

mdfa: Multi-Differential Fairness Auditor for Black Box Classifiers

Xavier Gitiaux* and Huzefa Rangwala

George Mason University

{xgitiaux, hrangwal}@gmu.edu

Abstract

Machine learning algorithms are increasingly involved in sensitive decision-making processes with adverse implications on individuals. This paper presents **mdfa**, an approach that identifies the characteristics of the victims of a classifier’s discrimination. We measure discrimination as a violation of multi-differential fairness. Multi-differential fairness is a guarantee that a black box classifier’s outcomes do not leak information on the sensitive attributes of a small group of individuals. We reduce the problem of identifying worst-case violations to matching distributions and predicting where sensitive attributes and classifier’s outcomes coincide. We apply **mdfa** to a recidivism risk assessment classifier and demonstrate that for individuals with little criminal history, identified African-Americans are three-times more likely to be considered at high risk of violent recidivism than similar non-African-Americans.

1 Introduction

Machine learning algorithms are increasingly used to support decisions that could impose adverse consequences on an individual’s life: for example in the judicial system, algorithms are used to assess whether a criminal offender is likely to recommit crimes; or within banks to determine the default risk of a potential borrower. At issue is whether classifiers are fair [Calsamiglia, 2009]: are classifiers’ outcomes independent of *exogenous irrelevant characteristics* or sensitive attributes like race and/or gender? Abundant examples of classifiers’ discrimination can be found in diverse applications ([Atlantic, 2016; ProPublica, 2016]). ProPublica ([ProPublica, 2016]) reported that COMPAS machine learning based recidivism risk assessment tool assigns disproportionately higher risk to African-American defendants than to Caucasian defendants.

Establishing contestability is challenging for potential victims of machine learning discrimination because (i) many assessment tools are proprietary and usually not transparent; and (ii) a precedent in United States case law places the

burden on the plaintiff to establish that characteristics irrelevant to the task affect the algorithm’s outcomes. Identifying the definitive characteristics of a classifier’s discrimination empowers the victims of such discrimination. Moreover, a classifier’s user needs warnings for individual instances in which suspicions of profiling/discrimination have been detected. In this paper, we present a theoretical framework (multi-differential fairness) and a practical tool (**mdfa**) to identify groups of individuals with similar characteristics but different classifier’s outcomes depending on their sensitive attribute.

We define a classifier as differentially fair if conditional on an individual’s features, the classifier’s outcomes does not depend on sensitive attributes. To be computationally operational, we relax differential fairness to multi-differential fairness by conditioning only on membership in balanced subsets that are parametrized by a collection of sub-populations. Multi-differential fairness guarantees that a classifier assigns identical outcome to similar individuals, similarity being controlled by the complexity of the collection of sub-populations and the requirement that the distribution of features is balanced across sensitive attributes.

mdfa audits for the worst-case violations of multi-differential fairness. Theoretically, our construction relies on a reduction to both matching distribution and agnostic learning problems. First, we re-balance the distribution of features by minimizing the maximum-mean discrepancy between distributions conditioned on different sensitive attributes. Secondly, we show that violations of differential fairness is a problem of finding correlations between sensitive attributes and classifier’s outcomes. Therefore, **mdfa** searches for instances of violation of multi-differential fairness by predicting where in the features space the binary values of sensitive attributes and classifier’s outcomes coincide. Lastly, worst-case violations are extracted by incrementally removing the least unfair instances.

This paper makes the following contributions:

- The proposed multi-differential fairness framework is the first attempt to efficiently identify groups of similar individuals to which a black box classifier assigns different outcomes depending on their sensitive attributes.
- We apply **mdfa** to a case study of a recidivism risk assessment in Broward County, Florida and find a sub-

*Contact Author

population of African-American defendants who are three times more likely to be considered at high risk of violent recidivism than similar individuals of other races.

- We apply **mdfa** to three other datasets related to crime, income and credit predictions and find that classifiers, even after being repaired for aggregate fairness, do discriminate against smaller sub-populations.

This paper belongs to a rapidly growing literature on algorithmic fairness [Chouldechova and Roth, 2018]. Our paper, as in [Hébert-Johnson *et al.*, 2017; Kim *et al.*, 2018; Kearns *et al.*, 2017] provides a definition of fairness that protects group of individuals as small as computationally possible. This is a step in the right direction since empirical observations in [Kearns *et al.*, 2017; Dwork *et al.*, 2012] show that aggregate level fairness cannot protect sub-populations against severe discrimination.

Prior contributions on algorithmic disparate treatment ([Zafar *et al.*, 2017]) have focused on whether sensitive attributes are used directly to train a classifier. This is a limitation when dealing with classifiers with unknown inputs. Multi-differential fairness addresses this limitation at the cost of a relaxation of disparate treatment compared to [Loftus *et al.*, 2018]: it marginalizes the classifier’s outcome over sub-population characteristics instead of individual characteristics; it is only an associational and not a causal notion of disparate treatment. However, results in theorem 3.4 show that because the classifier is a black box, a relaxation to sub-population is necessary to efficiently find violations of differential fairness.

Multi-differential fairness borrows from the literature in differential privacy [Dwork *et al.*, 2014]. Reinterpretations of fairness as a privacy problem can be found in [Jagielski *et al.*, 2018; Foulds and Pan, 2018], but those contributions marginalize only over sensitive attributes.

The relaxation of differential fairness to sub-population requires rebalancing the distribution of features across sensitive attributes. Our kernel matching technique borrows from domain adaptation (see e.g. [Mansour *et al.*, 2009]) and counterfactual analysis (e.g. [Johansson *et al.*, 2016]). Because **mdfa** balances the features across sensitive classes, its outcome cannot be explained by disparate impact (as in [Feldman *et al.*, 2015]). In fact, empirical results in section 4 show that even after removing disparate impacts, **mdfa** still detects violations of multi-differential fairness.

2 Individual and Multi-Differential Fairness

Preliminary. An individual i is defined by a tuple $((x_i, s_i), y_i)$, where $x_i \in \mathcal{X}$ denotes i ’s audited features; $s_i \in \mathcal{S}$ denotes the sensitive attributes; and $y_i \in \{-1, 1\}$ is a classifier f ’s outcome. The auditor draws m samples $\{((x_i, s_i), y_i)\}_{i=1}^m$ from a distribution D on $\mathcal{X} \times \mathcal{S} \times \{-1, 1\}$. Features in \mathcal{X} are not necessarily the ones used to train f , because the auditor may not have access to all features used to train f . Secondly, the auditor may decide to deliberately leave out some features used to train f because those features – e.g. small geography identifiers – correlate strongly with

sensitive attributes and that may break the following assumption of common support.

Assumptions. In our analysis we assume that the distributions of auditing features conditioned on sensitive attributes have common support.

Assumption 1. For all $x \in \mathcal{X}$, $Pr[S|X = x] > 0$.

2.1 Individual Differential Fairness

We define differential fairness as the guarantee that, conditioned on features relevant to the tasks, a classifier’s outcome is nearly independent of sensitive attributes:

Definition 2.1. (*Individual Differential Fairness*) For $\delta \geq 0$, a classifier f is δ -differential fair if $\forall x \in \mathcal{X}, \forall s \in \mathcal{S}, \forall y \in \{-1, 1\}$

$$e^{-\delta} \leq \frac{Pr[Y = y|S = s, x]}{Pr[Y = y|S \neq s, x]} \leq e^{\delta} \quad (1)$$

The parameter δ controls how much the distribution of the classifier’s outcome Y depends on sensitive attributes S given that auditing feature is x ; larger δ implies a less differentially fair classifier. δ -differential fairness bounds the maximum divergence between the distributions $Pr(Y|S = s, x)$ and $Pr(Y|S \neq s, x)$:

$$\max_{y \in Y} \ln \left(\frac{Pr[Y|S = s, x]}{Pr[Y|S \neq s, x]} \right) \leq \delta$$

Relation with Differential Privacy. Differential fairness re-interprets disparate treatment as a differential privacy issue [Dwork *et al.*, 2014] by bounding the leakage of sensitive attributes caused by Y given what is already leaked by the auditing features x . Formally, the fairness condition (1) is equivalent to bounding the maximum divergence between the distributions $Pr(S|Y, x)$ and $Pr(S|x)$ by δ .

Individual Fairness. Def. (2.1) is an individual level definition of fairness, since it conditions the information leakage on auditing features x . Compared to the notion of individual fairness [Dwork *et al.*, 2012], individual differential fairness does not require an explicit similarity metric. This is a strength of our framework since defining a similarity metric has been the main limitation of applying the concept of individual fairness [Chouldechova and Roth, 2018].

2.2 Multi-Differential Fairness.

Although useful, the notion of individual differential fairness cannot be computationally efficiently audited for. Looking for violations of individual differential fairness requires searching over a set of $2^{|\mathcal{X}|}$ individuals. Moreover, a sample from a distribution over $\mathcal{X} \times \mathcal{S} \times \{-1, 1\}$ has a negligible probability to have two individuals with the same auditing features x but different sensitive attributes s .

Therefore, we relax the definition of individual differential fairness and impose differential fairness for sub-populations. Formally, \mathbb{C} denotes a collection of subsets or group of individuals G in \mathcal{X} . The collection \mathbb{C}_α is α -strong if for $G \in \mathbb{C}$ and $y \in \{-1, 1\}$, $Pr[Y = y \& x \in G] \geq \alpha$.

Definition 2.2. (*Multi-Differential Fairness*) Consider a α -strong collection \mathbb{C}_α of sub-populations of \mathcal{X} . For $0 \leq \delta, \alpha$

classifier f is $(\mathbb{C}_\alpha, \delta)$ -multi differential fair with respect to \mathcal{S} if $\forall s \in \mathcal{S}, \forall y \in \{-1, 1\}$ and $\forall G \in \mathbb{C}_\alpha$:

$$e^{-\delta} \leq \frac{Pr[Y = y|S = s, G]}{Pr[Y = y|S \neq s, G]} \leq e^\delta \quad (2)$$

Multi-differential fairness guarantees that the outcome of a classifier f is nearly mean-independent of protected attributes within any sub-population $G \in \mathbb{C}_\alpha$. The fairness condition in Eq. 2 applies only to sub-populations with $Pr[Y = y \& x \in G] \geq \alpha$ for $y \in \{-1, 1\}$ to avoid trivial cases where $\{x \in G \& Y = y\}$ is a singleton for some y , implying $\delta = \infty$.

Collection of Indicators. We represent the collection of sub-populations \mathbb{C} as a family of indicators: for $G \in \mathbb{C}$, there is an indicator $c : \mathcal{X} \rightarrow \{-1, 1\}$ such that $c(x) = 1$ if and only if $x \in G$. The relaxation of differential fairness to a collection of groups or sub-population is akin to [Kim *et al.*, 2018; Kearns *et al.*, 2017; Hébert-Johnson *et al.*, 2017]. \mathbb{C}_α is the computational bound on how granular our definition of fairness is. The richer \mathbb{C}_α , the stronger the fairness guarantee offers by Def. 2.2. However, the complexity of \mathbb{C}_α is limited by the fact that we identify a sub-population G via random samples drawn from a distribution over $\mathcal{X} \times \mathcal{S} \times \{-1, 1\}$.

3 Fairness Diagnostics: Worst-Case Violations

The objective of **mdfa** is to find the sub-populations with the most severe violation of multi-differential fairness – that is to solve for $s \in \mathcal{S}$ and $y \in \{-1, 1\}$

$$\sup_{S \in \mathbb{C}_\alpha} \ln \left(\frac{Pr[Y = 1|G, S = s]}{Pr[Y = 1|G, S \neq s]} \right). \quad (3)$$

Our approach succeeds at tackling three challenges: (i) if \mathbb{C}_α is large, an auditing algorithm linearly dependent on $|\mathbb{C}|$ can be prohibitively expensive; (ii) the data needs to be balanced conditioned on sensitive attributes; (iii) finding efficiently the most-harmed sub-population implies that we can predict a function $c \in \mathbb{C}_\alpha$ for which we do not directly observe values $c(x)$. All proofs are in Appendix.

3.1 Reduction to Agnostic Learning

First, we reduce the problem of certifying for the lack of differential fairness to an agnostic learning problem.

In this section, we assume that for any $s \in \mathcal{S}$, the conditional distributions $p(x|S = s)$ and $p(x|S \neq s)$ are identical. This is not realistic for many datasets, but we will show how to handle unbalanced data in the next section. A balanced distribution does not leak any information on whether a sensitive attribute is equal to s : $Pr[S = s|x] = Pr[S \neq s|x]$. A violation of $(\mathbb{C}_\alpha, \delta)$ - multi differential fairness simplifies then to a sub-population $G \in \mathbb{C}_\alpha$, a $y \in \{-1, 1\}$ and $s \in \mathcal{S}$ such that

$$Pr[G, Y = y] \left\{ Pr[S = s|G, Y = y] - \frac{1}{2} \right\} \geq \gamma, \quad (4)$$

with $\gamma = \alpha (e^\delta / (1 + e^\delta) - 1/2)$. γ combines the size of the sub-population where a violation exists and the magnitude of the violation. We call a γ - unfairness certificate any triple (G, y, s) that satisfies Eq. (4). Further we postulate that

f is γ -unfair if and only if such certificate exists. Unfairness for balanced distributions is equivalent to the existence of sub-populations for which sensitive attributes can be predicted once the classifier’s outcomes are observed.

Searching for γ -unfairness certificate reduces to mapping the auditing features $\{x_i\}$ to the labels $\{s_i y_i\}$.

Lemma 3.1. *Let $s \in \mathcal{S}$. Suppose that the data is balanced. f is γ - multi-differential unfair for $y \in \{-1, 1\}$ if and only there exists $c \in \mathbb{C}_\alpha$ such that $Pr[rSY = c] \geq 1 - \rho(y) + 4\gamma$, where $r = \text{sign}(y)$ and $\rho(y) = Pr[S = rY]$.*

Lemma 3.1 allows us to reduce searching for a (G, y, s) unfairness certificate to predicting where sensitive attribute and outcomes of f (if $y = 1$) or outcomes of $\neg f$ (if $y = -1$) coincide. Our proposed approach is to solve the following empirical loss minimization:

$$\min_{c \in \mathbb{C}} \frac{1}{m} \sum_{i=1}^m l(c, a_i y_i) + \text{Reg}(c), \quad (5)$$

where $l(\cdot)$ is a 0 – 1 loss function and $\text{Reg}(\cdot)$ a regularizer.

The following result shows that (i) our reduction to a learning problem leads to an unbiased estimate of γ ; (ii) there is a computational limit on how granular multi-differential fairness can be, since for many concept classes \mathbb{C} agnostic learning is a NP-hard problem ([Feldman *et al.*, 2012]).

Theorem 3.2. *Let $\epsilon, \beta > 0$ and $\mathbb{C} \subset 2^{\mathcal{X}}$. Let $\gamma' \in (\gamma - \epsilon, \gamma + \epsilon)$.*

- (i) *There exists an algorithm that by using $O(\log(|\mathbb{C}|, \log(\frac{1}{\eta}), \frac{1}{\epsilon^2}))$ samples $\{(x_i, s_i), y_i\}$ drawn from a balanced distribution D outputs with probability $1 - \eta$ a γ' -unfairness certificate if y_i are outcomes from a γ -unfair classifier;*
- (ii) *\mathbb{C} is agnostic learnable: there exists an algorithm that with $O(\log(|\mathbb{C}|, \log(\frac{1}{\eta}), \frac{1}{\epsilon^2}))$ samples $\{x_i, o_i\}$ drawn from a balanced distribution D outputs with probability $1 - \eta$, $Pr_D[h(x_i) = o_i] + \epsilon \geq \max_{c \in \mathbb{C}} Pr_D[c(x_i) = o_i]$*

3.2 Unbalanced Data

Multi-differential fairness measures the max-divergence distance between the posterior distribution $Pr(S|Y, x)$ and the prior one $Pr(S|x)$. Therefore, it requires knowledge of $Pr(S|x)$. In the previous section, we circumvent the issue by assuming $Pr[S = s|x] = 1/2$. To generalize our approach, we propose to rebalance the data with the following weights: for $s \in \mathcal{S}$, $w_s(x, s) = Pr[S \neq s|x] / Pr[S = s|x]$ and for $s' \neq s$, $w_s(x, s') = 1$. Once reweighted, the conditional distributions $Pr_w(X|S = s)$ and $Pr_w(X|S \neq s)$ are identical and our learning reduction from the previous section applies.

However, in practice we do not have direct access to w_s . One approach is to directly estimate the density $P[S = s|x]$. This method is used in propensity-score matching methods ([Rosenbaum and Rubin, 1983]) in the context of counterfactual analysis. But, exact or estimated importance sampling results in large variance in finite sample ([Cortes *et al.*, 2010]). Instead, we use a kernel-based matching approach ([Gretton *et al.*, 2009] and [Cortes *et al.*, 2008]). Our method considers real-value classification $h : \mathcal{X} \rightarrow \mathbb{R}$ such that $c(x)$ is

equal to the sign of h^1 . The loss function $l(h, sy)$ in Eq. (5) is assumed to be convex. Our setting includes, for example, support vector machine and logistic classification. The following result bounds above the change in the solution of Eq. (5) when changing the weighting scheme from u to w .

Lemma 3.3. *Let ϕ be a feature mapping and k be its associated kernel with $k(x, x') = \langle \phi(x), \phi(x') \rangle$ and $\|k\|_\infty < \kappa < \infty$. Suppose that in Eq. (5), $Reg(h) = \lambda_c \|g\|_k^2$ and that for $x \in \mathcal{X}$, $h(x) = \langle h | k(x, \cdot) \rangle$. Suppose that l is σ -Lipchitz in its first argument. Denote h_u and h_w the solutions of the risk minimization Eq. (5) with weights u and w respectively. Then,*

$$\forall x \in \mathcal{X}, |h_u(x) - h_w(x)| \leq \kappa^2 \sigma \frac{\sqrt{\text{cond}(k)}}{\lambda_c} G_k(u, w),$$

where $\text{cond}(k)$ is the condition number of the Gram matrix of k and $G_k(u, w)$ is the maximum mean discrepancy between the distributions weighted by u and w :

$$G_k(u, w) = \left\| \sum_i u(x_i) \phi(x_i) - \sum_i w(x_i) \phi(x_i) \right\|.$$

Note that when the distribution is weighted with the importance sampling w_s , the maximum mean discrepancy of $Pr(X|S = s)$ and $Pr(X|S \neq s)$ is zero. By minimizing the maximum-mean discrepancy $G_k(u, w_s)$, we minimize an upper bound on the pointwise difference between h_u and h_{w_s} , that is the difference between the unfairness certificate we choose with weights u and the one we would have chosen if $Pr(X|S = s) = Pr(X|S \neq s)$. Therefore, **mdfa** solves:

$$\min_{\phi, u} \sum_i u_i(x_i) l(\phi(x_i), a_i y_i) + Reg(c) + \widehat{G}_k(u, a) \quad (6)$$

In our implementation, the feature representation ϕ is learned via a neural network that is then shared with both tasks of minimizing the re-weighted certifying risk and the empirical counterpart $\widehat{G}_k(u, s)$ of the maximum mean discrepancy between $Pr(X|S = s)$ and $Pr(X|S \neq s)$. The following result bounds above the bias in **mdfa**'s estimate of γ :

Theorem 3.4. *Let $\epsilon > 0$ and $\eta \in (0, 1)$. Suppose that \mathbb{C} is a concept class of VC dimension $d < \infty$. Solving for Eq. (6) finds a $\gamma - \epsilon$ -unfairness certificate if f is γ -unfair and at least $O\left(\frac{1}{\epsilon^2} \log(d) \log\left(\frac{1}{\eta}\right)\right)$ samples are queried.*

3.3 Worst-Case Violation

Solving the empirical minimization Eq. (6) allows certifying whether any black box classifier is multi-differential fair, but the solution of Eq. (6) does not distinguish a large sub-population S with low value of δ from a smaller sub-population with larger value of δ . For example, consider two sub-populations of same size G_0 and G_δ for $\delta > 0$. Assume that there is no violation of multi-differential fairness on G_0 ,

¹To prove our results, we will need $c(x) = g(h(x))$, where for small $\tau > 0$, $g(h) = \text{sign}(h)$ for $|h| > \tau/2$ and $g(h) = 1$ with probability h for $|h| < \tau/2$.

Algorithm 1 Worst Violation Algorithm (WVA)

Input: $\{(x_i, a_i), y_i\}_{i=1}^m$, $\mathbb{C} \subset 2^{\mathcal{X}}$, ξ, α , weights $u, y \in \{-1, 1\}$

- 1: $\alpha_0 = 1, u_{it} = u$
- 2: **while** $\alpha_t > \alpha$ **do**
- 3: $c_t = \underset{c \in \mathbb{C}}{\text{argmin}} \frac{1}{m} \sum_{i=1}^m u_{it}(x) l(a_i y_i, c(x_i)) + \lambda_c Reg(c)$
- 4: $\hat{\delta}_t \leftarrow \frac{\sum_{\substack{i=1, c(x_i)=1 \\ y_i=y, a_i=1}}^m u_i(x_i)}{\sum_{\substack{i=1, c(x_i)=1 \\ y_i=y, a_i=-1}}^m u_i(x_i)}$
- 5: $\hat{\alpha}_t \leftarrow \frac{\sum_{i=1, c(x_i)=1, y_i=1}^m u_i(x_i)}{\sum_{i=1}^m u_i(x_i)}$
- 6: $t \leftarrow t + 1, u_{it} \leftarrow u_{it} + u_i \xi$ if $a_i \neq y_i$ and $y_i = y$.
- 7: **end while**
- 8: **Return** $\ln(\hat{\delta}_t)$.

but a δ -violation on G_δ . The risk minimization Eq. 6 will pick indifferently G_δ and $G_\delta \cup G_0$ as unfairness certificates, although G_0 mixes the violation G_δ with a sub-population without any violation of differential fairness.

Worst-Case Violation Algorithm (WVA). At issue in the previous example is that for the sub-population G_0 , choosing $c = 1$ or $c = -1$ will lead to the same empirical risk Eq. 6. To force $c(x) = -1$ for $x \in G_0$, our approach is to put a slightly larger weight on samples whenever $s_i \neq y_i$. Now the empirical risk is smaller for $c = -1$ wherever there is no violation of multi differential fairness. More generally, our worst-violation algorithm 1 iteratively increases by $1 + \xi t$ the weight on samples whenever $s_i \neq y_i$, where $\xi > 0$. At iteration t , the solution c_t of the empirical risk minimization (5) identifies a sub-population $G_t = \{x | c_t(x) = 1\}$ with a $\delta(c_t)$ -violation of differential fairness, with $\delta \geq \ln((1 - h(\xi t))/h(\xi t))$, where h is an increasing function. The Algorithm 1 terminates whenever either $|G_t| \leq \alpha$. At the second to the last iteration T , theorem 3.5 guarantees that Algorithm 1 will identify a sub-population G_T with a δ_T -multi differential fairness violation and δ_T asymptotically close to δ_m .

Theorem 3.5. *Suppose $\xi > 0, \epsilon > 0, \eta \in (0, 1)$ and $\mathbb{C} \subset 2^{\mathcal{X}}$ is α -strong. Suppose that the classifier f has been certified with γ -multi-differential unfairness for $y \in \{-1, 1\}$. Denote δ_m the worst-case violation of multi differential fairness for \mathbb{C} as defined in (3). With probability $1 - \eta$, with $O\left(\frac{1}{\epsilon^2} \log(|\mathbb{C}|) \log\left(\frac{4}{\eta}\right)\right)$ samples and after $O\left(\frac{4(\gamma + \alpha)}{2\gamma + 3\alpha} \frac{2(4\gamma - 2\rho(y) + 1)}{\xi}\right)$ iterations, Algorithm 1 learns $c \in \mathbb{C}$ such that*

$$\left| \ln \left(\frac{Pr[Y = y | S = s, c(x) = 1]}{Pr[Y = y | S \neq s, c(x) = 1]} \right) - \delta_m \right| \leq \epsilon. \quad (7)$$

3.4 mdfa Auditor

Putting the building blocks together allows us to design a fairness diagnostic tool **mdfa** that identifies efficiently the most severe violation of differential unfairness.

Architecture. Inputs are a dataset with a classifier’s outcomes (labels ± 1) along with auditing features. **mdfa** first uses a neural network with four fully connected layers of 8 neurons to express the weights u as a function of the features x and minimizes the maximum-mean discrepancy function $\hat{G}_k(u, s)$. The outputs of the last hidden layer in the neural network are used as a feature mapping ϕ and serve along with the estimated weights u as an input to the empirical minimization Eq. (6), which outputs a certificate $(c, y, s) \in \mathbb{C} \times \{-1, 1\} \times \mathcal{S}$ of unfairness. The weights are re-adjusted until the identified worst-case violation has a size smaller than α . When terminating, **mdfa** outputs an estimate of the most-harmed sub-population c_m along with an estimate of δ_m .

Cross-Validation. The auditor chooses the minimum size α of the worst-case violation they would like to identify. The advantage of our approach is that, although we do not have ground truth for unfair treatment, we can propose heuristics to cross-validate our choice of regularization parameters used in Eq. (6). First, we split 70%/30% the input data into a train and test set. Using a 5-fold cross-validation, **mdfa** is trained on four folds and a grid search looks for regularization parameters that minimize the maximum-mean-discrepancy $\hat{G}_k(u, s)$ and the empirical risk on the fifth fold. Once **mdfa** is trained, the estimated δ_m and the corresponding characteristics of the most-harmed sub-population are computed on the test set.

4 Experimental Results

4.1 Synthetic Data

A synthetic data is constructed by drawing independently two features X_1 and X_2 from two normal distributions $N(0, 1)$. We consider a binary protected attribute $S = \{-1, 1\}$ drawn from a Bernoulli distribution with $S = 1$ with probability $w(x) = \frac{e^{\mu * (x_1 - x_2)}}{1 + e^{\mu * (x_1 + x_2)}}$. μ is the imbalance factor. $\mu = 0$ means that the data is perfectly balanced. The data is labeled according to the sign of $(X_1 + X_2 + e)^3$, where e is a noise drawn from $N(0, 0.2)$. The audited classifier f is a logistic regression classifier that is altered to generate instances of differential unfairness. For $x_1^2 + x_2^2 \leq 1$, if $S = -1$, the classifier’s outcomes Y is changed from -1 to 1 with probability $1 - \nu \in (0, 1]$; if $S = 1$, all $Y = -1$ are changed to $Y = 1$. For $\nu = 0$, the audited classifier is differentially fair; however, as ν increases, in the half circle $\{(x_1, x_2) | x_1^2 + x_2^2 \leq 1 \text{ and } y = -1\}$ there is a fraction ν of individuals with $S = 1$ who are not treated similarly as individuals with $S = -1$.

Results. First, we test whether Algorithm 1 identifies correctly the worst-case violation that occurs in the sub-space $\{(x_1, x_2) | x_1^2 + x_2^2 \leq 1 \text{ and } y = -1\}$. **mdfa** is trained using a support vector machine (RBF kernel) on a unbalanced data ($\mu = 0.2$) with value of δ_m varying from 0 to 3.0. Figure 1a plots the estimated $\hat{\delta}_m$ against the true one δ_m and shows that **mdfa**’s estimate $\hat{\delta}_m$ is unbiased. Figure 1b shows that at each iteration of the Algorithm 1, the estimated worst-case violation $\hat{\delta}_m$ progresses toward the true value δ_m . Secondly, we compare our balancing approach *MMD* to alter-

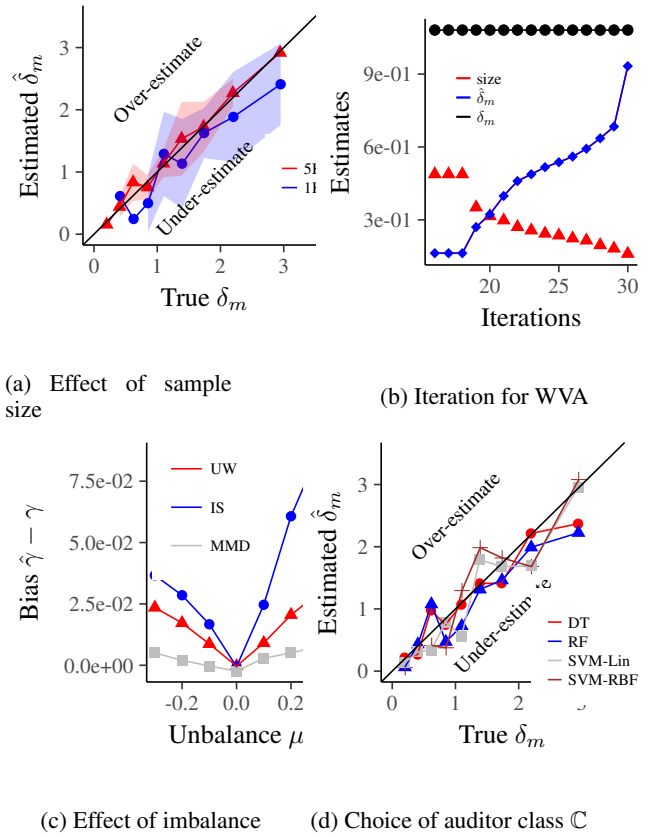


Figure 1: Performance of **mdfa** on synthetic data. If not precised otherwise, **mdfa** is trained with a support vector machine with $5K$ samples and imbalance factor $\mu = 0.2$

native re-balancing approaches: (i) uniform weights (*UW*) with $u(x) = 1/m$ for all x and (ii) importance sampling (*IS*) with exact weights $w(x)$. *UW* applies **mdfa** without re-balancing. *IS* uses oracle access to the importance sampling weights w , since they are known in this synthetic experiment. Figure 1c plots the bias $\hat{\gamma} - \gamma$ for each unfairness certificate obtained by **mdfa** with varying values of the imbalance factor μ . It shows that minimizing the maximum-mean discrepancy function is the only method that generates unbiased certificates regardless of data imbalance. Bias in estimates obtained with *UW* confirms that absent of a re-weighting scheme, **mdfa** cannot disentangle the information related to S leaked by the features x from the one leaked by the classifier’s outcomes y . Using importance sampling weights (*IS*) directly does not perform well: this confirms previous observations in the literature that in finite sample, the variance of the importance sample weights can be detrimental to a re-balancing approach. Lastly, Figure 1d shows that **mdfa**’s estimates of δ_m are robust to diverse classes of classifiers, including support vector machines with non-linear *SVM-RBF* or linear *SVM-Lin* kernels and random forest *RF*.

4.2 Case Study: COMPAS

We apply our method to the COMPAS algorithm, widely used to assess the likelihood of a defendant to become a recidivist

| Variable | Population | | Violation | |
|----------------------|----------------|----------------|----------------|----------------|
| | AA | Other | AA | Other |
| Prior Felonies | 4.44 (5.58) | 2.46 (3.76) | 0.79 (0.24) | 0.67 (0.17) |
| Charge Degree | 0.31 (0.46) | 0.4 (0.49) | 0.74 (0.23) | 0.74 (0.2) |
| Juvenile Felonies | 0.1 (0.49) | 0.03 (0.32) | 0.01 (0.02) | 0.0 (0.02) |
| Juvenile Misdemeanor | 0.14 (0.61) | 0.04 (0.3) | 0.01 (0.02) | 0.01 (0.01) |
| High Risk | 0.14 (0.35) | 0.05 (0.22) | 0.06 (0.04) | 0.02 (0.01) |

Table 1: Identifying the worst-case violation of differential fairness in the COMPAS risk score. The sensitive attribute is whether the individual is self-identified as African American (AA) or not (Other). () indicates standard deviation.

([ProPublica, 2016]). The research question is whether without knowledge of the design of COMPAS, **mdfa** can identify group of individuals that could argue for a disparate treatment. The data collected by ProPublica in Broward County from 2013 to 2015 contains 7K individuals along with a risk score and a risk category assigned by COMPAS. We transform the risk category into a binary variable equal to 1 for individuals assigned in the high risk category (risk score between 8 and 10). The data provides us with information related to the historical criminal history, misdemeanors, gender, age and race of each individual.

Worst Violations. We run **mdfa** on 100 different 70/30% train/test splits and report average value of auditing features and recidivism risk for the whole population and the worst-case subpopulation in Table 1. The first two columns show that the distribution of features in the whole population is disperse and differs between African American (AA) and Other. This is due to the data imbalance issue (c.f. Section 3). The probability of being classified as high risk is 0.14 for African-American, thereby 2.7 times higher than for non-African American. However, it is unclear whether that difference could be explained either by the distribution imbalance or by the classifier’s disparate treatment. The two last columns in Table 1 show that in the sub-population “violation” extracted by **mdfa**, the distribution of features is narrower and similar for African-American and non-African American: the sub-population is made of individuals with little criminal and misdemeanor history. However, African American are still three times more likely to be classified as high risk. A policy implication of **mdfa** findings is that a judge using COMPAS may discount its assessment for African-American with little criminal history.

4.3 Group Fairness vs. Multi-Differential Fairness

We evaluate whether previous fairness correcting approaches protect small group of individuals against violation of differential fairness. We consider two techniques: (i) [Feldman *et al.*, 2015]’s disparate impact with a logistic classification ($DI - LC$) and (ii) [Agarwal *et al.*, 2018]’s reduction with a logistic classification ($Red - LC$). We use **mdfa** to identify sub-population G with worst-case viola-

| Repair Technique | Adult | | German | | Crimes | |
|------------------|----------------|------|----------------|------|----------------|------|
| | DT_G | DI | DT_G | DT | DT_G | DI |
| LC | 1.88 (0.4) | 1.08 | 1.26 (0.14) | 1.07 | 5.76 (3.16) | 1.0 |
| DI-LC | 1.44 (0.32) | 0.99 | 1.1 (0.08) | 1.04 | 5.74 (2.19) | 1.0 |
| Red-LC | 1.6 (0.25) | 1.03 | 1.04 (0.21) | 1.01 | 5.24 (0.89) | 1.0 |

Table 2: Worst-case violations of multi-differential fairness identified by **mdfa** for classifiers trained with standard fairness repair techniques. () indicates standard deviation.

tions and measure disparate treatment as $DT_G = Pr[Y = 1|S = 1, G]/Pr[Y = 1|S = -1, G]$. We compare DT_G to its aggregate counterpart computed on the whole population $DI = Pr[Y = 1|S = 1]/Pr[Y = 1|S = -1]$.

Data. The experiment is carried on three datasets from [Friedler *et al.*, 2018; Kearns *et al.*, 2018]: **Adult** with 48840 individuals; **German** with 1000 individuals; and, **Crimes** with 1994 communities. In **Adult** the prediction task is whether an individual’s income is less than 50K and the sensitive attribute is gender; in **German**, the prediction task is whether an individual has bad credit and the sensitive attribute is gender; in **Crimes**, the task is to predict whether a community is in the 70th percentile for violent crime rates and the sensitive attribute is whether the percentage of African American is at least 20%. For each data, each repair technique produces a prediction; then, **mdfa** is trained on 70% of the data and computes estimates for disparate treatment DT_G on the remaining 30% of the data. The experiment is repeated with 100 train/test splits.

Results. In Table 2, even despite the fairness correction applied by $DI - LC$ and $Red - LC$, **mdfa** still finds sub-populations G for which DT_G is significantly larger than one. It indicates the existence of group of individuals who are similar but for their sensitive attributes and who are treated differently by the classifier trained by either $DI - LC$ or $Red - LC$. The repair techniques reduce the aggregate disparate impact compared to the baseline (LC), since DI is closer to one for $DI - LC$ and $Red - LC$ across all datasets. However, in the **Adult** dataset, DT_G remains between 1.44 and 1.6 after repair: **mdfa** identifies a group G of Females that are 44% – 60% more likely to be of low-income than Males with similar characteristics. In **Crimes** dataset, disparate treatment DT_G is around 5.7 for both $DI - LC$, $R - LC$: this means that there exist communities with dense African-American populations that are six times more likely to be classified at high risk than similar communities with lower percentages of African Americans.

5 Conclusion

In this paper, we present **mdfa**, a tool that measures whether a classifier treats differently individuals with similar auditing features but different sensitive attributes. We hope that **mdfa**’s ability to identify sub-populations with severe violations of differential fairness could inform decision-makers when to discount the classifier’s outcomes. It also provides

the victims with a framework to contest a classifier's outcomes.

Avenues for future research are to investigate (i) the properties of a classifier trained under a multi-differential fairness constraint; and, (ii) the possibility to extend our approach to re-balance distributions in order to make counterfactual inference [Johansson *et al.*, 2016] in the context of algorithmic fairness.

Acknowledgements

This work is partially supported by the National Science Foundation grant #1447489.

A Appendix: Proofs

Lemma 3.1.

Proof. Denote $\langle x, x' \rangle$ the inner product between x and x' . Observe that for $r = \pm$ the left-hand side in Eq. (4) can be written $\frac{1}{2} \langle \frac{c \pm 1}{2}, S \frac{1 \pm r Y}{2} \rangle$ since $\langle x, x' \rangle = Pr_w[x = x'] - 1$ for any $x, x' \in \{-1, 1\}$. The result from lemma 3.1 follows by remarking that $\langle S, 1 \rangle = \langle S, c \rangle = 2Pr_w[S = c] - 1 = 0$, since $Pr_w[S = s|x] = Pr_w[S \neq s|x]$. \square

Theorem 3.2.

Proof. (i) \Rightarrow (ii). Denote (x_i, s_i, o_i) a sample from a balanced distribution D over $\mathcal{X} \times \mathcal{S} \times \{-1, 1\}$. Denote $c^* \in \mathbb{C}$ such that $Pr[c^*(x_i) = o_i] = \max_{c \in \mathbb{C}} Pr[c(x_i) = o_i] = opt$. Construct a function f such that for (x_i, s_i, o_i) , $f(x_i, s_i) = s_i o_i$. Therefore, $f(x_i) s_i = o_i$ and $Pr[c^* = s_i f(x_i)] = Pr[c^* = o_i] = opt$: by lemma 3.1, c^* is a γ -unfairness certificate, with $\gamma = \frac{opt + \rho - 1}{4}$ and $\rho = Pr[o_i = 1]$. By (i), the certifying algorithm outputs a $(\gamma - \epsilon/4)$ -unfairness certificate $c \in \mathbb{C}$ with probability $1 - \eta$ and $O(\log(|\mathcal{C}|, \log(\frac{1}{\eta}), \frac{1}{\epsilon^2}))$ sample draws. Hence, by lemma 3.1, $Pr[c(x_i) = o_i] = Pr[c(x_i) = f(x_i) o_i] = 4(\gamma - \epsilon/4) + 1 - \rho = opt - \epsilon$, which concludes (i) \Rightarrow (ii).

(ii) \Rightarrow (i). Suppose that f is a γ -unfair. Denote $y_i = f(x_i, s_i)$. Samples $\{(x_i, s_i), y_i\}$ are drawn from a balanced distribution over $\mathcal{X} \times \mathcal{S} \times \{-1, 1\}$. By lemma 3.1, there exists $c \in \mathbb{C}$ such that $Pr[c(x_i) = s_i y_i] = 4\gamma + 1 - \rho_r$, with $r = \pm$. Assume, without loss of generality $r = +$. Then, since $\max_{c'} Pr[c(x_i) = s_i y_i] \geq 4\gamma + 1 - \rho_+$. By (ii), there exists an algorithm that outputs with probability $1 - \eta$ and $O(\log(|\mathcal{C}|, \log(\frac{1}{\eta}), \frac{1}{\epsilon^2}))$ sample draws $c \in \mathbb{C}$ such that $Pr[c(x_i) = s_i y_i] \geq \max_{c'} Pr[c(x_i) = s_i y_i] - \epsilon/4$. Therefore $Pr[c(x_i) = s_i y_i] \geq 4(\gamma - \epsilon) + 1 - \rho_+$. By lemma 3.1, c is a $(\gamma - \epsilon)$ -unfairness certificate for f , which concludes (ii) \Rightarrow (i). \square

Theorem 3.4. We first show the following lemma

Lemma A.1. *With the same assumption as in lemma 3.3, for any weights u, w , $\|u - w\| \leq G_k(u, w) / \sqrt{\lambda_{min}(k)}$, where $\lambda_{min}(k)$ is the smallest eigenvalue of the Gram matrix associated with k*

Proof. First, note that $G_k(u, w) = \sqrt{(u - w)^T k (u - w)}$ and by a standard bound on Rayleigh quotient, $\|u - w\| \leq G_k(u, w) / \sqrt{\lambda_{min}(k)}$. \square

Now we prove lemma 3.3:

Proof. The proof relies on the fact that the solution of Eq. (5) is distributionally stable as in [Cortes *et al.*, 2008]:

$$\|h_u - h_w\|_k \leq \kappa \sigma^2 \frac{\sqrt{\lambda_{max}(k)}}{2\lambda_c} \|u - w\|,$$

where $\lambda_{max}(k)$ is the largest eigenvalue of the Gram matrix associated to k (see [Cortes *et al.*, 2008], proof of theorem 1). Moreover, $\|h_u(x) - h_w(x)\| \leq \|h_u - h_w\|_k$. The result in lemma 3.3 follows from lemma A.1 and $cond(k) = \lambda_{max}(k) / \lambda_{min}(k)$. \square

The next result follows from [Gretton *et al.*, 2009] and bounds above $\hat{G}_k(u, w_s)$, the empirical counterpart of $G_k(u, w_s)$, where $w_s(x) = Pr[S \neq s|x] / (1 - Pr[S = s|x])$ are the importance sampling weights for $S = s$.

Lemma A.2. *Let $\eta > 0$. Denote $n_s = |\{i = 1, \dots, m | s_i = s\}|$ and $n_{\neq s} = |\{i = 1, \dots, m | s_i \neq s\}|$. Suppose that $\|u\|_\infty < B/n_s$ and that $E[u] < \infty$. There exists a constant $\kappa_1 > 0$ such that with probability $1 - \eta$,*

$$\hat{G}_k(u, w_s) \leq \kappa_1 \sqrt{2 \log \frac{2}{\eta} \left(\frac{B^2}{n_s} + \frac{1}{n_{\neq s}} \right)}$$

Proof. See [Gretton *et al.*, 2009] Lemma 1.5 \square

For $\tau > 0$, we construct $c \in \mathbb{C}$ from the real-valued function h as

$$c(x) = \begin{cases} \text{sign}(h) & \text{if } |h| > \tau \\ 1 \text{ w.p. } \frac{h + \tau}{2\tau} & \text{if } |h| \leq \tau \end{cases} \quad (8)$$

Lemma A.3. *Let $\tau, \eta > 0, \epsilon > 0$. Let \hat{c}_u, \hat{c}_w denote the certificates constructed from the solution of the empirical risk minimization \hat{h}_u (with weights u) and \hat{h}_w (with weights w_s). There exists $\kappa_2 > 0$ such that with probability $1 - \eta$,*

$$\frac{|\{i = 1, \dots, m | \hat{c}_u(x_i) \neq \hat{c}_w(x_i)\}|}{m} \leq \kappa_2 \sqrt{2 \log \frac{2}{\eta} \left(\frac{B^2}{n_s} + \frac{1}{n_{\neq s}} \right)}$$

Proof. Denote $\epsilon(m) = \frac{\kappa_2 \tau}{5} \sqrt{2 \log \frac{2}{\eta} \left(\frac{B^2}{n_s} + \frac{1}{n_{\neq s}} \right)}$, with

$\kappa_2 = \kappa_1 \kappa \sigma^2 \frac{\sqrt{cond(k)}}{2\lambda_c}$. By lemma 3.3 and lemma A.2, $\|\hat{h}_u - \hat{h}_w\|_2 \leq \epsilon_m$ with probability $1 - \eta$. Consider first the case $\hat{h}_w < -\tau$. Then, with probability $1 - \eta$, $\{x_i | \hat{c}_u(x_i) \neq \hat{c}_w(x_i) \wedge \hat{h}_w(x_i) < -\tau\} = \{x_i | (-\tau + \epsilon_m > \hat{h}_u(x_i) > -\tau) \wedge (\hat{c}_u(x_i) = 1) \wedge \hat{h}_w(x_i) < -\tau\}$ by lemma 3.3. Therefore, by construction of \hat{c}_u , $|\{x_i | \hat{c}_u(x_i) \neq \hat{c}_w(x_i) \wedge \hat{h}_w(x_i) < -\tau\}| = m \frac{\tau + \hat{h}_u}{2\tau} \leq m \frac{\epsilon_m}{5\tau}$. A similar result is obtained for $\hat{h}_w > \tau$.

Lastly, with probability $1 - \eta$, $|\{x_i | \hat{c}_u(x_i) \neq \hat{c}_w(x_i) \wedge \hat{h}_w(x_i) \in (-\tau, \tau)\}| \leq 2m \frac{\epsilon_m}{5\tau} + m \frac{|\hat{h}_u - \hat{h}_w|}{\tau} \leq 3 \frac{\epsilon_m}{5\tau}$. The first part of the inequality uses the results obtained in the previous paragraph for $|\hat{h}_u| > \tau$; the second part uses the construction of \hat{c}_u and \hat{c}_w . Therefore, with probability $1 - \eta$, $\frac{|\{i=1, \dots, m | \hat{c}_u(x_i) \neq \hat{c}_w(x_i)\}|}{m} \leq 5\epsilon(m)/\tau$. \square

Lemma A.4. Consider a random variable Z and \hat{c}_u, \hat{c}_w as in lemma A.3. Then $|Pr[Z = 1 \wedge \hat{c}_u = 1] - Pr[Z = 1 \wedge \hat{c}_w = 1]| \leq Pr[\hat{c}_u \neq \hat{c}_w]$

Proof. Note that $Pr[Z = 1 \wedge \hat{c}_u = 1] - Pr[Z = 1 \wedge \hat{c}_w = 1] = Pr[Z = 1 \wedge \hat{c}_u = 1 \wedge \hat{c}_w = -1] - Pr[Z = 1 \wedge \hat{c}_w = 1 \wedge \hat{c}_u = -1] \leq Pr[\hat{c}_u \neq \hat{c}_w]$. \square

The last result we need to prove theorem 3.4 is to link n_s and n_{-s} to sample size m

Lemma A.5. Denote $\alpha_s = Pr[S = s]$. Let $\epsilon, \eta > 0$. Therefore, if $m \geq \Omega\left(\frac{\log(1/\eta)}{\alpha_s^2 \epsilon^2}\right)$, with probability $1 - \eta$, $n_s \geq \alpha_s(1 - \epsilon)m$.

Proof. This is an application of a Hoeffding's inequality for Bernoulli random variable. \square

The proof of theorem 3.4 then follows from the fact that for a sample balanced with weights u , $\hat{\gamma}_u = Pr[\hat{c}_u = 1] \left(\frac{e^{\delta_u}}{e^{\delta_u} + 1} - \frac{1}{2}\right)$ with $e^{\delta_u} = Pr[Y = 1 | S = 1, \hat{c}_u = 1] / Pr[Y = 1 | S = -1, \hat{c}_u = 1]$. By lemmas A.4 and A.5, for any $\epsilon' > 0$ with $\Omega\left(\frac{1}{(\epsilon')^2 \alpha_s^2} \log \frac{2}{\eta}\right)$ samples, with probability $1 - \eta$, $|e^{\delta_u} - e^{\delta_w}| < \epsilon'$ and $|Pr[\hat{c}_w = 1 \wedge Y = 1] - Pr[\hat{c}_u = 1 \wedge Y = 1]| \leq \epsilon'$. There exists κ_3 such that for $\epsilon > 0$, with probability $1 - \eta$, $|\hat{\gamma}_u - \hat{\gamma}_w| \leq \kappa_3 \epsilon$. Moreover, by theorem 3.2, if \mathbb{C} has finite VC dimension, with probability $1 - \eta$ and $\Omega\left(\log(|\mathbb{C}|), \frac{1}{2\epsilon}, \log \frac{2}{\eta}\right)$ samples, $|\hat{\gamma}_w - \gamma_w| \leq \epsilon$. It follows that with probability $1 - \eta$ and $\Omega\left(\log(|\mathbb{C}|), \frac{1}{\epsilon^2 \alpha_s^2}, \log \frac{2}{\eta}\right)$ samples, $|\hat{\gamma}_u - \gamma_w| \leq (1 + \kappa_3)\epsilon$. That concludes the proof since by lemma 3.1, if f is γ -unfair, then $\gamma = \gamma_w$.

Theorem 3.5. Assume that the classifier f is γ -unfair for $Y = y \in \{-1, 1\}$. Let δ_m denote the worst-case violation. Note that by definition of γ and δ_m : $\gamma = \alpha \left(\frac{e^{\delta_m}}{e^{\delta_m} + 1} - \frac{1}{2}\right)$. At each iteration t , denote c_t the solution of the following optimization problem

$$\max_{c \in \mathbb{C}} E \left[\sum_{i=1}^m u_{it} \mathbb{1}_a(s_i y_i = c(x_i)) \right], \quad (9)$$

where u_{it} are the weights at iteration t and the expectation is taken over all the samples of size m drawn from D_f .

$$u_{it} = \begin{cases} u_i(1 + \nu t) & \text{if } y_i \neq s_i \wedge y_i = y \\ u_i & \text{otherwise.} \end{cases} \quad (10)$$

Lemma A.6. Let $s \in \mathcal{S}$. Assume that the classifier f is γ -unfair. At iteration t , denote $\delta_t = \ln(Pr[Y = 1 | c_t(x) = 1, S = s] / Pr[Y = 1 | c_t(x) = 1, S = -s])$. Then, $\frac{e^{\delta_t}}{1 + e^{\delta_t}} \geq 1 - h(\xi t)$, where $h(\xi t) = \frac{4\gamma + 1 - 2\rho_+}{\xi t}$.

Proof. Without loss of generality, we assume $y = 1$. Denote c_{-1} such that $c_{-1}(x) = -1$ for all x . By comparing the value

of the empirical risks for c_t and c_{-1} , if $c_t \neq c_{-1}$, then we can show that

$$Pr_u[c_t(x_i) = s_i y_i] \geq Pr_u[s_i \neq y_i] + \xi t E \left[\sum_{\substack{i=1, s_i \neq y_i \\ c_t(x_i)=1 \\ y_i=y}}^m u_i \right].$$

Moreover,

$$E \left[\sum_{\substack{i=1, s_i \neq y_i \\ c_t(x_i)=1 \\ y_i=y}}^m u_i \right] = Pr_u[s_i y_i \neq c_t(x_i) \wedge c_t(x_i) = 1 \wedge y_i = y].$$

It follows that if $c_t \neq c_{-1}$, since $Pr_u[c_t = 1, Y = y] \geq \alpha Pr_u[SY = c_t | c_t = 1, Y = y] \geq 1 - \frac{1}{\alpha \xi t} (Pr_u[c_t = SY]) - \rho_+$. Since the classifier f is γ -unfair for $y = 1$, we know that $\max_{c \in \mathbb{C}} Pr_u[SY = c] = 4\gamma + 1 - \rho_+$. Therefore, at iteration t , either $c_t = c_{-1}$ or

$$Pr_u[SY = c_t | c_t = 1, Y = y] \geq 1 - \frac{4\gamma + 1 - 2\rho_+}{\xi \alpha t} = 1 - h(\xi t).$$

Since $y = 1$, $Pr_u[Y = 1 | c_t = 1, S = 1] \geq 1 - h(\xi t)$. Without loss of generality, we can assume $Pr[Y = 1 | c_t = 1, S = 1] \geq 1 - h(\xi t)$. It follows that $e^{\delta_t} \geq \frac{1 - h(\xi t)}{h(\xi t)}$. \square

Lemma A.7. Denote $T = \frac{4\gamma + 1 - 2\rho(y)}{\xi \alpha} (e^{\delta_m} + 1)$. The algorithm stops for $t \geq T$ with $Pr[Y = y, C_t = 1] = \alpha$ and $\delta_t = \delta_m$.

Proof. First, note that $h(\xi T) = \frac{1}{e^{\delta_m} + 1}$ and thus that $\delta_T \geq \delta_m$. Moreover, at iteration t , c_t is chosen over c_m where c_m is the sub-population that corresponds to the worst-case violation δ_m . Comparing the expected risk for c_m and c_t at iteration t leads to

$$\begin{aligned} & 2Pr_u[SY = c_t] - 2Pr_u[SY = c_m] \\ & \geq \xi t E \left[\sum_{\substack{i=1 \\ s_i \neq y_i \\ c_t(x_i)=1 \\ y_i=1}}^m u_i - \sum_{\substack{i=1 \\ s_i \neq y_i \\ c_m(x_i)=1 \\ y_i=1}}^m u_i \right] \end{aligned} \quad (11)$$

$$\begin{aligned} & = (Pr_u[SY \neq c_t | c_t = 1, Y = 1] Pr_u[c_t = 1, Y = 1] \\ & - Pr[SY \neq c_m | c_m = 1, Y = 1] Pr[c_m = 1, Y = 1]) \end{aligned}$$

By definition of the worst-case violation for $y = 1$, $Pr_u[SY \neq c_t | c_t = 1, Y = 1] \geq Pr_u[SY \neq c_m | c_m = 1, Y = 1]$. Moreover, when the algorithm stops, $Pr_u[c_t = 1, Y = 1] = \alpha = Pr_u[c_m = 1, Y = 1]$. Therefore, the right-hand side of (11) is non-negative. It results that $Pr_u[SY = c_t] \geq Pr_u[SY = c_m]$. On the other hand, since f is γ -unfair, $Pr_u[SY = c_t]$ cannot be more than $4\gamma - \rho_+ + 1$. Therefore, $Pr_u[SY = c_t] = 4\gamma - \rho_+ + 1$.

It follows that at iteration t , when the algorithm stops

$$Pr_u[Y = 1, c_t = 1] \left(\frac{e^{\delta(c_t)}}{1 + e^{\delta(c_t)}} - \frac{1}{2} \right) = \gamma.$$

The algorithm stops when $Pr[Y = 1, c_t = 1] = \alpha$, which implies $\delta_t = \delta_m$. Therefore, $t \leq T$. \square

References

- [Agarwal *et al.*, 2018] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. *arXiv preprint arXiv:1803.02453*, 2018.
- [Atlantic, 2016] The Atlantic. How algorithms can bring down minorities credit scores? *The Atlantic*, 2016. <https://www.theatlantic.com/technology/archive/2016/12/how-algorithms-can-bring-down-minorities-credit-scores/509333/>.
- [Calsamiglia, 2009] Caterina Calsamiglia. Decentralizing equality of opportunity. *International Economic Review*, 50(1):273–290, 2009.
- [Chouldechova and Roth, 2018] Alexandra Chouldechova and Aaron Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018.
- [Cortes *et al.*, 2008] Corinna Cortes, Mehryar Mohri, Michael Riley, and Afshin Rostamizadeh. Sample selection bias correction theory. In *International Conference on Algorithmic Learning Theory*, pages 38–53. Springer, 2008.
- [Cortes *et al.*, 2010] Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weighting. In *Advances in neural information processing systems*, pages 442–450, 2010.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226. ACM, 2012.
- [Dwork *et al.*, 2014] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [Feldman *et al.*, 2012] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing*, 41(6):1558–1590, 2012.
- [Feldman *et al.*, 2015] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268. ACM, 2015.
- [Foulds and Pan, 2018] James Foulds and Shimei Pan. An intersectional definition of fairness. *arXiv preprint arXiv:1807.08362*, 2018.
- [Friedler *et al.*, 2018] Sorelle A Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P Hamilton, and Derek Roth. A comparative study of fairness-enhancing interventions in machine learning. *arXiv preprint arXiv:1802.04422*, 2018.
- [Gretton *et al.*, 2009] Arthur Gretton, Alex J Smola, Jiayuan Huang, Marcel Schmittfull, Karsten M Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. In *Dataset shift in machine learning*, pages 131–160. MIT Press, 2009.
- [Hébert-Johnson *et al.*, 2017] Ursula Hébert-Johnson, Michael P Kim, Omer Reingold, and Guy N Rothblum. Calibration for the (computationally-identifiable) masses. *arXiv preprint arXiv:1711.08513*, 2017.
- [Jagielski *et al.*, 2018] Matthew Jagielski, Michael Kearns, Jieming Mao, Alina Oprea, Aaron Roth, Saeed Sharifi-Malvajerdi, and Jonathan Ullman. Differentially private fair learning. *arXiv preprint arXiv:1812.02696*, 2018.
- [Johansson *et al.*, 2016] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029, 2016.
- [Kearns *et al.*, 2017] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. *arXiv preprint arXiv:1711.05144*, 2017.
- [Kearns *et al.*, 2018] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. An empirical study of rich subgroup fairness for machine learning. *arXiv preprint arXiv:1808.08166*, 2018.
- [Kim *et al.*, 2018] Michael P Kim, Omer Reingold, and Guy N Rothblum. Fairness through computationally-bounded awareness. *arXiv preprint arXiv:1803.03239*, 2018.
- [Loftus *et al.*, 2018] Joshua R Loftus, Chris Russell, Matt J Kusner, and Ricardo Silva. Causal reasoning for algorithmic fairness. *arXiv preprint arXiv:1805.05859*, 2018.
- [Mansour *et al.*, 2009] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- [ProPublica, 2016] ProPublica. How we analyzed the compas recidivism algorithm. *ProPublica*, 2016.
- [Rosenbaum and Rubin, 1983] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- [Zafar *et al.*, 2017] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1171–1180. International World Wide Web Conferences Steering Committee, 2017.