

Demographic-Agnostic Fairness without Harm

Zhongteng Cai, Mohammad Mahdi Khalili, Xueru Zhang

Department of Computer Science and Engineering
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210-1277 USA
{cai.1125, khaliligarekani.1, zhang.12807}@osu.edu

Abstract

As machine learning (ML) algorithms are increasingly used in social domains to make predictions about humans, there is a growing concern that these algorithms may exhibit biases against certain social groups. Numerous notions of fairness have been proposed in the literature to measure the unfairness of ML. Among them, one class that receives the most attention is *parity-based*, i.e., achieving fairness by equalizing treatment or outcomes for different social groups. However, achieving parity-based fairness often comes at the cost of lowering model accuracy and is undesirable for many high-stakes domains like healthcare. To avoid inferior accuracy, a line of research focuses on *preference-based* fairness, under which any group of individuals would experience the highest accuracy and collectively prefer the ML outcomes assigned to them if they were given the choice between various sets of outcomes. However, these works assume individual demographic information is known and fully accessible during training. In this paper, we relax this requirement and propose a novel *demographic-agnostic fairness without harm (DAFH)* optimization algorithm, which jointly learns a group classifier that partitions the population into multiple groups and a set of decoupled classifiers associated with these groups. Theoretically, we conduct sample complexity analysis and show that our method can outperform the baselines when demographic information is known and used to train decoupled classifiers. Experiments on both synthetic and real data validate the proposed method.

1 Introduction

Machine learning (ML) and automated decision-making systems trained with real-world data can have inherent bias and exhibit discrimination against certain social groups (Barocas, Hardt, and Narayanan 2023; Zhang et al. 2019; Zhang, Khalili, and Liu 2020; Zhang et al. 2020; Zhang and Liu 2021). One common approach to mitigating the issue is to impose certain fairness constraints while making ML predictions. Among all fairness notions proposed in the literature, a class that has received the most attention is based on parity (equality) in treatment or outcomes for different social groups, e.g., *demographic parity* (Dwork et al. 2012) and *equalized odds* (Hardt, Price, and Srebro 2016) that require (true/false) positive rates to be equalized across differ-

ent groups. However, enforcing these parity-based fairness notions often comes at the cost of lowering model accuracy, resulting in an inherent trade-off between fairness and accuracy (Menon and Williamson 2018; Pang et al. 2024; Khalili, Zhang, and Abroshan 2021, 2023; Zuo, Khalili, and Zhang 2023; Pham, Zhang, and Zhang 2023; Zuo et al. 2024; Zuo, Zhu, and Khalili 2025; Abroshan, Elliott, and Khalili 2024). In safety-critical domains such as healthcare, sacrificing predictive accuracy in exchange for fairness is undesirable as it violates both *beneficence* (i.e., do what is best for patients) and *non-maleficence* (i.e., do no harm) principles required in healthcare ethics (Beauchamp and Childress 1994).

To avoid harming model accuracy, a few studies focus on *preference-based* fairness (Zafar et al. 2017; Ustun, Liu, and Parkes 2019). The idea is to train decoupled ML models for different social groups so that each group of individuals would collectively prefer the model assigned to them to other models. In other words, if individuals were allowed to choose among a set of ML models, a majority of them would still stick to the assigned model. Unlike parity-based fairness, preference-based fairness aims to achieve *fairness without harm* and avoids inferior accuracy for each group.

In this paper, we focus on a preference-based fairness notion similar to the one proposed in (Ustun, Liu, and Parkes 2019), which requires that the decoupled ML models trained for different social groups satisfy two properties: *rationality* and *envy-freeness*. The former ensures that the accuracy of the decoupled model evaluated on the associated group is higher than that of the pooled model trained with data from all groups, while the latter ensures each group’s model has better accuracy than the models assigned to other groups. In addition to (Ustun, Liu, and Parkes 2019), some subsequent works proposed other variants of preference-based fairness and developed algorithms to train models satisfying the proposed fairness, which we discuss more in Section 2.

However, existing methods for achieving preference-based fairness typically assume that individual demographic information (i.e., sensitive attributes such as race and gender) is available during training or can be inferred from side information (Zhu et al. 2023), so that the decoupled models can be trained on the groups partitioned based on the sensitive attributes (Ustun, Liu, and Parkes 2019). Unfortunately, individuals’ sensitive attributes are not always accessible in practice, either because regulations mandate par-

Related Works	Preference-based	No harm	Without sensitive attributes	Intersectionality
Zafar et al. (2017)	✓	✓	✗	✗
Ustun, Liu, and Parkes (2019)	✓	✓	✗	✓
Lahoti et al. (2020)	✗	✗	✓	✓
Islam, Chen, and Cai (2024)	✗	✗	✓	✗
Pang et al. (2024)	✗	✓	✓	✗
DAFH (Ours)	✓	✓	✓	✓

Table 1: Comparison with related works: Zafar et al. (2017) and Ustun, Liu, and Parkes (2019) consider preference-based fairness to avoid harm on specific groups, but both require knowledge of sensitive attributes to generate group partition. Islam, Chen, and Cai (2024); Lahoti et al. (2020) consider parity-based fairness without accessing sensitive attributes, but they violate no harm constraint. Pang et al. (2024) introduces fairness-without-harm when sensitive attributes are inaccessible, but based on parity-based fairness metric. In contrast, our work aims to avoid harm without accessing sensitive attributes, and also deals with the intersectionality caused by the existence of multiple sensitive attributes.

ity of treatment, prohibiting the use of sensitive attributes in decision-making (Barocas and Selbst 2016), or due to user privacy concerns that prevent the disclosure of such attributes. For example, when ML models are trained on patient data from multiple hospitals, protected information must be de-identified to comply with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) (U.S. Congress 1996). Similarly, in applications like speech or face recognition, sensitive attributes such as race and gender may be unknown, yet it remains desirable to achieve high accuracy across all demographic groups.

Even when sensitive attributes are available, partitioning groups based on them may not lead to more accurate decoupled models. Because individuals are often associated with multiple sensitive attributes such as gender, race, and age, models that are fair with respect to one sensitive attribute can easily violate fairness when assessed on more fine-grained intersectional groups (Kearns et al. 2018). Training decoupled classifiers for all intersectional groups is also unrealistic because each intersectional can have limited data samples and the resulting decoupled classifier can easily overfit. More importantly, sensitive attributes alone cannot adequately capture the full heterogeneity of the population. This motivates a central question: *Can we identify a better group partition under which decoupled models satisfy preference-based fairness, without accessing sensitive attributes?*

To achieve this, we propose a *demographic-agnostic fairness without harm (DAFH)* optimization algorithm that jointly learns (i) a *group classifier* that finds the group partition for fairness without harm constraint (including both rationality and envy-freeness); and (ii) a set of *decoupled classifiers* with each a model assigned to a group partitioned based on the group classifier. Indeed, we can consider our approach as extending the previous methods of partitioning groups based on sensitive attributes to a more complex representation-based partition. Because our approach, by automatically learning individual representations, can more effectively capture population heterogeneity than fixed sensitive attributes. We theoretically show that our method can match or surpass the performance of trivial group partitioning based on sensitive attributes.

A comparison between our approach and representative

related works is summarized in Table 1. Our main contributions are as follows:

1. We propose the measure of fairness without harm when sensitive attributes are inaccessible, which is defined based on *rationality* and *envy-freeness* criteria.
2. We formulate a *demographic-agnostic fairness without harm (DAFM)* optimization algorithm with a novel objective function directly measuring the rationality and envy-freeness, which find group partition and decoupled classifiers that yield higher probability of satisfying fairness without harm.
3. We theoretically show that our algorithm can better capture the data heterogeneity and achieve the same and even better performance than trivially partitioning the dataset based on sensitive attribute.
4. We conduct experiments on synthetic and real datasets, and show that our algorithm can outperform baselines which have access to sensitive attributes.

2 Related Work

Preference-based fairness. Zafar et al. (2017) propose the first preference-based fairness notion by introducing the concept of envy-freeness from economics (Berliant, Thomson, and Dunz 1992; Varian 1973), which implies that no participant prefers the resources assigned to others. The resource is defined as the decoupled classifiers, and the utility of classifiers is often defined as the accuracy of the assigned classifier. Ustun, Liu, and Parkes (2019) propose using a decision tree to partition the group according to multiple sensitive attributes, until the divided group can best satisfy fairness without harm. Hossain, Mladenovic, and Shah (2020) defines similar concepts of group envy-freeness and equitability based on the utility of classification outcomes, and allow groups to be defined arbitrarily, not just based on sensitive attributes. Saito and Joachims (2022) formulate fairness in ranking as a resource allocation problem, where each individual must be assigned a position that satisfies two key criteria: envy-freeness and dominance over a uniform ranking, i.e., no individual will benefit by swapping positions with another or by switching to a uniform random ranking policy. Shah (2023) provides a survey of preference-based

notions of fairness rooted in the literature of computational social choice, including envy-freeness, and emphasize their relevance to various decision-making applications such as recommendation and classification.

Fairness without demographic data. Achieving fairness without accessing sensitive attributes is crucial when individual demographic data is not readily available (Andrus et al. 2021). However, the research on this topic has been largely focused on parity-based fairness (Ashurst and Weller 2023); this is fundamentally different from our work which aims to achieve preference-based fairness without accessing sensitive attributes. To attain fairness in the absence of demographic data, many studies attempt to infer sensitive attributes, e.g., by learning proxy models from a separate training dataset that includes demographic data, in order to predict the sensitive attributes in the target dataset for downstream tasks like debiasing (Grari, Lamprier, and Detyniecki 2021; Islam, Chen, and Cai 2024). However, empirical results have shown that relying on proxy variables to measure fairness can lead to biased outcomes (Zhu et al. 2023; Awasthi et al. 2021). Other approaches minimize the risk of disparity when group partitions are unknown (Lahoti et al. 2020; Hashimoto et al. 2018), akin to our method, which also does not require an additional dataset containing demographic information. Pang et al. (2024) proposed a parity-based fairness without harm, where new data samples without sensitive attributes are actively acquired during training to mitigate fairness parity. While they avoid harming accuracy by filtering fresh samples, we prevent harm by assigning decoupled classifiers to each group. Some works leverage properties of representations or performance on downstream tasks to identify group partition. Sohoni et al. (2020) increases the worst-case accuracy for subclasses with unknown labels by leveraging the separable representations in the feature space. Veldanda et al. (2024) uses incorrectly classified samples as proxies for disadvantaged groups. While these works focus on parity-based fairness, their method can be integrated with our framework to guide the design of our group classifier.

3 Problem Statement

Consider a dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ with n independent samples drawn from the same joint distribution $(x_i, y_i) \sim P(X, Y)$, where $x_i \in \mathcal{X}$ are individual’s features and $y_i \in \mathcal{Y} := \{\pm 1\}$ is the binary target variable to be predicted. We consider practical settings where individuals’ demographic information (i.e., sensitive attributes such as race, gender, and age) are unknown and not included in \mathcal{D} .

Since individuals’ sensitive attributes are unknown, we introduce a *group classifier* $\theta : \mathcal{X} \rightarrow \mathcal{K}$ from hypothesis class \mathcal{H}^G that maps the input features x_i to group label $k \in \mathcal{K} := [K]$. For samples assigned to each group k , let $h_k : \mathcal{X} \rightarrow \mathcal{Y}$ be the corresponding *decoupled classifier* from hypothesis class \mathcal{H}^D that predicts the target variable based on the input feature X . Given a classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$, its empirical risk $\hat{R}(h)$ and true risk $R(h)$ associated with 0-1

loss are defined respectively as follows:

$$\begin{aligned}\hat{R}(h) &= \frac{1}{n} \sum_{i=1}^n \mathbb{I}(h(x_i) \neq y_i); \\ R(h) &= \mathbb{E}_{X, Y} [\mathbb{I}(h(X) \neq Y)],\end{aligned}$$

where \mathbb{I} is the indicator function. Further define the *pooled classifier* h_0 as the classifier from \mathcal{H}^D that minimizes the true risk, i.e., $h_0 = \arg \min_h R(h)$. Given a classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$, define the group-specific empirical risk as the empirical risk evaluated on samples assigned to group k , i.e.,

$$\hat{R}_k(h) = \frac{1}{n_k} \sum_{i=1}^n [\mathbb{I}(\theta(x_i) = k) \cdot \mathbb{I}(h(x_i) \neq y_i)],$$

where n_k is the number of samples assigned to group k . Likewise, denote the group-specific true risk of group k as

$$R_k(h) = \mathbb{E}_{X, Y | \theta(X)=k} [\mathbb{I}(h(X) \neq Y)].$$

Our goal is to train the group classifier θ and a set of decoupled classifiers $\{h_k\}_{k \in [K]}$ that satisfy fairness without harm requirement, as defined below.

Definition 1 (Fairness without harm). *A group classifier θ and decoupled classifiers $\{h_k\}_{k \in [K]}$ satisfy fairness without harm if the resulting group-specific true risks satisfy both **rationality** and **envy-freeness** properties:*

1. **Rationality:** $R_k(h_k) \leq R_k(h_0), \forall k \in [K]$.
2. **Envy-freeness:** $R_k(h_k) \leq R_k(h_j), \forall k, j \in [K]$.

Definition 1 implies that for individuals in each group k , the true risk is lowest under the associated decoupled classifier h_k . In other words, if individuals are given opportunities to choose their own classifiers from h_0 and $\{h_k\}_{k \in [K]}$, a *majority* of individuals in each group will prefer their assigned decoupled classifiers to the pooled classifier (rationality) and the decoupled classifiers assigned to other groups (envy-freeness), given their preference to lower true risk. We adopt the same definitions of rationality and envy-freeness as presented in (Ustun, Liu, and Parkes 2019), with one key distinction: we define these notions over group partitions induced by a learned group classifier, rather than by sensitive attributes. It is worth noting that our definition does not require every individual to prefer their assigned classifier.

We approximate the degree of *rationality* violation by evaluating the difference in empirical risk between the pooled classifier and the assigned classifier for each group. Similarly, we estimate the degree of *envy-freeness* violation by examining the difference in empirical risk between the decoupled classifiers for each group. The goal is to find group classifier θ and decoupled classifiers $\{h_k\}_{k \in [K]}$ that maximize the probability that individuals satisfy rationality or envy-freeness, as defined in Definition 1. We formulate this as the following optimization:

$$\max_{\theta, \{h_k\}} \frac{1}{K} \sum_{k=1}^K \left\{ \underbrace{\left[\hat{R}_k(\hat{h}_0) - \hat{R}_k(h_k) \right]}_{\text{Rationality}} + \underbrace{\frac{1}{K} \sum_{j=1}^K \left[\hat{R}_k(h_j) - \hat{R}_k(h_k) \right]}_{\text{Envy-freeness}} \right\}, \quad (1)$$

where $\hat{h}_0 = \arg \min_h \hat{R}(h)$ is the classifier minimizing the empirical risk over all samples. Although the objective function (1) is a direct combination of two criteria (rationality

and envy-freeness), optimizing it is non-trivial due to the following challenges: (i) Objective (1) is a function of parameters in group classifier and decoupled classifiers; both must be trained jointly since the samples used for training each decoupled classifier depend on the outputs of the group classifier. (ii) Objective (1) is a highly non-convex and non-differentiable function, as it involves differences in empirical risks across multiple pairs of groups.

4 Proposed Algorithm

Next, we introduce our approach that solves optimization (1). We first derive a lower bound on the objective function (1) that is easier to optimize using the training data. Then, we approximate this new objective function using differentiable functions and find the optimal group classifier θ and decoupled classifiers $\{h_k\}_{k \in [K]}$ iteratively using gradient ascents.

Simplify objective function (1) with a lower bound. Given group classifier $\theta : \mathcal{X} \rightarrow [K]$ and training dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, denote $\pi_i^k := \mathbb{I}(\theta(x_i) = k)$ as the indicator of group assignment for i -th sample. We have $\sum_{k=1}^K \pi_i^k = 1, \forall i$. Let $L_i^0 := \mathbb{I}(h_0(x_i) \neq y_i)$ and $L_i^k := \mathbb{I}(h_k(x_i) \neq y_i)$ denote the losses of the sample (x_i, y_i) under pooled classifier h_0 and decoupled classifier h_k , respectively. Then we can have the following proposition:

Proposition 1. *By replacing n_k (the number of samples assigned to group k by θ) in (1) with n (the total number of samples), we obtain the the following lower bound on the objective function (1):*

$$\underline{\text{obj}} + \text{constant terms}$$

where we define

$$\underline{\text{obj}} = \frac{1}{nK^2} \sum_{i=1}^n \sum_{k=1}^K (L_i^k - 2K\pi_i^k L_i^k) \quad (2)$$

Proof. The objective function in (1) can be equivalently written as:

$$\frac{1}{K} \sum_{k=1}^K \left[\frac{1}{n_k} \sum_{i=1}^n \pi_i^k (L_i^0 - L_i^k) + \frac{1}{K} \sum_{j=1}^K \frac{1}{n_k} \sum_{i=1}^n \pi_i^k (L_i^j - L_i^k) \right].$$

Since the decoupled classifiers in each group can achieve at least the same accuracy as the pooled classifier and decoupled classifiers assigned to other groups, assume each group satisfies rationale and envy-freeness, therefore it has the following lower bound:

$$\begin{aligned} & \frac{1}{K} \sum_{k=1}^K \left[\frac{1}{n} \sum_{i=1}^n \pi_i^k (L_i^0 - L_i^k) + \frac{1}{K} \sum_{j=1}^K \frac{1}{n} \sum_{i=1}^n \pi_i^k (L_i^j - L_i^k) \right] \\ &= \frac{1}{n} \sum_{i=1}^n \left\{ \frac{1}{K} \sum_{k=1}^K \left[\pi_i^k (L_i^0 - L_i^k) + \frac{1}{K} \sum_{j=1}^K \pi_i^k (L_i^j - L_i^k) \right] \right\} \\ &= \frac{1}{Kn} \sum_{i=1}^n \sum_{k=1}^K \left(\pi_i^k L_i^0 - 2\pi_i^k L_i^k + \frac{1}{K} \pi_i^k \sum_{j=1}^K L_i^j \right) \\ &= \frac{1}{Kn} \sum_{i=1}^n L_i^0 + \frac{1}{Kn} \sum_{i=1}^n \left(-2 \sum_{k=1}^K \pi_i^k L_i^k + \frac{1}{K} \sum_{j=1}^K L_i^j \right) \\ &= \underbrace{\frac{1}{Kn} \sum_{i=1}^n L_i^0}_{\text{constant term}} + \underbrace{\frac{1}{nK^2} \sum_{i=1}^n \sum_{k=1}^K (L_i^k - 2K\pi_i^k L_i^k)}_{\text{obj}} \end{aligned}$$

Then we can prove the proposition. \square

It is worth noting that we replace n_k in (1) with n to improve optimization stability. Since n_k appears in the denominator of (1), it introduces non-differentiability and requires approximation—an issue that becomes particularly unstable when group sizes are small. Therefore, rather than directly optimizing the objective (1), we optimize the surrogate objective (2), which serves as a valid and more stable proxy.

Discussion on objective function. It is worth noting that the objective function in (2) is related but different from the overall accuracy. Specifically, let the overall accuracy of decoupled classifiers on the assigned groups be defined as:

$$\text{acc} = \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K \pi_i^k (1 - L_i^k). \quad (3)$$

Then, maximizing (2) is equal to maximizing the following:

$$\underline{\text{obj}} = \frac{2}{K} \text{acc} + \frac{1}{nK^2} \sum_{i=1}^n \sum_{k=1}^K L_i^k - \frac{2}{K} \quad (4)$$

Compared with the overall accuracy defined in (3), our objective function (2) also considers the performance of decoupled classifiers on groups that they are not assigned to. As we will show in Section 5 and 6, the group classifier θ found under our method better captures the data heterogeneity and can result in more accurate decoupled classifiers.

Approximation using differentiable functions. We further use differentiable functions to approximate (2). Specifically, instead of hard group indicator $\pi_i^k = \mathbb{I}(\theta(x_i) = k)$, we consider soft assignment and define $\tilde{\pi}_i^k = \Pr(\theta(x_i) = k)$ as the predicted probability that a sample belonging to group k . Let group classifier θ be a neural network which softmax layer its outputs, then we use k -th output of its softmax layer (denoted as $\theta(x)[k]$) to compute $\tilde{\pi}_i^k$. Similarly, instead of 0-1 loss L_i^k , we use $h_k(x)[1]$ to denote the predicted probability

Algorithm 1: Stochastic gradient ascent to optimize $\widetilde{\text{obj}}$

1: **Input:** Training data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, differentiable function $\widetilde{\text{obj}}$, learning rate α , hyperparameter λ
2: Initialize group and decoupled classifiers $\theta, \{h_k\}_{k \in [K]}$
3: **while** obj not converge **do**
4: **for** $i = 1, \dots, n$ **do**
5: $\theta \leftarrow \theta + \alpha \nabla_{\theta} \widetilde{\text{obj}}(\theta, \{h_k\}_{k \in [K]}, x_i, y_i, \lambda)$
6: **for** $k = 1, \dots, K$ **do**
7: $h_k \leftarrow h_k + \alpha \nabla_{h_k} \widetilde{\text{obj}}(\theta, \{h_k\}_{k \in [K]}, x_i, y_i, \lambda)$
8: **end for**
9: **end for**
10: **end while**
11: **Output:** $\theta, \{h_k\}_{k \in [K]}$

that a sample x has label 1 and approximate 0-1 loss with a Sigmoid function \widetilde{L}_i^k as follows.

$$\begin{aligned} \widetilde{L}_i^k &= \text{Sigmoid} \left(h_k(x_i)[1] - \frac{1}{2} \right) - y_i, \\ \widetilde{\pi}_i^k &= \theta(x_i)[k]. \end{aligned}$$

During the empirical studies in Section 6, we observe that the group classifier learned based on $\widetilde{\pi}_i^k, \widetilde{L}_i^k$ tends to assign a majority of training samples to the same group. To enforce balanced group size, we add a penalty term to the objective function, which is calculated as the KL-divergence between the group classifier output and uniform distribution, denoted as $\Gamma(\theta)$, times a hyperparameter λ . Combine the above, we have the following:

$$\widetilde{\text{obj}} = \frac{1}{nK^2} \sum_{i=1}^n \sum_{k=1}^K \left(\widetilde{L}_i^k - 2K\widetilde{\pi}_i^k \widetilde{L}_i^k \right) + \lambda \Gamma(\theta). \quad (5)$$

Since $\widetilde{\text{obj}}$ is a differentiable function of parameters in θ and h_k , we can find optimal parameters that yield the highest degree of fairness without harm using stochastic gradient ascent as shown in Algorithm 1. Here, we slightly abuse the notation by using θ and h_k denote the model parameters associated with group classifier and decoupled classifier. Let $\widehat{\text{obj}}(\theta, \{h_k\}_{k \in [K]}, x_i, y_i, \lambda)$ be calculated as in (5).

5 Theoretical Analysis

This section theoretically analyzes the performance of the proposed method. While our algorithm is trained by maximizing empirical performance on training data, we conduct analysis to show the generalization ability. We also show that $\{h_k\}_{k \in [K]}$ trained under our method are guaranteed to outperform decoupled classifiers trained by partitioning groups based on sensitive attributes. Unless otherwise stated, we focus on hyperparameter $\lambda = 0$ in (5).

5.1 Generalization

We first show that the difference between objective functions calculated through empirical risks and true risks is upper bounded. Define the true degree of fairness without harm

as:

$$\text{obj} = \frac{1}{K} \sum_{k=1}^K \left\{ [R_k(h_0) - R_k(h_k)] + \frac{1}{K} \sum_{j=1}^K [R_k(h_j) - R_k(h_k)] \right\},$$

while the empirical degree of fairness is the objective function of (1) defined as:

$$\widehat{\text{obj}} = \frac{1}{K} \sum_{k=1}^K \left\{ [\widehat{R}_k(h_0) - \widehat{R}_k(h_k)] + \frac{1}{K} \sum_{j=1}^K [\widehat{R}_k(h_j) - \widehat{R}_k(h_k)] \right\},$$

Assume that the group size decided by the group classifier has a lower bound, i.e. $\forall k \in [K], n_k \geq \underline{n}$. We now show that the difference between obj and $\widehat{\text{obj}}$ is upper bounded.

Theorem 1. *Let n_k be the number of samples assigned to group k under group classifier θ and denote the lower bound of group size as \underline{n} . For any $\epsilon > 0$, with probability at least $1 - 2K(K+1) \exp\{-2\epsilon^2 \underline{n}/9\}$, we have*

$$|\text{obj} - \widehat{\text{obj}}| \leq \epsilon.$$

Proof. For any hypothesis $h \in \mathcal{H}^{\mathcal{D}}$, denote $|R_k(h) - \widehat{R}_k(h)|$ as $\Delta_k(h)$. Then we have:

$$|\text{obj} - \widehat{\text{obj}}| \leq \frac{1}{K} \sum_{k=1}^K \left[\Delta_k(h_0) + 2\Delta_k(h_k) + \frac{1}{K} \sum_{j=1}^K \Delta_k(h_j) \right]$$

According to Hoeffding's Inequality (Shalev-Shwartz and Ben-David 2014), for any $\epsilon > 0$, we have:

$$\Pr [\Delta_k(h) > \epsilon] \leq 2 \exp\{-2\epsilon^2 n_k\} \leq 2 \exp\{-2\epsilon^2 \underline{n}\}.$$

$\forall k \in [K], j \in \{0, \dots, K\}$, denote $\Delta_k(h_j) > \epsilon$ as event A_{kj} , then we have the following:

$$\begin{aligned} \Pr \left[\bigcap_{k,j} \overline{A_{kj}} \right] &= \Pr \left[\overline{\bigcup_{k,j} A_{kj}} \right] = 1 - \Pr \left[\bigcup_{k,j} A_{kj} \right] \\ &\geq 1 - \sum_{k=1}^K \sum_{j=0}^K \Pr [A_{kj}] \geq 1 - 2K(K+1) \exp\{-2\epsilon^2 \underline{n}\}. \end{aligned}$$

Event $\bigcap_{k,j} \overline{A_{kj}}$ means that $\forall k, j, \Delta_k(h_j) \leq \epsilon$, which indicates that $|\text{obj} - \widehat{\text{obj}}| \leq 3\epsilon$, hence we can know that:

$$\begin{aligned} \Pr \left[|\text{obj} - \widehat{\text{obj}}| \leq 3\epsilon \right] &\geq \Pr \left[\bigcap_{k,j} \overline{A_{kj}} \right] \\ &\geq 1 - 2K(K+1) \exp\{-2\epsilon^2 \underline{n}\}. \end{aligned}$$

Now we can prove the theorem by replacing ϵ with $\frac{\epsilon}{3}$. \square

5.2 Feasibility

Suppose each individual has a hidden sensitive attribute $S \in \mathcal{S}$ indicating the demographic information (e.g., race, gender, age). We then prove that our algorithm can achieve at least the same performance as directly using the sensitive attribute S to divide the dataset.

Specifically, denote $P_k := \Pr(X, Y | \theta(X) = k)$ as the data distribution of group k determined under group classifier θ , and let $P_k^* := \Pr(X, Y | S = k)$ be the distribution of group k partitioned based on sensitive attribute S . For any two classifiers $h, h' \in \mathcal{H}^D$, let the difference between their performance on P_k and P_k^* be defined as $D_k(h, h') := \mathbb{E}_{X \sim P_k} [\mathbb{I}(h(X) \neq h'(X))]$ and $D_k^*(h, h') := \mathbb{E}_{X \sim P_k^*} [\mathbb{I}(h(X) \neq h'(X))]$, respectively.

Suppose the population is partitioned based on the sensitive attribute S , and define $\{h_k^*\}_{k \in [K]}$ as the set of classifiers that attains the maximal true degree of fairness without harm obj_{P^*} under $\{P_k^*\}_{k \in [K]}$:

$$\text{obj}_{P^*} = \frac{1}{K} \sum_{k=1}^K \left\{ [R_k^*(h_0) - R_k^*(h_k^*)] + \frac{1}{K} \sum_{j=1}^K [R_k^*(h_j) - R_k^*(h_k^*)] \right\}.$$

Denote the maximal objective value of our classifiers $\{h_k\}_{k \in [K]}$ assessed on $\{P_k\}_{k \in [K]}$ derived by our group classifier as:

$$\text{obj}_P = \frac{1}{K} \sum_{k=1}^K \left\{ [R_k(h_0) - R_k(h_k)] + \frac{1}{K} \sum_{j=1}^K [R_k(h_j) - R_k(h_k)] \right\}$$

Theorem 2. *Assume the loss function ℓ is symmetric and obeys triangle inequality, i.e., $\forall y_1, y_2, y_3 \in \mathcal{Y}, \ell(y_1, y_2) \leq \ell(y_1, y_3) + \ell(y_3, y_2)$, and for every pair of P_k^* and P_k , there exists $h \in \mathcal{H}^D$, such that $R_k^*(h) = R_k(h)$, then we have:*

$$\text{obj}_{P^*} - \text{obj}_P \leq 3 \max_k \text{disc}(P_k, P_k^*),$$

where $\text{disc}(P_k, P_k^*)$ measures the discrepancy between two distributions P_k and P_k^* and is defined as:

$$\text{disc}(P_k, P_k^*) = \max_{h, h' \in \mathcal{H}^D} |D_k(h, h') - D_k^*(h, h')|.$$

Proof. Define $f : \mathcal{X} \rightarrow \mathcal{Y}$ as the labeling function, i.e., for each sample (x, y) , we have $f(x) = y$, then true risk $R_k(h)$ can be represented as $D_k(h, f)$. Define group-specific true risk on distribution P_k^* as $R_k^*(h) = D_k^*(h, f) = \mathbb{E}_{X, Y | S=k} [\ell(h(X), Y)]$. Without lossing generality, we can assume $D_k(h, f) \geq D_k^*(h, f)$. For any classifier $h \in \mathcal{H}^D$ and group index k , according to the triangle inequality of the loss function, we have:

$$\begin{aligned} |R_k(h) - R_k^*(h)| &= |D_k(h, f) - D_k^*(h, f)| \\ &\leq |D_k(h, h') + D_k(h', f) - [D_k^*(h', f) - D_k^*(h, h')]| \\ &\leq |D_k(h, h') - D_k^*(h, h')| + |D_k(h', f) - D_k^*(h', f)|, \end{aligned}$$

where h' can be any hypothesis in \mathcal{H}^D . Since there exist h' such that $D_k^*(h', f) = D_k(h', f)$, and according to the definition of discrepancy distance, we can have:

$$|D_k(h, f) - D_k^*(h, f)| \leq \text{disc}(P_k, P_k^*).$$

Let $\{h_k^*\}$ be the set of classifiers that attains the maximal true degree of fairness without harm obj_{P^*} under distribution $\{P_k^*\}_{k \in [K]}$. Let obj'_P be the true degree of fairness attained with the same set of classifiers $\{h_k^*\}$ under distribution $\{P_k\}_{k \in [K]}$ determined by our group classifier, i.e.,

Hyperparameter	Values
Batch size	256, 1,024
Training iterations	2, 3, 4
Learning rate for group classifier	$10^{-4}, 10^{-3}$
Learning rate for decoupled classifiers	$10^{-2}, 10^{-1}$
Momentum for decoupled classifiers	0.9
Penalty parameter λ	0, 1, 10, 100

Table 2: Hyper-parameters used in experiments

$$\text{obj}'_P = \frac{1}{K} \sum_{k=1}^K \left\{ [R_k(h_0) - R_k(h_k^*)] + \frac{1}{K} \sum_{j=1}^K [R_k(h_j) - R_k(h_k^*)] \right\}.$$

According to our definition, $\text{obj}_P \geq \text{obj}'_P$. Then we can know that the difference between obj_{P^*} and obj_P is upper bounded:

$$\begin{aligned} \text{obj}_{P^*} - \text{obj}_P &\leq \text{obj}_{P^*} - \text{obj}'_P \\ &\leq 3 \max_k \text{disc}(P_k, P_k^*) \end{aligned}$$

□

Given a hypothesis class, distance metric $\text{disc}(\cdot, \cdot)$ measures the difference between two distributions: with a higher distance, it is more likely that the same hypothesis has diverse outputs (Mansour, Mohri, and Rostamizadeh 2009).

Theorem 2 shows that the performance gap between our method and the optimal models trained by partitioning groups based on sensitive attributes is bounded by the difference between P_k and P_k^* . Suppose there is a ground truth function $\phi : \mathcal{X} \rightarrow \mathcal{S}$ such that $\phi(x) = s$ and $\phi \in \mathcal{H}^G$, the following result is straightforward.

Corollary 1. *If $\phi \in \mathcal{H}^G$ and we can find a group classifier θ such that $\text{disc}(P_k, P_k^*) = 0, \forall k$, then $\text{obj}_P \geq \text{obj}_{P^*}$.*

6 Experiment

6.1 Dataset

Synthetic dataset. We generate the synthetic dataset as follows. Each data point has two sensitive attributes $s_1, s_2 \in \{\pm 1\}$, two features x_1, x_2 , and label $y \in \{\pm 1\}$. The synthetic dataset comprises 20,000 samples, with each sample being randomly assigned with sensitive attributes and a label. Features are associated with both sensitive attributes and the label:

$$\begin{aligned} x_1 &\sim \mathcal{N}(s_1 + \Delta s_1 s_2 y, \sigma), \\ x_2 &\sim \mathcal{N}(s_2 + \Delta s_1 s_2 y, \sigma), \end{aligned}$$

where $\Delta = 0.4, \sigma = 0.3$. The synthetic dataset simulates a setting where features can act as proxy variables of sensitive attributes, and different subgroups divided by two sensitive attributes follow different distributions. A visualization of a randomly sampled subset of the dataset is depicted in Fig. 1, where points with positive (resp. negative) labels are denoted by circles (resp. crosses), and different colors correspond to groups divided based on sensitive attributes (or the group classifier in our method).

Hyperparameter	Adult	Arrest	Violent	German	Bank
Batch size	1,024	1,024	1,024	256	1,024
Training epochs	3	3	3	3	2
Learning rate for group classifier	10^{-3}	10^{-3}	10^{-3}	10^{-3}	10^{-3}
Learning rate for decoupled classifiers	10^{-2}	10^{-2}	10^{-2}	10^{-2}	10^{-2}
Momentum for decoupled classifiers	0.9	0.9	0.9	0.9	0.9
Penalty parameter λ	10	10	10	100	10

Table 3: Final hyper-parameters selected for each dataset

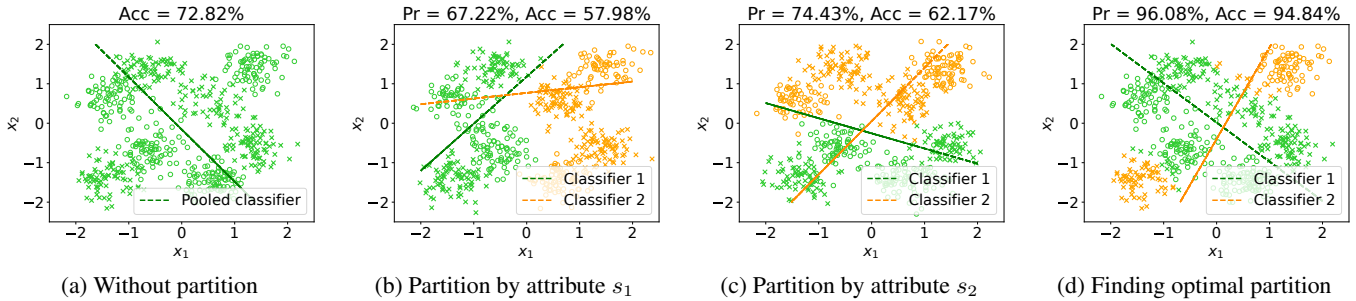


Figure 1: Illustration of different methods on synthetic data. Circles (respectively, crosses) represent positive (respectively, negative) points. Lime (respectively, orange) points are mapped to the group associated with classifier denoted by the green (respectively, dark orange) line. Features x_1 and x_2 are correlated with s_1 , s_2 , and the label, so using a pooled classifier or dividing the group based on a single attribute, cannot achieve desirable performance. In contrast, our method can identify the optimal group partition, achieving the highest probability of fairness (denoted by “Pr”) and accuracy (denoted by “Acc”).

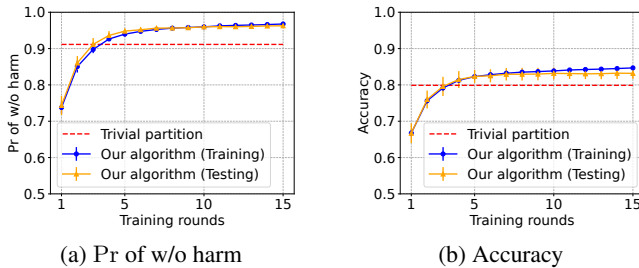


Figure 2: Convergence of our method on the Violent dataset: mean values of performance metrics during training, with error bars indicating standard deviation. Our algorithm gradually improves fairness without harm and accuracy, outperforming the *Trivial partition* in fewer than 15 steps.

Real datasets. We use five real datasets, as detailed below:

1. **Adult** (Becker and Kohavi 1996): Census data to predict whether an individual’s income exceeds \$50K per year. It includes 48,842 instances, each has 14 features and we choose *race* as the sensitive attribute.
2. **Arrest** (Angwin et al. 2022): Data used by the COMPAS algorithm to score the criminal defendants’ risk of recidivism. It includes 7,214 instances, each has 52 features and we choose *race* as the sensitive attribute.
3. **Violent** (Angwin et al. 2022): Data used by the COMPAS algorithm to predict the criminal defendants’ risk of violent recidivism. It includes 4,743 instances, each has 53 features and we choose *race* as the sensitive attribute.




Figure 3: Physical meaning of the learned group partition on Arrest dataset: the group classifier is used to divide the dataset into two sets. The probability of different subgroups defined by 3 attributes race, sex, and age is visualized. Results show that different groups have diverse distribution on features.

4. **German** (Hofmann 1994): Credit data used to predict individuals’ credit risks (as high or low). It includes 1,000 instances, each has 20 features and we choose *sex* as the sensitive attribute.
5. **Bank** (Moro, Rita, and Cortez 2012): The dataset focuses on bank marketing campaigns conducted via phone calls, aiming to predict whether clients will subscribe to a bank’s term deposit product. It includes 45,211 instances, each has 16 features and *age* is the sensitive attribute.

6.2 Baseline Methods

We compare our method with the following baselines:

Dataset	DAFH (Ours)		Trivial partition		Pooled classifier	
	Pr of w/o harm	Accuracy	Pr of w/o harm	Accuracy	Pr of w/o harm	Accuracy
Adult	90.70% \pm 0.88%	83.36% \pm 0.43%	91.64% \pm 0.86%	80.70% \pm 0.51%	N/A	81.29% \pm 0.41%
Arrest	87.05% \pm 1.21%	65.72% \pm 1.36%	78.60% \pm 1.54%	62.01% \pm 1.43%	N/A	63.84% \pm 1.30%
Violent	91.74% \pm 1.49%	81.91% \pm 1.53%	91.14% \pm 1.78%	79.88% \pm 1.94%	N/A	81.35% \pm 1.58%
German	90.64% \pm 1.65%	74.24% \pm 0.82%	78.88% \pm 2.42%	65.92% \pm 1.51%	N/A	70.72% \pm 0.93%
Bank	97.71% \pm 0.34%	90.20% \pm 0.10%	93.42% \pm 0.72%	88.77% \pm 0.55%	N/A	88.91% \pm 0.54%

Table 4: Comparison with baselines (*Trivial partition* and *Pooled classifier*) on five real datasets: “Pr of w/o harm” measures the probability that samples in the testing set satisfy both rationality and envy-freeness, given the group partition and decoupled classifiers (*Pooled classifiers* do not have group partition, hence we put “N/A”). “Accuracy” is the overall accuracy of classifiers on associated samples. Results show that we can outperform other baselines in both fairness and accuracy in most cases.

	Pr of w/o harm	Accuracy
Manual partition	81.39% \pm 2.39%	62.54% \pm 1.70%
Trivial partition	78.60% \pm 1.54%	62.01% \pm 1.43%
Our algorithm	85.72% \pm 1.14%	64.72% \pm 1.12%

Table 5: Comparison with *Manual partition* and *Trivial partition* on Arrest dataset: we manually partition the dataset into 2 groups according to 3 attributes race, sex, and age, which mimics the physical meaning we observe in the group partition generated by our group classifier, and train decoupled classifiers associated with each group. Results show that manually partitioning the dataset based on our group classifier can help derive decoupled classifiers with higher performance, compared with trivially partitioning the dataset based on the chosen sensitive attribute.

1. *Pooled classifier* is the classifier trained on the complete training set through structural risk minimization, and is applied to all samples in the testing set.
2. *Trivial partition* first splits each dataset according to the chosen sensitive attribute into two groups, and then trains decoupled classifiers accordingly through structural risk minimization. Typically it is restricted to the setting when the number of assigned groups $K = 2$.
3. *Clustering* first partitions the dataset through a clustering algorithm (e.g., K-means in our experiments), then trains decoupled classifiers based on each partitioned set through structural risk minimization.
4. *LR-All* (Ustun, Liu, and Parkes 2019) partitions the dataset by the intersection of multiple attributes, and trains decoupled classifiers for each group.
5. *TreeLR* (Ustun, Liu, and Parkes 2019) recursively grows a tree to find the group partition that best fulfills rationality and envy-freeness, with each group being assigned a decoupled classifier.

In all baselines, we use the logistic regression as decoupled classifiers, the same as our method for fair comparison.

6.3 Experimental Setup

Implementation details. Experiments are conducted on a personal computer with Intel Core i5-10210U CPU, 16

GB RAM, and Windows 11 operating system. Our algorithm and baselines are implemented with PyTorch 2.1.0 and scikit-learn 1.3.2. We use multi-layer perceptron with a hidden layer of 100 neurons as the group classifier, and logistic regression models as decoupled classifiers. Because our method does not require access to individuals’ demographic information, we remove the sensitive attributes from input features when training our models. We iterate over all samples in the training set to build the computation graph of the objective function, automatically calculate the gradients for classifier parameters using autograd, and update parameters using stochastic gradient ascent to maximize the objective.

Evaluation metrics. We measure the performance of our algorithm by the following metrics:

1. **Probability of fairness without harm:** It is the probability that a sample satisfies both rationality and envy-freeness. Given a group classifier θ , pooled classifier h_0 , and decoupled classifiers $\{h_k\}_{k \in [K]}$, a sample (x_i, y_i) satisfies both rationality and envy-freeness if the 0-1 loss of the assigned classifier $h_{\theta(x_i)}$ is not higher than the pooled classifier h_0 and any other decoupled classifier h_k , i.e., $\ell(h_{\theta(x_i)}(x_i), y_i) \leq \ell(h_k(x_i), y_i), \forall k \in \{0\} \cup [K]$. The probability of fairness without harm equals the number of samples that satisfy fairness without harm divided by the size of the dataset.
2. **Accuracy** of decoupled classifiers measures the proportion of samples in the dataset that are correctly classified by the assigned classifier, i.e, $\Pr(h_{\theta(X)}(X) = Y)$, given a group classifier θ and decoupled classifiers $\{h_k\}_{k \in [K]}$.
3. **# of violations** is the number of groups that violates rationality, plus the number of pairwise comparison among all groups that violate envy-freeness.
4. **max gain** is the maximal difference between the true risk of pooled classifier and decoupled classifiers, $\max_{k \in [K]} (R_k(h_0) - R_k(h_k))$, which is the higher the better.
5. **min envy** is the maximal difference between the true risk of decoupled classifiers and the assigned classifier, $\max_{k, k' \in [K]} (R_k(h_{k'}) - R_k(h_k))$, which is the higher the better.
6. **Δ disparity** is the disparity when using decoupled classifiers minus the disparity of using pooled classifier only,

Dataset	Algorithm	$K = 2$		$K = 3$		$K = 4$	
		Pr of w/o harm	Accuracy	Pr of w/o harm	Accuracy	Pr of w/o harm	Accuracy
Adult	DAFH (Ours)	90.70% \pm 0.88%	83.36% \pm 0.43%	94.55% \pm 1.31%	82.82% \pm 0.17%	93.98% \pm 0.47%	80.98% \pm 0.54%
	Clustering	88.79% \pm 0.58%	82.90% \pm 0.28%	87.43% \pm 0.53%	82.83% \pm 0.23%	86.00% \pm 1.19%	82.59% \pm 0.52%
Arrest	DAFH (Ours)	87.05% \pm 1.21%	65.72% \pm 1.36%	88.06% \pm 1.06%	65.82% \pm 0.98%	88.78% \pm 1.18%	65.77% \pm 0.80%
	Clustering	84.23% \pm 1.43%	64.86% \pm 0.89%	74.52% \pm 2.43%	65.38% \pm 0.81%	72.12% \pm 1.84%	65.42% \pm 0.38%
Violent	DAFH (Ours)	96.28% \pm 0.53%	83.18% \pm 1.53%	95.38% \pm 1.24%	83.16% \pm 1.69%	96.42% \pm 0.68%	84.54% \pm 0.74%
	Clustering	87.96% \pm 2.68%	83.82% \pm 1.36%	85.27% \pm 1.82%	83.58% \pm 1.56%	84.02% \pm 1.58%	83.64% \pm 1.54%
German	DAFH (Ours)	90.64% \pm 1.65%	74.24% \pm 0.82%	90.00% \pm 2.16%	76.00% \pm 3.27%	89.12% \pm 1.32%	73.52% \pm 1.42%
	Clustering	84.40% \pm 1.12%	73.16% \pm 2.38%	78.96% \pm 2.43%	71.92% \pm 1.51%	79.04% \pm 2.05%	73.36% \pm 1.69%
Bank	DAFH (Ours)	97.71% \pm 0.34%	90.20% \pm 0.10%	97.34% \pm 0.32%	90.06% \pm 0.28%	97.09% \pm 0.23%	90.11% \pm 0.30%
	Clustering	90.18% \pm 0.21%	88.92% \pm 0.25%	89.67% \pm 0.25%	89.07% \pm 0.16%	89.29% \pm 0.43%	88.84% \pm 0.24%

Table 6: Comparison with *Clustering* baseline on five real datasets when the number of decoupled classifiers $K = 2, 3, 4$: results show that we can outperform *Clustering* on both fairness and accuracy in most cases, indicating that our learned classifier can extract more information than merely the similarity in distance between different groups to improve performance.

where disparity is the maximal difference in accuracy between different groups, which is the lower the better.

We randomly split each dataset into 75%-25% training-testing set for 5 times. For the first two metrics, we report the mean and standard deviation of metrics. For the following four metrics, we report the results when the optimal accuracy is achieved on the testing set.

Hyperparameter selection. We test all combinations of hyperparameter values shown in Table 2, and choose the combination that can attain the highest probability of fairness without harm on each dataset to show the results. Notably, we choose different learning rates to update parameters of group classifier and decoupled classifiers. We use the same learning rate to update parameters of the pooled classifier and decoupled classifiers in both our algorithm and other baselines. The chosen hyper-parameters on each dataset are as shown in Table 3.

6.4 Results

Synthetic data. Fig. 1 evaluates our method on synthetic data. The performance metrics for each method are shown in the title of each figure. In Fig. 1a, a pooled classifier (green line) is applied to all samples. Since the distribution between different groups differs significantly, a single classifier cannot achieve desirable accuracy. In Fig. 1b and Fig. 1c, we partition the dataset based on sensitive attributes s_1 and s_2 , respectively, and train decoupled classifiers on the divided groups. In particular, lime (resp. orange) points represent points belonging to Group 1 (resp. Group 2) and are assigned with the classifier denoted by the green (resp. dark orange) line. Since the difference in distribution results from both s_1 and s_2 , partitioning groups based on only one of them leads to inferior fairness and accuracy. In contrast, by using an additional group classifier to search for the optimal partition, our algorithm shown in Fig. 1d can attain the best fairness and accuracy. This further indicates that our method can effectively tackle intersectional group fairness and has greater capacity to explore data heterogeneity to enhance fairness and accuracy.

Real data. Table 4 compares our method with *Trivial partition* and *Pooled classifier* when the number of decoupled classifiers $K = 2$. Results show that we can achieve a comparable or even higher probability of fairness without harm compared with *Trivial partition* on most datasets, indicating that our algorithm can find better group partition to capture the data heterogeneity. Because the *Pooled classifier* is applied to the whole dataset without any partitioned groups, we put ‘‘N/A’’ in Table 4. Moreover, our method can also achieve higher accuracy than *Trivial partition* and *Pooled classifier*, which means that our algorithm can increase the overall accuracy by exploring the data heterogeneity, even when we did not directly optimize for the accuracy metric.

Fig. 2 further shows the convergence of our algorithm on Violent dataset, where we use dashed horizontal lines to indicate the *Trivial partition* and use solid polylines to denote ours. We report the mean values with errorbars denoting standard deviation of 5 repeated experiments. Results show that our algorithm improves fairness without harm and accuracy on both training and testing sets by maximizing the objective function, and can outperform *Trivial partition* in less than 15 steps of mini-batch gradient ascent operations, with each batch containing 1,024 samples.

We explore the group partition generated by the learned group classifier based on the Arrest dataset, to find whether there are any physical meaning, as depicted in Fig. 3. We evaluate the proportions of different subgroups of people in the two generated groups, with each subgroup being defined by the intersection of three attributes: race (‘White’ or ‘Non-White’), sex (‘Female’ or ‘Male’), and age (‘Less than 25’, ‘25 - 45’, or ‘Greater than 45’). Notably, the proportion of non-white male less than 45 in Group 0 is especially higher than that in Group 1. The proportion of non-white female from 25 to 45 in Group 1 is especially higher than that in Group 0. To test whether these insights can enhance fairness and accuracy, we manually partition the Arrest dataset into two groups. We assign all non-white males under 45 to one group, all non-white females aged 25 to 45 to another group, and randomly distribute the remaining individuals between

Dataset	Metrics	Trivial Partition	LR-All	TreeLR	DAFH (Ours)
Adult	# violations	1	1	0	0
	max gain	2.9%	17.6%	4.1%	11.11%
	min envy	2.9%	33.3%	11.8%	5.56%
	Δ disparity	0.6%	0.4%	-0.4%	-2.91%
	# models	2	12	7	2
Arrest	# violations	0	3	1	0
	max gain	7.7%	11.5%	9.0%	10%
	min envy	7.7%	15.4%	12.8%	30%
	Δ disparity	-0.1%	1.0%	-2.7%	-7.9%
	# models	2	6	4	2
Violent	# violations	1	0	2	0
	max gain	9.3%	13.6%	9.3%	10.61%
	min envy	7.6%	17.0%	7.6%	13.64%
	Δ disparity	-9.7%	-14.0%	-9.8%	-10.29%
	# models	2	6	2	2

Table 7: Comparison with algorithms proposed in Ustun, Liu, and Parkes (2019) which rely on sensitive attributes to generate group partition. Metrics that are higher than those of other algorithms achieving the same # of violations on each dataset are boldfaced. Results show that we can ensure no rationality or envy-freeness violation across all three datasets. We can also achieve lower Δ disparity, indicating that our algorithm can satisfy parity-based fairness criteria even without access to sensitive attributes.

these groups. We compare this *Manual partition* algorithm with *Trivial partition* baseline and also our algorithm. Results in Table 5 show that *Manual partition* achieves a higher probability of fairness without harm and accuracy compared with *Trivial partition*, which indicates that using a single sensitive attribute is insufficient for capturing the nuances between different groups. Our *Manual partition* approach, which also considers age and sex in addition to race, underscores the importance of the intersectionality of attributes. While *Manual Partition* does not outperform our algorithm, which leverages a broader set of attributes through the group classifier, it still shows significant benefits.

We also test the performance of our algorithm on the same datasets when the number of groups K exceeds 2. Partitioning datasets into more than two groups based on sensitive attributes can result in some training sets being too small, potentially leading to overfitting. Therefore, we use *Clustering* baseline to generate a more balanced group partition. Results in Table 6 show that our algorithm outperforms the *Clustering* baseline in most cases in terms of probability of fairness without harm and accuracy, indicating that our algorithm can extract more information than merely similarity between different samples, thereby forming better group partition and increasing fairness and accuracy.

Finally, we compare our algorithm with three algorithms proposed in Ustun, Liu, and Parkes (2019): *Trivial Partition* (referred as *LR (One attribute)* in Ustun, Liu, and Parkes (2019)), *LR-All*, and *TreeLR*. We record the values of four metrics when each algorithm achieves optimal testing accuracy, evaluated across different numbers of models. The results in Table 7 demonstrate that our approach ensures no violations of rationality or envy-freeness across all datasets, i.e., every group satisfies fairness without harm, a standard

that other algorithms fail to meet. Additionally, compared to algorithms with the same number of violations, our DAFH algorithm achieves lower Δ disparity, highlighting its capability to satisfy parity-based fairness criteria. While other algorithms may achieve higher maximum gain or lower minimum envy, this is likely due to their ability to leverage sensitive attributes, enabling them to better capture groups with differing data distributions.

7 Conclusion

In this work, we propose a novel approach to achieve fairness without harm, a preference-based fairness notion, without accessing demographic data. Differs from parity-based fairness concept, we do not have to trade accuracy for fairness. Instead, we can train decoupled classifiers for divided groups and ensure each group prefer their own assigned classifier in terms of accuracy. Previous works require access to sensitive attributes to form group partition, which is not always feasible. Therefore, we propose the first approach that can achieve fairness without harm while not knowing sensitive attributes. Experiments show that our algorithm can achieve superior fairness and accuracy compared with baselines which have access to sensitive attributes.

Acknowledgments

This material is based upon work supported by the U.S. National Science Foundation under award IIS2202699, IIS-2416895, IIS-2301599, and CMMI-2301601, and by OSU President’s Research Excellence Accelerator Grant, and grants from the Ohio State University’s Translational Data Analytics Institute and College of Engineering Strategic Research Initiative.

References

- Abroshan, M.; Elliott, A.; and Khalili, M. M. 2024. Imposing fairness constraints in synthetic data generation. In *International Conference on Artificial Intelligence and Statistics*, 2269–2277. PMLR.
- Andrus, M.; Spitzer, E.; Brown, J.; and Xiang, A. 2021. What we can’t measure, we can’t understand: Challenges to demographic data procurement in the pursuit of fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 249–260.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2022. Machine bias. In *Ethics of data and analytics*, 254–264. Auerbach Publications.
- Ashurst, C.; and Weller, A. 2023. Fairness without demographic data: A survey of approaches. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–12.
- Awasthi, P.; Beutel, A.; Kleindessner, M.; Morgenstern, J.; and Wang, X. 2021. Evaluating fairness of machine learning models under uncertain and incomplete information. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 206–214.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.
- Barocas, S.; and Selbst, A. D. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104: 671.
- Beauchamp, T. L.; and Childress, J. F. 1994. *Principles of biomedical ethics*. Edicoes Loyola.
- Becker, B.; and Kohavi, R. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- Berliant, M.; Thomson, W.; and Dunz, K. 1992. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3): 201–216.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Grari, V.; Lamprier, S.; and Detyniecki, M. 2021. Fairness without the sensitive attribute via causal variational autoencoder. *arXiv preprint arXiv:2109.04999*.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- Hashimoto, T.; Srivastava, M.; Namkoong, H.; and Liang, P. 2018. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, 1929–1938. PMLR.
- Hofmann, H. 1994. Statlog (German Credit Data). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5NC77>.
- Hossain, S.; Mladenovic, A.; and Shah, N. 2020. Designing fairly fair classifiers via economic fairness notions. In *Proceedings of The Web Conference 2020*, 1559–1569.
- Islam, R.; Chen, H.; and Cai, Y. 2024. Fairness without Demographics through Shared Latent Space-Based Debiasing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12717–12725.
- Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International conference on machine learning*, 2564–2572. PMLR.
- Khalili, M. M.; Zhang, X.; and Abroshan, M. 2021. Fair sequential selection using supervised learning models. *Advances in Neural Information Processing Systems*, 34: 28144–28155.
- Khalili, M. M.; Zhang, X.; and Abroshan, M. 2023. Loss balancing for fair supervised learning. In *International Conference on Machine Learning*, 16271–16290. 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.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*.
- Menon, A. K.; and Williamson, R. C. 2018. The cost of fairness in binary classification. In *Conference on Fairness, accountability and transparency*, 107–118. PMLR.
- Moro, S.; Rita, P.; and Cortez, P. 2012. Bank Marketing. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5K306>.
- Pang, J.; Wang, J.; Zhu, Z.; Yao, Y.; Qian, C.; and Liu, Y. 2024. Fairness without Harm: An Influence-Guided Active Sampling Approach. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Pham, T.-H.; Zhang, X.; and Zhang, P. 2023. Fairness and Accuracy under Domain Generalization. In *The Eleventh International Conference on Learning Representations*.
- Saito, Y.; and Joachims, T. 2022. Fair ranking as fair division: Impact-based individual fairness in ranking. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1514–1524.
- Shah, N. 2023. Pushing the Limits of Fairness in Algorithmic Decision-Making. In *IJCAI*, 7051–7056.
- Shalev-Shwartz, S.; and Ben-David, S. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- Sohoni, N.; Dunnmon, J.; Angus, G.; Gu, A.; and Ré, C. 2020. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *Advances in Neural Information Processing Systems*, 33: 19339–19352.
- U.S. Congress. 1996. Health Insurance Portability and Accountability Act of 1996.
- Ustun, B.; Liu, Y.; and Parkes, D. 2019. Fairness without harm: Decoupled classifiers with preference guarantees. In *International Conference on Machine Learning*, 6373–6382. PMLR.
- Varian, H. R. 1973. Equity, envy, and efficiency.

- Veldanda, A. K.; Brugere, I.; Dutta, S.; Mishler, A.; and Garg, S. 2024. Hyper-parameter Tuning for Fair Classification without Sensitive Attribute Access. *Transactions on Machine Learning Research*.
- Zafar, M. B.; Valera, I.; Rodriguez, M.; Gummadi, K.; and Weller, A. 2017. From parity to preference-based notions of fairness in classification. *Advances in neural information processing systems*, 30.
- Zhang, X.; Khalili, M. M.; and Liu, M. 2020. Long-Term Impacts of Fair Machine Learning. *Ergonomics in Design: The Quarterly of Human Factors Applications*, 28(3): 7–11.
- Zhang, X.; Khaliligarekani, M.; Tekin, C.; and liu, m. 2019. Group Retention when Using Machine Learning in Sequential Decision Making: the Interplay between User Dynamics and Fairness. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Zhang, X.; and Liu, M. 2021. Fairness in learning-based sequential decision algorithms: A survey. In *Handbook of Reinforcement Learning and Control*, 525–555. Springer.
- Zhang, X.; Tu, R.; Liu, Y.; Liu, M.; Kjellstrom, H.; Zhang, K.; and Zhang, C. 2020. How do fair decisions fare in long-term qualification? *Advances in Neural Information Processing Systems*, 33: 18457–18469.
- Zhu, Z.; Yao, Y.; Sun, J.; Li, H.; and Liu, Y. 2023. Weak proxies are sufficient and preferable for fairness with missing sensitive attributes. In *International Conference on Machine Learning*, 43258–43288. PMLR.
- Zuo, Z.; Khalili, M.; and Zhang, X. 2023. Counterfactually fair representation. *Advances in neural information processing systems*, 36: 12124–12140.
- Zuo, Z.; Xie, T.; Tan, X.; Zhang, X.; and Khalili, M. M. 2024. Lookahead Counterfactual Fairness. *Transactions on Machine Learning Research*.
- Zuo, Z.; Zhu, D.; and Khalili, M. M. 2025. Post-processing for Fair Regression via Explainable SVD. In *International Conference on Artificial Intelligence and Statistics*, 1711–1719. PMLR.