

Unveiling the Impact of Local Homophily on GNN Fairness: In-Depth Analysis and New Benchmarks

Donald Loveland*

Danai Koutra*

Abstract

Graph Neural Networks (GNNs) struggle to generalize when graphs exhibit both homophily (same-class connections) and heterophily (different-class connections). Specifically, GNNs tend to underperform for nodes with local homophily levels that differ significantly from the global homophily level. This issue poses a risk in user-centric applications where *underrepresented* homophily levels are present. Concurrently, fairness within GNNs has received substantial attention due to the potential amplification of biases via message passing. However, the connection between local homophily and fairness in GNNs remains underexplored. In this work, we move beyond global homophily and explore how local homophily levels can lead to unfair predictions. We begin by formalizing the challenge of fair predictions for underrepresented homophily levels as an out-of-distribution (OOD) problem. We then conduct a theoretical analysis that demonstrates how local homophily levels can alter predictions for differing sensitive attributes. We additionally introduce three new GNN fairness benchmarks, as well as a novel semi-synthetic graph generator, to empirically study the OOD problem. Across extensive analysis we find that two factors can promote unfairness: (a) OOD distance, and (b) heterophilous nodes situated in homophilous graphs. In cases where these two conditions are met, fairness drops by up to 24% on real world data, and 30% in semi-synthetic data. Collectively, our theoretical insights, empirical analysis, and algorithmic contributions unveil a previously overlooked source of unfairness rooted in the graph's homophily information.

1 Introduction

Graph Neural Networks (GNNs) have achieved success in graph machine learning by enabling representation learning over graph-structured data [32]. However, it has been shown that GNNs do not always outperform non-graph baselines [33, 37]. To understand the conditions under which GNNs succeed, recent studies have investigated trends between graph structure and performance [19, 22, 30, 35]. In particular, the presence of homophily and heterophily – where nodes tend to connect

to similar and different classes, respectively – has proved challenging to simultaneously model. Thus, extensive work has explored the conditions under which GNNs perform well for both homophilous and heterophilous graphs, often focusing on degree and distinguishability [22, 36].

Originally used to describe the overall graph connectivity patterns, homophily has been expanded to characterize the local neighborhoods around nodes [20, 24]. This local perspective has helped refine the notion that GNNs perform poorly under heterophily, instead highlighting that a major difficulty lies in predicting for underrepresented local homophily levels, regardless of the global level [9, 21, 24]. We refer to the challenge of applying trained GNNs to settings with underrepresented local homophily levels as the out-of-distribution (OOD) local homophily problem. While existing research has primarily explored this issue from a performance perspective, we propose that the correlation, either causal or spurious, between sensitive attributes and class labels can allow the OOD problem to potentially compromise fairness in GNNs [10, 11]. From a fairness standpoint, most studies have investigated how GNNs might exploit sensitive user attributes [2, 3, 31], with *group fairness* measuring treatment differences between sensitive groups [7]. However, these studies often assess disparity on a global scale, overlooking disparities induced by local graph structures [4, 6, 17, 27, 29]. Recognizing that local patterns in graphs may be influenced by harmful processes, such as forced segregation in social networks [25], we emphasize the need to formally link local homophily and fairness to ensure fair predictions. Without local characterization, GNNs risk misusing underrepresented sensitive attributes, constituting a novel source of unfairness rooted in the local structures of the graph.

Moving beyond a global interpretation of group fairness and homophily, we investigate *local homophily-induced unfairness*, arising from a GNN's ability to exploit a node's sensitive attribute. Focusing on node classification, we motivate the problem through theoretical analysis and show that predictions for nodes with differing sensitive attributes can diverge in OOD settings. To

*University of Michigan, Ann Arbor {dlovelan, dkoutra}@umich.edu

generalize and validate our theoretical findings, we conduct empirical analyses using an OOD training framework, controlling the discrepancy in homophily level between the training and test sets. For datasets, we identify inadequacies in current benchmarks for studying the OOD problem due to limited homophily variability, insufficient preprocessing information, and weakly defined sensitive attributes. Thus, we introduce three new GNN fairness benchmarks with diverse local homophily levels, easy reproducibility, and meaningful fairness tasks. To further explore various local homophily patterns, we propose a semi-synthetic graph generation strategy based on graph re-wiring for precise control. Our theoretical insights, novel datasets, and algorithmic contributions provide evidence that GNNs induce systematic group unfairness for individuals whose local homophily differs from the majority. These findings highlight a fundamental risk in GNNs where *underrepresentation of an exogenous characteristic, the graph structure, can cause disparities related to an endogenous characteristic, the sensitive attribute*, presenting a risk beyond typical fair modeling concerns like sensitive attribute imbalance. Our contributions are outlined below:

- **Theory for Local Homophily-Induced Unfairness:** We establish a relationship between fairness and local homophily, showing that node treatment varies with the severity of the OOD problem.
- **New GNN Fairness Benchmarks:** We introduce three benchmarks that feature natural sensitive attributes closely tied to the learning problem and a wide range of local homophily levels.
- **Semi-synthetic Data Generation:** We propose a re-wiring algorithm to adjust a graph's local homophily distribution using optimal transport between local homophily subgroups.
- **Extensive Empirical Analysis of OOD Problem:** Using our real and semi-synthetic data, we show up to a 24% increase in unfairness as a byproduct of local homophily.

2 Preliminaries

2.1 Graphs Let $G = (V, E, \mathbf{X}, \mathbf{Y})$ denote a simple graph, where V is the node set, E is the edge set, $\mathbf{X} \in \mathbb{R}^{|V| \times f}$ is the feature matrix with f features per node, and $\mathbf{Y} \in \{0, 1\}^{|V| \times c}$ is the label matrix with c classes in one-hot encoded form. A specific node $i \in G$ has a feature vector \mathbf{x}_i , a class label $y_i \in \{1, \dots, c\}$, and a one-hot encoded class vector \mathbf{y}_i . The edge set is often represented using an adjacency matrix, $\mathbf{A} \in \{0, 1\}^{|V| \times |V|}$, where $\mathbf{A}_{i,j} = 1$ indicates an edge between nodes i and j . We use E when discussing the edges set of G and \mathbf{A} when referring to matrix computations

involving the edges. The k -hop neighborhood of node $i \in V$, denoted $N_k(i)$, comprises the nodes that can be reached from i within k hops.

2.2 Node Classification with GNNs We focus on the node classification task, where the objective is to learn a mapping from the features \mathbf{X} to the labels \mathbf{Y} . In a k -layer GNN, this process involves message passing over k -hop neighborhoods. The general steps consist of applying a non-linear transformation to the feature matrix \mathbf{X} , parameterized by a weight matrix \mathbf{W} , and aggregating the features within each node's neighborhood. The update rule for a node i 's representation \mathbf{r}_i is given by:

$$(2.1) \quad \mathbf{r}_i^{l+1} = \text{ENC}(\mathbf{r}_i^l, \text{AGGR}(\{\mathbf{r}_v^l \mid v \in N(i)\})),$$

where l is the current layer of the GNN, AGGR is an aggregation function over the neighboring nodes' representations, and ENC is an encoding function that synthesizes the representations of node i and its aggregated neighbors. After k updates, the final class prediction for node i is obtained by taking the arg max of its final representation \mathbf{r}_i^k . The initial representations $\mathbf{r}_i^0 = \mathbf{x}_i$.

2.3 Homophily and Heterophily We focus on edge homophily and note that while there are various ways to express homophily, many are global and not suitable for localized analysis. First, the global homophily ratio, h , describes the overall homophily level of the graph, where $h = 0$ indicates a fully heterophilous graph, and $h = 1$ indicates a fully homophilous graph [25]. It is defined as:

DEFINITION 2.1. Global Homophily Ratio. *The global homophily ratio h of a graph's edge set E is the fraction of edges in E that connect nodes u and v with the same label y_u and y_v :*

$$(2.2) \quad h = \frac{|\{(u, v) \in E \mid y_u = y_v\}|}{|E|}.$$

To analyze homophily at a more granular level, we use the local homophily ratio of a node t , h_t , defined as:

DEFINITION 2.2. Local Homophily Ratio. *The local homophily ratio h_t of a node t is the fraction of edges in its 1-hop neighborhood $N_1(t)$ that connect t to a neighbor u with the same class:*

$$(2.3) \quad h_t = \frac{|\{(u, t) : u \in N_1(t) \mid y_u = y_t\}|}{|N_1(t)|}.$$

Since local homophily ratios can vary throughout a graph, we define the local homophily distribution as:

DEFINITION 2.3. Local Homophily Distribution. *The local homophily distribution P_G of a graph G is the likelihood of observing a particular local homophily ratio h_t within G .*

We estimate the local homophily distribution by computing the local homophily levels of each node and discretizing these levels into bins to create a probability mass function. We then specify the OOD problem as:

DEFINITION 2.4. Out-of-Distribution Homophily Levels. *Given a training set of nodes with local homophily distribution P_G and a testing set of nodes with local homophily distribution $P_{G'}$, the OOD problem exists when $D(P_G, P_{G'}) > \epsilon$, where $D(\cdot, \cdot)$ is a measure of dissimilarity between the train and test set distributions.*

When ϵ is sufficiently large, the train and test nodes are significantly different, constituting an OOD scenario. In our experiments, we use the Earth Mover's Distance (EMD) as D to characterize the distance between the train and test homophily distributions.

2.4 Fairness We assume each node i has a sensitive attribute s_i , where s_i may be implicitly or explicitly encoded within \mathbf{x}_i . Under the group fairness paradigm, we use statistical parity (SP) as our notion of fairness with respect to pairs of sensitive attributes, defined as:

DEFINITION 2.5. Statistical Parity. *Given two sensitive attributes $s = 0$ and $s = 1$, statistical parity measures the difference in probability of attaining the preferred class c between the two sensitive attributes.*

$$(2.4) \quad |P(y = c|s = 0) - P(y = c|s = 1)|.$$

For a given model, the probabilities in the SP calculation are estimated empirically using a held-out dataset.

3 Theoretical Relationship between Local Homophily and Group Fairness

Despite independent theoretical studies on either fairness in GNNs [2, 4, 31] or the impact of local homophily on GNNs [21, 24], there are no theoretical analyses to bridge the two concepts. Thus, we present a theoretical analysis on how nodes with similar labels, but differing sensitive attributes, experience disparate treatment when their *local* homophily levels are underrepresented relative to the training set. We first outline the assumptions and setup for our theoretical framework, then demonstrate how changes in predictions can arise for nodes with differing sensitive attributes.

3.1 Theoretical Setup To assess the impact of OOD local homophily levels on GNN predictions, we assume a typical theoretical learning setup used in many

previous works [9, 21, 24, 31, 37]. Specifically, we focus on a binary classification task with binary sensitive attributes and consider a GNN where the representations \mathbf{r}_i for a node i , with degree d and local homophily level h_i , are computed using an aggregation function $\text{AGGR}(i) = \sum_{v \in N(i)} \mathbf{r}_v$ and an update function $\text{ENC}(i) = (\mathbf{r}_i + \text{AGGR}(i))\mathbf{W}$. In matrix form, the GNN can be represented as $(\mathbf{A} + \mathbf{I})\mathbf{X}\mathbf{W}$.

A node i 's features are designed to encode both the label y_i and sensitive attribute s_i . Specifically, i 's feature vector is the concatenation of the label and sensitive attribute features, $\mathbf{x}_i = [x_i^{(l)} \parallel x_i^{(s)}]$, where (1) the label is encoded as $x_i^{(l)} = -p_i$ if $y_i = 0$, and $x_i^{(l)} = p_i$ if $y_i = 1$, with $p_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_l, \sigma_l)$; (2) the sensitive attribute is encoded as $x_i^{(s)} = -q_i$ when $s_i = 0$ and $x_i^{(s)} = q_i$ when $s_i = 1$, with $q_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_s, \sigma_s)$. When $\mu_l = 0$, the variable $x_i^{(l)}$ follows the same distribution when $y_i = 0$ or 1. In contrast, a larger μ_l encourages the features to deviate and provide discriminability on the label. Similar logic holds for μ_s and $x_i^{(s)}$.

3.2 Impact of Out-of-Distribution Problem on Fairness

We study a biased learning setup where a GNN is trained on k data points with $y = s = 0$ and $n - k$ data points with $y = s = 1$, creating a correlation between the class and sensitive attribute. After solving for the expected weight matrix $\mathbb{E}[\mathbf{W}]$, we attain predictions for two test nodes u and v , with the same label (WLOG let $y_u = y_v = 0$), but differing sensitive attributes. We also assume that u and v have the same local homophily $h + \alpha$, where α is the relative shift in local homophily level from the global homophily level h . We show that as α increases in magnitude, denoting increase in OOD distance, the gap in expected predictions of the two test nodes also increases.

THEOREM 3.1. *Consider test nodes u and v with local homophily ratios $h + \alpha$, labels $y_u = 0$ and $y_v = 0$, and sensitive attributes $s_u = 0$ and $s_v = 1$. The difference in their expected logit associated with the correct label, $\mathbb{E}[\mathbf{p}_u]_{y_u}$ and $\mathbb{E}[\mathbf{p}_v]_{y_v}$, is:*

$$(3.5) \quad \mathbb{E}[\mathbf{p}_u]_{y_u} - \mathbb{E}[\mathbf{p}_v]_{y_v} = \frac{\mu_s^2 k (1 + d(2h + 2\alpha - 1))}{n(1 + d(2h - 1))((\mu_l^2 + \mu_s^2)}.$$

The proof for this theorem is provided in the Appendix. A large difference, either positive or negative, increases the likelihood of differing predictions for nodes u and v , suggesting disparate treatment of individuals with different sensitive attributes. To further assess the criteria that disparate treatment arises, we must determine when $|\mathbb{E}[\mathbf{p}_u]_{y_u} - \mathbb{E}[\mathbf{p}_v]_{y_v}| > 0$, i.e. there is preferred

treatment towards a sensitive attribute. This inequality holds when $\mathbb{E}[\mathbf{p}_u]_{y_u} - \mathbb{E}[\mathbf{p}_v]_{y_v} \neq 0$. Assuming $\mu_s \neq 0$ and $k > 0$, $\mathbb{E}[\mathbf{p}_u]_{y_u} - \mathbb{E}[\mathbf{p}_v]_{y_v} \neq 0$ when $\alpha \neq 0$. These results indicate that drifting away from the global homophily level can increase the distance between predicted logits. If $c = 0$ in our SP definition, our theorem can also be interpreted as characterizing the increase in SP for different α values. In the following sections, we validate our findings on semi-synthetic and real-world datasets.

4 New Real World Datasets

In this section, we outline current challenges with fairness benchmarks in node classification and present our improved datasets. Our datasets possess diverse homophily patterns while ensuring reproducibility and well-motivated fairness implications. These datasets are also necessary to study fair GNNs for node classification given most fairness datasets are tailored for link prediction [3]. Through these real-world datasets, we expand the scenarios for studying GNNs, thereby minimizing unexpected behavior upon deployment.

4.1 Limitations of Current Datasets While the number of GNNs continues to expand, the datasets to study GNN fairness remain limited [7]. Moreover, existing datasets exhibit fundamental issues, preventing systematic evaluation of how local homophily impacts fairness. In this section, we provide a high-level overview of the identified failure modes of current GNN fairness datasets, with thorough discussion and example datasets of each failure mode provided in the Appendix.

First, we focus on diversity with respect to homophily patterns. Given prominent fairness datasets are built around similarity networks [2, 3, 8], their respective distributions tend to be homophilous, making it difficult to study diverse OOD settings. Second, many fairness datasets are ambiguous given they have multiple variants with missing information [3, 4, 7]. Finally, a significant number of fairness benchmarks are datasets without sensitive attributes or fairness contexts [3]. Instead, structural attributes are arbitrarily chosen as sensitive attributes despite not being representative of disparate treatment [7, 34]. Together, these three challenges hinder the progression of GNN fairness research, particularly as it pertains to the OOD problem.

4.2 Proposed Real-World Datasets The proposed datasets are fully labeled with both class and sensitive attribute information. We focus on datasets with diverse local homophily patterns to study different OOD scenarios. To mitigate remote nodes, we extract the largest connected component from each dataset which satisfies the above criteria. We give key statistics in

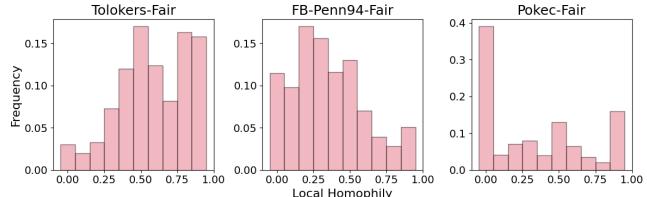


Figure 1: **Local Homophily Distributions for Proposed Datasets.** Each dataset displays a wide array of local and global homophily levels.

Table 1, and the homophily distributions in Figure 1.

Tolokers-Fair: The Tolokers graph, derived from users of the Tolokers crowdsourcing platform, was introduced by Lim et al. [26] as a large-scale benchmark for GNNs. We introduce a fairness-based learning task: predicting whether an individual contributor should be banned from the platform, using dominant language (English vs. non-English) as the sensitive attribute. This task is motivated by potential language biases. The features are retained based on their original implementation. Tolokers-Fair is skewed homophilous ($h = 0.58$).

FB-Penn94-Fair: The FB graphs, based on Facebook social networks, were introduced by Leskovec et al. [16]. Despite their use as benchmarks [18], the FB datasets have not been directly used for fairness. To study fairness, we propose predicting the major of a student (instead of gender in previous studies) and using gender as a sensitive attribute, addressing gender inequity in college majors. As the major has high cardinality, we extract the five most common majors and subsample nodes with these majors and a gender attribute. The largest component is taken as the dataset. The features are retained based on their original implementation. FB-Penn94-Fair is heterophilous ($h = 0.39$).

Pokec-Fair: Originally proposed by [4], Pokec is a fairness dataset with artificially scrubbed subgraphs retaining only 1% of the labels and sensitive attributes. No information is provided to recover the scrubbed information, nor the subgraphs. To perform local analyses, we create Pokec-Fair with the task of predicting occupation while using gender as the sensitive attribute. We do not perform major region sampling as done in [4], given region is a high cardinality feature with many missing values. Similar to FB-Penn94-Fair, we extract the 5-most common jobs, excluding unemployment, leaving classes “service work”, “construction”, “finance”, “healthcare”, and “transport”, and subsample the users that work in these fields and possess a gender attribute. As noted in [18], the pre-processing of features is highly variable, with [4] reporting 59 features, and [18] reporting 65 features. As open source access to the feature processing is not provided, we reduce inconsistencies by removing

Table 1: Proposed Dataset Statistics and Details

Dataset	Nodes	Edges	Task	Sens. Attr.
Tolokers-Fair	11,758	519K	Banned(2)	Language
FB-Penn94-Fair	7,016	59,845	Major(5)	Gender
Pokec-Fair	69,949	130K	Job(5)	Gender

the free text features, and only retain the categorical and numerical features. The Pokec-Fair dataset is heterophilous ($h = 0.38$).

5 Semi-Synthetic Data Generation

In addition to our proposed real-world datasets, we introduce a semi-synthetic graph generator that enables the exploration of diverse local homophily distributions. Semi-synthetic data generation is commonly used to assess GNN behavior under varied homophily conditions [13, 24, 37], but existing methods only control global homophily and often inflate node degrees by adding edges. To address these limitations, we propose a strategy that manages homophily at the node level without requiring a full node-class interaction matrix [23]. At a high level, we achieve this by defining a goal homophily distribution, parameterized by the beta probability distribution, and re-wiring a small set of edges around select nodes. This minimally disrupts the original graph, adjusting only necessary nodes to match the goal distribution. Although we do not explicitly consider feature information during re-wiring, we adhere to the common assumption that features encode the label information, allowing the underlying signal to be extracted during message passing. We describe our method below.

Parameterizing Goal Homophily Distribution:

We parameterize the goal homophily distribution as a beta distribution, where $X \sim Beta(\alpha, \beta)$ with probability density function $C(x)^{\alpha-1}(1-x)^{\beta-1}$, and normalization constant C . Specifying α and β determines the likelihood of sampling a particular local homophily ratio, as dictated by $\mathbb{E}[X] = \frac{\alpha}{\alpha + \beta}$. Thus, when $\alpha > \beta$, the distribution will skew homophilous, while $\alpha < \beta$ will skew heterophilous. The spread of the distribution is additionally controlled by $Var[X] = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$, where larger values create well peaked distributions.

Determining Edges to Re-wire: To achieve the goal homophily distribution for a graph, we start by specifying the target homophily distribution Q (coming from $Beta(\alpha, \beta)$), and computing the empirical homophily histogram P_G for graph G with b bins. We solve an optimal transport problem to align P_G with Q , yielding a transportation matrix $\mathbf{T} \in [0, 1]^{b \times b}$. Each

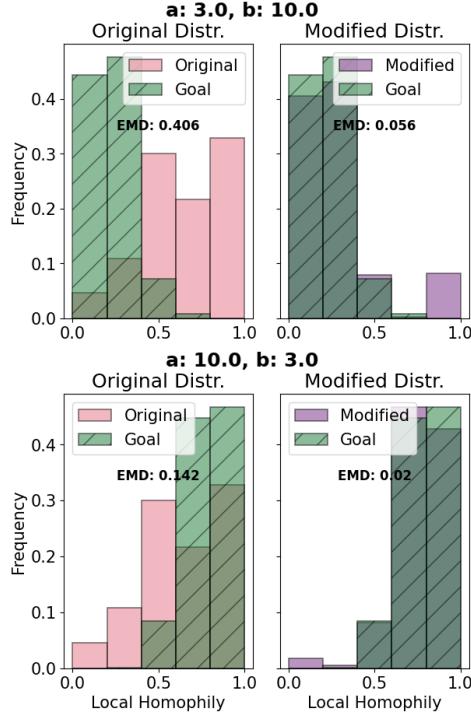


Figure 2: **Comparison of Original, Generated, and Goal Distributions for Tolokers-Fair.** Red and purple denote the original and modified distributions, respectively. The green (hatching) denotes the goal distribution. Overlap between the distributions indicates better re-wiring.

matrix entry $\mathbf{T}_{i,j}$ denotes the proportion of nodes whose local homophily ratio falls within bin i and needs to be modified to fall within bin j . Nodes that already meet the goal homophily level are not modified and retain their original structure. For the other nodes, we precompute the number of edges that need to be modified to reach their goal homophily level. For a node i with local homophily h_i , goal homophily h_g , and degree d_i , the total number of edges to move, e_i , can be bounded as $|(h_g - h_i)d_i| \leq e_i \leq |\frac{(h_g - h_i)d}{1 - h_g}|$ when $h_i < h_g$ and $|(h_g - h_i)d_i| \leq e_i \leq |\frac{(h_i - h_g)d}{h_g}|$ when $h_i > h_g$. The lower and upper bounds denote the number of edges that are moved when purely re-wiring (no degree change) or adding edges (degree change).

Our Minimally Disruptive Re-wiring Model:

Our generation begins by sampling nodes from G based on the matrix \mathbf{T} . These are denoted as source nodes. The number of edges to be modified in each source node's neighborhood is determined using the lower bounds above, with the criteria that removing an edge optimizes both the source and neighboring nodes homophily levels. Then, edges removed from the source

node are re-wired to candidate nodes outside the neighborhood, again optimizing both the source and candidate nodes' local homophily. As an example, if a node is to become more homophilous, a heterophilous edge will be removed from a neighboring node that also needs to become more homophilous and converted into a homophilous edge. If re-wiring alters a neighboring or candidate node's degree, the lower bound is recalculated.

As re-wiring is required to strictly decrease the distance between the current and goal homophily level for any nodes, it may not be possible to attain the goal homophily ratio under such constraint. Thus, we provide a secondary refinement phase where edges are added to further adjust the local homophily ratio of the node, at the expense of modify the degree distribution slightly. Specifically, the upper bounds above are computed for each node, and edges are added between source and candidate nodes to mutually benefit their local homophily level. This process continues until all nodes achieve their targets or no further beneficial modifications are possible. Note that this process can accommodate both binary and multi-class settings.

Evaluation: We assess the re-wiring quality using the EMD between the original/goal and generated/goal distributions. We provide distributions for the Tolokers-Fair dataset in Figure 2. For each case, the modified and goal distributions are in strong agreement, demonstrating that the semi-synthetic data can act as a promising tool to explore the space of local homophily distributions. Through these capabilities, GNN fairness can be comprehensively assessed under a variety of connectivity patterns beyond what is available in benchmarks.

6 Empirical Analysis

This section studies the relationship between OOD local homophily levels and unfairness through empirical analysis. We leverage the proposed real-world and semi-synthetic datasets and focus on two fundamental questions: (RQ1) *How does the global and local homophily levels of a graph impact fairness?*, and (RQ2) *How does GNN design mitigate unfairness in the OOD setting?*

6.1 Train/Val/Test Processing To study how the OOD problem impacts fairness, we create two test sets, one in-distribution and one OOD. Given a dataset's local homophily distribution P_G , we generate these two scenarios by concentrating the mass of P_G and sampling from the new distribution. Specifically, each bin $b \in B$ is raised to the γ power, and then normalized, resulting in P_G^γ . We then compute an inverted distribution \bar{P}_G^γ , where $\bar{P}_{G,b}^\gamma = (P_{G,b}^\gamma)^{-1} / \sum_{b \in B} (P_{G,b}^\gamma)^{-1}$. Parameter $\gamma \in [0, \infty)$ allows us to choose the concentration of mass within $P_{G,b}^\gamma$ and $\bar{P}_{G,b}^\gamma$, changing the OOD severity.

When $\gamma = 0$, the test set is in-distribution, while when $\gamma = 3$ the test set is OOD. The proportion of training and testing nodes in each homophily range bin b is determined by $P_{G,b}^\gamma / (P_{G,b}^\gamma + \bar{P}_{G,b}^\gamma)$ and $\bar{P}_{G,b}^\gamma / (P_{G,b}^\gamma + \bar{P}_{G,b}^\gamma)$, respectively. In both $\gamma = 0$ (standard splitting protocol), and $\gamma = 3$ (our stratified protocol), the train and test set are randomly sampled with an 80/20 split. Then, the validation set is sampled as a subset from the training set with an 80/20 train/val ratio. Figure 3 shows splits for different γ values.

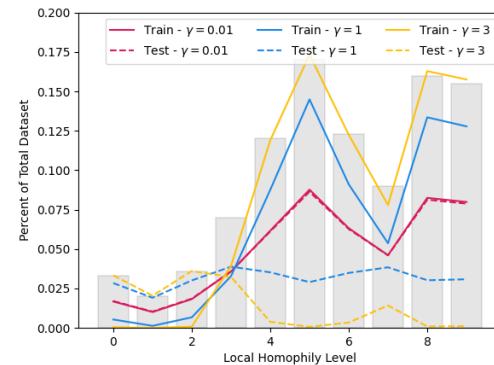


Figure 3: **Example Train/Test Splits with differing γ .** Gray bars denote the P_G for the Tolokers-Fair dataset. As γ increases, the train and test sets become more disjoint, leading to more OOD samples in the test set. When $\gamma \approx 0$, train and test are identical.

6.2 Models and Evaluation To understand how (un)fairness can arise, we evaluate three classes of GNNs: (1) Classic GNNs, (2) GNNs designed to accommodate homophily and heterophily, and (3) GNNs built with fairness considerations. Our goal is to address whether local homophily induced fairness degradation can be addressed by heterophilous learning mechanisms. We use GCN [14] to represent our classic GNN, LINKX [18] as our heterophilous model, and GraphSAGE [12] as a midpoint between the two. Specifically, GCN offers a simple degree weighted aggregation, while GraphSAGE and LINKX both de-couple the encoding of the ego and neighbor embeddings – a key design shown to help learn over heterophily. Additionally, LINKX separates the structure and feature learning to further help in heterophilous settings. For fairness-specific GNNs, we employ Nifty [2] and FairGNN [4], each with GCN and GraphSAGE backbones. An MLP is trained as a graph-agnostic baseline.

Hyperparameters are chosen via cross-validation. Each experiment is repeated three times. Performance and fairness are measured using the micro-F1 score and SP, respectively. As the common formulation of SP relies on binary properties [1, 2], we also consider a multi-class variant which computes the maximum SP for

Table 2: Model Performance Across Real World Datasets. The data includes F1 and statistical parity (SP) for $\gamma = 0$ and $\gamma = 3$, along with changes in F1 and SP. Rank is ordered by increasing unfairness, i.e. ΔSP . The OOD problem produces a 6.3% increase in SP on average, denoting increased unfairness.

Dataset	Model	$\gamma = 0$		$\gamma = 3$		Δ		Rank
		F1	SP	F1	SP	F1	SP	
Tolokers-Fair	GCN	0.06 (0.02)	-0.02 (0.02)	-0.03 (0.02)	0.07 (0.02)	-0.09	0.09	3
	SAGE	0.10 (0.06)	-0.05 (0.01)	-0.11 (0.02)	-0.03 (0.01)	-0.21	0.02	2
	LINKX	0.08 (0.05)	0.02 (0.03)	-0.12 (0.04)	0.01 (0.02)	-0.20	-0.01	1
	FairGCN	0.15 (0.04)	-0.17 (0.01)	-0.03 (0.02)	-0.06 (0.01)	-0.18	0.11	4
	FairSAGE	0.16 (0.04)	-0.19 (0.01)	-0.03 (0.02)	-0.07 (0.01)	-0.19	0.12	5
	NiftyGCN	-0.63 (0.04)	-0.20 (0.01)	-0.59 (0.02)	-0.07 (0.01)	0.04	0.13	6
Pokec-Fair	NiftySAGE	-0.62 (0.04)	-0.20 (0.01)	-0.56 (0.02)	-0.06 (0.01)	0.06	0.14	7
	GCN	0.02 (0.01)	-0.34 (0.02)	0.02 (0.02)	-0.22 (0.03)	0.0	0.12	4
	SAGE	0.04 (0.01)	-0.15 (0.03)	0.02 (0.03)	0.08 (0.03)	-0.02	0.23	6
	LINKX	-0.02 (0.02)	-0.41 (0.0)	-0.08 (0.03)	-0.35 (0.03)	-0.06	0.06	1
	FairGCN	-0.01 (0.01)	-0.71 (0.02)	0.08 (0.02)	-0.54 (0.04)	0.09	0.17	5
	FairSAGE	-0.01 (0.02)	-0.70 (0.02)	0.06 (0.02)	-0.46 (0.06)	0.07	0.24	3
FB-Penn94-Fair	NiftyGCN	-0.02 (0.01)	-0.76 (0.01)	0.06 (0.02)	-0.70 (0.03)	0.08	0.06	1
	NiftySAGE	-0.02 (0.01)	-0.73 (0.01)	0.05 (0.01)	-0.66 (0.02)	0.07	0.07	3
	GCN	0.02 (0.01)	0.06 (0.01)	0.0 (0.01)	0.01 (0.01)	-0.02	-0.05	3
	SAGE	0.07 (0.01)	0.05 (0.01)	0.07 (0.01)	0.02 (0.01)	0.0	-0.03	5
	LINKX	0.08 (0.01)	0.10 (0.03)	0.07 (0.02)	0.04 (0.01)	-0.01	-0.06	1
	FairGCN	0.01 (0.01)	0.03 (0.0)	0.02 (0.02)	0.01 (0.02)	0.01	-0.02	6
FB-Penn94-Fair	FairSAGE	0.0 (0.01)	0.07 (0.01)	-0.05 (0.02)	0.01 (0.01)	-0.05	-0.06	1
	NiftyGCN	0.0 (0.04)	0.04 (0.02)	0.01 (0.05)	0.0 (0.01)	0.01	-0.04	4
	NiftySAGE	0.03 (0.01)	0.03 (0.01)	-0.01 (0.0)	0.05 (0.01)	-0.04	0.02	7

all pairs of classes [5, 28]. For the real-world results in Table 2, we provide F1 and SP for both γ values, while for the semi-synthetic results in Table 3 we provide the change in F1 and SP between $\gamma = 3$ and $\gamma = 0$. For all results, the MLP F1 and SP scores are subtracted from GNN scores to control for spurious feature relationships, helping to isolate the impact of homophily.

6.3 Results In this section, we present the results of our empirical analysis, revealing insights into how homophily can induce unfairness and the effectiveness of different GNN designs in mitigating this. We begin with a high-level analysis to establish that the OOD problem exists, and then move to granular dataset- and model-level analysis to explain drivers for unfairness.

(RQ1) The Impact of Homophily on Fairness: Table 2 summarizes the performance and fairness across different datasets. On average, we observe a 6.3% increase in SP, indicating increased unfairness. Notably, the FB-Penn94-Fair dataset exhibits a 3.4% average decrease in SP, suggesting it is generally unaffected by the OOD problem. In contrast, the Tolokers-Fair and Pokec-Fair datasets show increases in SP of 8.6% and 13.6%, respectively. When we quantify the average EMD between the train and test sets for each dataset, we find the following values: 0.037 for FB-Penn94-Fair, 0.082 for Tolokers-Fair, and 0.179 for Pokec-Fair, *elucidating a near-linear relationship between OOD severity and unfairness*. This additionally corresponds with our theoretical analysis which relates the homophily shift level α to unfairness. We further characterize the im-

pact of global homophily through our semi-synthetic setting where EMD values remain roughly constant. This allows us to isolate changes in SP as a factor of homophily level, rather than EMD. We find that globally homophilous datasets tend to induce unfairness on heterophilous nodes, increasing average SP by 8.8%. Conversely, globally heterophilous datasets tend to improve fairness for homophilous nodes, reducing average SP by 5.4%. This pattern agrees and further explains FB-Penn94-Fair’s fairness levels, given it is global heterophilous with low OOD severity.

Overall, these results suggest a nuanced interplay between global and local homophily levels. Specifically, globally heterophilous datasets are beneficial for improving fairness in OOD settings until the OOD data points become significantly far from the training set. This finding implies that both local and global homophily levels must be considered together to fully understand fairness patterns. Previous works have addressed the challenge of fairness when applying GNNs to globally homophilous graphs but have overlooked (a) the local homophily level and (b) the shift of local homophily level relative to the global level [29].

(RQ2) The Impact of GNN Design on Fairness: We now analyze different GNN designs as outlined in the experimental setup. As opposed to RQ1, we look at two properties, the initial SP when $\gamma = 0$, as well as ΔSP . To ease comparisons, in Table 2 each model is ranked across datasets, with lower rank indicating less of a change in SP. Interestingly, LINKX is ranked among the best across each dataset, often producing smaller drops in

Table 3: Performance Across Semi-Synthetic Datasets with Parameters $\alpha = 3.0, \beta = 10.0$ and $\alpha = 10.0, \beta = 3.0$. The data includes changes in F1 and SP between $\gamma = 0$ and $\gamma = 3$. Homophilous datasets produce an 8.8% increase in SP, showing heterophilous nodes experience amplified unfairness in globally homophilous graphs.

Model	Tolokers-Fair Semi-Synth				Pokec-Fair Semi-Synth				FB-Penn94-Fair Semi-Synth			
	$\alpha = 3.0, \beta = 10.0$		$\alpha = 10.0, \beta = 3.0$		$\alpha = 3.0, \beta = 10.0$		$\alpha = 10.0, \beta = 3.0$		$\alpha = 3.0, \beta = 10.0$		$\alpha = 10.0, \beta = 3.0$	
	$\Delta F1$	ΔSP										
GCN	-0.13	0.09	-0.21	0.08	0.02	-0.05	-0.04	0.08	-0.01	0.12	-0.09	0.06
SAGE	-0.16	-0.09	-0.28	0.11	-0.01	-0.07	-0.12	0.10	-0.01	0.14	-0.35	0.27
LINKX	-0.12	-0.09	-0.20	0.07	0.0	-0.03	-0.04	0.22	0.03	-0.08	-0.21	0.22
FairGCN	-0.01	-0.14	-0.18	0.11	0.05	-0.15	-0.03	-0.06	0.01	-0.01	-0.03	0.05
FairSAGE	0.02	-0.10	-0.18	0.10	0.05	-0.12	-0.02	-0.04	-0.02	-0.11	0.02	-0.10
NiftyGCN	-0.05	-0.14	-0.08	0.11	0.05	-0.16	0.0	0.05	0.05	0.06	-0.04	0.15
NiftySAGE	-0.32	-0.15	-0.06	0.11	0.05	-0.16	0.0	0.30	-0.03	0.10	-0.05	-0.14

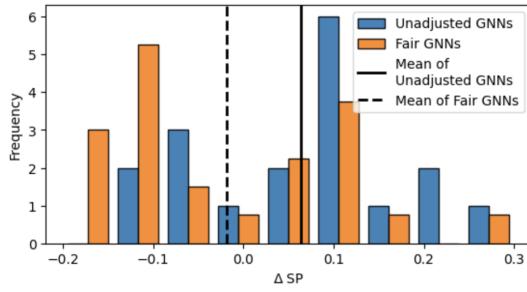


Figure 4: **Distribution of ΔSP for Real and Semi-Synthetic Datasets.** Results indicate that unadjusted GNNs (GCN, SAGE, LINKX) produce more frequent SP increases compared to fairness-based GNNs (Nifty, FairGNN).

fairness as opposed to the fair GNN architectures. Note that across each dataset the initial SP when $\gamma = 0$ is significantly lower for these fair models, highlighting *fair GNNs are beneficial for in-distribution data points, but are more susceptible to OOD fairness degradation*. This is further seen in the semi-synthetic data as when the OOD severity is controlled through similar homophily distributions, the GNNs without fairness adjustment do significantly worse on average. The comparison between unadjusted and fairness-adjusted GNNs is provided in 4 across both the real and semi-synthetic datasets, demonstrating an 8.2% difference in average ΔSP .

7 Related Work

While many related works are introduced in Sections 1 and 4, here we provide additional fairness methods not specifically related to homophily. Many works aim to ensure that the representations learned by GNNs do not heavily rely on sensitive attributes [3, 7]. There are two common paradigms: (1) pre-processing, where the original graph is altered to promote fairness, and (2) in-processing, where the computation graph is adjusted to facilitate fair representations. FairWalk [27] proposes a random walk technique that ensures nodes from different sensitive groups are sampled at equal rates. Simi-

larly, FairAdj [17] and FairDrop [29] modify the adjacency matrix to create more balanced neighborhoods, reducing the influence of sensitive attributes. For in-processing, adversarial debiasing has been utilized to mitigate the GNN’s reliance on sensitive attributes [4]. Nifty [2] approaches the problem from a representation perspective by promoting layer-wise re-normalization and neighborhood augmentation. Re-weighting has also inspired methods like FairGAT and FairVGNN, which tend to work well for in-distribution nodes [15, 31].

8 Conclusion

In this work, we addressed the relationship between homophily and fairness, particularly regarding the OOD problem in the local homophily distribution of a graph. We provided a theoretical analysis of how the OOD problem affects predictions for nodes with different sensitive attributes. Then, using our three proposed fairness benchmarks and semi-synthetic graph generator, we conducted an empirical study to identify conditions that lead to unfairness in GNNs. Our findings showed the OOD problem is prevalent and most severe in graphs with high global homophily, making it challenging to integrate heterophilous users into homophilous settings. Although fair GNN models can partially address this issue, significant gaps persist. This work shifts the focus from addressing sensitive attributes to also considering the structural under-representation in GNN fairness. Future research will aim to enhance fairness for all users across the network, rather than just the majority.

Acknowledgements

We thank the reviewers for their constructive feedback. We also thank Jiong Zhu for his contributions during initial discussions. This material is based upon work supported by the National Science Foundation under a Graduate fellowship, IIS 2212143 and CAREER Grant No. IIS 1845491. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or other funding parties.

References

[1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudik, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *ICML*. PMLR, 2018.

[2] Chirag Agarwal, Himabindu Lakkaraju, and Marinka Zitnik. Towards a unified framework for fair and stable graph representation learning. In *UAI*. PMLR, 2021.

[3] April Chen, Ryan A Rossi, Namyong Park, Puja Trivedi, Yu Wang, Tong Yu, Sungchul Kim, Franck Dernoncourt, and Nesreen K Ahmed. Fairness-aware graph neural networks: A survey. *ACM TKDD*, 18(6), 2024.

[4] Enyan Dai and Suhang Wang. Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information. *WSDM*, 2021.

[5] Christophe Denis, Romuald Elie, Mohamed Hebiri, and François Hu. Fairness guarantee in multi-class classification. *arXiv preprint arXiv:2109.13642*, 2021.

[6] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. Edits: Modeling and mitigating data bias for graph neural networks. In *ACM WebConf*, 2022.

[7] Yushun Dong, Jing Ma, Song Wang, Chen Chen, and Jundong Li. Fairness in graph mining: A survey. *IEEE TKDE*, 2023.

[8] Yushun Dong, Binchi Zhang, Hanghang Tong, and Jundong Li. Elegant: Certified defense on the fairness of graph neural networks. *arXiv:2311.02757*, 2023.

[9] Lun Du, Xiaozhou Shi, Qiang Fu, Xiaojun Ma, Hengyu Liu, Shi Han, and Dongmei Zhang. Gbk-gnn: Gated bi-kernel graph neural networks for modeling both homophily and heterophily. In *ACM WebConf*, 2022.

[10] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. Fairness in deep learning: A computational perspective. *IEEE Intell. Syst.*, 36(4), 2021.

[11] Chris Dulhanty and Alexander Wong. Auditing imagenet: Towards a model-driven framework for annotating demographic attributes of large-scale image datasets. *arXiv preprint arXiv:1905.01347*, 2019.

[12] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *NeurIPS*, 2017.

[13] Fariba Karimi, Mathieu Génois, Claudia Wagner, Philipp Singer, and Markus Strohmaier. Homophily influences ranking of minorities in social networks. *Scientific reports*, 8(1), 2018.

[14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[15] O Deniz Kose and Yanning Shen. Fairgat: Fairness-aware graph attention networks. *ACM TKDD*, 2023.

[16] Jure Leskovec and Julian Mcauley. Learning to discover social circles in ego networks. In *NeurIPS*, volume 25. Curran Associates, Inc., 2012.

[17] Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. On dyadic fairness: Exploring and mitigating bias in graph connections. In *ICLR*, 2020.

[18] Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and Ser Nam Lim. Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods. In *NeurIPS*, volume 34. Curran Associates, Inc., 2021.

[19] Zemin Liu, Trung-Kien Nguyen, and Yuan Fang. On generalized degree fairness in graph neural networks. In *AAAI*, 2023.

[20] Donald Loveland, Jiong Zhu, Mark Heimann, Ben Fish, Michael T Schaub, and Danai Koutra. On graph neural network fairness in the presence of heterophilous neighborhoods. *KDD Deep Learning on Graphs Workshop*, 2022.

[21] Donald Loveland, Jiong Zhu, Mark Heimann, Benjamin Fish, Michael T Schaub, and Danai Koutra. On performance discrepancies across local homophily levels in graph neural networks. In *LoG*, volume 231. PMLR, 27–30 Nov 2024.

[22] Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. Is homophily a necessity for graph neural networks? In *ICLR*, 2022.

[23] Seiji Maekawa, Yuya Sasaki, George Fletcher, and Makoto Onizuka. Gencat: Generating attributed graphs with controlled relationships between classes, attributes, and topology. *Information Systems*, 115, 2023.

[24] Haitao Mao, Zhikai Chen, Wei Jin, Haoyu Han, Yao Ma, Tong Zhao, Neil Shah, and Jiliang Tang. Demystifying structural disparity in graph neural networks: Can one size fit all? In *NeurIPS*, volume 36, 2023.

[25] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*.

[26] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenko. A critical look at the evaluation of gnns under heterophily: Are we really making progress? In *ICLR*, 2022.

[27] Tahleen Rahman, Bartłomiej Surma, Michael Backes, and Yang Zhang. Fairwalk: Towards fair graph embedding. In *IJCAI*, 2019.

[28] Julien Rouzot, Julien Ferry, and Marie-José Huguet. Learning optimal fair scoring systems for multi-class classification. In *IEEE ICTAI*, 2022.

[29] Indro Spinelli, Simone Scardapane, Amir Hussain, and Aurelio Uncini. Biased edge dropout for enhancing fairness in graph representation learning. *IEEE TAI*, 2021.

[30] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *ICLR*, 2022.

[31] Yu Wang, Yuying Zhao, Yushun Dong, Huiyuan Chen, Jundong Li, and Tyler Derr. Improving fairness in graph neural networks via mitigating sensitive attribute leakage. In *SIGKDD*. ACM, 2022.

[32] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.*, 32(1), 2021.

- [33] Yujun Yan, Milad Hashemi, Kevin Swersky, Yaoqing Yang, and Danai Koutra. Two sides of the same coin: Heterophily and oversmoothing in graph convolutional neural networks. In *IEEE ICDM*, 2022.
- [34] Tianxiang Zhao, Enyan Dai, Kai Shu, and Suhang Wang. Towards fair classifiers without sensitive attributes: Exploring biases in related features. In *WSDM*, 2022.
- [35] Jiong Zhu, Gaotang Li, Yao-An Yang, Jing Zhu, Xuehao Cui, and Danai Koutra. On the impact of feature heterophily on link prediction with graph neural networks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [36] Jiong Zhu, Yujun Yan, Mark Heimann, Lingxiao Zhao, Leman Akoglu, and Danai Koutra. Heterophily and graph neural networks: Past, present and future. *IEEE Data Eng. Bull.*, 46(2):12–34, 2023.
- [37] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. In *NeurIPS*, volume 33. Curran Associates, Inc., 2020.