A Comparative Analysis of Graph Construction Methods for Fairness in Graph-Based Semi-Supervised Learning

Willian Dihanster Gomes de Oliveira [1], **Lilian Berton**[1]

[1]Institute of Science and Technology
Federal University of São Paulo - UNIFESP
São José dos Campos – SP – Brazil

{*dihanster.willian*, *lberton*}@unifesp.br

Abstract. *Graph-based semi-supervised learning (GSSL) relies on the quality of the input graph, often constructed from tabular data. We compare KNN, Mutual KNN (M-KNN), and Sequential KNN (S-KNN) on fairness benchmark datasets. While KNN achieves the highest AUC-ROC, M-KNN and S-KNN, which promote more balanced degree distributions, yield fairer outcomes as measured by Disparate Impact and Average Absolute Odds Difference, with only a small loss in AUC-ROC. Additionally, our analysis reveals that degree disparities between protected groups are linked to unfair outcomes, and that more balanced connectivity can mitigate these effects, highlighting the importance of fairness-aware graph construction in GSSL.*

1. Introduction

Graph-based machine learning methods, such as Label Propagation Algorithms (LPAs) and Graph Neural Networks (GNNs), have gained increasing attention due to their ability to incorporate relational structure into learning processes [Wu et al. 2020, Song et al. 2022]. While originally popularized in domains with inherent graph structures, such as citation networks, social networks, and molecular data, recent works have demonstrated the feasibility of applying graph-based learning techniques to tabular and other types of data by constructing artificial graphs [Berton et al. 2018, de Oliveira et al. 2020, Liao and Li 2023].

In these cases, each data instance is represented as a node, and edges are introduced based on similarity functions, often through k-nearest neighbors (KNN) or ε-radius methods, typically using Euclidean distance metrics. These constructed graphs allow graph-based models to operate on otherwise non-relational data. However, despite their practical utility, such construction processes introduce inductive biases that can significantly influence classification outcomes [Wang et al. 2024, de Oliveira and Berton 2024]. In particular, the choice of similarity metric, the value of k, and the method used to connect nodes may amplify existing patterns in the data that relate to sensitive attributes such as gender, race, or age, for example.

The potential for bias is especially concerning in high-stakes applications, such as healthcare, finance, or criminal justice, where tabular data is ubiquitous and fairness is critical. For example, individuals from underrepresented groups may have fewer or noisier neighbors in the constructed graph, leading to worse model performance for those subpopulations. This issue has been observed in various fairness audits of GNNs and similar algorithms [Chen et al. 2024]. Moreover, fairness-aware learning for graph-based

models remains a nascent field, and many fairness mitigation strategies developed for conventional models cannot be directly transferred to the graph setting.

In this work, we investigate how different graph construction strategies applied to tabular data can introduce or exacerbate bias in classification tasks using graph-based learning methods. We focus on neighborhood-based methods such as KNN, analyzing their impact on both performance and fairness metrics across sensitive subgroups. By conducting systematic experiments and subgroup analysis, we aim to highlight the consequences of graph induction choices and provide guidelines for more equitable graph-based learning on tabular data.

Our contributions are:

- We evaluate how graph construction via KNN and its variants affects subgroup performance and fairness in tabular classification.
- We demonstrate that standard graph induction techniques can lead to more discriminatory graph topologies that harm unprivileged groups.
- We propose recommendations for practitioners on constructing fairer graphs from tabular datasets.

The remainder of the paper is structured as follows: Section 2 presents related work; Section 3 presents theoretical concepts of graph theory, graph construction methods, graph semi-supervised learning, and graph fairness; Section 4 presents the experimental methodology of the current work; Section 5 presents the results obtained by the experiments performed and a discussion of them; and finally, Section 6 presents the conclusions of the work.

2. Related Work

Previous studies have investigated bias in graphs; a chronological selection of key contributions is outlined below. [Bose and Hamilton 2019] propose training models to make predictions that are both accurate and invariant to sensitive attributes via adversarial losses. [Spinelli et al. 2021] introduces FairDrop, a biased edge dropout strategy to counteract homophily-based bias amplified by graph structure in GNN tasks like link prediction. [Dong et al. 2022] demonstrates that biases in graph structure itself (edge patterns) are a major cause of unfair predictions, and offers a framework to identify which edges drive bias. [Wang et al. 2022] shows how feature propagation over graph edges can leak sensitive information, exacerbating bias, and proposes feature-masking methods to mitigate this. [Kose and Shen 2022] reveals that both node features and graph topology contribute to bias and proposes data augmentation strategies that modify the graph during training to boost subgroup fairness. [Dong et al. 2023] reviews how underlying biased graph, whether from KNN, homophily, or relational data, leads to amplification of class or attribute bias in GNN outputs. [Chen et al. 2024] highlights that biased graph structures are a core challenge, and recommends pre-/post-processing solutions like edge rewiring or augmentation during graph construction to improve fairness.

This literature underscores the importance of treating graph structure decisions (number of neighbors, metric, edge selection) as fairness-critical choices. However, most of the works focused on GNN, here we aim to analyze the impact of the graph in label propagation like algorithms.

3. Background

3.1. Graph Theory

A graph $G = (V, E)$ is a mathematical structure used to represent relational data. It consists of a set of vertices (or nodes) V and a set of edges E connecting pairs of nodes. For machine learning classification tasks, we define the graph as an adjacency matrix $A \in R^{n \times n}$, where n is the number of vertices and $A_{ij} = 1$ if there is an edge between v_i and v_j, and 0 otherwise (or a real-valued similarity score in the weighted case).

To build a graph for classification, the following steps must be considered:

- Node Representation: Each data instance corresponds to a node.
- Feature Extraction: Each node is described by a feature vector, forming a feature matrix.
- Edge Construction: Edges are created based on similarity between feature vectors, typically using k-nearest neighbors (KNN) or ϵ-neighborhood approaches with Euclidean distance.
- Label Assignment: Some nodes are labeled with class labels Y, while the goal is to infer labels for the unlabeled nodes using the graph structure.

3.2. Neighborhood-Based Graph Construction Methods

In Graph-based Semi-supervised Learning (GSSL), the initial construction of the graph is a critical step that significantly influences the subsequent performance of the learning model. Among the unsupervised graph construction methods, neighborhood-based approaches are widely employed to generate sparse subgraphs by connecting data points deemed "proximal" within the feature space. These methods are crucial for capturing local data structures and facilitating information propagation in GSSL.

The k-Nearest Neighbor (KNN) graph is one of the most prevalent methods for constructing similarity graphs in GSSL. Its fundamental principle involves connecting each data element to its k closest neighbors. The definition of closest is determined by a pre-configured distance metric, such as Euclidean distance. The primary parameter for KNN graph construction is k, which specifies the number of neighbors to be connected. A notable characteristic of KNN graphs is their tendency to result in irregular structures, where different nodes can exhibit varying degrees (i.e., number of connections). This irregularity can lead to the formation of hubs, which are nodes with very high degrees, and can negatively impact the performance of downstream classification tasks.

Mutual k-Nearest Neighbor (M-KNN) graph is an extension of the traditional KNN approach designed to address the issues of irregular connectivity and hubs [Jebara et al. 2009, Ozaki et al. 2011a]. In an M-KNN graph, an edge is established between two nodes, v_i and v_j, only if v_i is among the k nearest neighbors of v_j and v_j is simultaneously among the k nearest neighbors of v_i. This stricter criterion for connectivity promotes a more symmetric and balanced graph structure compared to the standard KNN graph. By requiring mutual proximity, M-KNN aims to mitigate the hub problem, leading to potentially more robust graphs for label propagation in GSSL.

Finally, Sequential k-Nearest Neighbor (S-KNN) graph represents another variant developed to overcome the limitations of classic KNN graphs, specifically focusing on the problems of irregularity and the presence of hubs [Vega-Oliveros et al. 2014]. This

method incrementally builds the graph by prioritizing the connection of less-connected nodes first, thereby preventing the emergence of highly connected hubs and promoting a more regular degree distribution across the network. By aiming for a more balanced structure, S-KNN contributes to a more effective and stable propagation of labels in semi-supervised learning tasks, even if it does not necessarily guarantee a perfectly regular graph in terms of every node having the exact same degree.

In summary, while the KNN graph provides a straightforward and computationally efficient starting point for graph construction, its extensions, such as M-KNN and S-KNN, are crucial for addressing inherent limitations like irregular connectivity and the presence of hubs. These advanced neighborhood-based methods strive to produce more robust and balanced graph structures, which can optimize the label propagation process and model performance, especially in fairness-sensitive scenarios.

3.3. Graph-Based Semi-Supervised Learning

Graph-Based Semi-Supervised Learning (GSSL) is a prominent paradigm in machine learning where both labeled and unlabeled data are represented as nodes in a graph, and learning is guided by the assumption that neighboring nodes are likely to share similar labels [Zhu et al. 2003]. The central premise is the *smoothness assumption*, which states that if two points are close in the intrinsic geometry of the data manifold, their corresponding labels should also be similar.

Label Propagation (LP) is a classic algorithm for graph-based SSL, introduced by [Zhu et al. 2002]. It iteratively propagates labels from labeled nodes to their neighbors based on edge weights, under the constraint that labeled nodes retain their original labels. The Label Spreading (LS) algorithm, introduced by [Zhou et al. 2003] unlike Label Propagation that strictly adhere to initial labels, LS incorporates a regularization mechanism that allows for a degree of flexibility in the initially labeled nodes, enhancing its robustness against potential label noise. At its core, LS iteratively propagates label information across the graph, driven by a normalized similarity derived from the graph's adjacency and degree matrices.

Let $Y \in R^{n \times c}$ be the label matrix, where n is the number of vertices, c is the number of classes, and each row is a one-hot encoding for labeled nodes or initialized as zero for unlabeled nodes. Let W be the symmetric affinity matrix of the graph, and D be the diagonal degree matrix with $D_{ii} = \sum_j W_{ij}$. The propagation rule is defined as:

$$F^{(t+1)} = \alpha S F^{(t)} + (1 - \alpha) Y,$$

where $S = D^{-1/2} W D^{-1/2}$ is the normalized similarity matrix, F is the label distribution, and $\alpha \in (0, 1)$ is a clamping factor. The algorithm converges to a fixed point where labels are smoothly distributed across the graph.

Label Spreading is computationally efficient and often yields strong performance with minimal labeled data. However, its effectiveness heavily depends on the quality of the graph construction. In particular, if the graph structure fails to capture the true underlying data manifold or introduces spurious connections (e.g., due to outliers or imbalanced data [de Oliveira and Berton 2023]), the propagated labels may become biased or inaccurate.

3.4. Graph Fairness

Fairness in graph learning is complicated by the interconnected nature of the data. Nodes may represent individuals with sensitive attributes (e.g., gender, race, age), and edges encode relationships that may not be independent of those attributes. Biases can emerge due to:

- **Homophily Bias:** Individuals with similar sensitive attributes are more likely to be connected. This can cause over-smoothing or label leakage in GNNs and propagate group-specific biases.
- **Structural Bias:** The graph topology may reflect historical or societal inequalities (e.g., underrepresentation of marginalized groups), leading to disparate performance across groups.
- **Attribute-Topology Correlation:** Sensitive attributes may correlate with node degrees or community structures, allowing models to infer protected information even if it is not explicitly used.

Graph fairness definitions often draw from standard machine learning fairness notions but adapt them to the relational setting. Common criteria include:

- **Group Fairness:** Ensuring equal treatment or performance across demographic groups (e.g., equal opportunity or statistical parity).
- **Individual Fairness:** Similar nodes (based on features or graph structure) should receive similar predictions.
- **Counterfactual Fairness:** Predictions should remain unchanged under plausible counterfactual changes to sensitive attributes.

Despite recent advances, ensuring fairness in graph learning remains an open challenge. Key obstacles include the lack of standardized benchmarks, trade-offs between fairness and accuracy, and the dynamic nature of real-world graphs. Moreover, ensuring fairness under graph perturbations, adversarial attacks, and evolving node/edge features is still an active area of research.

4. Methodology

We considered 3 popular benchmark datasets for fairness, with more information is provided in Table 1, including the percentage of favorable labels, the definition of protected groups, their representativeness, and the disparate impact in each dataset.

In the Adult Census Income dataset [Becker and Kohavi 1996] the task is to predict income with sex and race as a protected attribute, in this case we consider only a sample of 10000 examples, for computational reasons; Compas Recidivism [Angwin et al. 2016] the task is to predict crime recidivism and has sex and race as protected attributes; and German Credit [Hofmann 1994] the task is to predict credit risk, sex and age are the protected attributes. For simplicity, these datasets will be referred to as "Adult", "Compas", and "German" in the following sections.

Figure 1 details the methodology used in this comparative study. We repeated the splitting, pseudo-labeling, and classification process 30 times to calculate the mean and standard deviation of the results, thus ensuring greater robustness of the results obtained. We use the datasets provided and processed by the AIF360 library [Bellamy et al. 2018],

Table 1. Statistics, definitions and description of the protected attributes of the datasets considered in the experimentation

Dataset	Rows	Label	% Favorable Label	Privileged Group	Unprivileged Group	Disparate Impact
Adult Census Income	10000 (Sample)	Income $>$ 50K or Income $\leq$ 50K	25%	Sex = Male (68%)	Sex = Female (32%)	0.36
				Race = White (86%)	Race $\neq$ White (14%) [Asian-Pac-Islander, Amer-Indian-Eskimo, Other, Black]	0.60
Compas Recidivism	6172	Did Not Reoffend or Reoffended	54%	Sex = Female (19%)	Sex = Male (81%)	0.80
				Race = Caucasian (34%)	Race $\neq$ Caucasian (66%) [African-American, Asian, Hispanic, Native American)	0.84
German Credit	1000	Good Credit or Bad Credit	70%	Sex = Male (69%)	Sex = Female (31%)	0.90
				Age $>$ 25 and $<$ 60 (76%)	Age $\leq$ 25 or $\geq$ 60 (24%)	0.84

where all categorical variables were one-hot encoded and rows with missing values are removed. For the numerical features we apply a normalization to mean 0 and variance 1. In detail:

- For each dataset and each protected attribute the data was divided into training (80%) and testing (20%).
- The training data is divided into a labeled (10%) and unlabeled (90%) part.
- We constructed the graph using the KNN, M-KNN and S-KNN methods. The Euclidean distance was used and the number of neighbors was set as $k = \sqrt{n}$, where n is the total of examples (labeled and unlabeled), following a common rule of thumb [Nadkarni 2016].
  - To avoid a low number of neighbors and disconnected components in M-KNN, we connect nodes with less than $k/2$ with their k closest neighbors until we have at least $k/2$ neighbors. We also connect disconnected components using the maximum spanning tree method [Ozaki et al. 2011b].
- The training data is then used in the Label Spreading algorithm, with default parameters other than k, where pseudo-labels are generated for the unlabeled data.
- The 10% labeled data and 90% pseudo-labels are used to train a LightGBM classifier [Shi et al. 2025] with $class\_weight = `balanced'$.
- We used the classifier on the initially defined test set and then the metrics are calculated:
  - For classification we calculated AUC-ROC metric.
  - For Fairness we consider Disparate Impact and Average Absolute Odds Difference metrics.

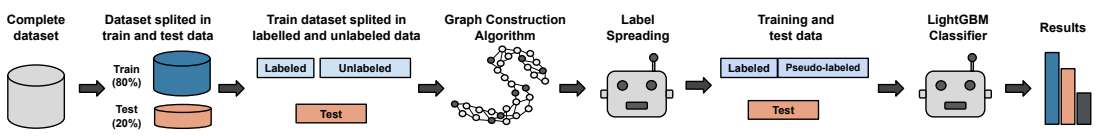

Figure 1. Overview of the experimental pipeline, including data splitting, graph construction, pseudo labeling, and classification stages

5. Results

Figure 2 details all the classification and fairness results obtained using the graph construction methods for all datasets and protected attributes considered in the methodology. Each bar represents the average of 30 repetitions of each experiment. The AUC-ROC is equal across the protected groups of a dataset.

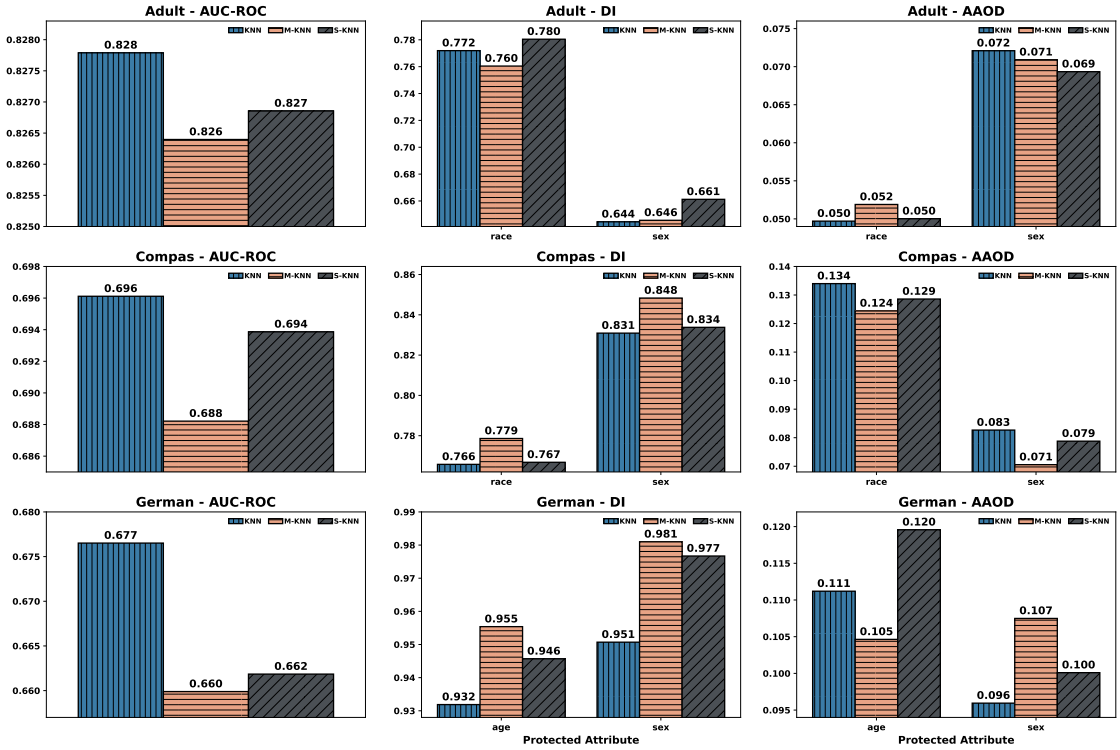

Figure 2. Overall classification and fairness results among graph construction algorithms for all datasets and protected attributes

In terms of classification performance measured by AUC-ROC, the KNN algorithm achieved the best results, reaching 0.828, 0.696, and 0.677 for the Adult, Compas, and German datasets, respectively, whereas the M-KNN algorithm yielded the lowest results.

For DI there is no clear winner in all the results. For the Adult dataset, S-KNN performed better, reaching 0.780, in relation to the race attribute, and 0.661 for the sex attribute. For Compas, M-KNN achieved the best results, with 0.779 for the race attribute and 0.848 for the sex attribute. Finally, for the German dataset, M-KNN achieved the best results in both fairness metrics, with a DI of 0.955 for age and 0.981 for sex.

Regarding AAOD there is also no algorithm that consistently performs better for all cases. For the Adult dataset, with the race attribute, KNN and S-KNN tied for the best results (0.050), closely to M-KNN (0.052), while for the sex attribute, S-KNN achieved the best result (0.069). For the race and sex attribute of Compas, M-KNN achieved the best results, 0.124 for race and 0.071 for sex. Finally, for the German dataset, M-KNN achieved the best result (0.105) for the age attribute, and KNN achieved the best result (0.096) for sex.

Figure 3 presents the results obtained by the graph construction methods in a global way, that is, grouping all the results obtained for the different datasets and protected attributes. For AUC-ROC the results show, globally, exactly what was previously found in the individual results, where KNN was the algorithm that consistently performed best, with 0.733. For DI, the M-KNN and S-KNN algorithms present better average DI results, with 0.828 and 0.827, respectively. In AAOD, the M-KNN algorithm presented the best result (0.088), with KNN and S-KNN having very close results.

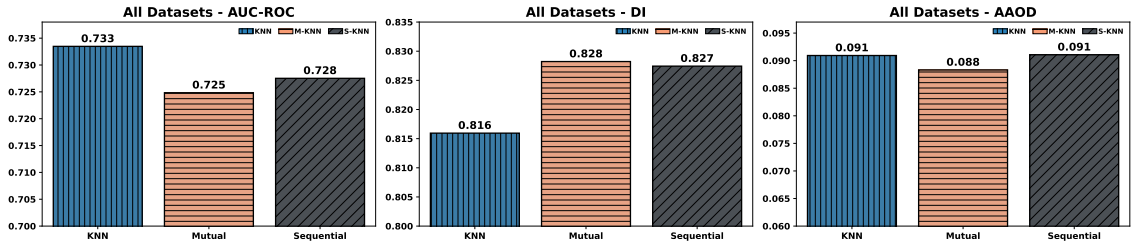

Figure 3. Overall classification and fairness results considering the graph construction algorithms

5.1. Graph Metrics Analysis

We also conducted an analysis of the graph topology in order to understand the results obtained. Thus, Table 2 details the metrics (averages) obtained by the graphs generated for each dataset and protected attribute. The table presents the values of average degree, average degree for the protected groups, the absolute difference between the degree of the groups and the percentage of intragroup edges in the graph.

By correlating the topological metrics with the fairness outcomes, we observe that the KNN algorithm consistently generates the largest disparity in node degree between protected groups. This is due to the nature of the algorithm, which does not seek to build a regular network and may be creating hubs between protected groups. In addition, a negative relationship is noted between this difference and the DI metric. Except for the race attribute of the Adult dataset, in all other datasets and attributes the KNN algorithm exhibits the greatest difference in degree between groups and the worst DI.

This pattern appears to be more related to the difference in degree between groups than to the network average degree itself: while KNN has the highest average degree, the M-KNN graph shows the second largest degree difference between groups (greater than that of S-KNN), even though its average degree is the lowest. These findings strongly suggest that the disparity in node degree influences the label propagation process, allowing the group with higher connectivity to exert a stronger influence.

5.2. Discussion

The results demonstrate that graph construction strategies significantly impact both classification performance and fairness in semi-supervised learning with tabular data. Overall, the KNN algorithm achieved the highest AUC-ROC across datasets, while M-KNN and S-KNN performed better in fairness metrics such as Disparate Impact (DI) and Average Absolute Odds Difference (AAOD). Among the alternatives, S-KNN offered the most balanced trade-off between performance and fairness.

Table 2. Graph topology metrics with the average degree of protected groups and proportion of intragroup edges for all datasets and protected attributes

Dataset	Protected Attribute	Graph Algorithm	Average Degree	Average Degree (Unprivileged Group)	Average Degree (Privileged Group)	Difference Average Degree Absolute	Intragroup Edges Ratio
Adult	Race	KNN	116,33	116,83	116,24	0,592	0,749
		M-KNN	68,97	70,49	68,70	**1,786**	**0,746**
		S-KNN	93,98	94,01	93,97	0,042	0,750
	Sex	KNN	116,33	115,67	116,65	**0,979**	**0,596**
		M-KNN	68,97	68,77	69,07	0,298	**0,596**
		S-KNN	93,98	93,72	94,10	0,383	0,597
Compas	Race	KNN	84,81	85,26	83,94	**1,317**	0,584
		M-KNN	57,24	57,05	57,61	0,564	**0,579**
		S-KNN	72,24	72,33	72,05	0,279	0,582
	Sex	KNN	84,81	85,10	83,58	**1,517**	0,703
		M-KNN	57,24	57,02	58,20	1,181	**0,695**
		S-KNN	72,24	72,27	72,11	0,155	0,699
German	Age	KNN	42,85	44,16	42,43	**1,723**	**0,653**
		M-KNN	19,53	19,97	19,39	0,578	0,660
		S-KNN	31,47	31,77	31,37	0,393	0,661
	Sex	KNN	42,85	43,24	42,68	**0,560**	**0,598**
		M-KNN	19,53	19,76	19,43	0,337	0,604
		S-KNN	31,47	31,43	31,48	0,052	0,604

A deeper investigation of the generated graph structures helps explain these outcomes. The KNN method, while effective at optimizing classification performance, tends to create irregular topologies characterized by highly connected hubs and uneven node degree distribution. This can lead to disproportionate influence during label propagation, favoring nodes from groups that are more central in the graph. This phenomenon was confirmed by the graph analysis: KNN consistently exhibited the highest average difference in node degree between protected and unprotected groups.

This structural imbalance appears directly linked to worse fairness outcomes. In most cases, the higher the degree gap between groups, the lower the DI score. This suggests that nodes in the dominant group exert greater influence during label diffusion, leading to label distributions that disadvantage the other group. These findings align with the literature on structural bias in graphs, where unequal connectivity can systematically reduce the influence of less connected groups.

In contrast, M-KNN and S-KNN produced more balanced graph topologies, with significantly smaller degree differences between protected groups and more uniform node connectivity overall. The more structured and symmetric nature of M-KNN and S-KNN graphs may implicitly lead to a propagation dynamic that is less dominated by local hubs, thus reducing bias leakage from overrepresented regions of the graph. Future work could further explore clustering and centrality metrics as additional fairness indicators.

6. Conclusions

This study highlights the critical role of graph construction strategies in shaping both the predictive performance and fairness of graph-based semi-supervised learning on tabular data. While KNN graphs provided higher classification accuracy, they also exhibited greater topological imbalance, reflected in uneven degree distributions and lower fairness metrics. In contrast, M-KNN and S-KNN offered more balanced connectivity across

protected groups, resulting in fairer outcomes, particularly in terms of Disparate Impact and AAOD, even at the cost of slightly lower overall AUC-ROC.

Our analysis shows that fairness disparities can be strongly linked to topological properties of the graph, such as degree variance between protected groups and intragroup edge density. These findings emphasize the importance of topology-aware fairness analysis in GSSL pipelines, and suggest that graph construction is not a neutral step, but one that can either mitigate or amplify bias.

In line with these observations, the results indicate that the choice of graph construction method directly affects the trade-off between accuracy and fairness. While KNN tends to achieve higher AUC-ROC, its irregular topology may contribute to disparities across protected groups. In contrast, methods like M-KNN and S-KNN, which promote more balanced connectivity, may help mitigate such effects. In fairness-sensitive contexts, it may be beneficial to test alternative graph constructions and monitor topological properties, such as degree distribution and intra/inter-group connectivity, to better understand their impact on subgroup outcomes. Evaluating fairness metrics alongside performance can also provide a more comprehensive view of model behavior.

As future work, we plan to explore adaptive and fairness-aware graph construction methods, such as balancing the number of intra/inter-group edges, extend our evaluation to other learning paradigms such as GNNs, include more granular subgroup analyses, and investigate the impact of other topological factors. We also envision employing hyperparameter optimization, including the k of the graph construction method, the parameters of Label Spreading, and the hyperparameters of a subsequent LightGBM model, to improve both performance and fairness.

Acknowledgements

The authors would like to thank the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) and the Conselho Nacional de Desenvolvimento Científico (CNPq).

References

Angwin, J., Larson, J., Mattu, S., and Kirchner, L. (2016). How we analyzed the compas recidivism algorithm.

Becker, B. and Kohavi, R. (1996). Adult. UCI Machine Learning Repository. DOI: https://doi.org/10.24432/C5XW20.

Bellamy, R. K. E., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilovic, A., Nagar, S., Ramamurthy, K. N., Richards, J., Saha, D., Sattigeri, P., Singh, M., Varshney, K. R., and Zhang, Y. (2018). AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias.

Berton, L., de Andrade Lopes, A., and Vega-Oliveros, D. A. (2018). A comparison of graph construction methods for semi-supervised learning. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Bose, A. and Hamilton, W. (2019). Compositional fairness constraints for graph embeddings. In *International conference on machine learning*, pages 715–724. PMLR.

Chen, A., Rossi, R. A., Park, N., Trivedi, P., Wang, Y., Yu, T., Kim, S., Dernoncourt, F., and Ahmed, N. K. (2024). Fairness-aware graph neural networks: A survey. *ACM Transactions on Knowledge Discovery from Data*, 18(6):1–23.

de Oliveira, W. D. G. and Berton, L. (2023). A systematic review for class-imbalance in semi-supervised learning. *Artificial Intelligence Review*, 56(Suppl 2):2349–2382.

de Oliveira, W. D. G. and Berton, L. (2024). A comparative study of pre-processing algorithms for fair classification in few labeled data context. *AI and Ethics*, pages 1–21.

de Oliveira, W. D. G., Penatti, O. A., and Berton, L. (2020). A comparison of graph-based semi-supervised learning for data augmentation. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 264–271. IEEE.

Dong, Y., Ma, J., Wang, S., Chen, C., and Li, J. (2023). Fairness in graph mining: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 35(10):10583–10602.

Dong, Y., Wang, S., Wang, Y., Derr, T., and Li, J. (2022). On structural explanation of bias in graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 316–326.

Hofmann, H. (1994). Statlog (German Credit Data). UCI Machine Learning Repository. DOI: https://doi.org/10.24432/C5NC77.

Jebara, T., Wang, J., and Chang, S.-F. (2009). Graph construction and b-matching for semi-supervised learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 441–448.

Kose, O. D. and Shen, Y. (2022). Fair node representation learning via adaptive data augmentation. *arXiv preprint arXiv:2201.08549*.

Liao, J. C. and Li, C.-T. (2023). Tabgsl: Graph structure learning for tabular data prediction. *arXiv preprint arXiv:2305.15843*.

Nadkarni, P. (2016). Chapter 10 - core technologies: Data mining and "big data". In Nadkarni, P., editor, *Clinical Research Computing*, pages 187–204. Academic Press.

Ozaki, K., Shimbo, M., Komachi, M., and Matsumoto, Y. (2011a). Using the mutual k-nearest neighbor graphs for semi-supervised classification on natural language data. In *Proceedings of the fifteenth conference on computational natural language learning*, pages 154–162.

Ozaki, K., Shimbo, M., Komachi, M., and Matsumoto, Y. (2011b). Using the mutual k-nearest neighbor graphs for semi-supervised classification on natural language data. In Goldwater, S. and Manning, C., editors, *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 154–162, Portland, Oregon, USA. Association for Computational Linguistics.

Shi, Y., Ke, G., Soukhavong, D., Lamb, J., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.-Y., Titov, N., and Cortes, D. (2025). *lightgbm: Light Gradient Boosting Machine*. R package version 4.6.0.99.

Song, Z., Yang, X., Xu, Z., and King, I. (2022). Graph-based semi-supervised learning: A comprehensive review. *IEEE Transactions on Neural Networks and Learning Systems*, 34(11):8174–8194.

Spinelli, I., Scardapane, S., Hussain, A., and Uncini, A. (2021). Fairdrop: Biased edge dropout for enhancing fairness in graph representation learning. *IEEE Transactions on Artificial Intelligence*, 3(3):344–354.

Vega-Oliveros, D. A., Berton, L., Eberle, A. M., de Andrade Lopes, A., and Zhao, L. (2014). Regular graph construction for semi-supervised learning. *Journal of physics: Conference series*, 490(1):012022.

Wang, Y., Zhao, Y., Dong, Y., Chen, H., Li, J., and Derr, T. (2022). Improving fairness in graph neural networks via mitigating sensitive attribute leakage. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 1938–1948.

Wang, Z., Qiu, M., Chen, M., Salem, M. B., Yao, X., and Zhang, W. (2024). Toward fair graph neural networks via real counterfactual samples. *Knowledge and Information Systems*, 66(11):6617–6641.

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*.

Zhou, D., Bousquet, O., Lal, T., Weston, J., and Schölkopf, B. (2003). Learning with local and global consistency. *Advances in neural information processing systems*, 16.

Zhu, X., Ghahramani, Z., and Lafferty, J. (2002). Learning from labeled and unlabeled data with label propagation. In *CMU CALD Tech Report*.

Zhu, X., Ghahramani, Z., and Lafferty, J. D. (2003). Semi-supervised learning using gaussian fields and harmonic functions. *ICML*, 3(2):912–919.