

Physics-Guided Fair Graph Sampling for Water Temperature Prediction in River Networks

Erhu He¹, Declan Kutscher¹, Yiqun Xie², Jacob Zwart³, Zhe Jiang⁴, Huaxiu Yao⁵, Xiaowei Jia¹

¹University of Pittsburgh

²University of Maryland

³U.S. Geological Survey

⁴University of Florida

⁵University of North Carolina at Chapel Hill

{erh108, dtk28}@pitt.edu, xie@umd.edu, jzwart@usgs.gov, zhe.jiang@ufl.edu, huaxiu@cs.unc.edu, xiaowei@pitt.edu

Abstract

This work introduces a novel graph neural networks (GNNs)-based method to predict stream water temperature and reduce model bias across locations of different income and education levels. Traditional physics-based models often have limited accuracy because they are necessarily approximations of reality. Recently, there has been an increasing interest of using GNNs in modeling complex water dynamics in stream networks. Despite their promise in improving the accuracy, GNNs can bring additional model bias through the aggregation process, where node features are updated by aggregating neighboring nodes. The bias can be especially pronounced when nodes with similar sensitive attributes are frequently connected. We introduce a new method that leverages physical knowledge to represent the node influence in GNNs, and then utilizes physics-based influence to refine the selection and weights over the neighbors. The objective is to facilitate equitable treatment over different sensitive groups in the graph aggregation, which helps reduce spatial bias over locations, especially for those in underprivileged groups. The results on the Delaware River Basin demonstrate the effectiveness of the proposed method in preserving equitable performance across locations in different sensitive groups.

Introduction

Water quality continues to degrade due to the increasing demand on water-related services (e.g., drinking water supply and energy consumption) and the changing climate (Hoekstra et al. 2012; Carr et al. 2013). Timely monitoring of water temperature, which serves as a master factor of water quality, can help inform water management decisions (e.g., water filtering, water release from reservoirs, enhanced management of nearby farmlands), maintain the desired habitat for aquatic life, and better understand ecological processes leading to the changes in water quality.

To address this problem, scientists from multiple disciplines have built physics-based models to simulate the internal water cycles and their interactions with weather, soil, and other water bodies based on general physical relationships such as energy and mass conservation (Hipsey et al. 2014; Markstrom 2012; Dugdale et al. 2017). However, these models still rely on qualitative parameterizations (i.e., approxi-

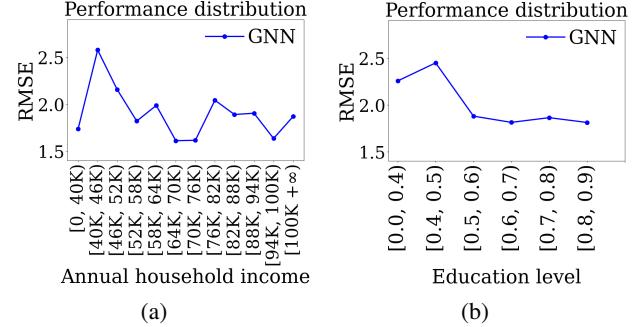


Figure 1: The distribution of RMSE for a GNN model's predictions in different sensitive values, (a) annual household income in USD, (b) education level as the proportion of the population that has attended college relative to the total population, in our study region in the Delaware River Basin.

mations) based on soil and surficial geologic classification along with topography, land cover, and climate input, leading to degraded predictive accuracy. Recently, graph neural networks (GNNs) have been widely adopted as a data-driven solution to model stream networks (Jia et al. 2021b; Moshe et al. 2020; Chen et al. 2022; Jia et al. 2021a; Chen et al. 2021; Jia et al. 2023) because they can learn to capture complex interactions amongst different stream segments (e.g., through mass advection and diffusion).

Although GNNs have shown promising results in improving the predictive accuracy in initial small-scale tests, their application to large-scale stream network modeling frequently results in substantial bias related to spatial locations. Such spatial bias can lead to harm to society by amplifying discrimination and social inequity in decision and policy making, as reported in a recent National Institute of Standards and Technology (NIST) document (Schwartz et al. 2022). For example, Fig. 1 shows the distribution of temperature prediction errors (in root-mean-square errors (RMSE)) by a GNN model over locations (i.e., stream segments) with different sensitive attribute values in the Delaware River Basin. In this example, we consider two attributes, annual household income and education level. Here the education level is quantified numerically as the proportion of the pop-

ulation that has attended college relative to the total population within the area. It can be seen that the model has relatively poor performance in low-income and low-education communities. Such biases may unintentionally induce unfair distribution of social resources (e.g., subsidies and assistance) and treatment with respect to environmental policies, especially for low-income and low-education regions.

Prior works (Xie et al. 2022; He et al. 2022, 2023, 2024b,a; Chen et al. 2024) have developed fairness formulations for spatial datasets by prioritizing training over locations with suboptimal performance, but these are not based on GNNs and fail to consider sensitive groups of different locations. Recent studies have also started investigating the bias of GNNs based on the connections amongst different nodes (Tang et al. 2020; Cong et al. 2023; Dong et al. 2022). In particular, GNNs learn node embeddings by aggregating and transforming node feature vectors over the graph topology. It is possible that a node aggregates information from neighbors that are mostly sampled from only a few sensitive groups (e.g., high-income regions). Therefore, the embeddings learned on the graph tend to be correlated with sensitive attributes and lead to discriminatory predictions. Most existing methods of addressing GNN-incurred model bias aim to balance the contribution of nodes from different sensitive groups (e.g., high-income and low-income regions). However, they do not explicitly consider the influence amongst different nodes for quantifying the node contribution in the graph aggregation process. This is especially important and challenging in scientific modeling tasks due to the strong variability in node interactions driven by complex physical processes.

In this paper, we propose a Physics-Guided Fair Graph (PGFG) model to enhance fairness by leveraging the physical knowledge about node interactions. In particular, we define the influence amongst nodes in GNNs and then quantify such influence based on known physical theories about water dynamics and advection. We also utilize existing physics-based models to simulate variables that are not available or not measurable but required in computing the influence. Once we pre-compute the influence over the graph based on underlying physics, we introduce an edge modifier to adjust the edge sampling in the graph, with the aim to balance the influence each node receives from different sensitive groups (e.g., locations with different income levels). The edge modifier is also designed to directly balance the influence density when considering continuous sensitive attributes.

We evaluate our method in predicting stream water temperature using data from the past 40 years in the Delaware River Basin, which covers large areas in New York, Pennsylvania, New Jersey, Delaware, and Maryland. Our experiments demonstrate the superiority of the proposed method over a diverse set of baselines in preserving fairness over two different sensitive attributes, annual household income and educational level, selected from the U.S. Census data (Manson 2022), without compromising predictive performance.

Related Work

Extensive research has focused on fairness regarding pre-defined categorical attributes (e.g., race, gender), includ-

ing regularization (Zafar et al. 2017; Kamishima et al. 2011; Serna et al. 2020), sensitive category de-correlation (Sweeney et al. 2020; Zhang et al. 2021; Alasadi et al. 2019), data collection or filtering strategies (Jo and Gebru 2020; Yang et al. 2020), and more (e.g., a survey (Mehrabi et al. 2021)). Recent developments have targeted fairness formulations for spatial locations (Xie et al. 2022; He et al. 2022, 2023) to promote unbiased decision-making in natural resource management, which is critical for ensuring environmental justice. Most of these methods aim to reduce the performance disparity (e.g., variance of performance) across locations, but neglect the relationship between bias and sensitive attributes (e.g., income, education levels, business types) over different locations. Moreover, they focus on addressing data-incurred bias (e.g., data heterogeneity and imbalance of training data), but are not designed for addressing model-incurred bias.

GNNs have been widely used in many scientific disciplines, including hydrology (Moshe et al. 2020), freshwater science (Jia et al. 2021b), agriculture (Fan et al. 2022), and climate science (Cachay et al. 2020). Recent studies have highlighted that the learning process of GNNs may inadvertently create societal biases (Dong et al. 2022; Cong et al. 2023; He et al. 2024c). The inherent feature aggregation mechanisms in GNNs pose risks of propagating these biases, which affects model predictions. Traditional fairness approaches designed for independently and identically distributed (IID) data are not directly applicable to GNNs due to their inability to address biases embedded in graph structures (Li et al. 2021). To enhance fairness, an approach is to modify the input graph to reduce the graph-embedded bias, typically by adjusting the adjacency matrix (Li et al. 2021; Laclau et al. 2021; Dong et al. 2022). Research along this direction aims to create balanced node representation from diverse sensitive groups, striving for equitable embeddings. However, these methods primarily focus on balancing the number of sampled edges or adjacency weights amongst different groups, and may not adequately capture node interaction dynamics in complex scientific modeling tasks over a large space, potentially leading to aggregation bias.

Prior research has highlighted the importance of incorporating physical knowledge to guide machine learning in scientific modeling tasks (Willard et al. 2022). Most works on this topic modify the loss function (Raissi et al. 2017; Jia et al. 2019) or model structures (Muralidhar et al. 2020; Liu et al. 2024) based on domain-specific physical knowledge. Some works also explored integrating physical relationships with the information propagation process in GNNs (Jia et al. 2021b), which can improve the model predictive performance and generalizability even when the model is trained with limited observations. Incorporating physics could also provide opportunities for mitigating spatial biases, which are often caused by inadequate modeling of many complex system behaviors (e.g., dynamics and interactions). However, it remains largely under-explored how to effectively leverage such physical knowledge to promote fairness.

Problem Formulation and Preliminaries

Problem definition: The objective of this study is to predict water temperature in an interconnected river network. We consider N stream segments, each indexed by i . Each stream segment i is characterized by a set of input features over time, which is represented as $\mathbf{x}_i = \{\mathbf{x}_{i,1}, \mathbf{x}_{i,2}, \dots, \mathbf{x}_{i,T}\}$, collected over T time steps (e.g., dates). These features include meteorological variables like daily average solar radiation, air temperature, precipitation, local wind speed, date of the year, and some modeled local variables. Additionally, observed target variables $\{y_i^t\}$ (i.e., water temperature) are available for certain segments $i \in \{1, \dots, N\}$ at certain time steps $t \in \{1, \dots, T\}$.

This study also focuses on promoting fairness in model performance. Fairness is measured with respect to a specific sensitive census attribute \mathbf{S} (e.g., income level), which can be either categorical or continuous. This attribute is associated with regions (e.g., subcounties) through which each stream segment flows. For each stream segment i , we aggregate the sensitive attribute values from all overlapping and adjacent regions, and we represent this aggregated sensitive attribute value as s_i . Our objective is to ensure that the model's predictive accuracy is equitable over locations with different attribute values (e.g., low-income and high-income regions).

Graph representation of stream networks: Graphs are widely used for representing the interactions amongst stream segments in a stream network (Jia et al. 2021b; Chen et al. 2022). Here a graph is defined as $G = \{\mathcal{V}, \mathcal{E}\}$, where the nodes in the set \mathcal{V} represent stream segments, and the edge set \mathcal{E} represents connections between upstream-downstream segment pairs. We create an edge between node i and node j if node i is anywhere upstream from node j . The edges in the graph can be assigned weights based on the inverse of the stream distance between two stream segments (refer to details in the dataset section). These edge weights are stored in a node adjacency matrix \mathbf{A} . GNNs leverage this graph structure for predictive analysis with two key benefits: (1) They aggregate contextual information from connected segments to enhance the accuracy in predicting target variables, which are often highly influenced by water dynamics in upstream segments. (2) GNNs enable the information diffusion of observational data from data-rich to data-scarce stream segments, which helps improve the overall performance.

Fairness formulation: A variety of fairness formulations have been proposed, broadly categorized into three types: (i) individual fairness (Lahoti et al. 2019a,b) that treats similar individuals similarly, (ii) group fairness (Hardt et al. 2016; Dai et al. 2021) that expects parity of statistical performance across groups, and (iii) fairness formulations that aim to improve per-group performance, such as Pareto-fairness (Ge et al. 2022) and Rawlsian Max-Min fairness (Lahoti et al. 2020; Zhang et al. 2014). In this work, we extend the formulations (ii) and (iii) to the spatial data of a large number of stream segments with respect to a specific continuous sensitive attribute (e.g., annual household income).

To evaluate fairness in the context of a continuous sensitive attribute, one approach would be to first discretize the sensitive values into distinct groups, e.g., different income intervals. This allows us to apply group fairness concepts

effectively. Here we introduce the fairness measure M_{fair} , defined on a set of groups \mathcal{P} . This metric is designed to balance model performance across all groups within \mathcal{P} , as represented by the following equation:

$$M_{fair}(\mathcal{F}_\Theta, M_{\mathcal{F}}, \mathcal{P}) = \sum_{p \in \mathcal{P}} \frac{d(M_{\mathcal{F}}(\mathcal{F}_\Theta, p), E_p)}{|\mathcal{P}|}, \quad (1)$$

where $M_{\mathcal{F}}$ represents a performance metric (e.g., RMSE) for model \mathcal{F}_Θ on a specific group p , E_p denotes the average model performance across groups, and $d(\cdot, \cdot)$ is a distance function. Intuitively, a significant deviation of a group's performance $M_{\mathcal{F}}(\mathcal{F}_\Theta, p)$ from the overall average E_p indicates potential unfairness in model \mathcal{F}_Θ . In our test, we calculate E_p as the overall performance of the model \mathcal{F}_Θ across all groups, which is represented as $E_p = M_{\mathcal{F}}(\mathcal{F}_\Theta, \{p \mid p \in \mathcal{P}\})$; and $d(\cdot, \cdot)$ is measured as the absolute distance.

In this work, we also propose new algorithms to directly improve the fairness with respect to the continuous attributes. In the evaluation phase, we assess the model's effectiveness in reducing spatial bias by measuring its performance in worst-case scenarios. Ideally, the model should perform consistently across all ranges of the sensitive attribute values. We use a sliding window with a predetermined size to traverse the entire range of attribute values. As it moves, we evaluate the model's predictive performance for stream segments within the current window. Our fairness objective is to improve the model performance in the worst-case scenarios identified by any window.

Proposed Method

Model Architecture

We develop a machine learning model architecture \mathcal{F} to simulate water temperature in stream segments by taking into account their spatial and temporal dependencies (Fig. 2). For each stream segment, its thermal status and water quantity change gradually based on the current weather input and its historical state. We use a long short-term memory (LSTM) network layer to capture this temporal dependency. This layer generates the hidden presentation $\{\mathbf{h}_{i,t}\} = \text{LSTM}(\{\mathbf{x}_{i,t}\})$ for each segment i at each time step t by integrating the current input $\mathbf{x}_{i,t}$ with the previous LSTM state and hidden representation $\mathbf{h}_{i,t-1}$.

The water temperature in a stream segment is also affected by its neighbors, e.g., through water advected from upstream segments. Hence, after gathering the hidden representations for all the segments, we use L graph convolutional layers to capture interactions between neighboring stream segments. Formally, the l -th layer of graph convolution can be expressed as follows:

$$\begin{aligned} \mathbf{z}_{i,t}^{(l)} &= g_a^{(l)}([\mathbf{z}_{i,t}^{(l-1)}, \mathbf{a}_{i,t}^{(l)}; \theta_a^{(l)}], \text{ for } l \in \{1, \dots, L\}, \\ \mathbf{a}_{i,t}^{(l)} &= \text{Pooling}(\{\mathbf{z}_{j,t}^{(l-1)}, \forall j \in \mathcal{N}(i)\}, \mathbf{A}_i), \end{aligned} \quad (2)$$

where $\mathbf{z}_{i,t}^{(0)} = \mathbf{h}_{i,t}$, $\mathcal{N}(i)$ denotes the set of neighboring stream segments to i , and $\theta_a^{(l)}$ represents the parameters in the l -th graph convolution layer. The latent representation

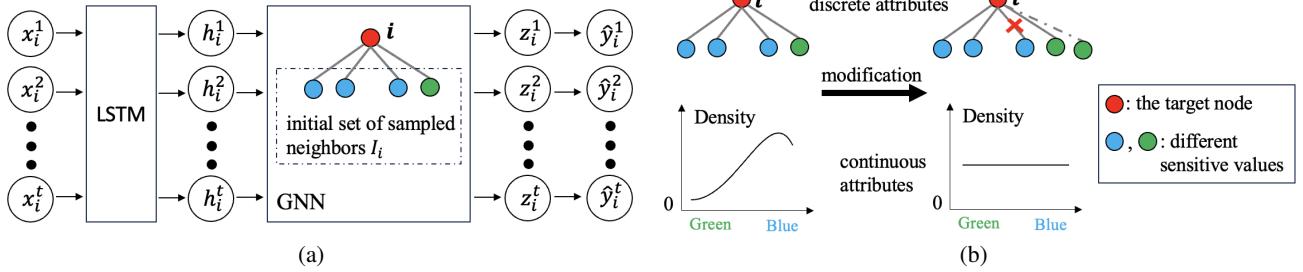


Figure 2: (a) A diagram of the proposed model. For each node i and each time t , the long short-term memory (LSTM) network extracts an embedding $\mathbf{h}_{i,t}$. Then we apply a graph neural network (GNN) to refine each time's embedding by aggregating information from neighboring nodes (highlighted in blue and green), producing a new embedding $\mathbf{z}_{i,t}$. Finally, the fully connected layers output the prediction \hat{y}_i^t . (b) Fair edge sampling for discrete and continuous sensitive attributes.

$\mathbf{a}_{i,t}^{(l)}$ embeds the information from neighboring stream segments, and is obtained through a weighted pooling of embeddings from the neighbors based on the weights from the adjacency matrix \mathbf{A} . We concatenate $\mathbf{a}_{i,t}^{(l)}$ with the last GNN layer's embedding of the target node $\mathbf{z}_{i,t}^{(l-1)}$ before the transformation using the function $g_a(\cdot)$. In this work, we adopt the GraphSAGE method (Hamilton et al. 2017) to implement the function $g_a(\cdot)$. Finally, a series of fully connected output layers are stacked to transform the aggregated embeddings $\mathbf{z}_{i,t}$ into the predicted output $\hat{y}_{i,t}$.

The model is updated by minimizing the distance between the predictions and true observations, as follows:

$$\mathcal{L} = \sum_{i,t \in \mathcal{D}} \|\hat{y}_{i,t} - y_{i,t}\|^2 / |\mathcal{D}|, \quad (3)$$

where \mathcal{D} is the set of $\{i, t\}$ for which the observed $y_{i,t}$ is available.

Physical Knowledge Integration

Despite the promise of GNNs in capturing spatial dependencies, the aggregation process (Eq. 2) can introduce learning bias over different groups. When training the GNN model, the prediction errors (Eq. 3) on each node i will be propagated to both node i and its neighbors because the hidden representation \mathbf{z}_j for any node $j \in \mathcal{N}(i)$ is involved in computing the latent representation \mathbf{a}_i for node i . Hence, the errors from the node i could also contribute to the training of its neighboring nodes. However, this aggregation process could lead to a group-biased learning process when the neighbors' influence is unevenly distributed over different groups. Here we first introduce a novel physics-guided graph approach to quantify the influence between neighboring nodes. The proposed influence measure will then be used to create a new fairness-driven aggregation method, which will be discussed later in the next section.

Based on the graph aggregation process, for each node i , we define the influence it receives from its neighbor node j at time t as follows:

$$\text{Influence}_t(j, i) = \left\| \frac{\partial \hat{y}_{i,t}}{\partial \mathbf{z}_{j,t}} \right\|, \text{ for } j \in \mathcal{N}(i), \quad (4)$$

where $\|\cdot\|$ represents the L2 norm of a vector.

We could estimate the gradients over the GNNs to compute the influence during the training process. However, this may degrade the training performance and stability because the influence estimate is dependent on aggregation parameters θ_a and thus can be affected by an immature model. The updated aggregation using the influence estimate will also in turn affect the learning of aggregation parameters θ_a . More importantly, true observations $\{y_{i,t}\}$ are often sparse over space and time in scientific modeling tasks, e.g., some low-income or remote regions have limited budgets in collecting high-quality observations, which leads to initial bias in the gradient estimates and θ_a . Moreover, directly utilizing these gradient estimates requires additional computational costs.

To address these challenges, we propose to leverage underlying physical knowledge to estimate the influence over the graph. Most dynamical systems are governed by partial differential equations (PDEs), which describe the change of target variables in response to external drivers and spatial interactions. When modeling a set of locations as different nodes in a graph, we can often estimate the influence between a pair of nodes based on the spatial interaction terms (represented as spatial derivatives) in the governing PDE. In stream networks, the dynamics of stream water temperature are governed by the heat transfer process following the law of energy conservation (Dugdale et al. 2017). The heat transfer process involves different heat fluxes, including shortwave and longwave radiation, latent heat flux (e.g., evaporation), and advective and convective heat flux. In particular, the advective heat flux describes the interaction between different stream segments. Through some derivations and approximations from the heat transfer PDE, we can obtain the relation of water temperature between neighboring segments (Dugdale et al. 2017; Boyd 1996), as follows:

$$y_{i,t} \approx \frac{y_{i,t-1}q_{i,t} + \sum_j y_{j,t}q_{j,t}}{q_{i,t} + q_{j,t}}, \text{ for } j \in \mathcal{N}(i), \quad (5)$$

where $q_{i,t}$ represents the streamflow at the stream segment i at time t . Then we can obtain $\frac{\partial y_{i,t}}{\partial y_{j,t}}$ as

$$\frac{\partial y_{i,t}}{\partial y_{j,t}} = \frac{q_{j,t}}{q_{j,t} + q_{i,t}}, \text{ for } j \in \mathcal{N}(i) \quad (6)$$

Finally, we can approximate the influence from node j to node i as follows:

$$\begin{aligned} \text{Influence}_t(j, i) &\approx \left\| \frac{\partial y_{i,t}}{\partial y_{j,t}} \frac{\partial \hat{y}_{j,t}}{\partial \mathbf{z}_{j,t}} \right\| \\ &= \frac{q_{j,t}}{q_{j,t} + q_{i,t}} \left\| \frac{\partial \hat{y}_{j,t}}{\partial \mathbf{z}_{j,t}} \right\|, \text{ for } j \in \mathcal{N}(i) \end{aligned} \quad (7)$$

When we compare the influence between different pairs of nodes, the term $\left\| \frac{\partial y_{i,t}}{\partial y_{j,t}} \right\|$ is constant because all the nodes use the same transformation layers to convert \mathbf{z} to the output \hat{y} . Hence we just need to consider $\frac{q_{j,t}}{q_{j,t} + q_{i,t}}$ for the influence from node j to node i .

As the streamflow observations q are very sparse for most stream segments, we complement observed streamflow with physical simulations. We leverage a physics-based PRMS-SNTemp model (Markstrom 2012) to simulate streamflow for all the stream segments and all the time steps. We use the simulated value when the streamflow $q_{i,t}$ is not available for certain segments or certain dates. Then we can use such combined streamflow data to estimate the influence between each pair of neighboring nodes.

Another issue is that the physical relationship described by Eq. 5 only applies to a node i and its direct upstream segment j . Most stream segments only have a few number of upstream segments, which makes the size of neighborhood relatively small. To enhance the power of GNN in leveraging spatial dependencies, prior work in this domain often extends the neighborhood to multi-hops away (Jia et al. 2021b; Chen et al. 2022), i.e., edges are created between node j and node i as long as node j is anywhere upstream from node i . This often leads to better performance because each node i can still be affected by water flow from a distant upstream segment j . When segment j to segment i are not geographically adjacent, we estimate the influence from node j to node i following the stream path between j and i , which is represented as $\{j_0 = j, j_1, j_2, \dots, j_M = i\}$. The influence can be estimated as follows:

$$\text{Influence}_t(j, i) = \prod_{m=1}^M \frac{q_{j_{m-1},t}}{q_{j_{m-1},t} + q_{j_m,t}} \quad (8)$$

This approach can be efficiently implemented as the influence over the graph can be pre-computed using the streamflow data. This approach is also generally applicable to many other scientific modeling tasks (e.g., hydrology, climate simulation) because we can often derive the spatial dependencies of target variables $\frac{\partial y_i}{\partial y_j}$ from the governing PDEs.

Fairness-driven Edge Sampling

Prior works on mitigating bias in GNNs (Li et al. 2021; Cong et al. 2023; Laclau et al. 2021; Dong et al. 2022) are focused on balancing the number of edges or the sum of adjacency weights across different sensitive groups, but they do not sufficiently consider the node influence and how that will bring bias to the learning process.

We introduce an edge modifier that leverages physics-based influence estimates to update the edge connections amongst nodes so as to improve group fairness. For each

node i , we randomly sample a fixed number of initial neighbors following the GraphSAGE (Hamilton et al. 2017). The edge modifier then either injects or removes edges to balance the physical influence across different sensitive groups within the sampled neighborhood of node i . Within the candidate neighbors for each sensitive group, we prioritize selecting neighbors that have similar features \mathbf{x} with the target node (i.e., the locations under similar weather conditions). This help mitigate noise in \mathbf{x} gathered through aggregation and improve the model robustness.

After these edge modification steps, we slightly rescale the edge weights to balance the influence from different groups. For each node i , we first measure the sum of influence from all the neighbors in each sensitive group \mathcal{G}_k as $\text{SI}_k = \sum_{j \in \mathcal{G}_k} \text{Influence}_t(j, i)$. Then we rescale the weight of each edge (j, i) such that the influence from each group matches the highest group influence, as follows:

$$\mathbf{A}'_{ji} = \mathbf{A}_{ji} \frac{\max_{k'} \text{SI}_{k'}}{\text{SI}_k}, \text{ for } j \in \mathcal{G}_k. \quad (9)$$

To train fair GNNs with continuous sensitive attributes, the edge modifier aims to generate a more balanced density distribution of influence over sensitive attribute values from each node's neighborhood. The central idea is to modify the edge connections to maximize the variance of sensitive values. Similar to the case with discrete attributes, we first sample an initial set of neighbors \mathcal{I}_i for each node i , and represent other candidate neighbors as $\mathcal{C}(i)$. Then for each of other candidate neighbors $j \in \mathcal{C}(i)$, we quantify the density at j by aggregating the influence of initial neighbors based on their similarity to the candidate neighbor, as follows:

$$\begin{aligned} \text{Density}(j) &= \sum_{j'} \text{sim}(s_j, s_{j'}) \cdot \text{Influence}(j', i), \\ &\text{for } j \in \mathcal{C}(i), j' \in \mathcal{I}_i, \end{aligned} \quad (10)$$

where $\text{sim}(a, b)$ is the similarity function and measured as $1 - \text{Norm}(|a - b|)$, which normalizes and inverts to ensure lower $|a - b|$ indicate greater similarity.

We prioritize adding node j as neighbors and including it in the aggregation process if it has a lower density value. Intuitively, if a candidate node j has a sensitive attribute value highly similar to many existing neighbors and these neighbors have already exerted significant influence, we opt not to include node j as a neighbor to avoid amplifying biases in the aggregation phase. Instead, we prefer to include candidate nodes that have their sensitive attribute values different from existing highly influential neighbors.

Experiments

Dataset

The dataset used in our evaluation is collected from the Delaware River Basin (DRB), sourced from the U.S. Geological Survey's National Water Information System (USGS 2016) and the Water Quality Portal (Read et al. 2017). Observations at specific latitudes and longitudes are matched to stream segments that vary in length from 48 to 23,120 m. These segments are defined by the geospatial fabric used for

the National Hydrologic Model (Regan et al. 2018), and are split up to have roughly a 1-day water travel time. We match observations to stream segments by snapping observations to the nearest stream segment within a tolerance of 250 m. Observations farther than 5,000 m along the river channel to the outlet of a segment are omitted from our dataset. Segments with multiple observation sites are aggregated to a single mean daily water temperature value. Refer to (Oliver et al. 2021) for the full observational dataset.

DRB contains 456 stream segments with input features at the daily scale from Jan 01, 1980, to Jul 31, 2020 (14,823 dates). In addition, amongst 456 segments, the number of streamflow observations only available for the 183 segments ranges from 16 to 14,774 with a total of 1,928,445 streamflow observations across all dates and segments. For those segments or specific dates where streamflow data are not available, we utilize simulated data produced by the physics-based PRMS-SNTemp model (Markstrom 2012).

We take sensitive attributes from the U.S. Census data (Manson 2022) and process them at the subcounty level. Here we focus on two sensitive attributes, median annual household income and education level. These attributes are considered in the context of two types of fairness: group fairness and fairness with continuous sensitive attributes. For median annual household income, we categorize it into three distinct groups: low income (0-50,000), middle income (50,000-100,000), and high income (>100,000). In this categorization, we identify 43 streams in the low-income group, 329 streams in the middle-income group, and 84 streams in the high-income group. Education level is quantified numerically based on the proportion of the population that has attended college relative to the total population within each subcounty. These data are then divided into two categories: a low-education group, where the proportion is less than 0.5, and a high-education group, where the proportion exceeds 0.5. Within these education categories, 69 streams are classified within the low-education group and 387 streams within the high-education group. Finally, we assign a sensitive attribute value to each stream segment based on the average sensitive values of all the subcounties it flows through.

Implementation Details

We generate the adjacency matrix \mathbf{A} based on the river distance between each pair of stream segment outlets, represented as $\text{dist}(i, j)$. We standardize the stream distance and then compute the graph edge weights as $A_{ij} = 1/(1 + \exp(\text{dist}(i, j)))$.

We implement different candidate approaches with the initial adjacency matrix for 100 epochs (converged), using the Adam optimizer ($\alpha = 0.001$). In the following experiments, we use data from the first 27 years (Jan 01, 1980, to Jan 20, 2007) for training and then test in the next 13 years (Jan 21, 2007, to Jul 31, 2020). Within the training data, the last 1/3 of the time steps are further separated as validation data to halt training so as to prevent overfitting.

Candidate Methods

We compare the performance and fairness of the proposed method with several representative baselines.

- GraphSAGE (Hamilton et al. 2017) is a widely used GNN model without fairness constraints, which only optimizes model predictive performance by randomly sampling an equal number of nodes for embedding aggregation. It also serves as the base model for all the following methods.
- REG (Serna et al. 2020) is a fairness-preserving method that adds the fairness objective as a regularization term in the loss function. For discrete attributes, REG includes the fairness metric by Eq. 1 as the regularization term. For continuous attributes, the regularization term is defined as the variance of performance across nodes.
- FairGNN (Dai et al. 2021) adopts an adversarial discriminator to rectify the final layer representations in the GNN model. Given the representation \mathbf{z} , the adversary aims to predict the node's sensitive group for discrete attributes and the specific sensitive values for continuous attributes.
- Bi-Level (Xie et al. 2022) introduces a bi-level optimization that adjusts learning rates via a global referee based on predictive performance. When considering discrete attributes or continuous attributes, it assigns relatively higher learning rates for underperforming nodes.
- Both DSGNN and SLDSGNN use degree-specific parameters to improve fairness (Tang et al. 2020). Specifically, SLDSGNN adds self-supervised learning for pseudo-labeling unlabeled nodes. To adapt these models for discrete attributes, we extend them by including group-specific parameters (using a similar method as in (Tang et al. 2020) for learning degree-specific parameters). We denote the variants of DSGNN and SLDSGNN with group-specific adjustments as DSGNN_g and SLDSGNN_g .
- PGFG is the proposed method that integrates physical knowledge and the edge modifier.

We include two variants of PGFG to show the effectiveness of using physical knowledge: (1) FairEdge, which only balances the number of edges across different sensitive values, and (2) FairAdj, which replaces the physical influence with the edge weights in the initial adjacency matrix \mathbf{A} .

Results

Overall accuracy and fairness evaluation: Table 1 presents an overview of the predictive performance and fairness of our proposed method and other baselines for two selected sensitive attributes. We use the RMSE as the evaluation metric. The fairness performance is measured using two metrics, i.e., performance disparity (discrete attribute) and worst-case performance (continuous attribute), as discussed in previous section. For the sensitive attribute of annual household income, we test three sliding window sizes: 1000, 3000, and 5000, and report the worst-case RMSE for each one. For the sensitive attribute of education level, we also try three sliding window sizes: 0.003, 0.006, and 0.01. For all metrics, a lower value indicates better performance.

Metrics in Table 1 indicate that the proposed PGFG method outperforms the baselines in terms of fairness for both discrete and continuous attribute scenarios. Moreover, PGFG effectively promotes the fairness without compromising the predictive performance. We can also observe

Method	Annual household income						Education level					
	Discrete scenario		Continuous scenario				Discrete scenario		Continuous scenario			
	RMSE	Fairness	RMSE	1000	3000	5000	RMSE	Fairness	RMSE	0.003	0.006	0.010
GraphSAGE	1.774	0.206	1.774	5.057	4.384	3.268	1.774	0.253	1.774	4.910	4.331	3.959
REG	1.785	0.181	1.765	4.802	4.040	3.073	1.786	0.226	1.765	4.625	4.323	4.007
FairGNN	1.774	0.157	1.784	4.864	3.807	3.022	1.788	0.215	1.767	4.477	4.201	3.760
Bi-Level	1.776	0.183	1.822	4.778	3.555	2.908	1.790	0.206	1.822	4.490	4.317	3.815
DSGNN _g	1.832	0.189	—	—	—	—	1.829	0.225	—	—	—	—
DSGNN	1.850	0.141	1.850	5.025	3.474	3.026	1.850	0.210	1.850	4.684	4.311	3.720
SLGSGNN _g	1.794	0.175	—	—	—	—	1.792	0.227	—	—	—	—
SLDSGNN	1.791	0.185	1.791	4.679	3.453	2.767	1.791	0.240	1.791	4.252	3.988	3.988
FairEdge	1.783	0.196	1.780	4.751	3.975	2.901	1.773	0.220	1.774	4.478	3.965	3.846
FairAdj	1.767	0.170	1.784	4.709	3.671	2.715	1.780	0.206	1.759	4.219	3.991	3.654
PGFG	1.774	0.134	1.748	4.456	3.333	2.573	1.783	0.189	1.777	3.999	3.817	3.597

Table 1: The fairness and overall RMSE with two continuous sensitive attributes. The lower value indicates the better result. Bold indicates the best-performing model for a given metric and attribute, and ‘—’ indicates that the model is not applicable.

that FairEdge has comparable fairness performance with other explicit fairness-preserving methods (e.g., REG and FairGNN), which shows the promise in promoting fairness through the modification of neighborhood in graphs. PGFG performs much better than both FairEdge and FairAdj. This confirms that merely balancing the number of neighbors or the sum of weights from different groups is not sufficient; it is also necessary to consider their physical influence.

Detailed fairness evaluation: We compare the proposed PGFG model with the base GraphSAGE model, the FairGNN model, and the SLDSGNN model in terms of the group fairness on the two sensitive attributes, as shown in Fig. 3. For each sensitive group p (x-axis), we report the absolute distances between RMSE achieved on each group p and the overall performance across all the groups $\{p|p \in \mathcal{P}\}$. Here the low-income and low-education regions have larger distances because they cover a relatively small number of locations (i.e., stream segments) compared to other groups and the models in general make larger errors on these locations. According to Fig. 3, PGFG has substantially reduced deviations for all groups compared to all other methods, which shows the effectiveness of balancing the physical influence across different sensitive groups.

Fig. 4 presents the distributions of RMSE achieved by different models within different ranges of sensitive values for the two sensitive attributes when they are considered as continuous variables. We evaluate the predictive performance for stream segments within a specific range of sensitive values. It is observed that PGFG improves mostly on poorly-performing segments in low-income and low-education communities (i.e., the left part of the distribution) while maintaining the competitive performance in well-performing segments in relatively high-income and high-education communities (the right part of the distribution).

Conclusion

We introduce PGFG, a GNN-based method for predicting stream water temperature while ensuring fairness across regions with varying sensitive attributes. PGFG leverages physical knowledge of stream dynamics to quantify influ-

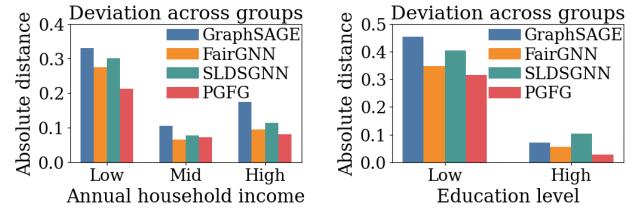


Figure 3: Group fairness comparison amongst PGFG, SLDSGNN, FairGNN, and GraphSAGE over the two different sensitive attributes. The lower value indicates less bias.

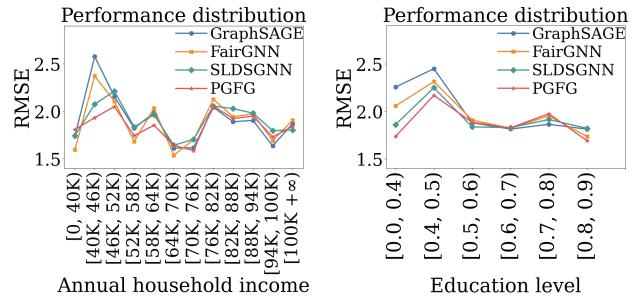


Figure 4: Continuous fairness comparison amongst PGFG, SLDSGNN, FairGNN, and GraphSAGE over the two different sensitive attributes. The lower value indicates less bias.

ence among graph nodes. We then introduce an edge modifier to refine the graph structure through edge sampling based on the obtained node influence. Evaluations on a large, heterogeneous river basin show that PGFG outperforms various baselines in enforcing fairness without compromising predictive accuracy. Although PGFG is developed and evaluated in the context of stream monitoring, it is applicable to many other important societal applications, such as weather and climate modeling. Future work will explore multiple sensitive attributes from U.S. Census data and identify the attribute combinations that lead to greater spatial bias.

Acknowledgments

This work was supported by the National Science Foundation (NSF) under grants 2239175, 2316305, 2147195, 242584, 2425845, 2430978, and 2126474, the USGS awards G21AC10564 and G22AC00266, the NASA grant 80NSSC24K1061, the Google’s AI for Social Good Impact Scholars program. This research was also supported in part by the University of Pittsburgh Center for Research Computing through the resources provided. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

Alasadi, J.; et al. 2019. Toward fairness in face matching algorithms. In *Proceedings of the 1st International Workshop on Fairness, Accountability, and Transparency in MultiMedia*, 19–25.

Boyd, M. S. 1996. Heat source: stream, river and open channel temperature prediction.

Cachay, S. R.; Erickson, E.; Bucker, A. F. C.; Pokropek, E.; Potosnak, W.; Osei, S.; and Lütjens, B. 2020. Graph Neural Networks for Improved El Niño Forecasting. *arXiv preprint arXiv:2012.01598*.

Carr, J. A.; D’Odorico, P.; Laio, F.; and Ridolfi, L. 2013. Recent history and geography of virtual water trade. *PloS one*, 8(2): e55825.

Chen, S.; Appling, A.; Oliver, S.; Corson-Dosch, H.; Read, J.; Sadler, J.; Zwart, J.; and Jia, X. 2021. Heterogeneous stream-reservoir graph networks with data assimilation. In *2021 IEEE International Conference on Data Mining (ICDM)*, 1024–1029. IEEE.

Chen, S.; et al. 2022. Physics-guided graph meta learning for predicting water temperature and streamflow in stream networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2752–2761.

Chen, W.; Xie, Y.; Jia, X.; He, E.; Bao, H.; An, B.; and Zhou, X. 2024. Referee-Meta-Learning for Fast Adaptation of Locational Fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 21949–21957.

Cong, Z.; Shi, B.; Li, S.; Yang, J.; He, Q.; and Pei, J. 2023. FairSample: Training Fair and Accurate Graph Convolutional Neural Networks Efficiently. *IEEE Transactions on Knowledge and Data Engineering*.

Dai, E.; et al. 2021. 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*, 680–688.

Dong, Y.; Liu, N.; Jalaian, B.; and Li, J. 2022. Edits: Modeling and mitigating data bias for graph neural networks. In *Proceedings of the ACM Web Conference 2022*, 1259–1269.

Dugdale, S. J.; et al. 2017. River temperature modelling: A review of process-based approaches and future directions. *Earth-Science Reviews*, 175: 97–113.

Fan, J.; Bai, J.; Li, Z.; Ortiz-Bobea, A.; and Gomes, C. P. 2022. A GNN-RNN approach for harnessing geospatial and temporal information: application to crop yield prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 11873–11881.

Ge, Y.; Zhao, X.; Yu, L.; Paul, S.; Hu, D.; Hsieh, C.-C.; and Zhang, Y. 2022. Toward Pareto efficient fairness-utility trade-off in recommendation through reinforcement learning. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 316–324.

Hamilton, W.; et al. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, 1024–1034.

Hardt, M.; et al. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.

He, E.; Chen, W.; Xie, Y.; and Jia, X. 2024a. 18 Fairness-Aware Deep Learning in Space. *Big Data: Techniques and Technologies in Geoinformatics*, 293.

He, E.; Xie, Y.; Chen, W.; Skakun, S.; Bao, H.; Ghosh, R.; Ravirathinam, P.; and Jia, X. 2024b. Learning With Location-Based Fairness: A Statistically-Robust Framework and Acceleration. *IEEE Transactions on Knowledge and Data Engineering*.

He, E.; Xie, Y.; Jia, X.; Chen, W.; Bao, H.; Zhou, X.; Jiang, Z.; Ghosh, R.; and Ravirathinam, P. 2022. Sailing in the location-based fairness-bias sphere. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, 1–10.

He, E.; Xie, Y.; Liu, L.; Chen, W.; Jin, Z.; and Jia, X. 2023. Physics Guided Neural Networks for Time-Aware Fairness: An Application in Crop Yield Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14223–14231.

He, E.; Xie, Y.; Sun, A.; Zwart, J.; Yang, J.; Jin, Z.; Wang, Y.; Karimi, H.; and Jia, X. 2024c. Fair Graph Learning Using Constraint-Aware Priority Adjustment and Graph Masking in River Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22087–22095.

Hipsey, M.; et al. 2014. GLM-General Lake Model: Model overview and user information.

Hoekstra, A. Y.; et al. 2012. The water footprint of humanity. *Proceedings of the national academy of sciences*, 109(9): 3232–3237.

Jia, X.; Chen, S.; Zheng, C.; Xie, Y.; Jiang, Z.; and Kalanat, N. 2023. Physics-guided Graph Diffusion Network for Combining Heterogeneous Simulated Data: An Application in Predicting Stream Water Temperature. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, 361–369. SIAM.

Jia, X.; Willard, J.; Karpatne, A.; Read, J.; Zwart, J.; Steinbach, M.; and Kumar, V. 2019. Physics Guided RNNs for Modeling Dynamical Systems: A Case Study in Simulating Lake Temperature Profiles. In *Proceedings of SIAM International Conference on Data Mining*.

Jia, X.; Xie, Y.; Li, S.; Chen, S.; Zwart, J.; Sadler, J.; Appling, A.; Oliver, S.; and Read, J. 2021a. Physics-Guided Machine Learning from Simulation Data: An Application in

Modeling Lake and River Systems. In *2021 IEEE International Conference on Data Mining (ICDM)*, 270–279. IEEE.

Jia, X.; Zwart, J.; Sadler, J.; Appling, A.; Oliver, S.; Markstrom, S.; Willard, J.; Xu, S.; Steinbach, M.; Read, J.; et al. 2021b. Physics-guided recurrent graph model for predicting flow and temperature in river networks. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, 612–620. SIAM.

Jo, E. S.; and Gebru, T. 2020. Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 306–316.

Kamishima, T.; et al. 2011. Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, 643–650. IEEE.

Laclau, C.; Redko, I.; Choudhary, M.; and Larggeron, C. 2021. All of the fairness for edge prediction with optimal transport. In *International Conference on Artificial Intelligence and Statistics*, 1774–1782. PMLR.

Lahoti, P.; Beutel, A.; Chen, J.; Lee, K.; Prost, F.; Thain, N.; Wang, X.; and Chi, E. 2020. Fairness without demographics through adversarially reweighted learning. *Advances in neural information processing systems*, 33: 728–740.

Lahoti, P.; et al. 2019a. ifair: Learning individually fair data representations for algorithmic decision making. In *2019 ieee 35th international conference on data engineering (icde)*, 1334–1345. IEEE.

Lahoti, P.; et al. 2019b. Operationalizing individual fairness with pairwise fair representations. *arXiv preprint arXiv:1907.01439*.

Li, P.; Wang, Y.; Zhao, H.; Hong, P.; and Liu, H. 2021. On dyadic fairness: Exploring and mitigating bias in graph connections. In *International Conference on Learning Representations*.

Liu, L.; Zhou, W.; Guan, K.; Peng, B.; Xu, S.; Tang, J.; Zhu, Q.; Till, J.; Jia, X.; Jiang, C.; et al. 2024. Knowledge-guided machine learning can improve carbon cycle quantification in agroecosystems. *Nature communications*, 15(1): 357.

Manson, S. M. 2022. IPUMS national historical geographic information system: version 17.0.

Markstrom, S. L. 2012. *P2S-coupled Simulation with the Precipitation-Runoff Modeling System (PRMS) and the Stream Temperature Network (SNTemp) Models*. US Department of the Interior, US Geological Survey.

Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6): 1–35.

Moshe, Z.; Metzger, A.; Elidan, G.; Kratzert, F.; Nevo, S.; and El-Yaniv, R. 2020. Hydronets: Leveraging river structure for hydrologic modeling. *arXiv preprint arXiv:2007.00595*.

Muralidhar, N.; Bu, J.; Cao, Z.; He, L.; Ramakrishnan, N.; Tafti, D.; and Karpatne, A. 2020. Phynet: Physics guided neural networks for particle drag force prediction in assembly. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, 559–567. SIAM.

Oliver, S. K.; et al. 2021. Predicting water temperature in the Delaware River Basin. U.S. Geological Survey Data Release.

Raissi, M.; et al. 2017. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*.

Read, E. K.; et al. 2017. Water quality data for national-scale aquatic research: The Water Quality Portal. *Water Resources Research*.

Regan, R. S.; et al. 2018. Description of the national hydrologic model for use with the precipitation-runoff modeling system (PRMS). Technical report, US Geological Survey.

Schwartz, R.; Vassilev, A.; Greene, K.; Perine, L.; Burt, A.; Hall, P.; et al. 2022. Towards a standard for identifying and managing bias in artificial intelligence. *NIST special publication*, 1270(10.6028).

Serna, I.; Morales, A.; Fierrez, J.; Cebrian, M.; Obradovich, N.; and Rahwan, I. 2020. Sensitiveloss: Improving accuracy and fairness of face representations with discrimination-aware deep learning. *arXiv preprint arXiv:2004.11246*.

Sweeney, C.; et al. 2020. Reducing sentiment polarity for demographic attributes in word embeddings using adversarial learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 359–368.

Tang, X.; Yao, H.; Sun, Y.; Wang, Y.; Tang, J.; Aggarwal, C.; Mitra, P.; and Wang, S. 2020. Investigating and mitigating degree-related biases in graph convolutional networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 1435–1444.

USGS. 2016. US Geological Survey. National water information system data available on the world wide web (USGS water data for the nation). doi: 10.5066/F7P55KJN.

Willard, J.; Jia, X.; Xu, S.; Steinbach, M.; and Kumar, V. 2022. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4): 1–37.

Xie, Y.; He, E.; Jia, X.; Chen, W.; Skakun, S.; Bao, H.; Jiang, Z.; Ghosh, R.; and Ravirathinam, P. 2022. Fairness by “Where”: A Statistically-Robust and Model-Agnostic Bi-Level Learning Framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Yang, K.; Qinami, K.; Fei-Fei, L.; Deng, J.; and Russakovsky, O. 2020. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagnet hierarchy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 547–558.

Zafar, M. B.; Valera, I.; Gomez Rodriguez, M.; and Gummadi, K. P. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, 1171–1180.

Zhang, C.; et al. 2014. Fairness in multi-agent sequential decision-making. *Advances in Neural Information Processing Systems*, 27.

Zhang, H.; et al. 2021. Towards Fair Deep Anomaly Detection. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 138–148.