-agnostic* which means they require non-trivial modifications to the base model for adaption. Second, most existing works are unable to incorporate fairness constraints towards *multiple* sensitive attributes. Third, some methods cannot handle *non-binary* sensitive attributes. Table 1 summarizes these approaches together with our proposed work FairMILE. Compared with these approaches, our work can simultaneously (1) accommodate the base model easily while (2) achieve fairness towards multiple non-binary sensitive attributes. Most importantly, we will demonstrate that (3) FairMILE significantly outperforms these baselines in terms of efficiency.

3.3 Scalable Graph Embedding

Many methods for graph embedding have been proposed in recent years, including NetMF [42], DeepWalk [40], and node2vec [17]. Despite their excellent performance on various machine learning tasks, their lack of scalability prohibits them from processing large

Algorithm 1 Graph coarsening for fair embedding

Input: Graph $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i)$

Output: Coarsened graph \mathcal{G}_{i+1}

- 1: Sort \mathcal{V}_i in an increasing order of node degrees
- 2: **for** unmatched $u \in \mathcal{V}_i$ **do**
- 3: **if** All neighbors of u are matched **then**
- 4: $u' \leftarrow \{u\}$
- 5: **else**
- 6: Find unmatched v s.t. $(u, v) \in \mathcal{E}_i$ maximizing Equation (3),
then let $u' \leftarrow \{u, v\}$
- 7: Add supernode u' to \mathcal{V}_{i+1}
- 8: Connect supernodes in \mathcal{V}_{i+1} based on \mathcal{E}_i
- 9: Build \mathcal{G}_{i+1} from \mathcal{V}_{i+1} and \mathcal{E}_{i+1}

datasets. Recent research addressed the scalability issue of graph embedding using different methodologies. Some studies leverage high-performance computing techniques [33, 41, 54]. Another group of studies adopts the multi-level framework for better scalability. This framework is widely used for various graph problems [6, 26, 35] and the essential idea is to solve the problem from a smaller coarsened graph. However, none of these methods are fairness-aware. We are the first study that considers fairness based on this framework.

4 METHODOLOGY

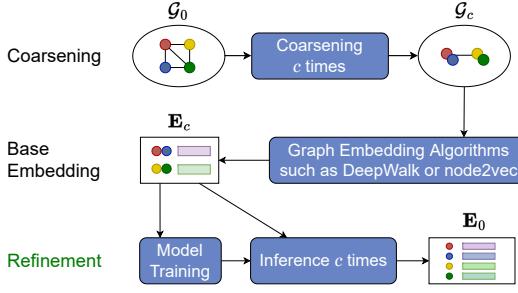
We propose a fairness-aware graph embedding framework FairMILE (shown in Figure 2a), which consists of three modules: graph coarsening, base embedding, and refinement. The idea is to first coarsen the original graph into a smaller one, then learn the embeddings of the coarsened graph using the base model, and eventually refine them into the embeddings of the original graph.

There exist several definitions of fairness [13, 31] - all with the shared principle that all subgroups should receive the positive outcome at the same level in a given measure (e.g., positive rate for Δ_{DP} , true positive rate for Δ_{EO}). Since this paper studies unsupervised representation learning, **our key idea** is to minimize the variance among the representations of different groups. Intuitively, the downstream models trained with such representations (or embeddings) will lead to a decrease in bias. Next, we introduce the functionality of each module and explain how FairMILE enforces fairness in the graph representations while improving efficiency and versatility.

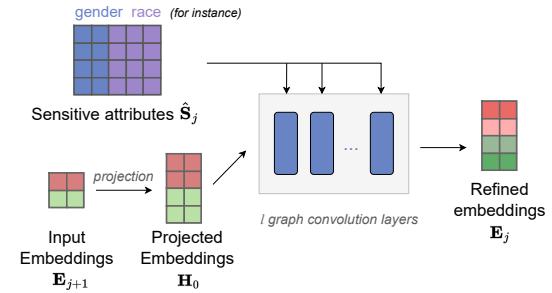
4.1 Graph Coarsening

We develop a new fairness-aware graph coarsening algorithm (shown in Algorithm 1). Given the initial graph $\mathcal{G}_0 = \mathcal{G}$, it shrinks the graph size by collapsing a group of nodes into a supernode in the output graph \mathcal{G}_1 . If two nodes are connected in \mathcal{G}_0 , there exists an edge between their supernodes in \mathcal{G}_1 . As a result, the numbers of nodes and edges are reduced in \mathcal{G}_1 . After repeating this process c times, we can finally get the coarsened graph \mathcal{G}_c .

The key challenges in graph coarsening are two-fold: retaining the structural information for better utility while incorporating fairness toward the sensitive attributes. For utility, we adopt a high-utility coarsening approach Normalized Heavy-Edge Matching (NHEM) [26], which merges two nodes if their normalized edge weight is maximum. Formally, given a node $u \in \mathcal{V}$, NHEM



(a) An illustrative pipeline of FairMILE



(b) Architecture of the refinement model.

Figure 2: Overview of FairMILE

computes the normalized weight of the edge $(u, v) \in \mathcal{E}$ defined as:

$$w(u, v) = \frac{A_{u,v}}{\sqrt{\delta(u)\delta(v)}} \quad (1)$$

Intuitively, it encourages nodes with fewer connections to match other nodes and penalizes the hub nodes. Node matching allows a supernode to have structural properties from different nodes. However, due to the homophily in many real-world networks, the bias in graph structure can be reinforced if two nodes from the same group are merged.

In light of this, we add a new term in the node matching function to reflect the divergence of the sensitive attribute distributions of two nodes for fairness consideration. In FairMILE, we use a matrix $S \in \mathbb{R}^{N \times M}$ to denote the sensitive attribute values, where each row s_u is the attribute value distribution vector for node u , and M is the dimensionality of sensitive attribute values. We use these attribute distribution vectors to quantify the divergence between nodes.

Initially, each node in the original graph \mathcal{G}_0 has a concatenation of one-hot vectors for the attributes. For example, \mathcal{F} has two attributes: ‘gender’ (female, male) and ‘race’ (African, Asian, white). The attributes of a female African u can be modeled as $s_u = [1, 0, 1, 0, 0]$. For a node v in a coarser graph $\mathcal{G}_i (i > 0)$, we use s_v to denote the distribution of all nodes in \mathcal{G}_0 merged into it. For example, $s_v = [1, 1, 1, 1, 0]$ in \mathcal{G}_1 indicates that the supernode v contains one female and one male on attribute ‘gender’, and one African and one Asian on ‘race’. To measure the difference between the sensitive attribute distributions of two nodes, we define the following function based on Kullback–Leibler divergence [30]:

$$\phi(u, v) = 1 - \left(1 + \sum_{j=1}^M \frac{s_{u,j}}{\|s_u\|} \log \left(\frac{s_{u,j}/\|s_u\|}{s_{v,j}/\|s_v\|} \right) \right)^{-1} \quad (2)$$

It essentially maps the divergence of two normalized attribute distributions to $[0, 1]$. The higher the score is, the more different their sensitive attributes are. Finally, given a node u of \mathcal{G}_i , we formulate the node matching policy as follows:

$$\max_{v:(u,v) \in \mathcal{E}_i} (1 - \lambda_c) w(u, v) + \lambda_c \phi(u, v) \quad (3)$$

where λ_c is the weight of fairness in node matching. The objective here is to find the neighbor of node u that maximizes the edge weight and attribute divergence together. Intuitively, a large value

of λ_c generates more inter-group matching in graph coarsening. Section 5 will empirically show that Algorithm 1 can improve the fairness in graph representations.

4.2 Base Embedding

Like other multi-level frameworks [26, 35], FairMILE applies the base model on the coarsest graph in an agnostic manner. Since the input is the coarsened graph and the output is its node embeddings, it is straightforward that FairMILE accommodates any unsupervised graph embedding method such as DeepWalk [40] or node2vec [17] with no modification required. This step generates the embeddings E_c on the coarsened graph.

4.3 Refinement

In the last phase of FairMILE, we seek to learn the representations of graph \mathcal{G}_0 from the embeddings of the coarsest graph \mathcal{G}_c . Generally, we train a fairness-aware refinement model based on graph convolution networks (GCN) [29] to infer the embeddings E_{c-1} of \mathcal{G}_{c-1} from E_c . Then we iteratively apply it until we get E_0 .

4.3.1 Model Architecture. Figure 2b shows the architecture of our refinement model. Without loss of generality, the refinement model has a projection layer followed by l GNN layers where the input and output of layer $i \in [1, l]$ are denoted as H_{i-1} and H_i . Given two graphs \mathcal{G}_j and \mathcal{G}_{j+1} , we initialize H_0 by projecting the embeddings of supernodes in \mathcal{G}_{j+1} to its associated nodes in \mathcal{G}_j . Note that if two nodes that have different sensitive attribute values are merged into \mathcal{G}_{j+1} , they share the same initial embeddings in \mathcal{G}_j , which mitigates the potential bias in learned representations.

In each layer, we use a normalized adjacency matrix $\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$ for message passing. Here we drop the notation referring to a specific graph for clarity. To take the sensitive attributes into consideration, we concatenate the input of each layer with the row-normalized sensitive attribute matrix \tilde{S} . Formally, the i -th convolution layer in our refinement model can be formulated as

$$H_i = \tanh \left(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} (H_{i-1} \parallel \tilde{S}) \Theta_i \right) \quad (4)$$

where Θ_i is the trainable linear transformation matrix of layer i . Finally, we can infer \mathcal{G}_j 's node representations $E_j = H_l / \|H_l\|_2$.

4.3.2 Training objectives. To achieve a trade-off between utility and fairness, we have two distinct objectives for each of them. For utility, we expect the refined representations to be close to the input so that matched nodes are still close to each other in the embedding space after refinement. Therefore, we minimize the difference between the projected embeddings and the predicted ones generated by the refinement model, which is defined as

$$L_u = \frac{1}{|\mathcal{V}_c|} \|H_0 - H_l\|^2 \quad (5)$$

L_u is the utility loss in our framework. To improve fairness, we encourage nodes with different sensitive attribute values to be closer in the embedding space. Specifically, we create a subset of edges \mathcal{E}'_c that consists of links between two nodes with significantly different sensitive attributes, which can be formulated as

$$\mathcal{E}'_c = \{(u, v) : (u, v) \in \mathcal{E}_c \wedge \phi(u, v) \geq \gamma\}$$

where γ is a threshold parameter for attribute divergence. Then we use the learned representations to reconstruct an adjacency matrix that represents the node distances in the embedding space. To reward the links between diverse attribute groups, we use the Hadamard product of the reconstruction matrix and the adjacency matrix of \mathcal{E}'_c (denoted as \mathcal{A}'_c) as our fairness objective. As a result, the trained model generates similar embeddings for inter-group nodes. The fairness loss (L_f) is formally defined as

$$L_f = -\frac{1}{|\mathcal{E}'_c|} [\text{sigmoid}(\mathbf{H}_l \mathbf{H}_l^\top) \odot \mathcal{A}'_c] \quad (6)$$

The multiplication of the embedding matrix with a sigmoid activation (denoted by $\text{sigmoid}(\mathbf{H}_l \mathbf{H}_l^\top)$) reconstructs the entire $N \times N$ matrix. Its element-wise product with \mathcal{A}'_c performs the masking operation to select the node pairs with diverse sensitive attributes. The negative minimization of L_f ensures that the similarity of the selected node pairs – measured by the dot products of their embeddings – is increased. This loss ensures that the embeddings learned by the refinement model of the node pairs with diverse sensitive attributes are *proximal* to each other.

Combining the utility and fairness loss functions, the overall training objective of our refinement model is

$$\min_{\{\Theta_i, \forall i \in [1, I]\}} (1 - \lambda_r) L_u + \lambda_r L_f \quad (7)$$

where $\lambda_r \in [0, 1]$ controls the weight of fairness objective.

4.4 Theoretical analysis

Time complexity: The time complexity of FairMILE depends on the selected embedding approach. Note that such approaches typically have a time complexity of at least $O(d|\mathcal{V}_0|)$ [8], for example, the time complexity of DeepWalk is $O(d|\mathcal{V}_0| \log |\mathcal{V}_0|)$. Considering that the number of nodes can be reduced by up to half after each time of coarsening (observed in Section 5.4), the efficiency is significantly improved by embedding the coarsened graph. Apart from embedding the coarsened graph, FairMILE spends additional $O(cl(d + M)(|\mathcal{E}_0| + d|\mathcal{V}_0|))$ time on coarsening and refinement. Given that $d, M \ll |\mathcal{V}_0|$, the additional time of these two phases is typically much less than the reduced time of base embedding – empirically verified in Section 5.4. Additional details on complexity analysis are included in our supplement.

Fairness: We prove that the difference between the mean representations of different demographic groups is bounded depending on the network topology.

Theorem 1. *When L_f is minimized, the 2-norm of the difference between the mean embeddings of any two demographic groups regarding a given sensitive attribute is bounded by*

$$\|\mu_p - \mu_q\|_2 \leq 2(1 - \min(\beta_p, \beta_q))$$

where p, q are any two different values of the given sensitive attribute. For $i \in \{p, q\}$, μ_i denotes the mean embedding values of nodes from group i , and β_i denotes the ratio of nodes from group i that have at least one inter-group edge.

Theorem 1 shows that the difference between the mean embeddings of two groups depends on the ratio of inter-group connected nodes in each group, which is typically large. For example, among the datasets used in our experiments, the minimum β is 0.676 in Credit and 0.958 in German, respectively. When the mean embeddings of different demographic groups are close to each other, they have similar representations and therefore with a high likelihood, they will receive similar outcomes in the downstream task. The full proof of Theorem 1 is provided in our supplement.

5 EXPERIMENTS

5.1 Experiment Setup

Datasets: We examine the performance of FairMILE on both node classification and link prediction tasks. Our experiments are conducted on seven real-world datasets from different application scenarios widely used in the fairness literature [1, 10, 11, 34, 46]. Statistics of the datasets are shown in Table 2, where $|\mathcal{Y}|$ denotes the number of predicted labels, and $|\mathcal{S}|$ denotes the number of sensitive attribute values. For details, please refer to our supplement.

Metrics: To quantify the prediction performance in node classification, we use AUROC, F1-score (for binary class problems), and Micro-F1 (for multi-class problems) as our utility metric. To measure the group fairness, we use Δ_{DP} and Δ_{EO} described in Section 3 as our fairness metrics. We also report the end-to-end running time in seconds to show the efficiency of all methods.

For link prediction, following prior works [34, 46], we use AUROC, Average Precision (AP), and accuracy as the utility metrics, and compute the disparity in expected prediction scores between intra-group and inter-group links. Specifically, the fairness metrics are formulated as:

$$\begin{aligned} \Delta_{DP, LP} &= |\mathbb{E}_{(u, v) \sim \mathcal{V} \times \mathcal{V}} [\hat{Y}|S(u) = S(v)] \\ &\quad - \mathbb{E}_{(u, v) \sim \mathcal{V} \times \mathcal{V}} [\hat{Y}|S(u) \neq S(v)]| \\ \Delta_{EO, LP} &= |\mathbb{E}_{(u, v) \sim \mathcal{V} \times \mathcal{V}} [\hat{Y}|(u, v) \in \mathcal{E}, S(u) = S(v)] \\ &\quad - \mathbb{E}_{(u, v) \sim \mathcal{V} \times \mathcal{V}} [\hat{Y}|(u, v) \in \mathcal{E}, S(u) \neq S(v)]| \end{aligned}$$

where $\hat{Y} \in [0, 1]$ is the link prediction score.

Baselines: Our baselines include 1) *Specialized approaches*: For node classification, we use the vanilla GCN [29] and three state-of-the-art fair node classification methods (NIFTY [1], FairGNN [10], and EDITS [11]) with GCN as their base model. For link prediction, we use VGAE [28] and FairAdj [34] with VGAE as the base predictor in our comparative experiments. In addition, we adapt CFGE [5] for both

Table 2: Statistics of datasets (NC denotes Node Classification, and LP denotes Link Prediction).

Dataset	Task	# Nodes	# Edges	# Features	Label ($ \mathcal{Y} $)	Sensitive Attributes ($ \mathcal{S} $)
German	NC	1,000	22,242	27	credit risk (2)	gender (2)
Recidivism	NC	18,876	321,308	18	crime (2)	race (2)
Credit	NC	30,000	1,436,858	13	payment default (2)	age (2)
Pokec-n	NC	66,569	729,129	59	field of work (4)	region (2), gender (2)
Cora	LP	2,708	5,278	1433	citation (2)	paper category (7)
Citeseer	LP	3,312	4,660	3703	citation (2)	paper category (6)
Pubmed	LP	19,717	44,338	500	citation (2)	paper category (3)

tasks which is the only baseline that accommodates multiple sensitive attributes. 2) *Graph embedding approaches*: We choose three popular unsupervised graph embedding techniques: NetMF [42], DeepWalk [40], and Node2vec [17]. We also include FairWalk [43], which is essentially a fairness-aware adaption of Node2vec. Note that these approaches can be used in both downstream tasks. 3) *Our framework*: We let FairMILE run with the three embedding approaches above for evaluation.

Parameters and environments: All hyperparameters of specialized approaches are set following the authors' instructions. In particular, unless otherwise specified, we set $\lambda = 50$ in CFGE for a better tradeoff between fairness and utility. The dimensionality d of representations for graph embedding approaches and our work is set to 128. In FairMILE, we set $\lambda_c = 0.5$ for graph coarsening. For refinement, We train a two-layer model for 200 epochs with $\lambda_r = 0.5$, $\gamma = 0.5$, and the learning rate of 1×10^{-3} on all datasets. To evaluate unsupervised approaches for node classification and all approaches for link prediction, we train a linear classifier on the learned embeddings (for node classification) or the Hadamard products of embeddings of sampled node pairs (for link predictions). For the task of link prediction, we randomly sample 10% of edges to build test sets and remove them from the training data, then we add the same number of negative samples as positive edges in the training and test sets, respectively. All methods are evaluated on the same test sets and trained on CPUs for fair comparisons. Experiments are conducted on a Linux machine with a 28-core Intel Xeon E5-2680 CPU and 128GB RAM. We report the average results of 5 runs with different data splits. For reproducibility, our codes and data are available¹.

5.2 Results for Node Classification

We first compare our work with specialized GNN-based approaches and unsupervised graph embedding baselines on three datasets with a single binary sensitive attribute. Specifically, we address the following questions: **Q1:** Does FairMILE improve the fairness of a base embedding method? How is the fairness of FairMILE compared with specialized approaches? **Q2:** Does FairMILE outperform other baselines in terms of efficiency? **Q3:** Does FairMILE retain the embedding's utility while improving its fairness and efficiency? Table 3 only presents the results on two datasets (German and Credit). The results on Recidivism are included in our supplement. Methods are categorized into different groups by their base models, and the optimal performance in each group is highlighted in bold.

¹<https://github.com/heyuntian/FairMILE>

A1: FairMILE achieves better fairness scores. Compared with graph embedding approaches without fairness consideration (i.e., NetMF, DeepWalk, and Node2vec), FairMILE on top of them always has lower Δ_{DP} and Δ_{EO} . In German, FairMILE decreases the Δ_{DP} scores of Node2vec and DeepWalk by 83.7% and 90.7%, respectively. In terms of Δ_{EO} , FairMILE improves the fairness of Node2vec and DeepWalk by 84.0% and 96.6%. When choosing NetMF as the base model, FairMILE is optimal (zero) on both fairness metrics, which indicates that the learned representations lead to a perfectly fair classification. Compared with FairWalk (the fair adoption of Node2vec), FairMILE-Node2vec improves Δ_{DP} by 49.7% on Credit while FairWalk only improves by 4.8%. Similar results are also observed on other datasets, which reveals that FairMILE successfully mitigates the bias in general-purpose graph embedding baselines.

We also evaluate specialized methods (i.e., FairGNN, NIFTY, EDITS, and CFGE) and observe that FairMILE achieves comparable or better fairness with respect to these approaches. FairGNN outperforms the other specialized methods on all datasets by gaining the largest improvements on both Δ_{DP} and Δ_{EO} with respect to Vanilla GCN. It means that their methodology of adversarial training is more effective than the regularization adopted by NIFTY, and EDITS suffers from its task agnosticity. On German and Credit, the best fairness metrics of FairMILE are comparable to or better than FairGNN and other specialized techniques. These results demonstrate that FairMILE is effective in reducing the bias in graph representations compared to the state-of-the-art models.

A2: FairMILE is more efficient than other baselines. First, FairMILE outperforms the GNN-based specialized approaches on all datasets in terms of efficiency. In German, they take up to 2 minutes for training, while FairMILE can finish in only 6 seconds. The difference becomes more significant on larger datasets. When CFGE needs more than 3 hours on Credit, FairMILE finishes within 2 minutes which is 110.6-137.7x faster. On the other hand, FairMILE also improves the efficiency of all base embedding methods. In Credit, FairMILE on top of NetMF saves up to 80% of the running time of vanilla NetMF. The improvement in German is sometimes invisible because German is a small graph. But FairMILE can still finish within seconds. Compared with FairWalk, FairMILE-Node2vec is always faster.

A3: FairMILE learns quality graph representations. We observe the quality of learned representations through the AUROC and F1 scores. With respect to the base embedding methods, FairMILE has a similar performance on both utility metrics which is fairly remarkable given that FairMILE significantly improves fairness. An interesting observation is that while Vanilla GCN, graph embedding approaches, and FairMILE have a similar performance

Table 3: Comparison on utility, fairness, and efficiency metrics in node classification between FairMILE and other baselines.

Dataset	Method	AUROC (\uparrow)	F1 (\uparrow)	Δ_{DP} (\downarrow)	Δ_{EO} (\downarrow)	Time (\downarrow)
German	NetMF	65.16 \pm 2.45	80.63 \pm 1.10	5.71 \pm 2.89	3.66 \pm 2.11	2.48
	FairMILE-NetMF	61.93 \pm 3.38	82.35 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	6.31
	DeepWalk	58.54 \pm 4.43	75.78 \pm 1.49	7.22 \pm 3.86	7.69 \pm 3.26	16.99
	FairMILE-DeepWalk	63.31 \pm 3.63	82.40 \pm 0.33	0.67 \pm 0.88	0.26 \pm 0.39	7.84
	Node2vec	63.37 \pm 3.77	78.69 \pm 1.25	3.69 \pm 2.60	2.75 \pm 1.34	12.76
	Fairwalk	63.98 \pm 2.07	77.64 \pm 1.62	3.67 \pm 2.74	3.28 \pm 2.50	11.93
	FairMILE-Node2vec	62.00 \pm 2.59	82.32 \pm 0.20	0.60 \pm 0.96	0.44 \pm 0.41	8.29
	Vanilla GCN	64.75 \pm 7.20	77.93 \pm 3.53	16.27 \pm 5.86	13.28 \pm 5.06	23.75
	FairGNN	53.12 \pm 5.73	82.35 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	136.29
	NIFTY	56.65 \pm 6.84	81.35 \pm 1.54	1.20 \pm 1.45	0.83 \pm 1.20	91.05
Credit	EDITS	64.93 \pm 2.90	79.64 \pm 2.27	5.00 \pm 3.38	2.76 \pm 2.26	84.24
	CFGE	64.38 \pm 0.77	81.59 \pm 0.33	4.54 \pm 2.91	4.30 \pm 1.84	3900.11
	NetMF	74.93 \pm 0.43	88.36 \pm 0.08	2.66 \pm 0.55	1.34 \pm 0.93	240.38
	FairMILE-NetMF	74.69 \pm 0.43	88.31 \pm 0.08	0.68 \pm 0.50	0.68 \pm 0.66	90.28
	DeepWalk	75.09 \pm 0.39	88.36 \pm 0.15	2.50 \pm 0.54	1.81 \pm 0.71	570.35
	FairMILE-DeepWalk	74.60 \pm 0.53	88.31 \pm 0.11	1.34 \pm 0.57	1.05 \pm 0.70	112.39
	Node2vec	74.95 \pm 0.36	88.26 \pm 0.18	2.92 \pm 0.41	1.93 \pm 0.87	249.92
	Fairwalk	74.88 \pm 0.49	88.25 \pm 0.15	2.78 \pm 0.47	1.58 \pm 0.81	261.23
	FairMILE-Node2vec	74.01 \pm 0.47	87.99 \pm 0.16	1.47 \pm 0.74	1.33 \pm 0.64	105.43
	Vanilla GCN	72.83 \pm 2.83	82.79 \pm 1.76	5.77 \pm 0.32	4.36 \pm 0.59	2265.58
Pokec-n	FairGNN	71.14 \pm 4.58	83.29 \pm 3.27	1.40 \pm 0.97	1.21 \pm 0.50	3201.72
	NIFTY	72.43 \pm 0.68	81.80 \pm 0.37	5.52 \pm 0.33	4.35 \pm 0.68	7735.00
	EDITS	73.22 \pm 0.77	81.55 \pm 0.30	5.36 \pm 0.42	4.17 \pm 0.65	9078.12
	CFGE	70.51 \pm 1.31	87.85 \pm 0.13	2.99 \pm 0.47	1.47 \pm 0.18	12430.37

in terms of utility, the supervised GCN-based fair approaches have lower AUROC or F1 scores on one or both datasets, which reveals that these approaches require fine-tuning to perform well in terms of utility while enforcing fairness.

In summary, FairMILE in conjunction with popular base embedding approaches can compete or improve on the fairness criteria with various specialized methods while outperforming them significantly in terms of efficiency and retaining comparable utility.

5.3 Fairness towards Multiple Sensitive Attributes

To explore how FairMILE learns fair representations towards multiple sensitive attributes, we conduct an experiment on the Pokec-n dataset. Pokec-n has a multi-class predicted label and two sensitive attributes. Most baselines cannot process such datasets since they restrictively only cater to the case of a single binary label or sensitive attribute. CFGE [5] is the only baseline that accommodates multiple sensitive attributes. Therefore we compare FairMILE with CFGE [5] and graph embedding methods, and we set $c = 4$ for our work. Note that FairWalk can consider only one sensitive attribute at a time, thus we run it with each one of the two sensitive attributes and show the results of both runs. In Table 4, we show the results of all methods on Pokec-n in terms of utility, efficiency, and fairness with respect to two sensitive attributes. First of all, with respect to standard embedding methods, FairMILE improves the efficiency and the fairness towards both sensitive attributes while the utility remains competitive. For example, FairMILE reduces the Δ_{DP}

of DeepWalk by 65.8% on ‘region’ and 51.1% on ‘gender’, respectively. Second, although FairWalk also fulfills the fairness towards the assigned attribute, FairMILE-Node2vec has a better fairness score on both attributes. Third, FairMILE significantly outperforms CFGE on efficiency given that they have comparable utility and fairness performance. To study the performance of CFGE, we tune the hyperparameter λ which controls the strength of fairness. We pick λ from 1 to 100. Going beyond 100 we find the drop in utility exceeds 20% which is often unacceptable. When the constraint is strict ($\lambda = 100$), CFGE has better fairness outcomes at a significant cost to the utility. For $\lambda = 1, 10$ and 50 , CFGE and FairMILE have competitive performance in terms of fairness and utility tradeoff. However, while CFGE takes around 9500 seconds, FairMILE finishes in only 200-300 seconds, which is up to $46\times$ faster.

5.4 Ablation Study

In the ablation study, we showcase the impact of coarsen level c on FairMILE’s performance and the effectiveness of its modules in the fairness of learned representations. In our supplement, we also play with hyperparameters λ_c and λ_r to study the trade-off between fairness and effectiveness.

5.4.1 Impact of coarsening. We vary the coarsen level c to observe its impact on graph sizes and model performance. Table 5 shows the results with NetMF on the Credit dataset. Other results are similar and can be found in our supplement. Specifically, we study:

P1: How the graph changes after each time of coarsening. We observe that increasing c exponentially reduces the numbers of nodes and edges, which corroborates the analysis in Section 4.4.

Table 4: Case study on Pokec-n consisting of multiple sensitive attributes.

Method	AUROC (\uparrow)	Micro-F1 (\uparrow)	Region		Gender		Time (\downarrow)
			Δ_{DP} (\downarrow)	Δ_{EO} (\downarrow)	Δ_{DP} (\downarrow)	Δ_{EO} (\downarrow)	
NetMF	73.12 ± 0.46	62.18 ± 0.37	0.26 ± 0.15	1.82 ± 0.48	0.39 ± 0.24	1.04 ± 0.58	1542.78
FairMILE-NetMF	68.86 ± 0.91	55.96 ± 0.67	0.17 ± 0.07	1.18 ± 0.29	0.31 ± 0.06	0.56 ± 0.30	204.99
DeepWalk	72.58 ± 0.67	61.44 ± 0.68	0.38 ± 0.07	1.49 ± 0.63	0.45 ± 0.22	1.61 ± 0.47	1139.30
FairMILE-DeepWalk	67.43 ± 1.16	54.98 ± 0.88	0.13 ± 0.11	0.70 ± 0.33	0.22 ± 0.12	0.59 ± 0.24	352.38
Node2vec	72.71 ± 0.46	61.71 ± 0.32	0.36 ± 0.05	2.06 ± 0.51	0.44 ± 0.09	1.20 ± 0.49	798.01
Fairwalk (Region)	72.78 ± 0.44	61.99 ± 0.40	0.32 ± 0.09	1.61 ± 0.64	0.40 ± 0.15	1.61 ± 0.49	829.41
Fairwalk (Gender)	72.63 ± 0.49	61.79 ± 0.46	0.34 ± 0.12	1.50 ± 0.67	0.42 ± 0.13	0.98 ± 0.47	784.69
FairMILE-Node2vec	70.10 ± 0.61	57.77 ± 0.92	0.23 ± 0.06	1.22 ± 0.38	0.14 ± 0.08	0.87 ± 0.48	230.85
CFGE $_{\lambda=1}$	71.09 ± 0.76	60.21 ± 1.03	0.28 ± 0.18	0.96 ± 0.49	0.26 ± 0.14	0.68 ± 0.28	9572.75
CFGE $_{\lambda=10}$	68.93 ± 0.95	58.98 ± 0.59	0.26 ± 0.19	1.11 ± 0.59	0.31 ± 0.13	0.50 ± 0.51	9447.50
CFGE $_{\lambda=50}$	58.49 ± 1.25	52.08 ± 0.77	0.17 ± 0.10	0.54 ± 0.31	0.29 ± 0.21	0.40 ± 0.28	9521.35
CFGE $_{\lambda=100}$	56.01 ± 0.59	51.04 ± 0.31	0.16 ± 0.06	0.29 ± 0.25	0.14 ± 0.06	0.23 ± 0.13	9391.33

P2: How FairMILE improves the efficiency by coarsening.

Generally, the efficiency is significantly improved when c increases. A small c (e.g., $c = 1$) may make FairMILE slower because the time of coarsening and refinement outweighs the saved time of learning embeddings when the coarsened graph is not small enough.

P3: How the fairness evolves with varying c . In terms of fairness in the downstream task, we observe that increasing c can visibly improve the fairness of representations. For example, vanilla NetMF has $\Delta_{DP} = 2.66$ and $\Delta_{EO} = 1.34$, which is improved to $\Delta_{DP} = 0.68$ and $\Delta_{EO} = 0.68$ by FairMILE ($c = 4$).

P4: How the utility is impacted by the information loss. We find increasing c leads to a slight decrease in AUROC and F1 scores. The AUROC score only decreases by 0.3% after FairMILE coarsens the graph 4 times. In some cases, FairMILE achieves a better utility than the base embedding method (i.e., with $c = 1$). Given the little cost of utility, we suggest using a large c for the sake of fairness and efficiency.

5.4.2 Effectiveness of fairness-aware modules. To evaluate the effectiveness of our fairness-aware modules for graph coarsening and refinement, we observe the change in performance when we remove the fairness-aware designs. Additionally, we choose MILE [35] as a baseline of scalable embedding methods without fairness considerations. Table 6 reports the results on three datasets. The result shows that FairMILE outperforms MILE in terms of fairness, which demonstrates our effectiveness in fairness improvement. We also notice that **the fairness scores decline on all datasets when removing either fairness-aware design in FairMILE**, indicating that these designs effectively mitigate the bias in the learned embeddings. Furthermore, the little differences between the utility scores of these methods demonstrate that FairMILE is able to improve fairness without impacting utility when compared with MILE. However, this improved fairness does come at some cost to efficiency w.r.t MILE (which is always faster than FairMILE).

5.5 Results for Link Prediction

We next evaluate FairMILE in the context of link prediction. Table 7 only shows the results on Cora and Citeseer. Full results and detailed analysis are in our supplement. First, FairMILE effectively mitigates the bias in link prediction. Our framework has lower Δ_{DP} , LP and

Δ_{EO} , LP than the specialized approaches and graph embedding approaches. On the other hand, FairMILE is more efficient than other baselines, especially the fairness-aware competitors FairAdj and CFGE. Finally, FairMILE has a similar or slightly better utility compared with other approaches. This demonstrates that FairMILE consistently achieves the tradeoff between fairness, efficiency, and utility in a different downstream task.

6 CONCLUSION

In this paper, we study the problem of fair graph representation learning and propose a general framework FairMILE, which can incorporate fairness considerations with any unsupervised graph embedding algorithms and learn fair embeddings towards one or multiple sensitive attributes. We conduct comprehensive experiments to demonstrate that with respect to state-of-the-art techniques for fair graph representation learning, our work achieves similar or better performance in terms of utility and fairness, while FairMILE can significantly outperform them on the axis of efficiency (up to two orders of magnitude faster). Planned future work includes evaluating the use of FairMILE for real-world graph-based model auditing [37] in deployed online settings.

ACKNOWLEDGMENTS

This material is supported by the National Science Foundation (NSF) under grants OAC-2018627, CCF-2028944, SES-1949037, and CNS-2112471. Any opinions, findings, and conclusions in this material are those of the author(s) and may not reflect the views of the respective funding agency.

REFERENCES

- [1] 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] Taha Atahan, Alkyildiz, Amro Alabsi Aljundi, and Kamer Kaya. 2020. GOSH: Embedding big graphs on small hardware. In *ICPP*. 1–11.
- [3] Abolfazl Asudeh, HV Jagadish, Julia Stoyanovich, and Gautam Das. 2019. Designing fair ranking schemes. In *SIGMOD*. 1259–1276.
- [4] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *NeurIPS*, Vol. 29.
- [5] Avishhek Bose and William Hamilton. 2019. Compositional fairness constraints for graph embeddings. In *International Conference on Machine Learning*. PMLR,

Table 5: Impact of coarsen level on graph sizes and FairMILE's performance.

Metric	Vanilla NetMF	$c = 1$	$c = 2$	$c = 3$	$c = 4$
# Nodes	30000	15033	7544	3789	1899
# Edges	1.44M	550K	256K	138K	83K
Time (↓)	240.38	487.60	208.75	139.79	90.28
Δ_{DP} (↓)	2.66 ± 0.55	2.23 ± 0.38	2.07 ± 0.22	1.88 ± 0.37	0.68 ± 0.50
Δ_{EO} (↓)	1.34 ± 0.93	1.22 ± 0.80	1.14 ± 0.76	1.09 ± 0.69	0.68 ± 0.66
AUROC (↑)	74.93 ± 0.43	75.12 ± 0.52	74.95 ± 0.40	74.80 ± 0.41	74.69 ± 0.43
F1 (↑)	88.36 ± 0.08	88.34 ± 0.11	88.41 ± 0.11	88.34 ± 0.11	88.31 ± 0.08

Table 6: Ablation study of each module's effectiveness in fairness (with NetMF).

Dataset	Method	AUROC (↑)	F1 (↑)	Δ_{DP} (↓)	Δ_{EO} (↓)	Time (↓)
German	FairMILE	61.93 ± 3.38	82.35 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	6.31
	FairMILE w/o fair coarsening	63.61 ± 2.91	82.27 ± 0.58	2.20 ± 0.98	1.47 ± 1.54	5.98
	FairMILE w/o fair refinement	63.37 ± 2.91	81.08 ± 0.60	2.94 ± 1.58	2.02 ± 1.06	5.64
	MILE	63.02 ± 2.76	82.04 ± 0.76	3.28 ± 2.10	2.01 ± 0.97	5.06
Recidivism	FairMILE	89.52 ± 0.50	77.65 ± 0.47	2.81 ± 0.50	0.75 ± 0.55	29.66
	FairMILE w/o fair coarsening	90.23 ± 0.43	79.62 ± 0.64	3.18 ± 0.33	1.34 ± 0.61	22.36
	FairMILE w/o fair refinement	90.25 ± 0.23	79.94 ± 0.50	3.27 ± 0.36	1.30 ± 0.66	26.01
	MILE	90.75 ± 0.30	80.46 ± 0.64	3.29 ± 0.32	1.43 ± 0.65	18.42
Credit	FairMILE	74.69 ± 0.43	88.31 ± 0.08	0.68 ± 0.50	0.68 ± 0.66	90.28
	FairMILE w/o fair coarsening	74.65 ± 0.34	88.30 ± 0.12	1.50 ± 0.52	0.92 ± 0.74	57.21
	FairMILE w/o fair refinement	74.42 ± 0.33	88.25 ± 0.15	3.15 ± 0.49	1.71 ± 0.68	82.01
	MILE	74.65 ± 0.41	88.33 ± 0.12	2.54 ± 0.48	1.38 ± 0.81	49.98

Table 7: Comparison on utility, fairness, and efficiency metrics in link prediction between FairMILE and other baselines.

Dataset	Method	AUROC (↑)	AP (↑)	Accuracy (↑)	$\Delta_{DP, LP}$ (↓)	$\Delta_{EO, LP}$ (↓)	Time (↓)
Cora	VGAE	90.86 ± 0.79	92.81 ± 0.89	82.68 ± 1.19	47.51 ± 2.47	24.12 ± 3.29	15.17
	FairAdj $T_2=2$	89.56 ± 1.06	91.31 ± 1.21	81.57 ± 1.22	45.47 ± 2.52	20.44 ± 3.48	56.76
	FairAdj $T_2=5$	88.49 ± 1.30	90.34 ± 1.36	80.97 ± 0.83	42.39 ± 2.95	16.75 ± 3.26	70.37
	FairAdj $T_2=20$	85.97 ± 0.62	87.84 ± 0.53	77.68 ± 0.54	35.67 ± 1.45	11.26 ± 3.07	145.01
	CFGE	90.76 ± 1.27	92.59 ± 1.30	82.35 ± 1.26	32.06 ± 1.33	13.73 ± 1.63	531.79
	NetMF	91.33 ± 1.35	92.96 ± 1.16	86.41 ± 1.10	42.03 ± 2.57	17.24 ± 2.48	22.80
Citeseer	FairMILE-NetMF	92.51 ± 0.74	93.56 ± 1.03	84.42 ± 0.87	41.19 ± 1.47	16.09 ± 1.46	8.16
	DeepWalk	88.52 ± 1.01	89.84 ± 0.77	77.78 ± 0.99	42.53 ± 1.55	21.58 ± 2.19	47.30
	FairMILE-DeepWalk	92.00 ± 1.27	93.67 ± 0.93	82.26 ± 0.64	33.17 ± 1.02	14.31 ± 1.40	19.40
	Node2vec	90.87 ± 1.03	92.03 ± 0.77	80.15 ± 0.55	45.62 ± 1.80	23.37 ± 3.09	21.30
	Fairwalk	89.49 ± 1.24	91.00 ± 0.88	78.92 ± 0.99	42.21 ± 2.28	18.68 ± 2.68	22.71
	FairMILE-Node2vec	91.77 ± 0.89	93.13 ± 1.03	85.16 ± 0.74	28.99 ± 0.98	11.65 ± 0.81	13.71
Citeseer	VGAE	88.24 ± 1.83	90.53 ± 2.78	81.67 ± 3.39	24.82 ± 2.67	7.13 ± 3.32	20.33
	FairAdj $T_2=2$	87.90 ± 1.81	89.99 ± 2.81	81.61 ± 3.02	24.17 ± 2.65	6.49 ± 3.34	84.10
	FairAdj $T_2=5$	87.47 ± 1.74	89.48 ± 2.67	81.50 ± 3.01	23.65 ± 2.72	6.47 ± 3.47	109.07
	FairAdj $T_2=20$	86.90 ± 2.23	88.61 ± 3.23	80.41 ± 3.35	22.74 ± 3.23	6.47 ± 3.38	240.16
	CFGE	87.36 ± 3.29	90.45 ± 3.11	71.12 ± 13.42	13.60 ± 4.86	3.00 ± 1.39	478.18
	NetMF	87.62 ± 0.73	91.20 ± 0.46	84.53 ± 0.30	23.32 ± 1.48	4.66 ± 1.72	19.35
Citeseer	FairMILE-NetMF	89.22 ± 0.34	91.71 ± 0.69	85.30 ± 0.58	20.96 ± 1.65	2.79 ± 1.45	7.57
	DeepWalk	88.83 ± 0.26	90.96 ± 0.31	81.59 ± 1.54	23.99 ± 1.88	5.55 ± 3.14	48.10
	FairMILE-DeepWalk	89.50 ± 1.02	92.55 ± 0.70	86.39 ± 0.37	17.67 ± 1.18	2.62 ± 1.31	21.20
	Node2vec	89.00 ± 0.65	91.77 ± 0.35	83.88 ± 0.81	22.85 ± 1.57	3.52 ± 2.17	16.88
	Fairwalk	88.86 ± 0.85	91.71 ± 0.33	83.52 ± 0.61	23.59 ± 1.38	4.13 ± 2.43	16.78
	FairMILE-Node2vec	89.54 ± 0.95	92.31 ± 0.52	86.91 ± 0.78	12.49 ± 0.91	2.08 ± 1.08	11.72

715–724.

[6] Haochen Chen, Bryan Perozzi, Yifan Hu, and Steven Skiena. 2018. HARP: hierarchical representation learning for networks. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. 2127–2134.

[7] Lee Cohen, Zachary C. Lipton, and Yishay Mansour. 2020. Efficient Candidate Screening Under Multiple Tests and Implications for Fairness. In *1st Symposium on Foundations of Responsible Computing, FORC 2020*, Vol. 156. 1:1–1:20.

[8] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. 2018. A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering* 31, 5 (2018), 833–852.

[9] Sean Current, Yuntian He, Saket Gurukar, and Srinivasan Parthasarathy. 2022. FairEGM: Fair Link Prediction and Recommendation via Emulated Graph Modification. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–14.

[10] 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.

[11] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. 2022. EDITS: Modeling and mitigating data bias for graph neural networks. In *Proceedings of the ACM Web Conference 2022*. 1259–1269.

[12] Yushun Dong, Jing Ma, Chen Chen, and Jundong Li. 2022. Fairness in Graph Mining: A Survey. *arXiv preprint arXiv:2204.09888* (2022).

[13] 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.

[14] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *The World Wide Web Conference*. 417–426.

[15] Danilo Franco, Nicolò Navarin, Michele Donini, Davide Anguita, and Luca Oneto. 2022. Deep fair models for complex data: Graphs labeling and explainable face recognition. *Neurocomputing* 470 (2022), 318–334.

[16] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, et al. 2020. Fairness-aware explainable recommendation over knowledge graphs. In *SIGIR*. 69–78.

[17] Aditya Grover and Jure Leskovec. 2016. Node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 855–864.

[18] Saket Gurukar, Nikil Pancha, Andrew Zhai, Eric Kim, Samson Hu, Srinivasan Parthasarathy, Charles Rosenberg, and Jure Leskovec. 2022. MultiBiSage: A Web-Scale Recommendation System Using Multiple Bipartite Graphs at Pinterest. *Proceedings of the VLDB Endowment* 16, 4 (2022), 781–789.

[19] Saket Gurukar, Priyesh Vijayan, Balaraman Ravindran, Aakash Srinivasan, Goonmeet Bajaj, Chen Cai, Moniba Keymanesh, Saravana Kumar, Pranav Maneriker, Anasua Mitra, et al. 2022. Benchmarking and Analyzing Unsupervised Network Representation Learning and the Illusion of Progress. *Transactions on Machine Learning Research* (2022).

[20] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 1025–1035.

[21] Yuntian He, Saket Gurukar, Pouya Kousha, Hari Subramoni, Dhabaleswar K Panda, and Srinivasan Parthasarathy. 2021. DistMILE: A Distributed Multi-Level Framework for Scalable Graph Embedding. In *HiPC*. IEEE, 282–291.

[22] Yuntian He, Yue Zhang, Saket Gurukar, and Srinivasan Parthasarathy. 2022. WebMILE: democratizing network representation learning at scale. *Proceedings of the VLDB Endowment* 15, 12 (2022), 3718–3721.

[23] Zeinab S Jalali, Weixiang Wang, Myunghwan Kim, Hema Raghavan, and Sucheta Soundarajan. 2020. On the information unfairness of social networks. In *Proceedings of the 2020 SIAM International Conference on Data Mining*. SIAM, 613–521.

[24] Guangyin Jin, Qi Wang, Cunchao Zhu, Yanghe Feng, Jincai Huang, and Jianguiping 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.

[25] Jian Kang, Jingrui He, Ross Maciejewski, and Hanghang Tong. 2020. Inform: Individual fairness on graph mining. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 379–389.

[26] George Karypis and Vipin Kumar. 1998. Multilevel k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed computing* 48, 1 (1998), 96–129.

[27] Ahmad Khajehnejad, Moein Khajehnejad, Mahmoudreza Babaei, Krishna P Gummadi, Adrian Weller, and Baharan Mirzasoleiman. 2021. CrossWalk: Fairness-enhanced Node Representation Learning. *arXiv preprint arXiv:2105.02725* (2021).

[28] Thomas N. Kipf and Max Welling. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016).

[29] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.

[30] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.

[31] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. *Advances in Neural Information Processing Systems* 30 (2017).

[32] Preethi Lahoti, Krishna P Gummadi, and Gerhard Weikum. 2020. Operationalizing Individual Fairness with Pairwise Fair Representations. *Proceedings of the VLDB Endowment* 13, 4 (2020).

[33] Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. Pytorch-bigraph: A large scale graph embedding system. *Proceedings of Machine Learning and Systems* 1 (2019), 120–131.

[34] Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. 2021. On Dyadic Fairness: Exploring and Mitigating Bias in Graph Connections. In *9th International Conference on Learning Representations, ICLR 2021*.

[35] Jiongqian Liang, Saket Gurukar, and Srinivasan Parthasarathy. 2021. MILE: A Multi-Level Framework for Scalable Graph Embedding. *Proceedings of the International AAAI Conference on Web and Social Media* 15, 1 (2021), 361–372.

[36] Andrew Lohn and Micah Musser. 2022. AI and Compute: How Much Longer Can Computing Power Drive Artificial Intelligence Progress. *Center for Security and Emerging Technology (CSET)* (2022).

[37] Pranav Maneriker, Cody Burley, and Srinivasan Parthasarathy. 2023. Online Fairness Auditing through Iterative Refinement. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1665–1676.

[38] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.

[39] Galileo Namata, Ben London, Lise Getoor, Bert Huang, and U Edu. 2012. Query-driven active surveying for collective classification. In *10th International Workshop on Mining and Learning with Graphs*, Vol. 8. 1.

[40] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 701–710.

[41] Jiezhong Qiu, Laxman Dhulipala, Jie Tang, Richard Peng, and Chi Wang. 2021. LightNE: A Lightweight Graph Processing System for Network Embedding. In *SIGMOD*.

[42] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. 2018. Network embedding as matrix factorization: Unifying Deepwalk, LINE, PTE, and Node2vec. In *WSDM*. 459–467.

[43] Tableen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. 2019. Fairwalk: Towards Fair Graph Embedding. In *IJCAI*. AAAI Press, 3289–3295.

[44] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.

[45] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.

[46] Indro Spinelli, Simona Scardapane, Amir Hussain, and Aurelio Uncini. 2021. FairDrop: Biased Edge Dropout for Enhancing Fairness in Graph Representation Learning. *IEEE Transactions on Artificial Intelligence* (2021).

[47] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. *arXiv preprint arXiv:1906.02243* (2019).

[48] Lubos Taka and Michal Zabovsky. 2012. Data analysis in public social networks. In *International Scientific Conference and International Workshop Present Day Trends of Innovations*, Vol. 1.

[49] Sheridan Wall and Hilke Schellmann. 2021. LinkedIn's job-matching AI was biased. The company's solution? More AI. <https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>.

[50] Daixin Wang, Jianbin Lin, Peng Cui, Quanhui Jia, Zhen Wang, Yanming Fang, Quan Yu, Jun Zhou, Shuang Yang, and Yuan Qi. 2019. A semi-supervised graph attentive network for financial fraud detection. In *ICDM*. IEEE, 598–607.

[51] Jizhe Wang, Pipei Huang, Huan Zhao, Zhibo Zhang, Binqiang Zhao, and Dik Lun Lee. 2018. Billion-scale commodity embedding for e-commerce recommendation in alibaba. In *KDD*. 839–848.

[52] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *The 7th International Conference on Learning Representations*.

[53] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*. PMLR, 962–970.

[54] Zhaocheng Zhu, Shizhen Xu, Jian Tang, and Meng Qu. 2019. Graphvite: A high-performance CPU-GPU hybrid system for node embedding. In *The World Wide Web Conference*. 2494–2504.

A FULL THEORETICAL ANALYSIS

Corollary 1. *Algorithm 1 coarsens a graph $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i)$ into a smaller graph $\mathcal{G}_{i+1} = (\mathcal{V}_{i+1}, \mathcal{E}_{i+1})$ such that $\frac{1}{2}|\mathcal{V}_i| \leq |\mathcal{V}_{i+1}| \leq |\mathcal{V}_i|$ and $|\mathcal{E}_{i+1}| \leq |\mathcal{E}_i|$.*

PROOF. In the optimal case, all nodes in \mathcal{G}_i get matched and therefore $|\mathcal{V}_{i+1}| = \frac{1}{2}|\mathcal{V}_i|$. The worst case is that all nodes are isolated so the number of nodes does not decrease.

Our coarsening algorithm adds an edge (u, v) in \mathcal{G}_{i+1} if and only if there exists at least one edge in \mathcal{G}_i that connects one of u 's child nodes and one of v 's child nodes. Therefore $|\mathcal{E}_{i+1}| \leq |\mathcal{E}_i|$. \square

Lemma 1. *In the phase of graph coarsening, FairMILE consumes $O\left(M\left(\sum_{i=0}^{c-1}(|\mathcal{V}_i| + |\mathcal{E}_i|)\right)\right)$ time.*

PROOF. Without loss of generality, we assume the input graph is \mathcal{G}_i ($i < c$). For each edge, Algorithm 1 computes the fairness-aware edge weight in $O(M)$ time. Hence it takes $O(M(|\mathcal{E}_i| + |\mathcal{V}_i|))$ time to match and merge nodes. The time of creating a coarsened graph after node matching is also $O(M(|\mathcal{E}_i| + |\mathcal{V}_i|))$, which is mainly used for computing the attribute distribution of new nodes and the weights of new edges in the coarsened graph. Summing up the c coarsen levels, the total time complexity of graph coarsening is $O\left(M\left(\sum_{i=0}^{c-1}(|\mathcal{V}_i| + |\mathcal{E}_i|)\right)\right)$. \square

Corollary 2. *FairMILE can reduce the time of graph embedding exponentially in the optimal case since $|\mathcal{V}_c| \geq 2^{-c}|\mathcal{V}_0|$.*

PROOF. In Corollary 1, we analyze that the number of nodes can be reduced by up to half at each coarsen level. Thus the time complexity of base embedding can also be reduced exponentially when c increases. \square

Lemma 2. *If the refinement model has l layers, the time complexity of refinement is $O\left(l(d+M)\left[\sum_{i=0}^{c-1}(|\mathcal{E}_i| + d|\mathcal{V}_i|)\right]\right)$.*

PROOF. We again assume the input graph is \mathcal{G}_{i+1} ($i > 0$) without loss of generality. Before applying the model, FairMILE projects the embeddings from the supernodes to the child nodes of \mathcal{G}_i in $O(d|\mathcal{V}_i|)$ time. In each layer of the model, FairMILE needs $O((d+M)|\mathcal{V}_i|)$ to concatenate the input with the sensitive attributes. The following message passing process and the matrix multiplication take $O((d+M)|\mathcal{E}_i|)$ and $O(d(d+M)|\mathcal{V}_i|)$ time, respectively. Therefore the time complexity of each layer is $O((d+M)(|\mathcal{E}_i| + d|\mathcal{V}_i|))$. Finally, the total time complexity of applying the refinement model is $O\left(l(d+M)\left[\sum_{i=0}^{c-1}(|\mathcal{E}_i| + d|\mathcal{V}_i|)\right]\right)$. \square

Theorem 1. *When L_f is minimized, the 2-norm of the difference between the mean embeddings of any two demographic groups regarding a given sensitive attribute is bounded by*

$$\|\mu_p - \mu_q\|_2 \leq 2(1 - \min(\beta_p, \beta_q)) \quad (8)$$

where p, q are any two different values of the given sensitive attribute (e.g., gender or race). For $i \in \{p, q\}$, μ_i denotes the mean embedding values of nodes from group i , and β_i denotes the ratio of nodes from group i that have at least one inter-group edge.

PROOF. Our fairness loss L_f in Equation (6) is the average negative cosine similarity of embeddings of all connected node pairs $(u, v) \in \mathcal{E}'_c$ with diverse sensitive attributes, which means $\phi(u, v) \geq \gamma$. To reach the global minimum of L_f , the parameters $\{\Theta_i\}_{i=1}^l$ are optimized to completely focus on sensitive attributes and generate identical embeddings for the nodes, i.e., the similarity is 1. In general, using a smaller γ will have more edges impacted by the fairness objective. $\gamma = 0$ is the strictest value that always leads to all nodes with inter-group connections having the same representations (denoted as \bar{h}). A special case is when $c = 1$, $\phi(u, v) = 1$ holds for any $u, v \in \mathcal{V}_0$ such that $s_u \neq s_v$. Therefore any γ enforces fairness in the same strength.

Recall that p, q are any two different values of the given sensitive attribute (e.g., gender). For $i \in \{p, q\}$, let \mathcal{U}_i be the set of nodes with attribute value i (e.g., all nodes in \mathcal{U}_i share the same gender i). Then let \mathcal{U}_{ii} be the subset of \mathcal{U}_i in which the nodes are only connected to nodes in \mathcal{U}_i . Note that $\beta_i = 1 - |\mathcal{U}_{ii}|/|\mathcal{U}_i|$. The learned embedding of node u is denoted as \mathbf{h}_u . When L_f is minimized with $\gamma = 0$, we have

$$\begin{aligned} & \|\mu_p - \mu_q\|_2 \\ &= \|\beta_p \bar{h} + \frac{1 - \beta_p}{|\mathcal{U}_{pp}|} \sum_{u \in \mathcal{U}_{pp}} \mathbf{h}_u - \beta_q \bar{h} - \frac{1 - \beta_q}{|\mathcal{U}_{qq}|} \sum_{v \in \mathcal{U}_{qq}} \mathbf{h}_v\|_2 \\ &\leq |\beta_p - \beta_q| \|\bar{h}\|_2 + (1 - \beta_p) \|\frac{1}{|\mathcal{U}_{pp}|} \sum_{u \in \mathcal{U}_{pp}} \mathbf{h}_u\|_2 + (1 - \beta_q) \|\frac{1}{|\mathcal{U}_{qq}|} \sum_{v \in \mathcal{U}_{qq}} \mathbf{h}_v\|_2 \end{aligned}$$

Note that the output embeddings of our refinement model are L2-normalized, hence we finally have

$$\|\mu_p - \mu_q\|_2 \leq |\beta_p - \beta_q| + (1 - \beta_p) + (1 - \beta_q) = 2(1 - \min(\beta_p, \beta_q))$$

Therefore Equation (8) holds when L_f is minimized. \square

B DATASET DESCRIPTION

In *German* [1], each node is a client in a German bank, and two nodes are linked if their attributes are similar. The task is to classify a client's credit risk as good or bad, and the sensitive attribute is the client's gender.

Recidivism [1] is a graph created from a set of bail outcomes from US state courts between 1990-2009, where nodes are defendants and edges connect two nodes if they have similar past criminal records and demographic attributes. The task is to predict if a defendant will commit a violent crime or not, while the sensitive attribute is race.

Credit [1] consists of credit card applicants with their demographic features and payment patterns. Each node is an individual, and two nodes are connected based on feature similarity. In this dataset, age is used as the sensitive attribute, and the predicted label is whether the applicant will default on an upcoming payment.

Pokec-n [48] is collected from a Slovakia social network. We use both region and gender as the sensitive attributes, and choose each user's field of work as the predicted label. Note that Pokec-n has multiple sensitive attributes and a multi-class target, which FairMILE can handle by design. However, existing research [10, 11, 15] has only evaluated the use of this data with one sensitive attribute at-a-time with the target label binarized - a key limitation. We discuss how FairMILE can redress this limitation in Section 5.3.

Table 8: Comparison in node classification between FairMILE and other baselines on Recidivism dataset.

Dataset	Method	AUROC (\uparrow)	F1 (\uparrow)	Δ_{DP} (\downarrow)	Δ_{EO} (\downarrow)	Time (\downarrow)
Recidivism	NetMF	94.63 \pm 0.17	85.46 \pm 0.29	3.41 \pm 0.21	1.62 \pm 0.78	141.90
	FairMILE-NetMF	89.52 \pm 0.50	77.65 \pm 0.47	2.81 \pm 0.50	0.75 \pm 0.55	29.66
	DeepWalk	93.33 \pm 0.35	83.62 \pm 0.42	3.47 \pm 0.37	1.28 \pm 0.60	303.68
	FairMILE-DeepWalk	86.93 \pm 0.74	73.50 \pm 0.99	2.71 \pm 0.58	1.08 \pm 0.77	45.93
	Node2vec	92.56 \pm 0.26	83.31 \pm 0.36	3.61 \pm 0.56	1.57 \pm 0.97	136.33
	Fairwalk	92.43 \pm 0.43	82.99 \pm 0.51	3.32 \pm 0.24	1.48 \pm 0.66	133.62
	FairMILE-Node2vec	87.00 \pm 0.50	71.34 \pm 0.86	2.75 \pm 0.35	1.15 \pm 0.65	38.67
	Vanilla GCN	88.16 \pm 1.72	77.68 \pm 1.63	3.83 \pm 0.59	1.46 \pm 0.71	474.57
	FairGNN	67.26 \pm 7.80	44.63 \pm 14.87	0.67 \pm 0.45	1.24 \pm 0.40	1071.39
	NIFTY	77.89 \pm 4.21	64.44 \pm 6.11	1.34 \pm 1.01	0.63 \pm 0.42	1651.09

The remaining three datasets (namely, *Cora* [44], *Citeseer* [44], and *Pubmed* [39]) are citation networks widely evaluated in the graph representation learning literature [9, 34, 46]. In these data, each node denotes a paper, and each edge links two nodes if one paper cites the other. As in prior work [9, 34, 46], we treat the category of a paper as its sensitive attribute. The task is to predict whether a paper is cited by the other (or vice versa).

C NODE CLASSIFICATION ON RECIDIVISM

We conduct the experiments of node classification on another dataset Recidivism and revisit the questions in Section 5.2. Results are shown in Table 8.

A1) Fairness: FairMILE improves the fairness of all unsupervised graph embedding approaches. In Recidivism, FairMILE decreases the Δ_{DP} scores of NetMF and DeepWalk by 17.6% and 21.9%, respectively. In terms of Δ_{EO} , FairMILE improves the fairness of NetMF and DeepWalk by 53.7% and 15.6%. On top of Node2vec, FairMILE outperforms FairWalk in terms of both Δ_{DP} and Δ_{EO} . Among the specialized methods, FairGNN has the lowest Δ_{DP} score and NIFTY has the best Δ_{EO} score which is slightly better than FairMILE-NetMF (0.63% v.s. 0.75%). However, this is because these models trade too much utility for fairness (For example, in terms of AUROC, NIFTY 77.89% v.s. FairMILE-NetMF 89.52%).

A2) Efficiency: FairMILE is more efficient than other baselines. While GNN-based approaches take up to 2498.5 seconds, FairMILE on top of NetMF finishes in only 29.7 seconds, which is 84.2 \times faster. Compared with the unsupervised graph embedding approaches, FairMILE still improves the efficiency of graph embedding.

A3) Utility: Compared with the base embedding methods, the utility scores of FairMILE slightly drop which is fairly remarkable given that FairMILE significantly improves fairness and efficiency. Among the specialized approaches, all approaches except the vanilla GCN are outperformed by FairMILE in terms of AUROC and F1. This demonstrates that FairMILE achieves a better tradeoff between utility and fairness than these GNN-based approaches.

In summary, FairMILE on top of graph embedding approaches can compete or improve on fairness and utility with various specialized methods while outperforming them significantly in terms of efficiency.

D FULL ABLATION STUDY

D.1 Tuning the Coarsen Level

We vary the coarsen level c to observe its impact on utility, fairness, and efficiency. Results are shown in Figure 3. Note that when $c = 0$, FairMILE is performing the base embedding method on the original graph. Generally, increasing c leads to a slight decrease in AUROC and F1 scores. For example, the AUROC score of DeepWalk only decreases by 0.6% after FairMILE coarsens the graph 4 times. In some cases, FairMILE achieves a better utility than the base embedding method (e.g., FairMILE-Node2vec with $c = 1$ on German). While the decrease of utility is negligible, increasing c can visibly improve the fairness of representations. For example, vanilla DeepWalk has $\Delta_{DP} = 7.22$ and $\Delta_{EO} = 7.69$ on German, which is improved to $\Delta_{DP} = 0.67$ and $\Delta_{EO} = 0.26$ by FairMILE ($c = 2$). Last of all, increasing the coarsen level significantly improves the efficiency. Using a small c may make FairMILE slower because the time of coarsening and refinement outweighs the saved time of learning embedding when the coarsened graph is not small enough. Examples include $c = 1$ on Credit. Given the little cost of utility, we suggest using a large c for the sake of fairness and efficiency.

D.2 Trade-off between Utility and Fairness

To further explore the trade-off of FairMILE between utility and fairness, we choose the values of λ_c and λ_r from $\{0.1, 0.3, 0.5, 0.7, 0.9\}$ respectively to observe the impact on performance. Figure 4 shows the results of FairMILE-NetMF on Recidivism with $c = 4$ (We only report these results for one dataset since results on other datasets are similar). We use AUROC and Δ_{DP} as the metrics for utility and fairness. It is clear that there is a trade-off between the utility scores and the fairness of learned embeddings on this dataset. Increasing fairness (represented by lower Δ_{DP}) often causes a decrease in utility scores. We also observe that λ_r has a larger impact on this tradeoff than λ_c . We also find in general that our choice of $\lambda_c = \lambda_r = 0.5$ achieves a reasonable trade-off (applies to this dataset and the other datasets and tasks in our study). We do note of course that for different scenarios the designer may prefer to choose these parameters appropriately.

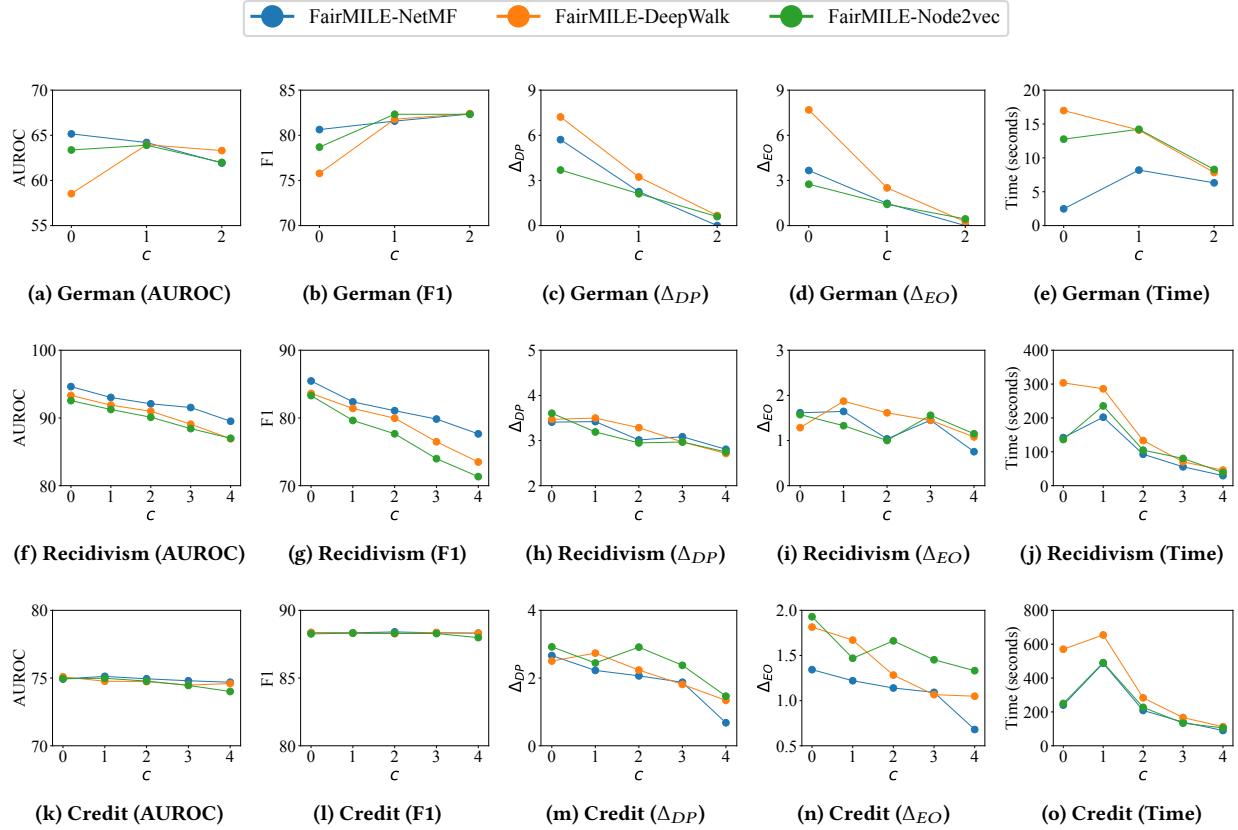


Figure 3: Impact of coarsen level c on FairMILE's utility, fairness, and efficiency.

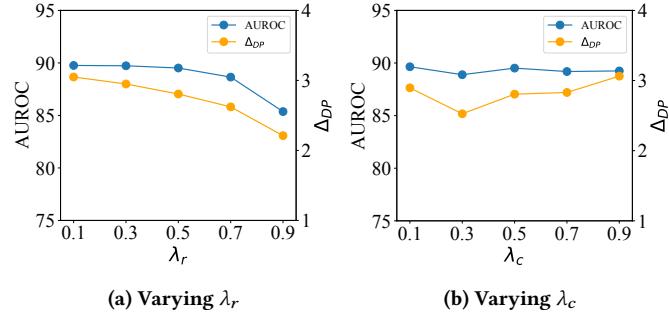


Figure 4: Impact of varying λ_c and λ_r on utility and fairness on Recidivism dataset.

E FULL RESULTS FOR LINK PREDICTIONS

We evaluate FairMILE in the context of link prediction on three datasets. For FairMILE, we set $c = 2$ on smaller datasets (Cora and Citeseer) and $c = 4$ on Pubmed. Table 9 shows the results on Pubmed. For the results on other datasets, please refer to Table 7 in Section 5.5. First, FairMILE makes fair predictions on all datasets. Our framework has an improvement of up to 45.3% on Δ_{DP} , LP compared with the base embedding approaches. In terms of Δ_{EO} , LP, while the performance of FairMILE declines on Pubmed very slightly (2.70% v.s. 2.04% in NetMF), it greatly reduces the unfair predictions on Cora and Citeseer. Combining the observations on both metrics,

FairMILE successfully enforces fairness in the task of link prediction. When compared with FairWalk, FairMILE-Node2vec always has a better fairness score (e.g., 12.49% v.s. 23.59% on Citeseer). In addition, we notice that FairAdj is less biased than VGAE, which demonstrates its effectiveness in debiasing. However, its best performance with $T2 = 20$ is still outperformed by FairMILE on all datasets. For example, the Δ_{DP} , LP score of FairMILE-Node2vec on Citeseer is 45.1% lower than that of FairAdj ($T2 = 20$). Compared with CFGE, FairMILE on top of Node2vec has a better performance in terms of fairness.

Table 9: Comparison in link prediction between FairMILE and other baselines on Pubmed dataset.

Dataset	Method	AUROC (\uparrow)	AP (\uparrow)	Accuracy (\uparrow)	$\Delta_{DP, LP} (\downarrow)$	$\Delta_{EO, LP} (\downarrow)$	Time (\downarrow)
Pubmed	VGAE	95.03 \pm 0.18	94.96 \pm 0.19	87.48 \pm 0.21	39.37 \pm 0.88	10.28 \pm 1.58	347.80
	FairAdj $T_2=2$	94.29 \pm 0.17	94.07 \pm 0.14	86.63 \pm 0.27	37.12 \pm 0.89	7.57 \pm 1.56	2218.41
	FairAdj $T_2=5$	93.57 \pm 0.19	93.21 \pm 0.13	85.69 \pm 0.16	35.06 \pm 1.01	5.60 \pm 1.55	2480.66
	FairAdj $T_2=20$	91.78 \pm 0.12	91.27 \pm 0.24	83.20 \pm 0.22	30.41 \pm 0.89	2.41 \pm 1.28	4532.99
	CFGE	91.25 \pm 5.32	92.08 \pm 5.22	82.78 \pm 5.79	33.03 \pm 6.19	9.55 \pm 2.35	4237.43
	NetMF	98.43 \pm 0.07	98.26 \pm 0.05	93.86 \pm 0.19	38.59 \pm 0.14	2.04 \pm 0.15	281.35
	FairMILE-NetMF	98.11 \pm 0.12	97.29 \pm 0.20	94.84 \pm 0.31	31.97 \pm 0.68	2.70 \pm 0.22	126.17
	DeepWalk	98.35 \pm 0.14	98.05 \pm 0.17	91.77 \pm 0.29	35.02 \pm 0.43	0.40 \pm 0.12	354.27
	FairMILE-DeepWalk	99.57 \pm 0.04	99.32 \pm 0.08	97.61 \pm 0.06	27.30 \pm 0.23	0.37 \pm 0.11	201.03
	Node2vec	99.52 \pm 0.04	99.44 \pm 0.04	93.11 \pm 0.21	40.28 \pm 0.41	0.21 \pm 0.13	249.52
	Fairwalk	99.50 \pm 0.05	99.43 \pm 0.05	92.86 \pm 0.24	38.58 \pm 0.35	0.65 \pm 0.12	225.99
	FairMILE-Node2vec	99.23 \pm 0.07	98.68 \pm 0.14	96.43 \pm 0.06	26.51 \pm 0.35	0.59 \pm 0.05	143.01

On the other hand, FairMILE also performs well in terms of utility. In comparison to the standard embedding approaches, FairMILE achieves a similar or better utility performance. For example, FairMILE **actually enhances the accuracy** of DeepWalk from 91.77% to 97.61% on Pubmed. Similar results can also be observed on the other metrics and datasets. Compared with VGAE-based methods, FairMILE outperforms them again on utility. Examples include that AUROC scores of VGAE and FairMILE-DeepWalk on Pubmed are 95.03% v.s. 99.57%, respectively.

Finally, FairMILE is more efficient than other baselines. For example, on the largest dataset Pubmed, FairMILE-NetMF takes around 2 minutes, while NetMF needs around 5 minutes, and FairAdj with $T_2 = 20$ even requires more than one hour to finish. In summary, FairMILE can flexibly generalize to the link prediction task improving over the state of the art on both counts of fairness and efficiency at a marginal cost to utility.