Reproducibility Study Of "Learning Fair Graph Representations Via Automated Data Augmentations"

NeurIPS 2024

Thijmen Nijdam, Taiki Papandreou-Lazos, Juell Sprott, Jurgen de Heus

Why fair graph representations?

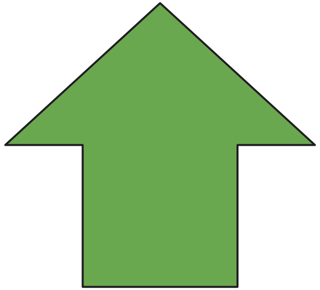

GNNs are increasingly popular…

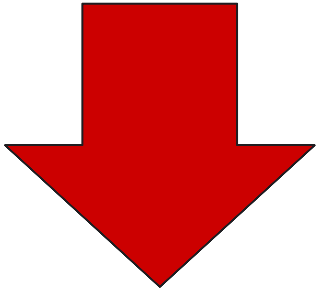

…but can inherit or even amplify bias in training data (Dai & Wang, 2021)

How to create fair graph representations?

The authors propose a graph augmentation framework, called Graphair, that is optimized for both fairness and informativeness (Ling et al., 2022)

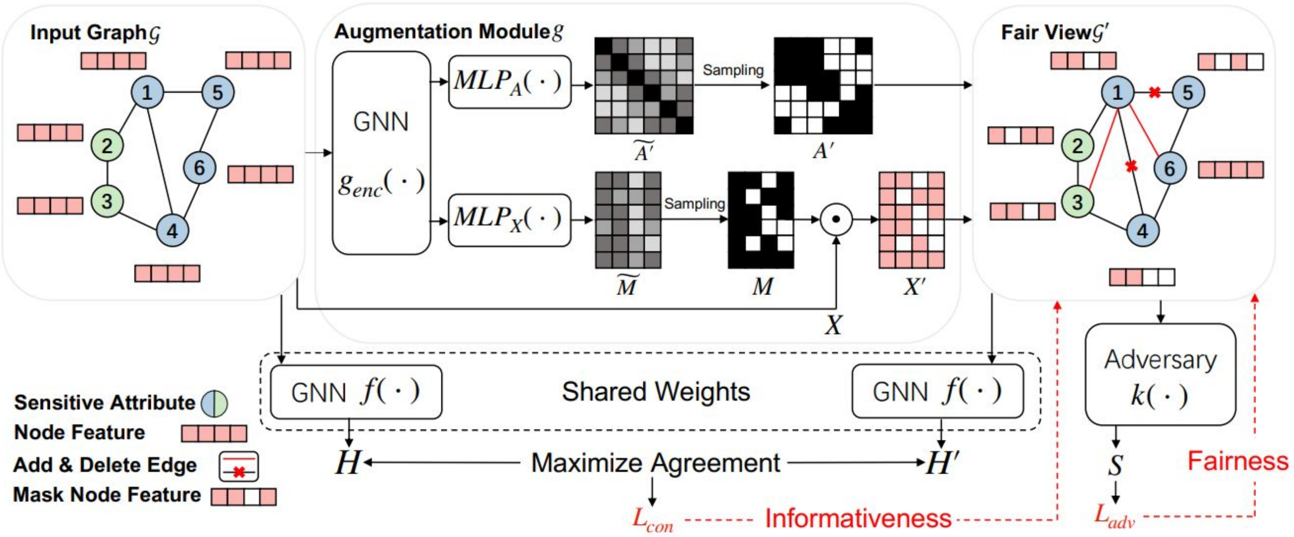

(Ling et al., 2022)

Claims

- Claim 1: Graphair consistently outperforms baselines in fairness-accuracy trade-off in node classification tasks

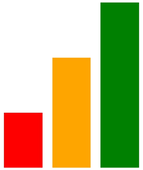

- Claim 2: Both the feature embeddings and the graph topology of the augmented graph contribute to mitigating prediction bias

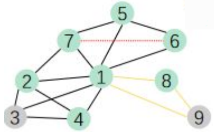

- Claim 3: Graphair can automatically learn to generate fair graph data without prior knowledge of fairness-relevant graph properties

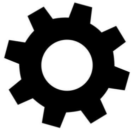

Claim 1: Outperform baselines in node classification

- NBA dataset is fully reproducible
- Pokec datasets are not fully reproducible with the available code[1]

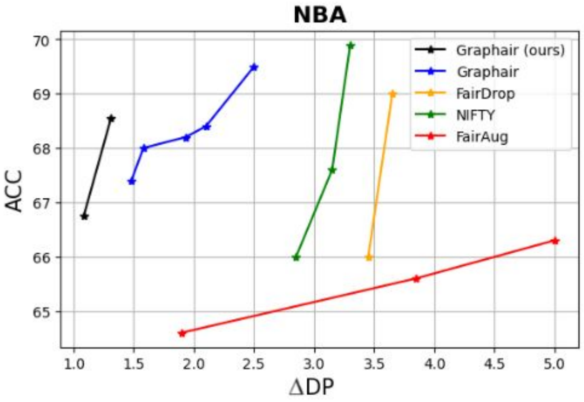

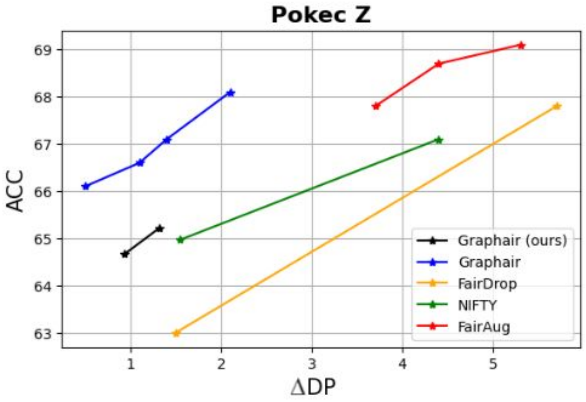

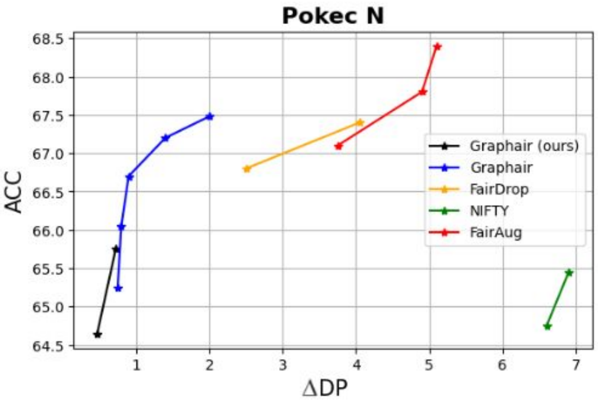

1: https://github.com/divelab/DIG

Extension: Link prediction

- Metrics
  - Mixed dyadic-level fairness
  - Subgroup dyadic-level fairness
- Baselines
  - FairDrop (Spinelli et al., 2021)
  - FairAdj (Li et al., 2021)

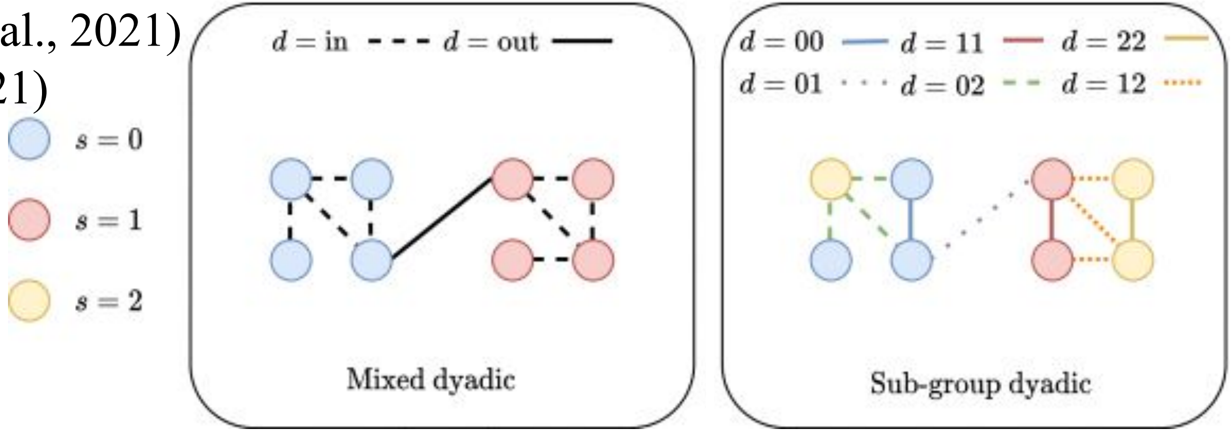

Indro Spinelli et al.[2]

Extension: Results

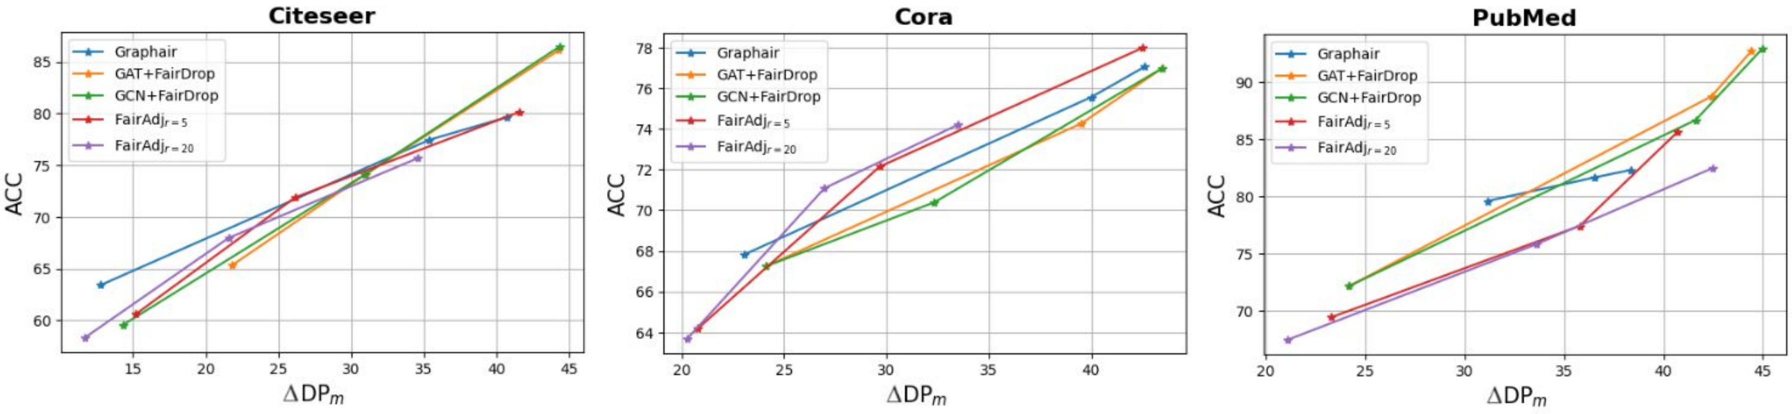

Graphair performs comparably against baselines on the mixed dyadic-level fairness metric

Extension: Results

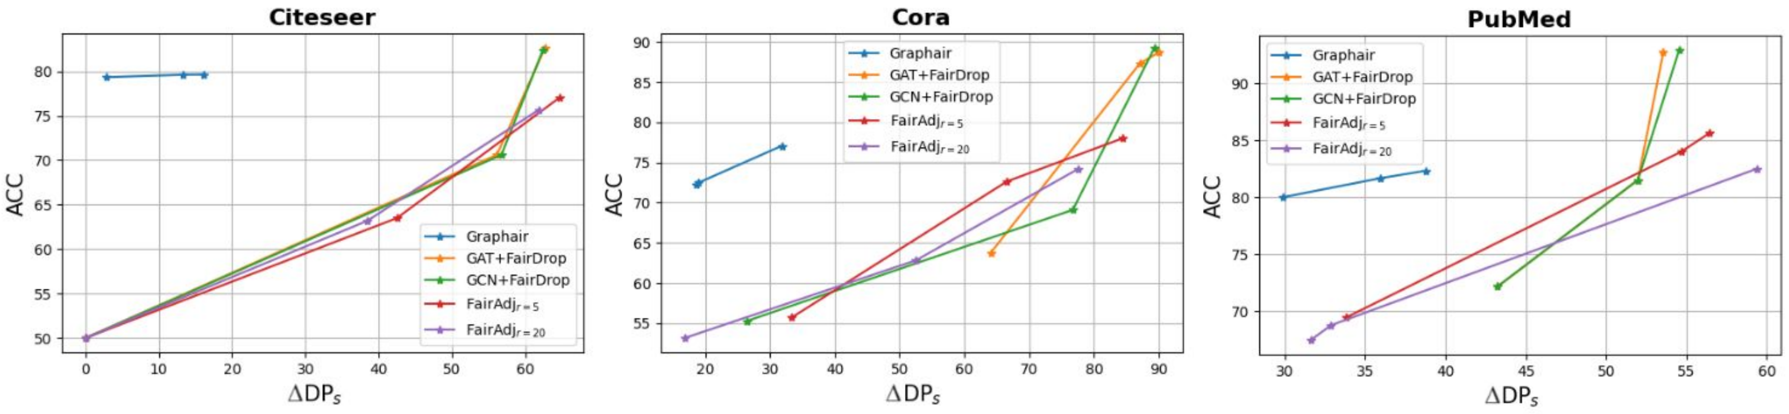

Graphair outperforms baselines on the subgroup dyadic-level fairness metric

Claim 2: Component contribution analysis

Graphair w/o feature masking

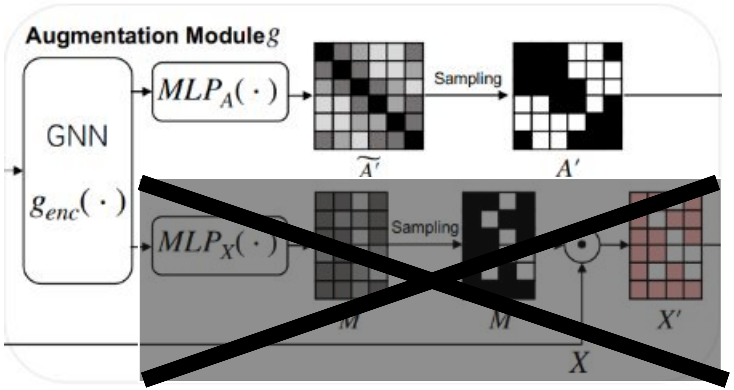

Graphair w/o edge perturbation

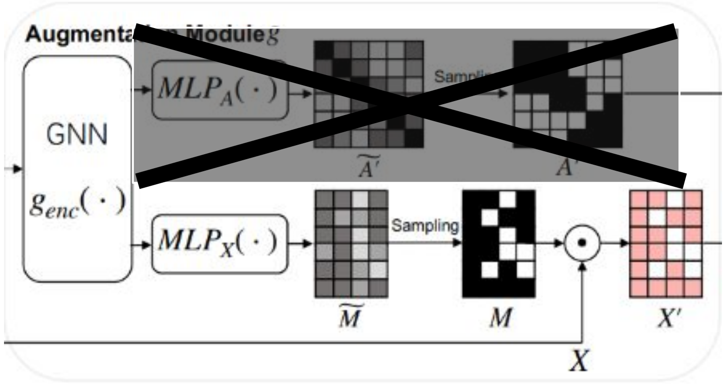

Fairness 

Conclusion: Both FM and EP are necessary!

Claim 3: Automated Fair Graph Generation

Fair Node Topology

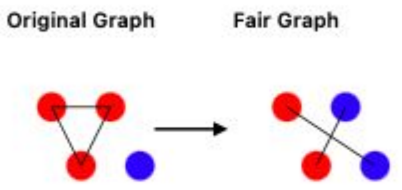

Homophily

Fairness

Fair Node Features

Spearman Correlation

Fairness

Discussion (Main Claims)

Claim 1	Claim 2	Claim 3

⚠️ **500 vs 10,000 Epochs**

Discussion (Extension)

- Link prediction performance on fixed dyadic-level fairness

Graphair ≈ Baselines

- Link prediction performance on subgroup dyadic-level fairness

Graphair > Baselines

Comments on reproducibility

- The code is part of a library and is well-documented
- Code and paper are closely aligned
- The paper is clear and well-written
- The paper and code provides a straightforward way to reproduce results

References

Hongyi Ling, Zhimeng Jiang, Youzhi Luo, Shuiwang Ji, and Na Zou. Learning fair graph representations via automated data augmentations. In The Eleventh International Conference on Learning Representations, 2022.

Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. On dyadic fairness: Exploring and mitigating bias in graph connections. In International Conference on Learning Representations, 2021.

Indro Spinelli, Simone Scardapane, Amir Hussain, and Aurelio Uncini. Fairdrop: Biased edge dropout for enhancing fairness in graph representation learning. IEEE Transactions on Artificial Intelligence, 3(3): 344–354, 2021.

Enyan Dai and Suhang Wang. Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information. In Proceedings of the 14th ACM International Conference on Web Search and Data Mining, pp. 680–688, 2021.