

Boosting Fair Classifier Generalization through Adaptive Priority Reweighing

ZHIHAO HU*, University of Science and Technology of China, China

YIRAN XU*, Warwick Business School, University of Warwick, United Kingdom

MENGNAN DU, New Jersey Institute of Technology, United States

JINDONG GU, University of Oxford, United Kingdom

XINMEI TIAN†, University of Science and Technology of China, China

FENGXIANG HE, University of Edinburgh, United Kingdom

With the increasing penetration of machine learning applications in critical decision-making areas, calls for algorithmic fairness are more prominent. Although there have been various modalities to improve algorithmic fairness through learning with fairness constraints, their performance does not generalize well in the test set. A performance-promising fair algorithm with better generalizability is needed. This paper proposes a novel adaptive reweighing method to eliminate the impact of the distribution shifts between training and test data on model generalizability. Most previous reweighing methods propose to assign a unified weight for each (sub)group. Rather, our method granularly models the distance from the sample predictions to the decision boundary. Our adaptive reweighing method prioritizes samples closer to the decision boundary and assigns a higher weight to improve the generalizability of fair classifiers. Extensive experiments are performed to validate the generalizability of our adaptive priority reweighing method for accuracy and fairness measures (i.e., equal opportunity, equalized odds, and demographic parity) in tabular benchmarks. We also highlight the performance of our method in improving the fairness of language and vision models. The code is available at <https://github.com/che2198/APW>.

CCS Concepts: • **Computing methodologies** → **Machine learning**.

Additional Key Words and Phrases: algorithmic fairness, reweighing method, trustworthy AI

ACM Reference Format:

Zhihao Hu, Yiran Xu, Mengnan Du, Jindong Gu, Xinmei Tian, and Fengxiang He. xxx. Boosting Fair Classifier Generalization through Adaptive Priority Reweighing. *ACM Trans. Knowl. Discov. Data.* 1, 1 (June xxx), 25 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

With the increasing penetration of machine learning applications in critical decision-making areas, such as mortgage approval [36], employee selection [38], recidivism prediction [11], and healthcare,

*These authors contributed to the work equally.

†Corresponding author.

Authors' addresses: Zhihao Hu, University of Science and Technology of China, Hefei, Anhui, 230027, China, huzhihao12@gmail.com; Yiran Xu, Warwick Business School, University of Warwick, Coventry, West Midlands, CV4 7AL, United Kingdom, Yiran.Xu@warwick.ac.uk; Mengnan Du, New Jersey Institute of Technology, Newark, NJ, 07102, United States, mengnan.du@njit.edu; Jindong Gu, University of Oxford, Oxford, OX1 3PJ, United Kingdom, jindong.gu@outlook.com; Xinmei Tian, University of Science and Technology of China, Hefei, Anhui, 230027, China, xinmei@ustc.edu.cn; Fengxiang He, University of Edinburgh, Edinburgh, EH8 9AB, United Kingdom, F.He@ed.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© xxx Association for Computing Machinery.

1556-4681/xxx/6-ART \$15.00

<https://doi.org/XXXXXXXX.XXXXXXX>




Fig. 1. Illustration of our adaptive priority reweighing method. Samples of different shapes indicate samples with different labels, different colors refer to different sensitive attributes, and darker color indicates higher weight. The dotted lines represent the decision boundaries for Empirical Risk Minimization(ERM) and Equal Reweighting, while the full line represents the decision boundary for Adaptive Priority Reweighting. In ERM, fairness is disregarded altogether; all samples have the same weight, which could easily lead to unfairness. In equal reweighing, (sub)groups are created and all points within one (sub)group are assigned the same weight, ensuring improved fairness on the training set. However, using equal reweighing methods induces many points clustered around the decision boundary, which limits the generalization of fair classifiers. In comparison, our method models the distance from the sample to the decision boundary, which improves the generalizability of fair classifiers.

calls for algorithmic fairness are more prominent. In reality, machine learning tools can behave in unintentional or even harmful ways [46]. For example, research on the COMPAS dataset [15] shows that a well-calibrated classification algorithm tends to classify African-American defendants with a higher risk of re-conviction while classifying white defendants with lower risk [5]. In this case, the machine learning algorithms perpetuated and amplified the long-standing discrimination against people of color [47]. Both categorizations produced during data collection and/or analysis and the data itself are persistent, which paves the way for algorithms to compound the impact of historically biased data. More importantly, bias exists in forms of distribution, thus almost impossible to address bias by discarding individual instances. Therefore, it is much more difficult to achieve fairness by removing sensitive attributes from the data. Machine learning researchers have been focusing on higher classification accuracy, yet our study suggests the need to shift away from the fixation on accuracy and be actively attentive to mitigating biases from training data and ensuring generalizable performance across real-world datasets.

In the recent decade, machine learning researchers have engaged in the conversation of algorithmic fairness by investigating the correlation between model output and sensitive attributes (e.g., gender and race). Diverse fairness metrics are proposed to measure the disparities in model output between different sensitive attributes. In order to improve model performance in measuring fairness, we identify fairness as a constraint term, enabling fairness to be improved by solving the constrained optimization problem. Existing methods use various regularization terms or constraints to address the problem, ranging from covariance [53], equalized odds [18], to Rényi correlation [7]. Although they provide a shared foundation for machine learning researchers to engage in the fairness conundrum, they all suffer from generalization issues. That is, although these methods minimize the value of constraints on the training set, they cannot always guarantee fairness on the test set (see Figure 2). Chuang and Mroueh (2021) [17] and Ramachandran and Rattani (2022) [40] address the generalization issues of algorithmic fairness using data augmentation. Chuang and Mroueh (2021) propose applying mixup path regularization to enhance the generalizability of the

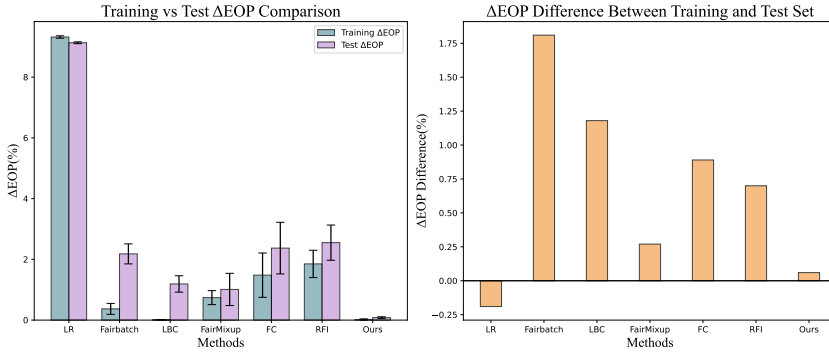


Fig. 2. Comparison of Training and Testing of Equal Opportunity Gap [26](Δ_{EOP}). We first use a histogram to illustrate the performance of six methods on the Equal Opportunity fairness metric on the training and test set, showcasing the range of Δ_{EOP} using error bars. In the histogram on the right, the difference in model performance is measured by subtracting the performance metrics on the training sets from those on the test set. The generalizability issue for fairness measures becomes particularly salient for algorithms with commendable fairness performance, such as Fairbatch [44] and Label Bias Correction(LBC) [29]. For example, these algorithms achieve considerable fairness performance on the training set. However, the same level of fairness performance does not necessarily translate to the test set. Our method demonstrates superior efficacy in ensuring fairness performance on both training and test sets, achieving optimal generalizability.

fair classifier. Ramachandran and Rattani (2022) further build on the conversation by engaging with the generative model to mitigate gender classification bias. Unfortunately, both studies require training data to conform to a specific distribution, leading to a limited representation of real-world datasets and reduced accuracy.

We aim to address the generalization issue of fairness improvement, but at the same time preserve accuracy. To do that, we propose an adaptive priority weight assignment strategy to improve fairness. As illustrated in Figure 1, we first divide the subgroups using sensitive attributes and predictions of the classifier. Data points closer to the decision boundary are more likely to be wrongly classified [30] – a phenomenon highlighting the critical role of decision boundaries in a model’s generalizability [3]. Therefore, our method models the distance from data points to the decision boundary to improve the generalizability of fair classifiers. This solves the generalizability issues associated with equal reweighing. Equal reweighing methods induce the cluster of sample points around the decision boundary. Distribution shifts could cause the model to classify sample points closer to the decision boundary incorrectly, thus inducing unfairness in the test set and limiting the generalization of fair classifiers. Additionally, we take the statistical independence of model output value and sensitive attribute as the termination condition and assign higher weights to samples whose model output is closer to the decision boundary in each subgroup.

We design experiments to validate the generalizability of our adaptive priority reweighing method for accuracy and fairness measures (i.e., equal opportunity, equalized odds, and demographic parity) on tabular, language, and vision benchmarks. We further highlight the performance of our method in improving the fairness of language and vision models. Our method shows promising results in improving the fairness of any pre-trained models simply via fine-tuning. The code is available at <https://github.com/che2198/APW>.

We summarize our contributions as follows:

1. We propose a novel reweighing method to mitigate data bias under different fairness measures, whose performance is validated through experiments on tabular, language, and vision benchmarks.

2. Our method provides a theoretical guarantee of generalization error bound.
3. Our method improves the fairness of any pre-trained model via fine-tuning by experimenting with both language and vision benchmarks.

2 RELATED WORK

Machine Learning Fairness. Most fairness-promoting algorithms focus on binary classification under binary-sensitive attributes. Generally, these methods can be divided into three categories: pre-processing, in-processing, and post-processing. Pre-processing methods tend to alter the sample distributions of protected variables to remove discrimination from the training data. Approaches to preprocessing data include sample selection [43, 44, 56], sample reweighing [13, 29, 30], fair representation learning [36, 54], fair data generation [28, 50, 51], etc. In-processing methods incorporate changes into the objective function or impose a constraint to remove discrimination during the model training process. Fairness constraints measure the correlation between model outputs and sensitive attributes [2, 10], such as adversarial layers [1, 55], contrastive learning loss [59], uncorrelation constraints [53], counterfactual logit pairing [24], and mutual information [7, 42]. Post-processing transforms model output in different demographic groups to achieve fairness [22, 26, 39, 49]. These methods can satisfy the fairness constraints on the training set but may not be generalized to the test set.

Recently, some methods based on data augmentation have been proposed to enhance the generalizability of fair classifiers, such as mix-up paths regularization [17], data generation [40]. However, these methods require training data to conform to a specific distribution or lead to computationally expensive algorithms.

Data Reweighing. The concept of data reweighing has gained considerable attention as an effective approach to improving the performance and fairness of machine learning classifiers. Various studies have explored different aspects of data reweighing techniques to address the challenges associated with biased training data and classifier fairness.

Many approaches introduced data reweighing as a means of improving classifier generalization [16, 60]. They highlighted the importance of reweighing samples in the training dataset to mitigate issues arising from class imbalance, leading to improved overall classifier performance. Meanwhile, [35, 57] demonstrated that reweighing techniques can be instrumental in reducing bias during training, and then promoting the fairness of classifiers across diverse demographic groups.

A subset of research has concentrated specifically on fairness improvement through data reweighing. Kamiran et al. [30] pioneered the concept of fixed reweighing, which involves assigning different weights to instances from underrepresented groups to balance their influence during training. Adaptive reweighing methods, as proposed by [29, 32, 33, 48], dynamically adjust instance weights based on the evolving characteristics of the dataset. These methods attempt to alleviate bias and enhance fairness during training.

However, these data reweighing methods on fairness minimize constraints on the training set. They cannot guarantee fairness on the test set. Though Chai and Wang [13] attempt to address this by assigning varying weights to individual samples within each subgroup, it may not adequately address the generalizability issue, for it only allocates higher weight on sample point with higher loss. Therefore, we propose a novel reweighing method that addresses the generalization issue of fairness improvement and preserves classification accuracy by allocating higher weight to sample points closer to the decision boundary.

3 PRELIMINARIES

We consider the binary classification task. A sample is $z = (x, a, y)$, where $x \in \mathcal{X}$ is the input feature, $a \in \mathcal{A} = \{0, 1\}$ is the sensitive attribute (e.g., gender, race, and age), and $y \in \{0, 1\}$ is the prediction

target. Let Z , X , A , and Y denote the corresponding random variables of z , x , a , and y , respectively. Then, the goal of fair machine learning is to learn a binary classifier $h_\theta : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ while ensuring a specific notion of fairness with respect to the sensitive attribute A , where θ is the parameter of the classifier. For simplicity, we denote $\hat{Y} := f(h_\theta(X, A)) \in \{0, 1\}$ as the prediction of the classifier h for variable $Z = (X, A, Y)$, where f is a label function.

3.1 Group fairness

A major family of fairness concepts is the *group fairness*, which aims to characterize discrimination across various groups of individuals. The most used definitions of group fairness are *demographic parity* [21], *equalized odds* [26], and *equal opportunity* [26].

DEFINITION 1 (FAIRNESS DEFINITIONS). *Given a data distribution p_z , a classifier h satisfies:*

- *Demographic parity if the prediction \hat{Y} is independent of the sensitive attribute A .*
- *Equalized odds if the prediction \hat{Y} and the sensitive attribute A are independent conditional on the target Y .*
- *Equal opportunity if $\mathbb{P}(\hat{Y} = 1|A = 0, Y = 1) = \mathbb{P}(\hat{Y} = 1|A = 1, Y = 1)$.*

The following notions are designed to assess the degree to which the classifier satisfies the fairness constraints presented in Definition 1.

DEFINITION 2 (FAIRNESS MEASURES). *Given a data distribution p_z , for a binary classifier h :*

- *The demographic parity Δ_{DP} is defined as:*

$$\Delta_{DP}(\hat{Y}, p_z) = |\mathbb{P}(\hat{Y} = 1|A = 0) - \mathbb{P}(\hat{Y} = 1|A = 1)|. \quad (1)$$

- *The equalized odds gap Δ_{EO} is defined as:*

$$\Delta_{EO}(\hat{Y}, p_z) = \max_{y \in \{0,1\}} |\mathbb{P}(\hat{Y} = 1|A = 0, Y = y) - \mathbb{P}(\hat{Y} = 1|A = 1, Y = y)|. \quad (2)$$

- *The equal opportunity gap Δ_{EOP} is defined as:*

$$\Delta_{EOP}(\hat{Y}, p_z) = |\mathbb{P}(\hat{Y} = 1|A = 0, Y = 1) - \mathbb{P}(\hat{Y} = 1|A = 1, Y = 1)|. \quad (3)$$

Combined with Definition 1, it is clear that a small value of the fairness measure in Definition 2 would demonstrate a strong non-discriminatory of the given classifier, vice versa. When the notion Δ_{DP} , Δ_{EO} , or Δ_{EOP} is equal to zero, the classifier h perfectly satisfies demographic parity, equalized odds, or equal opportunity, respectively.

3.2 Subgroup weights

We divided the training samples into subgroups according to the sensitive attributes and predictions of the classifier. Specifically, we define three types of subgroups: $G_{y,a} := \{(x_i, a_i, y_i) : y_i = y, a_i = a\}$, $G_{*,a} := \{(x_i, a_i, y_i) : a_i = a\}$, and $G_{y,*} := \{(x_i, a_i, y_i) : y_i = y\}$. Let m represent the total number of training samples. We assigned the subgroup weight $W_{y,a}$ to samples that satisfy $Y = y$ and $A = a$. To quantify the number of samples within a subgroup, we introduced the notation $m_{y,a}$, defined as $m_{y,a} := |\{i : y_i = y, a_i = a\}|$. Similarly, we defined $m_{y,*} := |\{i : y_i = y\}|$ and $m_{*,a} := |\{i : a_i = a\}|$. Additionally, we introduced the proportion $p_{y,a}$, which represents the fraction of samples in $G_{y,a}$ compared to the total number of samples: $p_{y,a} := \frac{m_{y,a}}{m}$. For the equal reweighting approach, each sample is assigned a weight equal to its subgroup weight.

3.3 Generalization error

When dealing with the problem of an unknown distribution D , the empirical risk is commonly employed as an approximation of the expected risk. The empirical risk is defined as follows:

$$\hat{R}_{L,D}(h) = \frac{1}{m} \sum_{i=1}^m L(y_i, h(x_i, a_i)), \quad (4)$$

where the corresponding expected risk is

$$R_{L,D}(h) = R[D, h, L] = E_{(X,A,Y) \sim D} [\hat{R}_{L,D}(h)], \quad (5)$$

and L represents a surrogate loss function, such as square loss, logistic loss, or hinge loss. The classifier is learned through the process of empirical risk minimization (ERM) [45]:

$$h' = \arg \min_{h \in H} \hat{R}_{L,D}(h). \quad (6)$$

The consistency of $\hat{R}_{L,D}(h')$ with respect to $\min_{h \in H} R_{L,D}(h)$ is of utmost importance when designing surrogate loss functions and learning algorithms. Let

$$h^* = \arg \min_{h \in H} R_{L,D}(h). \quad (7)$$

It has been demonstrated in previous work [6] that the following proposition holds true:

$$R_{L,D}(h') - R_{L,D}(h^*) \leq 2 \sup_{h \in H} |R_{L,D}(h) - \hat{R}_{L,D}(h)|. \quad (8)$$

The term on the right-hand side is commonly referred to as the generalization error, which measures the performance of the learned model on unseen data. The consistency of the model is ensured through the convergence of the generalization error.

4 PROPOSED METHOD

In this section, we present our Adaptive Priority Reweighting method. First, we introduce the updating rules for adaptive priority reweighting, followed by its application to prominent fairness measures, including equal opportunity, equalized odds, and demographic parity. Finally, we provide the theoretical guarantee of our method.

4.1 Updating Rules for Adaptive Priority Reweighting

For a given dataset $\{(x_i, a_i, y_i)\}_{i=1}^m$ of m samples, the vanilla training objective without reweighting can be formulated as:

$$\min_{\theta} \frac{1}{m} \sum_{i=1}^m L(y_i, h_{\theta}(x_i, a_i)). \quad (9)$$

Given that the unweighted training objective does not incorporate a fairness criterion, its minimizer often fails to satisfy the desired fairness requirement. To overcome this limitation, we introduce a weighting scheme for each sample to ensure the desired fairness guarantee. For example, assigning higher weights to a specific sensitive group can result in improved accuracy specifically for that group. Consequently, the batch gradient estimate provides an unbiased estimation of the reweighted empirical risk. In other words, if we draw a training example i with weight w_i , the training objective can be formulated as follows:

$$\min_{\theta} \frac{1}{m} \sum_{i=1}^m w_i L(y_i, h_{\theta}(x_i, a_i)). \quad (10)$$

This observation leads to the following bilevel optimization-based interpretation of how reweighting interacts with the inner optimization algorithm. Initially, batch SGD optimizes the unweighted

empirical risk. Subsequently, based on the results of the inner optimization, the outer optimizer iteratively refines the weights $\mathbf{w} := (w_1, w_2, \dots, w_m)$ assigned to each training example. The inner optimizer then operates on batches drawn from a new distribution, reoptimizing the inner objective function. This process continues until convergence is achieved. Consequently, the Adaptive Priority Reweighting algorithm can be perceived as a combination of an outer optimizer and an inner optimizer, solving the following bilevel optimization problem:

$$\min_{\mathbf{w}} \text{Cost}(\boldsymbol{\theta}_{\mathbf{w}}), \quad (11)$$

$$\boldsymbol{\theta}_{\mathbf{w}} = \arg \min_{\boldsymbol{\theta}} \frac{1}{m} \sum_{i=1}^m w_i L(y_i, h_{\boldsymbol{\theta}}(x_i, a_i)), \quad (12)$$

where $\text{Cost}(\cdot)$ captures the goal of the optimization. To simplify the discussion, we initially focus on the fairness metric Δ_{DP} , and then extend our analysis to include Δ_{EO} and Δ_{EOP} . The outer optimization problem, specifically for the Δ_{DP} metric, takes the following form:

$$\min_{\mathbf{w}} \left| \frac{1}{m_{1,*}} \sum_{z_i \in G_{1,*}} f(h_{\boldsymbol{\theta}_{\mathbf{w}}}(x_i, a_i)) - \frac{1}{m_{0,*}} \sum_{z_j \in G_{0,*}} f(h_{\boldsymbol{\theta}_{\mathbf{w}}}(x_j, a_j)) \right|. \quad (13)$$

To address this optimization problem, we begin by computing the subgroup weight $G_{y,a}$ for each subgroup. Subsequently, we determine the sample weight based on the distance of each sample from the decision boundary within its respective subgroup. Now, let us delve into the rationale behind the calculation of the subgroup weight $G_{y,a}$.

If the classifier $h_{\boldsymbol{\theta}}$ exhibits fairness, meaning that A and \hat{Y} are statistically independent, the expected probability $P_{\text{exp}}(\hat{Y} = 1, A = 1)$ can be expressed as follows:

$$P_{\text{exp}}(\hat{Y} = 1, A = 1) := \frac{|\{i : \hat{y}_i = 1\}|}{m} \times \frac{|\{i : a_i = 1\}|}{m}. \quad (14)$$

In reality, the observed probability is given by:

$$P_{\text{obs}}(\hat{Y} = 1, A = 1) := \frac{|\{i : \hat{y}_i = 1, a_i = 1\}|}{m}, \quad (15)$$

where m represents the total number of samples. It is possible that the expected probability differs from the observed probability. If the expected probability exceeds the observed probability, it indicates a bias towards the subgroup $G_{1,1}$.

To address the bias, we aim to assign lower weights to objects that have been either deprived or favored. For this purpose, we assign weights to each subgroup $G_{y,a}$ as follows:

$$W_{y,a} := \frac{P_{\text{exp}}(\hat{Y} = y, A = a)}{P_{\text{obs}}(\hat{Y} = y, A = a)} = \frac{|\{i : \hat{y}_i = y\}| \times |\{i : a_i = a\}|}{m \times |\{i : \hat{y}_i = y, a_i = a\}|}. \quad (16)$$

The weight assigned to each subgroup will be calculated as the ratio between the expected probability of observing an instance with its sensitive attribute value and prediction under the assumption of independence, and its corresponding observed probability.

Based on the aforementioned considerations, we propose the following algorithm for determining the subgroup weights:

$$\forall t \in \{1, 2, \dots, T\} : W_{y,a}^{(t)} = W_{y,a}^{(t-1)} \times \frac{|\{i : \hat{y}_i^{(t-1)} = y\}| \times |\{i : a_i = a\}| + \alpha}{m \times |\{i : \hat{y}_i^{(t-1)} = y, a_i = a\}| + \alpha}, \quad (17)$$

where α is a constant that regulates the magnitude of the weight adjustment, and T represents the training iteration. The weight assigned to each sample is calculated as follows:

$$\forall t \in \{1, 2, \dots, T\} : w_i^{(t)} = \frac{W_{y,a}^{(t)}}{\sum_y \sum_a W_{y,a}^{(t)}} \times \frac{p_{y,a} \exp(-\eta \phi_i^{(t)})}{\sum_{i=1}^{m_{y,a}} \exp(-\eta \phi_i^{(t)})}, z_i \in G_{y,a}, \quad (18)$$

where η represents the step size, and $\phi_i^{(t)}$ denotes the margin between the prediction of sample z_i and the decision boundary value d :

$$\forall t \in \{1, 2, \dots, T\} : \phi_i^{(t)} = |h_\theta^{(t)} - d|, i = 1, 2, \dots, m, \quad (19)$$

For a comprehensive implementation of the fair classifier h_θ targeting Demographic Parity, please refer to Algorithm 1.

Algorithm 1 Training a fair classifier for Demographic Parity.

- 1: **Input:** Training dataset $\mathcal{Z}_{[m]} = \{(x_i, a_i, y_i)\}_{i=1}^m$, training iteration T , subgroup learning rate α , step size η , value of decision boundary d , classification procedure H .
- 2: **Output:** Fair classifier h_θ for Demographic Parity.
- 3: Initialize the model parameters θ of the classifier h , subgroup weights $W_{1,1}^{(0)} = W_{1,0}^{(0)} = W_{0,1}^{(0)} = W_{0,0}^{(0)} = 1$, sample weights $w_1^{(0)} = w_2^{(0)} = \dots = w_m^{(0)} = \frac{1}{m}$, margin $\phi_1^{(0)} = \phi_2^{(0)} = \dots = \phi_m^{(0)} = 0$, respectively.
- 4: Let $h_\theta^{(0)} := H(\mathcal{Z}_{[m]}, \{w_i^{(0)}\}_{i=1}^m)$.
- 5: **for** t **in** $1, \dots, T$ **do**
- 6: Update the margin for each sample:

$$\phi_i^{(t)} = |h_\theta^{(t)} - d| \text{ for } i \in [m].$$

- 7: **for** $y \in \{0, 1\}$ **do**
- 8: **for** $a \in \{0, 1\}$ **do**
- 9: Calculate subgroup weight $W_{y,a}^{(t)}$:

$$W_{y,a}^{(t)} = W_{y,a}^{(t-1)} \times \frac{|\{i : \hat{y}_i^{(t-1)} = y\}| \times |\{i : a_i = a\}| + \alpha}{m \times |\{i : \hat{y}_i^{(t-1)} = y, a_i = a\}| + \alpha}.$$

- 10: Update sample weights:

$$w_i^{(t)} = \frac{W_{y,a}^{(t)}}{\sum_y \sum_a W_{y,a}^{(t)}} \times \frac{p_{y,a} \exp(-\eta \phi_i^{(t)})}{\sum_{i=1}^{m_{y,a}} \exp(-\eta \phi_i^{(t)})} \text{ for } z_i \in G_{y,a}.$$

- 11: Update $h_\theta^{(t)} := H(\mathcal{Z}_{[m]}, \{w_i^{(t)}\}_{i=1}^m)$.
 - 12: **return** $h_\theta^{(T)}$
-

Intuitively, if the observed positive prediction rate for a protected class $A = a$ is lower than the expected positive prediction rate, we adjust the weights of the positively labeled samples of $A = a$ by increasing them, while decreasing the weights of the negatively labeled samples of $A = a$. This approach encourages the classifier to improve its accuracy on the positively labeled samples within $A = a$, while reducing its accuracy on the negatively labeled samples. Moreover, the weight assigned to a sample is higher when it is closer to the decision boundary within each subgroup. These two mechanisms collectively contribute to an increase in the positive prediction rate for $A = a$ and enhance the generalizability of the fair classifier h_θ .

Algorithm 1 operates through a series of iterative steps as follows:

- Update the margin for each sample.
- Adjust subgroup weights by the observed probability and the expected probability.
- Compute the weights for each sample based on these margin and subgroup weights.
- Retrain the classifier given these weights.

Algorithm 1 takes as input a classification procedure H , which, given a dataset $\mathcal{Z}[m] = \{(x_i, a_i, y_i)\}_{i=1}^m$ and weights $w_{i=1}^m$, outputs a classifier. In practice, H can be any training procedure that minimizes a weighted loss function over a specific parameterized function class, such as logistic regression.

4.2 Extensions to Other Fairness Notions

Equalized Odds: Equalized odds requires that the prediction \hat{Y} and the sensitive attribute A are independent conditional on the target Y . Case 1, if the observed true positive rate for a protected class $A = a$ is lower than the expected true positive rate, we increase the weights of the positively labeled samples of $A = a$. This will encourage the classifier to increase its accuracy on the positively labeled samples in $A = a$. Case 2, if the observed true negative rate for a protected class $A = a$ is lower than the expected true negative rate, we increase the weights of the negatively labeled samples of $A = a$. This will encourage the classifier to increase its accuracy on the negatively labeled samples in $A = a$. In addition, the closer the sample is to the decision boundary in each subgroup, the higher the weight is assigned. See the full pseudocode of learning fair classifier h_θ for Equalized Odds in Algorithm 2.

Equal Opportunity: Equalized opportunity requires that prediction \hat{Y} and the sensitive attribute A are independent conditional on the target $Y = 1$. If the observed true positive rate for a protected class $A = a$ is lower than the expected true positive rate, we increase the weights of the positively labeled samples of $A = a$. This will encourage the classifier to increase its accuracy on the positively labeled samples in $A = a$. In addition, the closer the sample is to the decision boundary in each subgroup, the higher the weight is assigned. Either of these two events will cause the positive prediction rate on $A = a$ to increase and the generalizability of h_θ to improve. This forms the intuition behind Algorithm 3.

4.3 Theoretical Analysis

In this section, we present theoretical guarantees regarding the learned classifier h when employing the Adaptive Priority Reweighting technique.

PROPOSITION 4.1. *Given the proportion $p_{y,a}$, and considering $wL(y, h(x, a))$ to be upper bounded by b , we can state that for any $\delta > 0$, with a probability of at least $1 - \delta$, the following holds:*

$$\begin{aligned} & \sup_{h \in H} |R_{L,D}(h) - \hat{R}_{wL,D}(h)| \\ &= \sup_{h \in H} |E_{(X,A,Y) \sim D} [\hat{R}_{wL,D}(h)] - \hat{R}_{wL,D}(h)| \\ &\leq \max_{a,y} p_{a,y} \cdot \mathfrak{R}(L \circ H) + 2b \sqrt{\frac{\log(1/\delta)}{2m}}, \end{aligned} \quad (20)$$

where the Rademacher complexity $\mathfrak{R}(L \circ H)$ is defined by [9]

$$\mathfrak{R}(L \circ H) = E_{(X,A,Y) \sim D, \sigma} \left[\sup_{h \in H} \frac{2}{m} \sum_{i=1}^m \sigma_i L(y_i, h(x_i, a_i)) \right], \quad (21)$$

and $\sigma_1, \dots, \sigma_m$ are i.i.d. Rademacher variables.

Algorithm 2 Training a fair classifier for Equalized Odds.

- 1: **Input:** Training dataset $\mathcal{Z}_{[m]} = \{(x_i, a_i, y_i)\}_{i=1}^m$, training iteration T , subgroup learning rate α , step size η , value of decision boundary d , classification procedure H .
 - 2: **Output:** Fair classifier h_θ for Equalized Odds.
 - 3: Initialize the model parameters θ of the classifier h , subgroup weights $W_{1,1}^{(0)} = W_{1,0}^{(0)} = W_{0,1}^{(0)} = W_{0,0}^{(0)} = 1$, sample weights $w_1^{(0)} = w_2^{(0)} = \dots = w_m^{(0)} = \frac{1}{m}$, margin $\phi_1^{(0)} = \phi_2^{(0)} = \dots = \phi_m^{(0)} = 0$, respectively.
 - 4: Let $h_\theta^{(0)} := H(\mathcal{Z}_{[m]}, \{w_i^{(0)}\}_{i=1}^m)$.
 - 5: **for** t **in** $1, \dots, T$ **do**
 - 6: Update the margin for each sample:

$$\phi_i^{(t)} = |h_\theta^{(t)} - d| \text{ for } i \in [m].$$
 - 7: **for** $y \in \{0, 1\}$ **do**
 - 8: **for** $a \in \{0, 1\}$ **do**
 - 9: Calculate subgroup weight $W_{y,a}^{(t)}$:

$$W_{y,a}^{(t)} = W_{y,a}^{(t-1)} \times \frac{|\{i : \hat{y}_i^{(t-1)} = y, y_i = y\}| \times |\{i : y_i = y, a_i = a\}| + \alpha}{m_{y,*} \times |\{i : \hat{y}_i^{(t-1)} = y, y_i = y, a_i = a\}| + \alpha}.$$
 - 10: Update sample weights:

$$w_i^{(t)} = \frac{W_{y,a}^{(t)}}{\sum_y \sum_a W_{y,a}^{(t)}} \times \frac{p_{y,a} \exp(-\eta \phi_i^{(t)})}{\sum_{i=1}^{m_{y,a}} \exp(-\eta \phi_i^{(t)})} \text{ for } z_i \in G_{y,a}.$$
 - 11: Update $h_\theta^{(t)} := H(\mathcal{Z}_{[m]}, \{w_i^{(t)}\}_{i=1}^m)$.
 - 12: **return** $h_\theta^{(T)}$
-

The Rademacher complexity has a convergence rate of order $\mathcal{O}(\sqrt{1/n})$ [9]. When the function class satisfies appropriate conditions on its variance, the Rademacher complexity converges rapidly and is of order $\mathcal{O}(1/n)$ [8]. In Proposition 4.1, we utilize the Rademacher complexity method to derive the generalization bound. It is worth noting that alternative hypothesis complexities and methods can also be employed to establish the generalization bound.

Considering the inequality:

$$\begin{aligned}
 & R_{L,D}(h') - R_{L,D}(h^*) \\
 &= R_{wL,D}(h') - R_{wL,D}(h^*) \\
 &\leq 2 \sup_{h \in H} |R_{wL,D}(h) - \hat{R}_{wL,D}(h)|,
 \end{aligned} \tag{22}$$

we observe that the consistency rate is preserved when learning with sample weights.

Building upon Proposition 4.1, we are now ready to present our main result, which addresses classification in the presence of label bias using the framework of Adaptive Priority Reweighting.

THEOREM 4.2. *The Adaptive Priority Reweighting method allows for the utilization of any surrogate loss functions originally designed for traditional classification problems in the context of classification in the presence of label bias.*

To derive generalization bounds, we utilize the Rademacher complexity method [9].

Algorithm 3 Training a fair classifier for Equal Opportunity

- 1: **Input:** Training dataset $\mathcal{Z}_{[m]} = \{(x_i, a_i, y_i)\}_{i=1}^m$, training iteration T , subgroup learning rate α , step size η , value of decision boundary d , classification procedure H .
 - 2: **Output:** Fair classifier h_θ for Equal Opportunity.
 - 3: Initialize the model parameters θ of the classifier h , subgroup weights $W_{1,1}^{(0)} = W_{1,0}^{(0)} = W_{0,1}^{(0)} = W_{0,0}^{(0)} = 1$, sample weights $w_1^{(0)} = w_2^{(0)} = \dots = w_m^{(0)} = \frac{1}{m}$, margin $\phi_1^{(0)} = \phi_2^{(0)} = \dots = \phi_m^{(0)} = 0$, respectively.
 - 4: Let $h_\theta^{(0)} := H(\mathcal{Z}_{[m]}, \{w_i^{(0)}\}_{i=1}^m)$.
 - 5: **for** t **in** $1, \dots, T$ **do**
 - 6: Update the margin for each sample:

$$\phi_i^{(t)} = |h_\theta^{(t)} - d| \text{ for } i \in [m].$$
 - 7: **for** $a \in \{0, 1\}$ **do**
 - 8: Calculate subgroup weight $W_{1,a}^{(t)}$:

$$W_{1,a}^{(t)} = W_{1,a}^{(t-1)} \times \frac{|\{i : \hat{y}_i^{(t-1)} = 1, y_i = 1\}| \times |\{i : y_i = 1, a_i = a\}| + \alpha}{m_{1,*} \times |\{i : \hat{y}_i^{(t-1)} = 1, y_i = 1, a_i = a\}| + \alpha}.$$
 - 9: Update sample weights:

$$w_i^{(t)} = \frac{W_{1,a}^{(t)}}{\sum_y \sum_a W_{y,a}^{(t)}} \times \frac{p_{1,a} \exp(-\eta \phi_i^{(t)})}{\sum_{i=1}^{m_{1,a}} \exp(-\eta \phi_i^{(t)})} \text{ for } z_i \in G_{1,a},$$

$$w_i^{(t)} = \frac{1}{m \sum_y \sum_a W_{y,a}^{(t)}} \text{ for } z_i \in G_{0,a}.$$
 - 10: Update $h_\theta^{(t)} := H(\mathcal{Z}_{[m]}, \{w_i^{(t)}\}_{i=1}^m)$.
 - 11: **return** $h_\theta^{(T)}$
-

Let $\sigma_1, \dots, \sigma_m$ be independent Rademacher variables, and $(x_1, a_1), \dots, (x_m, a_m)$ be i.i.d. variables. Consider a real-valued function class H . The Rademacher complexity of the function class over these variables is defined as follows:

$$\mathfrak{R}(H) = E_\sigma \left[\sup_{h \in H} \frac{2}{m} \sum_{i=1}^m \sigma_i h(x_i, a_i) \right]. \quad (23)$$

THEOREM 4.3. ([9]) Let H be a real-valued function class on $\mathcal{X} \times \mathcal{A}$, $S = \{(x_1, a_1), \dots, (x_m, a_m)\} \in \mathcal{X}^m \times \mathcal{A}^m$ and

$$\Phi(S) = \sup_{h \in H} \left| \frac{1}{m} \sum_{i=1}^m E[h(x, a)] - h(x_i, a_i) \right|. \quad (24)$$

Then, $E_S[\Phi(S)] \leq \mathfrak{R}(H)$.

The following theorem, which leverages Theorem 4.3 and Hoeffding's inequality, plays a crucial role in establishing the generalization bounds.

THEOREM 4.4. ([9]) *Let H be an $[a, b]$ -valued function class on $X \times \mathcal{A}$, and $S = \{(x_1, a_1), \dots, (x_m, a_m)\} \in X^m \times \mathcal{A}^m$. For any $h \in H$ and any $\delta > 0$, there exists a high-probability bound such that, with probability at least $1 - \delta$, the following inequality holds:*

$$E_{(x,a)}[h(x, a)] - \frac{1}{m} \sum_{i=1}^m h(x_i, a_i) \leq \mathfrak{R}(H) + (b - a) \sqrt{\frac{\log(1/\delta)}{2m}}. \quad (25)$$

Based on Theorem 4.4, it can be readily proven that for any function class with values in the range $[-b, b]$ and any $\delta > 0$, there exists a high-probability bound such that, with probability at least $1 - \delta$, the following statement holds:

$$\sup_{h \in H} |E_{(x,a,y) \sim D} \hat{R}_{wL,D} - \hat{R}_{wL,D}| \leq \mathfrak{R}(w \circ L \circ H) + 2b \sqrt{\frac{\log(1/\delta)}{2n}}. \quad (26)$$

Considering that w is bounded from above by $\max_{a,y} p_{a,y}$, we can utilize the Lipschitz composition property of Rademacher complexity, commonly known as Talagrand's Lemma (refer to, for instance, Lemma 4.2 in [37]), to establish the following relationship:

$$\mathfrak{R}(w \circ L \circ H) \leq \max_{a,y} p_{a,y} \mathfrak{R}(L \circ H). \quad (27)$$

Proposition 4.1 can be demonstrated alongside the fact that $E_{(x,a,y) \sim D} \hat{R}_{wL,D} = R_{wL,D} = R_{L,D}$. Theorem 4.2 follows from Proposition 4.1.

5 EXPERIMENT DETAILS

This section presents our experiments' implementation details, including the adopted datasets, the investigated fairness-aware algorithms, and experiment settings.

5.1 Datasets

We employ five datasets that have been widely adopted in the context of fairness, including three tabular datasets, one image dataset, and one language dataset: (1) Adult dataset [34] consists of 48,842 samples with gender as its sensitive attribute. The assigned classification task is to predict whether a person earns over 50K per year. (2) COMPAS dataset [4] consists of 7,981 samples with race as its sensitive attribute. The assigned task is to predict whether a defendant will re-offend. (3) Synthetic dataset is generated by CTGAN [52], which consists of 70,000 samples with gender as its sensitive attribute. The assigned classification task is to predict whether an individual's income is above \$50,000. (4) UTKFace dataset [58] consists of over 20,000 samples with race as its sensitive attribute. The assigned classification task is to predict an individual's gender. (5) MOJI dataset [12] consists of over 150,000 samples with race as its sensitive attribute. The assigned classification task is to predict an individual's emotions. For tabular datasets, we follow the same data pre-processing procedures as that in IBM AI Fairness 360 [19], and for vision or language datasets, we follow [44], or [41].

5.2 Fairness-aware Algorithms

This subsection investigates ten existing fairness-aware algorithms, including four pre-processing methods, four in-processing methods, and two post-processing methods. For pre-processing methods, we investigate four different approaches: (1) **Cutting**, which equals the data size of different sensitive groups by removing data, (2) **Reweighting** [30], which mitigates discrimination by weighing samples with a fixed weigh, (3) **Label Bias Correction(LBC)** [29], which iteratively weighs examples towards a fair data distribution, (4) **Fairbatch** [44], which adaptively selects minibatch sizes to improve model fairness. Besides, for in-processing methods, we investigate four different

Table 1. Performance on the Adult, COMPAS, and Synthetic test sets w.r.t. equal opportunity. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EOP} , denoted by \downarrow . **Bold** values represent the best performance.

Method	Adult		COMPAS		Synthetic	
	Accuracy \uparrow	$\Delta_{EOP}\downarrow$	Accuracy \uparrow	$\Delta_{EOP}\downarrow$	Accuracy \uparrow	$\Delta_{EOP}\downarrow$
LR	84.50 \pm 0.03	9.13 \pm 0.03	68.58 \pm 0.11	23.92 \pm 0.61	83.56 \pm 0.02	32.31 \pm 0.02
Cutting	80.22 \pm 0.22	8.21 \pm 1.31	67.75 \pm 0.64	21.51 \pm 0.55	83.26 \pm 0.95	30.00 \pm 1.21
Reweighting [30]	83.41 \pm 0.10	15.01 \pm 0.85	68.55 \pm 0.08	15.70 \pm 0.07	82.67 \pm 0.89	8.74 \pm 1.01
Fairbatch [44]	83.91 \pm 0.11	2.18 \pm 0.33	66.13 \pm 0.12	4.23 \pm 0.08	83.19\pm0.04	0.67 \pm 0.31
LBC [29]	84.35\pm0.03	1.19 \pm 0.27	68.58\pm0.05	2.81 \pm 0.00	82.72 \pm 0.07	2.33 \pm 0.11
FairMixup [17]	81.49 \pm 0.12	1.01 \pm 0.53	65.46 \pm 0.21	2.14 \pm 0.35	80.54 \pm 0.65	1.07 \pm 0.54
FC [53]	84.00 \pm 0.22	2.37 \pm 0.85	66.68 \pm 0.12	5.92 \pm 1.82	81.95 \pm 0.61	1.46 \pm 0.32
AD [55]	84.14 \pm 0.36	4.25 \pm 0.66	66.39 \pm 0.46	6.71 \pm 2.14	80.84 \pm 1.11	3.51 \pm 1.04
RFI [7]	83.86 \pm 0.21	2.55 \pm 0.58	66.73 \pm 0.22	5.18 \pm 2.31	81.65 \pm 0.52	1.33 \pm 0.41
EO [26]	80.17 \pm 0.53	4.56 \pm 1.35	60.31 \pm 0.31	6.55 \pm 1.27	77.48 \pm 1.40	5.18 \pm 2.01
CEO [39]	79.69 \pm 0.45	4.56 \pm 1.02	59.35 \pm 0.52	7.58 \pm 2.03	77.35 \pm 1.49	5.23 \pm 1.88
Ours	84.34 \pm 0.04	0.08\pm0.03	68.34 \pm 0.02	0.06\pm0.01	82.59 \pm 0.00	0.02\pm0.00

Table 2. Performance on the Adult, COMPAS, and Synthetic test sets w.r.t. equalized odds. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EO} , denoted by \downarrow . **Bold** values represent the best performance.

Method	Adult		COMPAS		Synthetic	
	Accuracy \uparrow	$\Delta_{EO}\downarrow$	Accuracy \uparrow	$\Delta_{EO}\downarrow$	Accuracy \uparrow	$\Delta_{EO}\downarrow$
LR	84.50 \pm 0.03	9.13 \pm 0.03	68.58 \pm 0.11	23.92 \pm 0.61	83.56 \pm 0.02	32.31 \pm 0.02
Cutting	80.22 \pm 0.22	8.21 \pm 1.31	67.75 \pm 0.64	21.51 \pm 0.55	83.26\pm0.95	30.00 \pm 1.21
Reweighting [30]	83.41 \pm 0.10	15.01 \pm 0.85	68.55 \pm 0.08	15.70 \pm 0.07	82.67 \pm 0.89	8.74 \pm 1.01
Fairbatch [44]	84.21 \pm 0.10	5.28 \pm 0.43	68.63 \pm 0.12	5.21 \pm 0.07	82.85 \pm 0.13	1.92 \pm 0.10
LBC [29]	84.30 \pm 0.03	3.69 \pm 0.21	68.83 \pm 0.05	4.16 \pm 0.03	82.68 \pm 0.06	1.86 \pm 0.11
FairMixup [17]	80.42 \pm 0.68	5.31 \pm 1.36	65.43 \pm 0.21	3.99 \pm 0.22	80.45 \pm 0.22	1.18 \pm 0.34
FC [53]	84.10 \pm 0.23	4.33 \pm 0.65	67.74 \pm 0.16	5.92 \pm 1.82	81.25 \pm 0.53	1.88 \pm 0.42
AD [55]	84.33\pm0.30	4.35 \pm 0.42	66.39 \pm 0.46	6.71 \pm 2.14	79.88 \pm 1.99	2.62 \pm 0.97
RFI [7]	83.89 \pm 0.25	4.55 \pm 0.53	67.73 \pm 0.22	5.18 \pm 2.31	80.85 \pm 0.50	2.01 \pm 0.47
EO [26]	80.17 \pm 0.53	4.56 \pm 1.35	60.31 \pm 0.31	6.55 \pm 1.27	77.38 \pm 1.14	6.18 \pm 1.88
CEO [39]	79.69 \pm 0.45	5.12 \pm 1.08	59.35 \pm 0.52	7.58 \pm 2.03	76.92 \pm 1.49	6.83 \pm 2.13
Ours	81.49 \pm 0.12	0.87\pm0.10	68.99\pm0.11	0.99\pm0.15	81.39 \pm 0.01	1.09\pm0.03

approaches: (1) **Fairness Constraints(FC)** [53], which adopts decision boundary covariance as a regularization term, (2) **Adversarial debiasing(AD)** [55], which suppresses the dependence between the prediction and sensitive attribute in an adversarial learning manner, and (3) **Rényi fair inference(RFI)** [7], which uses Rényi correlation as a regularization term. (4) **FairMixup** [17], which employs mixup paths regularization. Finally, for post-processing methods, we investigate

Table 3. Performance on the Adult, COMPAS, and Synthetic test sets w.r.t. demographic parity. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{DP} , denoted by \downarrow . **Bold** values represent the best performance.

Method	Adult		COMPAS		Synthetic	
	Accuracy \uparrow	$\Delta_{DP}\downarrow$	Accuracy \uparrow	$\Delta_{DP}\downarrow$	Accuracy \uparrow	$\Delta_{DP}\downarrow$
LR	84.50 \pm 0.03	16.33 \pm 0.23	68.58 \pm 0.11	20.45 \pm 0.51	83.56 \pm 0.02	26.50 \pm 0.02
Cutting	80.22 \pm 0.22	15.22 \pm 1.14	67.75 \pm 0.64	9.19 \pm 0.53	83.26\pm0.95	23.70 \pm 0.55
Reweighting [30]	83.41\pm0.10	8.04 \pm 0.81	68.55 \pm 0.08	15.77 \pm 0.37	82.38 \pm 0.91	13.22 \pm 1.20
Fairbatch [44]	82.33 \pm 0.13	1.18 \pm 0.13	68.13 \pm 0.12	2.63 \pm 0.08	81.59 \pm 0.07	6.03 \pm 0.34
LBC [29]	82.35 \pm 0.03	1.01 \pm 0.07	68.74\pm0.09	1.27 \pm 0.06	80.36 \pm 0.00	0.06\pm0.00
FairMixup [17]	82.46 \pm 0.33	1.45 \pm 0.30	64.92 \pm 0.41	2.10 \pm 0.33	80.76 \pm 0.14	0.98 \pm 0.31
FC [53]	82.22 \pm 0.22	1.45 \pm 0.62	68.48 \pm 0.33	5.42 \pm 1.62	79.52 \pm 0.47	3.91 \pm 0.77
AD [55]	81.51 \pm 0.76	2.64 \pm 1.56	68.39 \pm 0.44	6.61 \pm 1.18	72.39 \pm 0.46	3.11 \pm 1.18
RFI [7]	81.92 \pm 0.41	1.53 \pm 0.56	68.66 \pm 0.25	3.46 \pm 1.11	79.20 \pm 0.28	4.83 \pm 0.78
EO [26]	80.17 \pm 0.53	5.31 \pm 0.65	60.31 \pm 0.31	5.02 \pm 1.26	78.47 \pm 1.50	6.48 \pm 0.77
CEO [39]	79.69 \pm 0.45	5.23 \pm 0.82	59.35 \pm 0.52	6.08 \pm 1.33	77.31 \pm 1.31	6.82 \pm 1.13
Ours	82.31 \pm 0.02	0.14\pm0.03	68.50 \pm 0.04	1.15\pm0.05	80.38 \pm 0.04	0.09 \pm 0.02

two different approaches: (1) **Equalized Odds** (EO) [26], which solves a linear program to optimize equalized odds, (2) **Calibrated Equalized Odds** (CEO) [39], which optimizes over calibrated classifier score outputs to find probabilities with which to change output labels with an equalized odds objective.

5.3 Experiment Settings

In order to enhance the fairness of the classifier, we employed algorithm 1, algorithm 2, or algorithm 3 for training with 200 epochs on the tabular data. Additionally, we performed fine-tuning with 40 epochs for the UTKFace dataset and 35 epochs for the MOJI dataset using the pre-training model. The default batch sizes were set as follows: 1000 for the Adult dataset, 200 for the COMPAS dataset, 2000 for the synthetic dataset, 32 for the UTKFace dataset, and 1024 for the MOJI dataset. In the experiment, we set the learning rate to 0.1, the subgroup learning rate α to a value within the range of [0, 10000], the step size η to [0.5, 3], and the decision boundary value d to 0.5. We used the cross-entropy loss function and updated the model parameters using the stochastic gradient descent (SGD) algorithm. Cross-validation was performed on the training set to select the optimal value that resulted in the highest accuracy with minimal fairness violations.

6 EXPERIMENT ANALYSIS

This section compares our method with the other approaches in the Adult, COMPAS, and synthetic test sets w.r.t. accuracy and three fairness metrics. Furthermore, the experiment shows that our method can improve the fairness of any pre-trained unfair model via fine-tuning.

6.1 Accuracy and Fairness

We compare our method with the baseline and three types of fairness-aware algorithms: (1) non-fair method: LR; (2) pre-processing methods: Cutting, Reweighting [30], LBC [29], and Fairbatch [44]; (3) in-processing methods: FC [53], AD [55], RFI [7], and FairMixup [17]; (4) post-processing methods:

Table 4. Comparison of model performance on training and test sets in terms of Accuracy and Δ_{EOP} on the Adult dataset. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EOP} , denoted by \downarrow .

Method	Training		Test		Difference	
	Accuracy \uparrow	$\Delta_{EOP}\downarrow$	Accuracy \uparrow	$\Delta_{EOP}\downarrow$	Accuracy Diff	Δ_{EOP} Diff
LR	84.90 \pm 0.02	9.32 \pm 0.04	84.50 \pm 0.03	9.13 \pm 0.03	-0.40	-0.19
Fairbatch [44]	84.69 \pm 0.03	0.37 \pm 0.18	83.91 \pm 0.11	2.18 \pm 0.33	-0.78	1.81
LBC [29]	84.40 \pm 0.02	0.01 \pm 0.01	84.35 \pm 0.03	1.19 \pm 0.27	-0.05	1.18
FairMixup [17]	82.43 \pm 0.22	0.74 \pm 0.23	81.49 \pm 0.12	1.01 \pm 0.53	-0.94	0.27
FC [53]	84.56 \pm 0.17	1.48 \pm 0.73	84.00 \pm 0.22	2.37 \pm 0.85	-0.56	0.89
RFI [7]	84.06 \pm 0.19	1.85 \pm 0.45	83.86 \pm 0.21	2.55 \pm 0.58	-0.20	0.70
Ours	84.81 \pm 0.04	0.02 \pm 0.02	84.34 \pm 0.04	0.08 \pm 0.03	-0.47	0.06

EO [26], and CEO [39]. Table 1 compares the performance of our method against the aforementioned SOTA methods on the Adult, COMPAS, and Synthetic test sets w.r.t. equal opportunity. Compared to the logistic regression (LR) method that does not engage with fairness intervention at all, post-processing algorithms (i.e., EO and CEO) improve fairness, but at the same time sacrifice greatly on the accuracy of the classification. In pre-processing methods, Fairbatch and LBC methods adaptively learn the sample probability or weight in each subgroup and are widely acknowledged for performing better than Cutting and Reweighting methods. However, these two methods suffer from poor generalizability, for they assign the same weight to all samples in each subgroup and neglect the differences among samples. Though our method lies in the category of pre-processing methods, we address a well-balanced trade-off between accuracy and fairness. More importantly, our method proves to be more generalizable. Similarly, in the in-processing algorithms, even though FC, AD, and RFI improve algorithmic fairness by training the algorithms with fairness constraints, their performance does not generalize well at the test set. FairMixup attentively addresses the generalizability problem for fairness measures. Sadly, it causes too much damage to the classification accuracy. In comparison, our method not only addresses the generalizability of fair classifiers, we also preserve a more than acceptable level of accuracy. That is, our method performs exceptionally well in improving equal opportunity on the test set without sacrificing much accuracy. To further evaluate our method, we ran experiments on two other most commonly used fairness measures - equalized odds Table 2 and demographic parity Table 3. The results strongly suggest that our method has optimal fairness performance on all three fairness measures across test sets.

6.1.1 Comparison of Model Performance on Training and Test Sets. The discrepancy in model performance, referred to as 'diff', is calculated by subtracting the performance metrics obtained on training sets from those on test sets. This discrepancy highlights the challenge of generalizing fairness measures, an issue that becomes pronounced in algorithms lauded for their fairness, such as LBC and Fairbatch. For instance, while these algorithms exhibit commendable fairness on training sets—particularly in terms of the equal opportunity metric—their fairness does not consistently extend to the test sets.

Our comparative analysis (using the Adult dataset as an example), focusing on Equal Opportunity (refer to Table 4) and Demographic Parity (refer to Table 5), reveals that our method excels in fairness generalizability, as also illustrated in Figure 2. When considering Equalized Odds (see Figure 6), we find no significant generalizability concerns among the six evaluated methods. This

Table 5. Comparison of model performance on training and test sets in terms of Accuracy and Δ_{DP} on the Adult dataset. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{DP} , denoted by \downarrow .

Method	Training		Test		Difference	
	Accuracy \uparrow	$\Delta_{DP}\downarrow$	Accuracy \uparrow	$\Delta_{DP}\downarrow$	Accuracy Diff	Δ_{DP} Diff
LR	84.90 \pm 0.02	9.32 \pm 0.04	84.50 \pm 0.03	9.13 \pm 0.03	-0.40	-0.19
Fairbatch [44]	83.11 \pm 0.17	0.39 \pm 0.23	82.33 \pm 0.13	1.18 \pm 0.13	-0.78	0.79
LBC [29]	83.00 \pm 0.02	0.01 \pm 0.00	82.35 \pm 0.03	1.01 \pm 0.07	-0.65	1.00
FairMixup [17]	82.61 \pm 0.41	1.01 \pm 0.52	82.46 \pm 0.33	1.45 \pm 0.30	-0.15	0.44
FC [53]	82.48 \pm 0.21	0.89 \pm 0.50	82.22 \pm 0.22	1.45 \pm 0.62	-0.26	0.56
RFI [7]	81.88 \pm 0.53	1.32 \pm 0.47	81.92 \pm 0.41	1.53 \pm 0.56	0.04	0.21
Ours	82.69 \pm 0.08	0.02 \pm 0.01	82.31 \pm 0.02	0.14 \pm 0.03	-0.38	0.12

Table 6. Comparison of model performance on training and test sets in terms of Accuracy and Δ_{EO} on the Adult dataset. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EO} , denoted by \downarrow .

Method	Training		Test		Difference	
	Accuracy \uparrow	$\Delta_{EO}\downarrow$	Accuracy \uparrow	$\Delta_{EO}\downarrow$	Accuracy Diff	Δ_{EO} Diff
LR	84.90 \pm 0.02	9.32 \pm 0.04	84.50 \pm 0.03	9.13 \pm 0.03	-0.40	-0.19
Fairbatch [44]	84.72 \pm 0.12	4.54 \pm 0.22	84.21 \pm 0.10	5.28 \pm 0.43	-0.51	0.74
LBC [29]	84.55 \pm 0.07	3.44 \pm 0.06	84.30 \pm 0.03	3.69 \pm 0.21	-0.25	0.25
FairMixup [17]	81.09 \pm 1.48	5.44 \pm 1.08	80.42 \pm 0.68	5.31 \pm 1.36	-0.67	-0.13
FC [53]	84.52 \pm 0.18	3.82 \pm 0.72	84.10 \pm 0.23	4.33 \pm 0.65	-0.42	0.51
RFI [7]	84.22 \pm 0.31	4.60 \pm 0.46	83.89 \pm 0.25	4.55 \pm 0.53	-0.33	-0.05
Ours	81.66 \pm 0.16	0.89 \pm 0.12	81.49 \pm 0.12	0.87 \pm 0.10	-0.17	-0.02

could be because none of the methods particularly excel in ensuring fairness, as measured by Equalized Odds, even on the training set.

Fairmixup closely follows our method in terms of fairness generalizability. This can be attributed to its inclusion of constraint terms in its algorithm, specifically addressing the critical issue of fairness generalization.

6.1.2 The trade-offs between accuracy and fairness. As depicted in Figure 3, We selected four other algorithms with notable fairness performance and conducted a comparative analysis of the trade-offs between accuracy and three fairness metrics on the Adult datasets. Our method demonstrates superior performance in balancing the trade-offs between accuracy and fairness in both Equal Opportunity and Demographic Parity metrics. In terms of the Equalized Odds metric, as illustrated, our method sacrifices a substantial amount of accuracy when attempting to enhance fairness. We highlight the trade-off effects between accuracy and fairness, calling for future research on algorithmic fairness to investigate these dynamics, especially on the fairness metric - Equalized Odds.




Fig. 3. The trade-offs between accuracy and fairness measures on Adult. The dashed line represents the accuracy and fairness metrics corresponding to the baseline (LR). The upper-right corner of each diagram indicates optimal performance for both accuracy and fairness.

Table 7. A selective case analysis on the Adult dataset. We examine the classification of income groups based on individual attributes. The analysis was specifically centered around cases where $Income > 50K$, employing the Equal Opportunity metric as the basis for the analysis.

Attributes	Case 1	Case 2	Case 3	Case 4
Age	49	40	45	47
Workclass	Private	Private	Private	Self-emp-not-inc
Education	HS-grad	HS-grad	Bachelors	5th-6th
Education-Num	9	9	13	3
Marital-Status	Married-civ-spouse	Married-civ-spouse	Divorced	Married-civ-spouse
Occupation	Craft-repair	Exec-managerial	Other-service	Sales
Relationship	Wife	Husband	Not-in-family	Husband
Race	White	White	White	White
Sex	Female	Male	Male	Male
Capital-Gain	0	0	0	0
Capital-Loss	0	0	0	0
Hours-Per-Week	40	40	45	55
Native-Country	United-States	United-States	?	Italy
Income	>50K	>50K	>50K	>50K
Prediction	<=50K	<=50K	<=50K	<=50K

6.1.3 *Case Study.* We conducted further investigation into instances of misclassification, delving into the reasons persisting despite the implementation of our method (Table 7). We present specific cases to demonstrate that our method achieves the purpose of debias by providing alternative explanations for the misclassification instances, highlighting that bias is not the root cause.

- Case 1: The individual’s education level (HS-Grad) and Occupation (Craft repair) are generally not directly associated with high income, which may contribute to the error of classifying it as " $\leq 50K$ ". Gender is usually a sensitive feature closely associated with high income, in this case, however, its impact on model prediction might be relatively minor compared to that of education and occupation.
- Case 2: Though the individual’s occupation is Exec-managerial, which is typically associated with higher income levels, it may not be typical for individuals in this occupation with an HS-grad education level to have higher income. This could result in the model incorrectly

Table 8. Fine-tuning the Performance of pre-trained models with our method. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EO} , denoted by \downarrow . **Bold** values represent the best performance.

Dataset	Pre-trained model	Method	Accuracy \uparrow	$\Delta_{EO}\downarrow$
UTKFace	ResNet18	Original	89.13 \pm 0.61	9.12 \pm 1.01
		Reweighting [30]	87.70 \pm 0.73	8.06 \pm 0.82
		Fairbatch [44]	89.08 \pm 0.20	6.46 \pm 0.18
		LBC [29]	89.11\pm0.13	6.03 \pm 0.10
		Ours	89.06 \pm 0.55	5.88\pm0.64
MOJI	DeepMoji	Original	72.11 \pm 0.11	28.56 \pm 0.41
		Reweighting [30]	73.29 \pm 0.36	18.03 \pm 0.39
		FairBatch [44]	75.03\pm0.33	10.22 \pm 0.61
		LBC [29]	74.99 \pm 0.21	8.89 \pm 0.30
		Gate [25]	74.55 \pm 0.23	8.63 \pm 0.51
		Ours	74.59 \pm 1.03	8.02\pm1.83

predicting their income as " \leq 50K", especially when high-income individuals in the dataset typically have higher education levels.

- Case 3: Despite being a 45-year-old male with a bachelor's degree, which is usually a positive indicator of higher income, his occupation is listed as "other-service". which is generally considered a low-paying job. In addition, the model is likely to associate a "divorced" marital status and "not-in-family" with lower income, especially when there lacks strong indicators supporting a high-income prediction. Also, the absence of nationality information may prove to be challenging for the model to accurately assess his income, especially when certain nationalities are positively associated with specific economic opportunities.
- Case 4: Despite working in sales and potentially having higher-income potential as a self-employed individual, this 47-year-old male has a lower education level (5th-6th grade), which typically does not align with a high-income group, leading to a potential underestimation of his income ability. Additionally, this individual works in Italy which indicates certain market and economic opportunities, nevertheless, the model may fail to fully understand income potential in different national contexts, especially for self-employed individuals.

6.2 Fine-Tuning Pretrained Unfair Models for Fairness

Having evaluated our method against state-of-the-art algorithms on tabular benchmarks (Table 1, Table 2, and Table 3), we extended our analysis to vision and language data, employing widely recognized models such as ResNet18 [27] and DeepMoji [20]. These models, despite their effectiveness, have raised fairness concerns in their respective domains[41, 44], motivating us to explore fine-tuning techniques to enhance fairness while maintaining model accuracy.

In the comprehensive comparison detailed in Table 8, we juxtaposed our method with the Reweighting approach, Fairbatch [44], and LBC [29]—methods previously validated in tabular contexts but here applied to vision and language tasks. Among these, our method distinguished itself by not just matching the accuracy levels of these established methods but also by leading in fairness metrics, achieving the lowest Δ_{EO} across the board. This singular achievement highlights the capability of our approach to enhance fairness effectively without detracting from the predictive performance of the models. We also explored the Gate [25] method, notable for its application of




Fig. 4. Results obtained for different α and different η on the Adult test set w.r.t. equal opportunity.




Fig. 5. Results obtained for different α and different η on the COMPAS test set w.r.t. equal opportunity.

reweighting and auxiliary networks to improve fairness in language processing, which demonstrated notable improvements as well.

The collective results from our evaluation illuminate the diverse landscape of fairness enhancement methods, each contributing to the reduction of fairness disparities in their own right. Among these, our approach stands out by striking a nuanced balance between fairness and accuracy, which is pivotal in nuanced and multifaceted contexts like AI. This balance not only underscores the adaptability of our method across various data types—from tabular to the more complex vision and language tasks—but also highlights its potential to thoughtfully address fairness concerns within these domains.

6.3 Evaluating Impacts of α and η Parameters

We examine the performance of our adaptive priority reweighing method for accuracy and three fairness measures - equal opportunity (Figures 4, 5, 6), equalized odds (Figures 7, 8, 9), and demographic parity (Figures 10, 11, 12 across three datasets - Adult, COMPAS, and Synthetic dataset). We illustrate the performance using color - the lighter the shade indicates better performance on accuracy and fairness. The adaptive priority reweighing algorithm is governed by two pivotal hyperparameters: the subgroup learning rate, α , and the step size, η . In this section, we delve into the effects of varying these hyperparameters on the algorithm’s performance.

From these figures, we validated that applying adjustment to η and prioritizing those closer to the decision boundary with higher weight improves the generalizability of fair classifiers. More specifically, when $\eta > 0$, our method proves to optimize fairness on the equal opportunity measure




Fig. 6. Results obtained for different α and different η on the Synthetic test set w.r.t. equal opportunity.




Fig. 7. Results obtained for different α and different η on the Adult test set w.r.t. equalized odds.

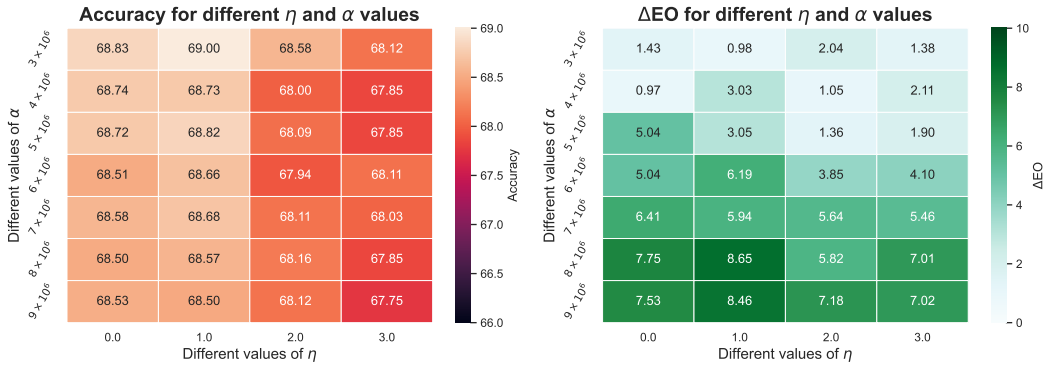


Fig. 8. Results obtained for different α and different η on the COMPAS test set w.r.t. equalized odds.

(Figures 4, 5, 6) as well as on the demographic parity measure (Figures 10, 11, 12) across all datasets. Though the performance of our model on improving fairness in accordance with the equalized odds measure was not evident, especially on the Synthetic dataset, it corresponds with the premises of our study that our method enables fairness improvement to fairness measures on the test dataset

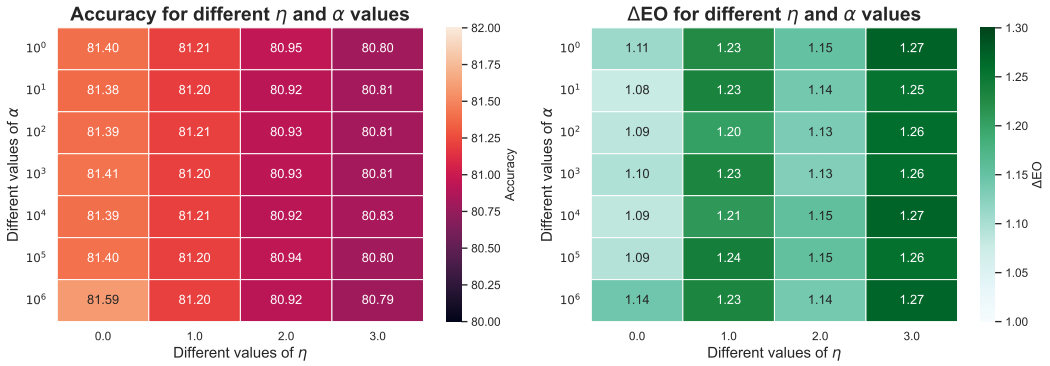


Fig. 9. Results obtained for different α and different η on the Synthetic test set w.r.t. equalized odds.




Fig. 10. Results obtained for different α and different η on the Adult test set w.r.t. demographic parity.




Fig. 11. Results obtained for different α and different η on the COMPAS test set w.r.t. demographic parity.

when they have already performed well on the training dataset. The equalized odds measure (see Table 6 and Figure 3) shows limitations on the training set, as we specified in section 6.1.1.

However, as illustrated in the figures, owing to the fact that increasing η allocates higher weight to samples closer to the decision boundary, augmenting it excessively can result in a decline in both accuracy and fairness performance. Specifically, in the case of equal opportunity, adjusting α shows minimal impact on both accuracy and fairness, whereas η , under appropriate settings, enhances




Fig. 12. Results obtained for different α and different η on the Synthetic test set w.r.t. demographic parity.

Table 9. Performance comparison of different baselines with our method. Our method aims to maximize accuracy, denoted by \uparrow and to minimize Δ_{EO} , denoted by \downarrow .

Dataset	Baselines	Method	Accuracy \uparrow	$\Delta_{EO}\downarrow$
Adult	LR	Original	84.50 \pm 0.03	9.13 \pm 0.03
		Ours	81.49 \pm 0.12	0.87 \pm 0.10
	GBDT	Original	86.53 \pm 0.11	8.73 \pm 0.09
		Ours	84.51 \pm 0.13	0.85 \pm 0.17
	XGBoost	Original	86.53 \pm 0.04	8.11 \pm 0.06
		Ours	84.03 \pm 0.07	0.82 \pm 0.43
UTKFace	ResNet18	Original	89.13 \pm 0.61	9.12 \pm 1.01
		Ours	89.06 \pm 0.55	5.88 \pm 0.64
	ResNet50	Original	89.25 \pm 0.43	8.53 \pm 0.88
		Ours	89.14 \pm 0.36	5.92 \pm 0.61
MOJI	DeepMoji	Original	72.11 \pm 0.11	28.56 \pm 0.41
		Ours	74.59 \pm 1.03	8.02 \pm 1.83
	Bert	Original	72.85 \pm 0.18	31.01 \pm 0.60
		Ours	75.01 \pm 1.22	7.95 \pm 1.46

model’s fairness performance on the test set substantially. Similarly, in cases of equalized odds and demographic parity, η holds a greater influence on accuracy and fairness. However, adjusting α within these fairness measures exhibits a more noticeable impact on fairness.

A cursory observation across all tables indicates that the performance metrics exhibit limited variation with increasing values of α , suggesting a relatively subdued influence of α on the algorithm. Conversely, the step size, η , appears to exert a more pronounced effect. For instance, Figure 4 reveals that optimal fairness is predominantly achieved at $\eta \approx 1.2$ for the Adult dataset. In the case of the COMPAS dataset, a consistent η of 0.4 yields optimal fairness across a spectrum of α values. This behavior underscores that α primarily serves to cushion abrupt shifts during weight iteration updates, while adjusting η accentuates the weights of samples near the decision boundary, bolstering fairness generalization.

Certain α and η combinations yield optimal fairness for specific datasets. A case in point is the Synthetic dataset in Figure 12, where an α of 4×10^8 coupled with η at 0.2 minimizes Δ_{DP} to a value of 0.09. A noteworthy observation is the heightened sensitivity of the COMPAS dataset to alterations in both hyperparameters, manifesting in substantial variations in its performance metrics compared to the Adult and Synthetic datasets. The intricacies in the interplay between α and η and their resultant effect on the metrics across datasets emphasize the nuanced relationship between these parameters and algorithmic performance. Certain parameter configurations manage to strike a harmonious balance, ensuring both high accuracy and commendable fairness.

6.4 Performance Comparison of Different Baselines

The experiments aim to validate the efficacy of our method in improving model fairness across various baselines (Table 9). In addition to the conventional models such as LR, ResNet18, and DeepMoj, we further explored higher-performing baselines in tabular (i.e., GBDT [23] and XGBoost [14]), vision (i.e., ResNet50 [27]), and language (i.e., Bert [31]) dataset. Leveraging these more advanced baselines, our method consistently achieves improvement in classification accuracy while maintaining fairness performance. Especially on tree models such as GBDT and XGBoost, our method demonstrates significant improvements compared to the original models. Maintaining equivalent fairness performance, our method outperforms LR by a notable three percent.

7 CONCLUSION

This paper responds to the calls for fair algorithms that generalize well across the test set by designing a novel reweighting method – the adaptive priority reweighting method. This method eliminates the impact of distribution shifts between training and testing data on model generalizability, a conundrum constantly faced with equal reweighting methods.

The performance of our method for accuracy and fairness measures (i.e., equal opportunity, equalized odds, and demographic parity) is evaluated through extensive experiments on tabular datasets. We further highlight the performance of our method in improving the generalizability of fair classifiers by experimenting with both language and vision benchmarks. We believe that our method performs extremely well in improving fairness and will show promising results in improving the fairness of any pre-trained models simply via fine-tuning.

REFERENCES

- [1] Tameem Adel, Isabel Valera, Zoubin Ghahramani, and Adrian Weller. 2019. One-network adversarial fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2412–2420.
- [2] Sina Aghaei, Mohammad Javad Azizi, and Phebe Vayanos. 2019. Learning optimal and fair decision trees for non-discriminative decision-making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1418–1426.
- [3] Matthew Almeida, Yong Zhuang, Wei Ding, Scott E Crouter, and Ping Chen. 2021. Mitigating class-boundary label uncertainty to reduce both model bias and variance. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 2 (2021), 1–18.
- [4] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. In *Ethics of data and analytics*. Auerbach Publications, 254–264.
- [5] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2022. Machine bias. In *Ethics of data and analytics*. Auerbach Publications, 254–264.
- [6] Martin Anthony, Peter L Bartlett, Peter L Bartlett, et al. 1999. *Neural network learning: Theoretical foundations*. Vol. 9. cambridge university press Cambridge.
- [7] Sina Baharlouei, Maher Nouiehed, Ahmad Beirami, and Meisam Razaviyayn. 2019. Rényi Fair Inference. In *International Conference on Learning Representations*.
- [8] Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. 2005. Local rademacher complexities. (2005).
- [9] Peter L Bartlett and Shahar Mendelson. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3, Nov (2002), 463–482.

- [10] Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2017. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409* (2017).
- [11] Alex Beutel, Jilin Chen, Tulse Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in recommendation ranking through pairwise comparisons. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2212–2220.
- [12] Su Lin Blodgett, Lisa Green, and Brendan O'Connor. 2016. Demographic dialectal variation in social media: A case study of African-American English. *arXiv preprint arXiv:1608.08868* (2016).
- [13] Junyi Chai and Xiaoqian Wang. 2022. Fairness with Adaptive Weights. In *International Conference on Machine Learning*. PMLR, 2853–2866.
- [14] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [15] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5, 2 (2017), 153–163.
- [16] Xu Chu, Yujie Jin, Wenwu Zhu, Yasha Wang, Xin Wang, Shanghang Zhang, and Hong Mei. 2022. DNA: Domain Generalization with Diversified Neural Averaging. In *International Conference on Machine Learning*. PMLR, 4010–4034.
- [17] Ching-Yao Chuang and Youssef Mroueh. 2021. Fair mixup: Fairness via interpolation. *arXiv preprint arXiv:2103.06503* (2021).
- [18] Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. 2018. Empirical risk minimization under fairness constraints. *Advances in neural information processing systems* 31 (2018).
- [19] Rachel K. E. Bellamy et al. 2018. AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias. <https://arxiv.org/abs/1810.01943>
- [20] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. *stat* 1050 (2017), 1.
- [21] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 259–268.
- [22] Benjamin Fish, Jeremy Kun, and Ádám D Lelkes. 2016. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM International Conference on Data Mining*. SIAM, 144–152.
- [23] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [24] David García-Soriano and Francesco Bonchi. 2021. Maxmin-fair ranking: individual fairness under group-fairness constraints. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 436–446.
- [25] Xudong Han, Timothy Baldwin, and Trevor Cohn. 2022. Balancing out Bias: Achieving Fairness Through Balanced Training. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 11335–11350.
- [26] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, Vol. 29.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [28] Taeuk Jang, Feng Zheng, and Xiaoqian Wang. 2021. Constructing a fair classifier with generated fair data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 7908–7916.
- [29] Heinrich Jiang and Ofir Nachum. 2020. Identifying and correcting label bias in machine learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 702–712.
- [30] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* 33, 1 (2012), 1–33.
- [31] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*. 4171–4186.
- [32] Emmanouil Kerasanakis, Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, and Yiannis Kompatsiaris. 2018. Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. In *Proceedings of the 2018 world wide web conference*. 853–862.
- [33] Peizhao Li and Hongfu Liu. 2022. Achieving fairness at no utility cost via data reweighting with influence. In *International Conference on Machine Learning*. PMLR, 12917–12930.
- [34] Moshe Lichman et al. 2013. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml> 40 (2013).
- [35] Tongliang Liu and Dacheng Tao. 2015. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence* 38, 3 (2015), 447–461.
- [36] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2018. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*. PMLR, 3384–3393.
- [37] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.

- [38] Dana Pessach and Erez Shmueli. 2023. Algorithmic fairness. In *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*. Springer, 867–886.
- [39] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On fairness and calibration. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 5684–5693.
- [40] Sreeraj Ramachandran and Ajita Rattani. 2022. Deep generative views to mitigate gender classification bias across gender-race groups. *arXiv preprint arXiv:2208.08382* (2022).
- [41] Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection. *arXiv preprint arXiv:2004.07667* (2020).
- [42] Yuji Roh, Kangwook Lee, Steven Whang, and Changho Suh. 2020. Fr-train: A mutual information-based approach to fair and robust training. In *International Conference on Machine Learning*. PMLR, 8147–8157.
- [43] Yuji Roh, Kangwook Lee, Steven Whang, and Changho Suh. 2021. Sample selection for fair and robust training. *Advances in Neural Information Processing Systems* 34 (2021), 815–827.
- [44] Yuji Roh, Kangwook Lee, Steven Euijong Whang, and Changho Suh. 2020. FairBatch: Batch Selection for Model Fairness. In *International Conference on Learning Representations*.
- [45] Vladimir Vapnik. 1999. *The nature of statistical learning theory*. Springer science & business media.
- [46] Mingyang Wan, Daochen Zha, Ninghao Liu, and Na Zou. 2023. In-processing modeling techniques for machine learning fairness: A survey. *ACM Transactions on Knowledge Discovery from Data* 17, 3 (2023), 1–27.
- [47] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. 2019. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF international conference on computer vision*. 5310–5319.
- [48] Hongyi Wen, Xinyang Yi, Tiansheng Yao, Jiayi Tang, Lichan Hong, and Ed H Chi. 2022. Distributionally-robust Recommendations for Improving Worst-case User Experience. In *Proceedings of the ACM Web Conference 2022*. 3606–3610.
- [49] Blake Woodworth, Suriya Gunasekar, Mesrob I Ohannessian, and Nathan Srebro. 2017. Learning non-discriminatory predictors. In *Conference on Learning Theory*. PMLR, 1920–1953.
- [50] Depeng Xu, Yongkai Wu, Shuhan Yuan, Lu Zhang, and Xintao Wu. 2019. Achieving causal fairness through generative adversarial networks. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- [51] Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. 2018. Fairgan: Fairness-aware generative adversarial networks. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 570–575.
- [52] Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2019. Modeling Tabular data using Conditional GAN. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [53] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P Gummadi. 2017. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*. PMLR, 962–970.
- [54] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning fair representations. In *International conference on machine learning*. PMLR, 325–333.
- [55] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 335–340.
- [56] Tao Zhang, Tianqing Zhu, Jing Li, Mengde Han, Wanlei Zhou, and S Yu Philip. 2020. Fairness in semi-supervised learning: Unlabeled data help to reduce discrimination. *IEEE Transactions on Knowledge and Data Engineering* 34, 4 (2020), 1763–1774.
- [57] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyang Shen. 2021. Deep stable learning for out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5372–5382.
- [58] Zhifei Zhang, Yang Song, and Hairong Qi. 2017. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5810–5818.
- [59] Chang Zhou, Jianxin Ma, Jianwei Zhang, Jingren Zhou, and Hongxia Yang. 2021. Contrastive learning for debiased candidate generation in large-scale recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3985–3995.
- [60] Xiao Zhou, Yong Lin, Renjie Pi, Weizhong Zhang, Renzhe Xu, Peng Cui, and Tong Zhang. 2022. Model agnostic sample reweighting for out-of-distribution learning. In *International Conference on Machine Learning*. PMLR, 27203–27221.

Received 14 September 2023; revised 12 February 2024; accepted 14 May 2024