Optimal Transport of Classifiers to Fairness

Anonymous Author(s)
Affiliation
Address
email

Abstract

In past work on fairness in machine learning, the focus has been on forcing the prediction of classifiers to have similar statistical properties for people of different demographics. To reduce the violation of these properties, fairness methods usually simply rescale the classifier scores, ignoring similarities and dissimilarities between members of different groups. Yet, we hypothesize that such information is relevant in quantifying the unfairness of a given classifier. To validate this hypothesis, we introduce Optimal Transport to Fairness (OTF), a method that quantifies the violation of fairness constraints as the smallest Optimal Transport cost between a probabilistic classifier and any score function that satisfies these constraints. For a flexible class of linear fairness constraints, we construct a practical way to compute OTF as a differentiable fairness regularizer that can be added to any standard classification setting. Experiments show that OTF can be used to achieve an improved trade-off between predictive power and fairness.

1 Introduction

Machine learning methods are increasingly being deployed for automated decision making due to the potential benefits in terms of efficiency and effectiveness. Yet, making automated decisions about *people* comes with substantial legal and ethical risks, as evidenced by more and more cases of undesirable algorithmic discrimination [27, 13, 10] with respect to the protected traits of individuals.

The research field of *fairness* in machine learning [29] therefore studies ways in which a model's discrimination with respect to a person's *sensitive features* can be reduced or removed. Many notions of fairness have been proposed for binary classification tasks [36]. A popular example is *demographic parity*, which is expressed as a constraint that enforces equality between the rates at which positive decisions are made for each protected group [6].

Motivation For a given fairness notion, we would like to tune a model that (to some extent) satisfies the constraint that expresses the notion. To this end, it can be practical to quantify the violation of this constraint as a fairness regularization term that can be added to any probabilistic classifier [26]. Such an approach provides an incentive to the model to learn parameters that result in fair probability scores, e.g. by learning to ignore highly biased features. A straightforward fairness regularizer is the norm of the difference between the quantities that should be equal according to the fairness notion [39, 30].

However, such a quantification of unfairness may be limiting, as it only considers the statistical properties of the model's output, while ignoring the input features for which those scores were computed are ignored. We argue that, to quantify the unfairness of a model as a fairness regularizer,

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

it can be beneficial to more strongly consider data points with similar features but significant unfair discrepancies in probability scores. To this end, we employ *Optimal Transport* (OT) [32] as a way to measure these discrepancies while taking feature similarity into account.

Contributions

- We propose to quantify the unfairness of a probabilistic classifier as the *Optimal Transport to Fairness* (OTF) cost, which is defined as the smallest OT cost between the score function of the classifier and any fair score function over the same data.
- We make novel derivations to compute OTF as a *differentiable fairness regularizer* that can be added to the training loss of any probabilistic classifier and is efficiently computed for popular notions of fairness such as *demographic parity* and *equalised odds*. Our approach is capable of handling multiple sensitive variables, that can be categorical or continuous.
- In experiments, our method shows its benefit of increased flexibility over other OT methods. In some cases, it also achieves a trade-off between predictive power and fairness that is significantly more effective for equalized odds.

2 Background

2.1 Fair Classification

Let $\mathcal{Z} \triangleq \mathcal{X} \times \mathcal{S} \times \{0,1\}$ denote the sample space, from which we draw samples $Z \triangleq (X, S, Y)$ and try to estimate the *output* label $Y \in \{0,1\}$ from *input* features $X \in \mathcal{X}$, without discriminating with respect to *sensitive* features $S \in \mathcal{S}$. It is assumed that $\mathcal{X} \subset \mathbb{R}^{d_\mathcal{X}}$ and $\mathcal{S} \subset \mathbb{R}^{d_\mathcal{S}}$.We sample binary predictions $\hat{Y} \in \{0,1\}$ from a probabilistic classifier $h : \mathcal{X} \to [0,1]$ that assigns a score $h(X)$ to the belief that a sample with features X belongs to the positive class. Our general goal in fair classification is to minimize a loss function $\mathcal{L}_Y(h)$, in our case the cross-entropy between $h(X)$ and Y, while also scoring well on a fairness measurement of $h(X)$.

2.1.1 Group Fairness

The fairness of classifiers is commonly enforced through some notion of independence from sensitive information $S \in \mathcal{S}$. In particular, *group fairness* is concerned with constraints over statistical measurements of this independence over all members of a protected group [29]. For now, assume that $\mathcal{S}$ consists of a single, categorical feature with $d_\mathcal{S}$ values. We denote $S \in \mathcal{S}$ as the one-hot encoding[1] values $S_k \in \{0,1\}$. This encoding is indexed by values from $[d_\mathcal{S}] = \{0, ..., d_\mathcal{S} - 1\}$.

Assuming that the predictions $\hat{Y} \mid X$ are randomly sampled from a probabilistic classifier $h(X)$, then the traditional fairness notion of *demographic parity* (DP) is equivalent to enforcing zero covariance between $h(X)$ and each group in S, which we refer to as Probabilistic Demographic Parity (PDP):

$$\text{(PDP)} \qquad \forall k \in [d_\mathcal{S}] : \mathbb{E}_Z\left[h(X)S_k\right] - \mathbb{E}_Z[h(X)]\,\mathbb{E}_Z[S_k] = 0. \tag{1}$$

Remark 1. If $\hat{Y}$ is not sampled from $h(X)$ but instead decided by a threshold (e.g. $\hat{Y} = 1$ if $h(X) > 0.5$), then Eq. (1) is a relaxation of the actual DP notion and may therefore admit unfair models [28]. However, such a threshold introduces a discontinuity with respect to the input, making it difficult to efficiently find fair classifiers in practice [35, 1]. Moreover, by considering the probabilities directly, we account for the uncertainty expressed by the probabilistic classifier.

The notation in Eq. (1) is easily extended to fairness notions that condition on other variables. For example, the *equalized odds* (EO) [21] fairness notion conditions DP on the actual label Y. We apply EO to probabilistic classifiers by conditioning the covariance constraint of PDP on the one-hot encoding of Y and retrieve the following expression for Probabilistic Equalized Odds (PEO):

$$\text{(PEO)} \quad \forall l \in \{0,1\}, \forall k \in [d_\mathcal{S}] : \mathbb{E}_Z\left[h(X)S_kY_l\right] - \mathbb{E}_Z[h(X)Y_l]\,\mathbb{E}_Z[S_kY_l] = 0. \tag{2}$$

[1]For example, if $\mathcal{S}$ only contains two categories, then we denote $\mathcal{S} = \{[1,0],[0,1]\}$ with $|\mathcal{S}| = d_\mathcal{S} = 2$.

Table 1: Mapping of binary fairness notions onto the $\mathbf{g}$ functions of their probabilistic version to be written as linear fairness notions.

	Predictions Constraint	Linear Fairness Notion
(P)DP	$\mathbb{P}(\hat{Y}=1 \mid S)=\mathbb{P}(\hat{Y}=1)$	$\mathbf{g}_k=\frac{S_k}{\mathbb{E}_Z[S_k]}-1$
(P)EO	$\mathbb{P}(\hat{Y}=1 \mid S,Y)=\mathbb{P}(\hat{Y}=1 \mid Y)$	$\mathbf{g}_{k+ld_{\mathcal{S}}}=Y_l\left(\frac{S_k}{\mathbb{E}_Z[S_kY_l]}-1\right)$

2.1.2 Linear Fairness

The PDP and PEO covariance constraints in Eq. (1) and Eq. (2) are practical, because they can be written as linear sums over the probabilities given by the classifier h. As was done in past work [1], we are generally interested in all fairness notions that can be enforced through such linear constraints.

Definition 1 (Linear fairness notion)**.** A notion of fairness is a *linear fairness notion* when the set $\mathcal{F}$ of all score functions $f:\mathcal{X}\to[0,1]$ that satisfy it is given by

$$\mathcal{F}\triangleq\{f:\mathcal{X}\to[0,1]:\mathbb{E}_Z\left[\mathbf{g}(Z)f(X)\right]=\mathbf{0}_{d_{\mathcal{F}}}\} \tag{3}$$

with $\mathbf{0}_{d_{\mathcal{F}}}$ a vector of $d_{\mathcal{F}}$ zeros and the vector-valued function $\mathbf{g}(Z):\mathcal{Z}\to\mathbb{R}^{d_{\mathcal{F}}}$ constant w.r.t. $f(X)$.

PDP and PEO are clearly linear fairness notions, with their $\mathbf{g}$ function listed in Table 1. Linear constraints are practical because the set of fair score functions $\mathcal{F}$ is convex. Moreover, multiple linear fairness notions can be combined into a new linear fairness notion by concatenating their constraint vectors. Sensitive information that is inherently continuous, like age, can be considered as a single dimension in S, thereby allowing for a mix of categorical and continuous sensitive attributes.

However, under linear fairness, a non-linear dependence between a score function $f(X)$ and *continuous* sensitive attributes S can remain undetected. Also, some well-known fairness notions are not linear. For example, notions related to *sufficiency* involve conditioning on $f(X)$, meaning that $f(X)$ also shows up in the $\mathbf{g}$ function. To this end, approaches have been proposed [9] that can approximate classifiers with non-linear fairness through a reduction to problems with linear constraints.

2.2 Optimal Transport of Score Functions

Optimal Transport (OT) [31] theory considers the problem of moving a mass from one measure to another at the smallest possible total cost. Here, we let every score function f correspond with the measure θ_f defined over the *input* space $\mathcal{X}$ endowed with the Borel σ-algebra: $\theta_f\triangleq\sum_{x\in\mathcal{D}_{\mathcal{X}}}f(x)\delta_x$, with δ_x the Dirac measure and $\mathcal{D}_{\mathcal{X}}$ the input features of samples in a collection $\mathcal{D}$. Note that the input space measure θ_f is not normalized, though this is not necessary to apply OT theory. In what follows, we implicitly consider the score functions h and f as their corresponding input space measures θ_h and θ_f when used in the OT problem. See the Appendix for further clarification.

For a collection $\mathcal{D}$ of n samples (e.g. a full dataset or only a batch from it), let $\mathbf{h}$ and $\mathbf{f}$ denote the n-dimensional vectors of score function values for all data points, i.e. $\mathbf{h}_i=h(x_i)$ and $\mathbf{f}_i=f(x_i)$. Furthermore, for a non-negative transport cost function c defined over $\mathcal{X}\times\mathcal{X}$, let $\mathbf{C}\in\mathbb{R}_+^{n\times n}$ represent the matrix of cost terms, i.e. $\mathbf{C}_{ij}=c(x_i,x_j)$. Similarly, with $\pi(x_i,x_j)$ the *coupling* that reflects how much score mass was transported from x_i to x_j, define the matrix $\mathbf{P}\in\mathbb{R}_+^{n\times n}$ with $\mathbf{P}_{ij}=\pi(x_i,x_j)$. The OT cost is then simplified to

$$\mathrm{OT}(h,f)=\min_{\mathbf{P}\in\Pi(h,f)}\langle\mathbf{C},\mathbf{P}\rangle \tag{4}$$

with $$\Pi(h,f)=\left\{\mathbf{P}\in\mathbb{R}_+^{n\times n}:\mathbf{P}\mathbf{1}_n=\mathbf{h},\mathbf{P}^T\mathbf{1}_n=\mathbf{f}\right\}.$$

where $\mathbf{1}_n$ is the n-dimensional vector of ones and $\langle\mathbf{C},\mathbf{P}\rangle=\sum_{ij}\mathbf{C}_{ij}\mathbf{P}_{ij}$.

3 Optimal Transport to Fairness

Assume that a set $\mathcal{F}$ is available that denotes all score functions which satisfy the required notion of fairness. We would like to quantify the unfairness of a classifier's score function h as the amount

of 'work' minimally required to make it a member of this set by measuring how far h is from its fair *projection* onto $\mathcal{F}$, i.e. the $f \in \mathcal{F}$ that is *closest* to h. We measure this closeness as the OT cost between h and f, as we then assign a higher unfairness cost to h if scores need to be transported between highly dissimilar individuals in order to reach a fair function $f \in \mathcal{F}$.

Thus, we propose to quantify unfairness as the Optimal Transport to Fairness (OTF) cost, i.e. the cost of the OT-based projection of h onto $\mathcal{F}$:

$$\text{OTF}(h) = \min_{f \in \mathcal{F}} \text{OT}(h, f). \tag{5}$$

In what follows, we expand on the OTF method by constructing it in three steps.

1. In Sec. 3.1, we directly express the OTF(h) objective as a linear programming problem by making the assumption that $\mathcal{F}$ is defined by a linear fairness notion as in Def. 1.
2. In Sec. 3.2, we add entropic smoothing to OTF(h), thereby making the resulting OTF$_\epsilon(h)$ cost a differentiable fairness regularizer that can be computed efficiently.
3. In Sec. 3.3, we address the fact that due to this smoothing, OTF$_\epsilon(h)$ does not necessarily equal zero for a fair h. To this end, we subtract the adjustment term OTFR$_\epsilon(h)$, resulting in the adjusted OTF$^0_\epsilon(h)$ cost.

Finally in Sec. 3.4, we show how OTF$^0_\epsilon(h)$ can be minimized w.r.t. (the parameters of) h.

For all derivations and proofs, we refer to the Appendix.

3.1 Optimal Transport to Linear Fairness

To compute the minimization in Eq. (5), first note that in the OT(h, f) cost in Eq. (4), f only shows up in the constraint $\mathbf{P}^T\mathbf{1}_n = \mathbf{f}$ on the column marginals of coupling matrix $\mathbf{P}$. It therefore suffices to weaken this constraint to $\mathbf{P}^T\mathbf{1}_n \in \mathcal{F}$. Next, recall that a *linear* fairness notion is enforced through a vector of constraints on the expectation of $\mathbf{g}(Z)f(X)$. For a collection $\mathcal{D}$ of n samples, let $\mathbf{G}_{cj} = \mathbf{g}_c(z_j)$, i.e. $\mathbf{G} \in \mathbb{R}^{d_\mathcal{F} \times n}$ is the constraints matrix with a row for every constraint and a column for every data point. It defines the linear fairness notion for $\mathcal{D}$.

Definition 2 (OTF)**.** For a linear fairness notion expressed through constraints matrix $\mathbf{G}$ for n samples and with non-negative cost matrix $\mathbf{C}$, the *Optimal Transport to Fairness* cost for score function $h : \mathcal{X} \to [0, 1]$ is computed as

$$\text{OTF}(h) = \min_{\mathbf{P} \in \Pi^\mathcal{F}(h)} \langle \mathbf{C}, \mathbf{P} \rangle \tag{6}$$

with $$\Pi^\mathcal{F}(h) = \left\{ \mathbf{P} \in \mathbb{R}^{n \times n}_+ : \mathbf{P}\mathbf{1}_n = \mathbf{h}, \mathbf{G}\mathbf{P}^T\mathbf{1}_n = \mathbf{0}_{d_\mathcal{F}} \right\}.$$

The optimal coupling computed for OTF(h) implicitly transports the scores of $\mathbf{h}$ to the fair vector $\mathbf{f}$ given the coupling's column marginals (i.e. $\mathbf{f} = \mathbf{P}^T\mathbf{1}_n$). Note that for any coupling matrix that satisfies the first constraint ($\mathbf{P}\mathbf{1}_n = \mathbf{h}$), we have that $\mathbf{h}^T\mathbf{1}_n = \mathbf{1}_n^T\mathbf{P}^T\mathbf{1}_n = \mathbf{f}^T\mathbf{1}_n$, meaning that $\mathbf{h}$ and this implicitly found vector $\mathbf{f}$ sum to the same value. Scores of h are therefore only *transported* to a fair score function and no extra scores are created or destroyed.

For $\mathbf{P}_{ij} \geq 0$, the optimization problem in Eq. (6) can be solved through linear programming. However, its solution is not differentiable with respect to the score function h [12, 17]. It is thus impractical as a fairness regularization term.

3.2 Entropic Smoothing

A well-known trick to make OT costs differentiable is to apply *entropic regularization* or *smoothing* [12, 32], which is done by maximizing the entropy of the coupling while optimizing the OT cost. We apply the same principle to make OTF differentiable.

Definition 3 (OTF_ϵ)**.** For a linear fairness notion expressed through constraints matrix $\mathbf{G}$ for n samples and with non-negative cost matrix $\mathbf{C}$, the *smooth OTF* cost with smoothing strength $\epsilon > 0$ for score function $h : \mathcal{X} \to (0, 1]$ is computed as

$$\text{OTF}_\epsilon(h) = \min_{\mathbf{P} \in \Pi^{\mathcal{F}}(h)} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}) \tag{7}$$

with $H(\mathbf{P}) = -\sum_{ij} \mathbf{P}_{ij} \left(\log \left(\mathbf{P}_{ij} \right) - 1 \right)$ the entropy of coupling $\mathbf{P}$ and $\epsilon > 0$ a hyperparameter that regulates the strength of entropic smoothing.

The use of an entropy term in Eq. (7) requires some justification, since we do not assume that $\mathbf{h}$ represents a normalized probability distribution (i.e. sums to one). Therefore, $\mathbf{P}$ also does not represent a normalized joint distribution. However, the $H(\mathbf{P})$ term still provides the practical advantage that it only admits couplings with $\mathbf{P}_{ij} > 0$, which is why we require that $\mathbf{h}_i > 0$. Furthermore, the addition of $H(\mathbf{P})$ makes the objective cost strongly convex.

3.2.1 Duality

The OTF problem for a linear fairness notion has $n + d_{\mathcal{F}}$ constraints, which is usually far fewer in number than the n^2 variables that make up the couplings in the primal problem. Because OTF_ϵ is also strongly convex, we will derive and then maximize the dual function instead.

Proposition 1. *If $\mathbf{G}$ has a non-empty null-space (i.e. $\exists \mathbf{f} \in \mathbb{R}^n : \mathbf{f} \neq \mathbf{0}_n \wedge \mathbf{G}\mathbf{f} = \mathbf{0}_{d_{\mathcal{F}}}$), then the $\text{OTF}_\epsilon(h)$ problem has a unique solution for any $h : \mathcal{X} \to (0, 1]$. Moreover, the OTF_ϵ problem is strongly convex and then enjoys strong duality. Its dual function is*

$$L(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_i \boldsymbol{\lambda}_i \mathbf{h}_i - \epsilon \sum_{ij} \exp \left(\frac{1}{\epsilon} \left[-\mathbf{C}_{ij} + \boldsymbol{\lambda}_i + \sum_c \boldsymbol{\mu}_c \mathbf{G}_{cj} \right] \right) \tag{8}$$

with $\boldsymbol{\lambda} \in \mathbb{R}^n$ and $\boldsymbol{\mu} \in \mathbb{R}^{d_{\mathcal{F}}}$ the dual variable vectors for the marginalization and fairness constraints.

3.2.2 Optimization

Though Eq. (8) could be maximized directly, we follow standard OT approaches [32] and perform our optimization with *exact* coordinate ascent. This strategy is particularly useful here, because the $\boldsymbol{\lambda}_i$ variable that maximizes $L(\boldsymbol{\lambda}, \boldsymbol{\mu})$ while $\boldsymbol{\mu}$ is fixed, is found independently of other variables in $\boldsymbol{\lambda}$. All variables in $\boldsymbol{\lambda}$ can therefore be updated at the same time:

$$\boldsymbol{\lambda}_i \leftarrow \epsilon \log \mathbf{h}_i - \epsilon \log \sum_j \exp \left(\frac{1}{\epsilon} \left[-\mathbf{C}_{ij} + \sum_c \boldsymbol{\mu}_c \mathbf{G}_{cj} \right] \right) \tag{9}$$

where we can use the stabilized log-sum-exp operation.

Unfortunately, there is no closed form expression for the $\boldsymbol{\mu}_c$ that maximizes $L(\boldsymbol{\lambda}, \boldsymbol{\mu})$. Instead, we preprocess $L(\boldsymbol{\lambda}, \boldsymbol{\mu})$ and numerically solve for each $\boldsymbol{\mu}_c$:

$$\boldsymbol{\mu}_c \leftarrow \operatorname*{argmin}_{\boldsymbol{\mu}_c} \sum_j \boldsymbol{\eta}_j(\boldsymbol{\lambda}) \exp \left(\frac{1}{\epsilon} \sum_{k \neq c} \boldsymbol{\mu}_k \mathbf{G}_{kj} \right) \exp \left(\frac{1}{\epsilon} \boldsymbol{\mu}_c \mathbf{G}_{cj} \right) \tag{10}$$

with $\boldsymbol{\eta}_j(\boldsymbol{\lambda}) = \sum_i \exp \left(\frac{1}{\epsilon} \left[-\mathbf{C}_{ij} + \boldsymbol{\lambda}_i \right] \right)$.

3.2.3 Complexity

The dual problem involves two variable vectors: $\boldsymbol{\lambda}$ with length n the number of samples and $\boldsymbol{\mu}$ with length $d_{\mathcal{F}}$ the number of linear fairness constraints. It is commonly the case that $d_{\mathcal{F}} \ll n$, because the number of distinguished protected groups is limited. This is in contrast with the traditional OT_ϵ problem, where the dual problem involves two dual variable vectors of length n.

In terms of computational complexity, the update of the full $\boldsymbol{\lambda}$ has complexity $O(n(n + d_{\mathcal{F}}))$. However, we stress that each $\boldsymbol{\lambda}_i$ can be updated in parallel with complexity $O(n + d_{\mathcal{F}})$. When

keeping the $\boldsymbol{\lambda}$ vector fixed, we can perform a $O(n^2)$ operation to precompute the $\boldsymbol{\eta}_j(\boldsymbol{\lambda})$ values for the $\boldsymbol{\mu}$ updates. Though these updates no closed form, the inner loop of the update for μ_c only has complexity complexity $O(n + d_{\mathcal{F}})$ for each inner loop, and should only be performed for $d_{\mathcal{F}}$ variables. By permitting a high tolerance on the convergence of $\boldsymbol{\lambda}$, the necessary number of updates is also limited. In our experiments, we found that a single update is often enough.

We can achieve a memory complexity of $O(n + d_{\mathcal{F}})$. However, we can perform efficient matrix computations by storing the cost matrix C and the fairness constraints matrix G. The memory required for practical use is therefore $O(n(n + d_{\mathcal{F}}))$.

3.3 Adjusted OTF_ϵ

A cause for concern with the smooth OTF_ϵ cost is that, as opposed to the standard OTF, it does not necessarily equal zero if h is already a fair score function. Indeed, while it may be feasible to achieve $\langle \mathbf{C}, \mathbf{P} \rangle = 0$, it is generally the case that the entropy $H(\mathbf{P}) > 0$. We would thus like to find an interesting lower bound on OTF_ϵ that is tight when $h \in \mathcal{F}$.

Definition 4 (OTFR_ϵ)**.** The OTF_ϵ cost in Def. 3 is relaxed to the OTFR_ϵ cost as

$$\text{OTFR}_\epsilon(h) = \min_{\mathbf{P} \in \Pi^{\mathcal{F}}(h)} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}) \tag{11}$$

with $$\Pi^{\mathcal{F}}(h) = \left\{ \mathbf{P} \in \mathbb{R}_+^{n \times n} : \mathbf{P}\mathbf{1}_n = \mathbf{h}, |\mathbf{G}\mathbf{P}^T \mathbf{1}_n| \leq |\mathbf{G}\mathbf{h}| \right\}.$$

Definition 5 (OTF_ϵ^0)**.** For a score function $h : \mathcal{X} \to (0, 1]$, the *adjusted* OTF_ϵ *cost* is computed as

$$\text{OTF}_\epsilon^0(h) = \text{OTF}_\epsilon(h) - \text{OTFR}_\epsilon(h). \tag{12}$$

Whereas the $\text{OTF}_\epsilon(h)$ problem involved transport to fair score vectors ($\mathbf{G}\mathbf{P}^T\mathbf{1}_n = \mathbf{0}_{d_{\mathcal{F}}}$), the $\text{OTFR}_\epsilon(h)$ problem relaxes this constraint by allowing unfairness up to the unfairness already present in h. It is thus easily seen that $\text{OTF}_\epsilon(h) \geq \text{OTFR}_\epsilon(h)$. Interestingly, because this relaxation is bounded by the unfairness of h, we have that $\text{OTFR}_\epsilon(h)$ is a tight lower bound if h is already fair.

Proposition 2. $h \in \mathcal{F} \implies \text{OTF}_\epsilon^0(h) = 0$.

In the Appendix, we provide empirical intuition for this theoretical result by using $\text{OTF}_\epsilon^0(h)$ as a postprocessing approach and illustrate that $\text{OTF}_\epsilon(h)$ gradually converges to $\text{OTFR}_\epsilon(h)$ as the fairness of h improves.

The adjustment term OTFR_ϵ can be computed by using strong duality in a similar way as $\text{OTF}_\epsilon(h)$, though now with two dual variables $\boldsymbol{\phi} \in \mathbb{R}^{d_{\mathcal{F}}}$ and $\boldsymbol{\psi} \in \mathbb{R}^{d_{\mathcal{F}}}$ that enforce the fairness inequality, in addition to the dual variable vector, here denoted by $\boldsymbol{\kappa} \in \mathbb{R}^n$, that enforces the row constraint on the coupling. The dual function of OTFR_ϵ is

$$L(\boldsymbol{\kappa}, \boldsymbol{\phi}, \boldsymbol{\psi}) = \sum_i \boldsymbol{\kappa}_i \mathbf{h}_i + \sum_c \boldsymbol{\gamma}_c(\boldsymbol{\phi}_c + \boldsymbol{\psi}_c) - \epsilon \sum_{ij} \exp\left(\frac{1}{\epsilon} \left[-\mathbf{C}_{ij} + \boldsymbol{\kappa}_i + \sum_c (\boldsymbol{\phi}_c - \boldsymbol{\psi}_c) \mathbf{G}_{jc} \right] \right)$$

with $\boldsymbol{\phi}_c \leq 0$, $\boldsymbol{\psi}_c \leq 0$ and $\boldsymbol{\gamma}_c = |\sum_j \mathbf{G}_{cj}\mathbf{h}_j|$. For a full discussion, we refer to the Appendix.

3.4 Minimising Unfairness

Given an efficient way to compute the parameters of the adjusted, smooth $\text{OTF}_\epsilon^0(h)$ cost, we can now jointly minimize the weighted sum of the classifier's loss and unfairness cost with respect to h.

Let $\mathcal{L}_Y(h)$ denote the cross-entropy loss of $h(X)$ for the output labels Y. To add our $\text{OTF}_\epsilon^0(h)$ cost to this objective, we can pose the following optimization problem:

$$\min_h (1 - \alpha)\mathcal{L}_Y(h) + \alpha \text{OTF}_\epsilon^0(h) \tag{13}$$

with $0 \leq \alpha \leq 1$ a hyperparameter that regulates the importance of the unfairness cost term $\text{OTF}_\epsilon^0(h)$.

So far, we considered the scores vector of h as constant with respect to the $\text{OTF}_\epsilon^0(h)$ problem. However, the solutions for $\text{OTF}_\epsilon(h)$ and $\text{OTFR}_\epsilon(h)$ clearly depend on h. Let $\boldsymbol{\lambda}^*$, $\boldsymbol{\mu}^*$, $\boldsymbol{\kappa}^*$, $\boldsymbol{\phi}^*$ and $\boldsymbol{\psi}^*$ denote the fully converged dual variables. Then

$$\frac{\partial}{\partial h}\text{OTF}_\epsilon^0(h) = \frac{\partial}{\partial h}\left(L(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) - L(\boldsymbol{\kappa}^*, \boldsymbol{\phi}^*, \boldsymbol{\psi}^*)\right) \tag{14}$$

Even if we did not update the dual variables until convergence, we can still use intermediate values to approximate the true gradient [18, 16]. We further approximate the true $\text{OTF}_\epsilon^0(h)$ cost by only computing it for a subset of the dataset, e.g. when batching is already used to optimize $\mathcal{L}_Y(h)$. The $\text{OTF}_\epsilon^0(h)$ can be optimized on a batch by only using the values of $\mathbf{C}$ and the columns of $\mathbf{G}$ that relate to these data points. For large datasets (including those we use in experiments), it is then unnecessary to precompute the $\mathbf{C}$ matrix between all pairs of data points, since most pairs will never actually co-occur in randomly sampled batches.

4 Related Work

Fairness in machine learning is commonly understood in two (extreme) ways [29]. First, the notion of *individual* fairness states that similar individuals should be treated similarly [15]. Second, *group* fairness requires that as a group, people are treated equally. The primary goal of our work is to achieve group fairness by measuring the distance of a score function to a set of group-fair functions. Yet, to an extent we also approach individual fairness because our unfairness quantification takes the features of each individual into account in its cost function c. By appropriately choosing c, it may be possible to further improve upon individual fairness in future work.

Methods that aim to improve fairness can be categorized into three types: those that perform *preprocessing* where the data is made more fair [25, 8], as opposed to those that do *postprocessing* where the model's predictions are modified to fit a notion of fairness [21]. As a fairness regularization term, our work clearly fits in the class of *inprocessing* methods, i.e where the machine learning model or its training is directly altered to reduce unfairness [7, 39, 38, 1, 9]. However, we highlight that the OTF cost could also be used as a postprocessing approach by setting $\alpha = 1$.

Our work draws inspiration from two kinds of inprocessing methods. First, with those methods where a (still unfair) model is explicitly projected unto the set of fair models. A popular definition of distance for this projection is the Kullback-Leibler divergence [2, 23, 5]. Second, our work uses Optimal Transport (OT) [32] theory. Several works have applied OT to fairness, though they almost all do so by finding barycenters [34]: measures that are 'equally far' from the measures belonging to each demographic. This type of method has been applied to preprocessing [19], classification [23, 40] and regression [11, 20], with the overarching goal being to move each group's measure closer to the barycenter. Such methods require fairness notions to be encoded in the format of explicit equality, between a limited amount of groups. They do not allow for more complex constraints such as those involving continuous or multiple sensitive attributes, and make strong assumption that the cost function is a distance metric.

Our proposed method instead combines the idea of projecting the model to a fairness-constrained set and the use of OT to measure the cost of this projection. In this manner, our work is most similar to [33], which computes the OT projection onto such a fair set. However, they only use the method as a group fairness test. Our method has the advantage that it computes a differentiable cost that can be added to the loss of any probabilistic classifier.

5 Experiments

Our experiments were performed on two staple datasets from fair binary classification literature [14] with different kinds of sensitive information. First, the UCI *Adult* Census Income dataset with two binary sensitive attributes SEX and RACE. Second, the UCI *Bank* Marketing dataset, where we measure fairness with respect to the original continuous AGE values and the binary, quantized version AGE_BINNED. Additional information on datasets and hyperparameters is given in the Appendix.

5.1 Methods

All experiments were performed using a logistic regression classifier. To achieve fairness, we jointly optimize fairness regularizers with the cross-entropy loss as in Eq. (13), and compute the gradient of their sum with respect to the parameters of the classifier.

The OTF cost is the adjusted version from Def. 5, with $\epsilon = 10^{-3}$. For the cost function c, we use the Euclidean distance between non-protected features. The fairness constraints for both sensitive features in the Adult dataset were combined by concatenating their $\mathbf{G}$ matrices.

As a first baseline, we use the L^1 norm of the PDP and PEO fairness constraints and will refer to this method as the *Norm* cost. This simple type of fairness regularizer is frequently proposed in literature [26, 39, 37, 30]. In our experiments, it thus represents the class of methods that can accommodate a flexible range of fairness constraints while only considering the output of a classifier and not the input features. Our second baseline is an OT *Barycenter*[2] approach [24]. This method directly minimizes the OT cost between the score distributions of two sensitive groups. However, the implementation can only minimize PDP, and does not admit a composition of linear fairness notions or continuous attributes. We therefore only minimize unfairness with respect to SEX and AGE_BINNED. To show the drawbacks of this limitation, we also report PDP unfairness with respect to RACE and the continuous AGE. The third baseline is *Unfair*, which was a classifier trained without a fairness cost term.

5.2 Evaluation

To measure the predictive power of the methods, the ROC AUC score was measured. Violation of PDP is quantified as the absolute value of the maximal Pearson correlation between the scores $h(X)$ and each dimension of the sensitive attributes S_k. For PEO, the violation is computed as the maximal PDP for the data points of each one-hot encoded output label Y_l. These measures mirror the requirements in Eq. (1) and Eq. (2) that the covariances ought to equal zero, though we choose the Pearson correlation instead such that the violation is normalized between 0 and 1. Thus we avoid that a method may improve on linear fairness notions by squeezing the variance of its scores.

All methods were tested for various strengths of α (see Eq. (13)). Each configuration of each method (i.e. each α value and fairness notion) was tested for five train-test splits with proportions 80/20. We report mean test set scores in Fig. 1 and show the confidence ellipse[3] of their estimator for the first standard deviation. The train set results are reported in the Appendix and show the same trends.

Note that for the OTF and Norm regularizers, we omit fairness scores for notions that were not optimized in that configuration. While SEX and RACE are minimized at the same time, the AGE and AGE_BINNED attributes were minimized in separate training runs.

All experiments were conducted using half the hyperthreads on an internal machine equipped with a 12 Core Intel(R) Xeon(R) Gold processor and 256 GB of RAM. In total, all experiments, including preliminary and failed experiments, used this machine for approximately 200 hours.

5.3 Results

The results in the Adult dataset in Fig. 1a show that our OTF method makes a similar trade-off between AUC and PDP fairness as the baselines for various α values. However, OTF clearly outperforms Norm for both sensitive attributes when they are trained to minimize PEO, especially in the lower violation range. This advantage may be due to the fact that OTF has the incentive to achieve fairness by exchanging classifier scores between similar individuals. Conditioning on the output label rewards such exchanges, as individuals with similar labels can be expected to have similar features. Note that the Barycenter approach was not trained to improve fairness with respect to RACE, as the implementation can not handle multiple sensitive attributes. Our proposed OTF method therefore displays its advantage in both an effective AUC-fairness trade-off, and a flexible capacity for fairness.

[2] Retrieved from `https://github.com/deepmind/wasserstein_fairness` (Apache License 2.0)

[3] We use the variance of the estimator of the mean, which is smaller than that of individual scores

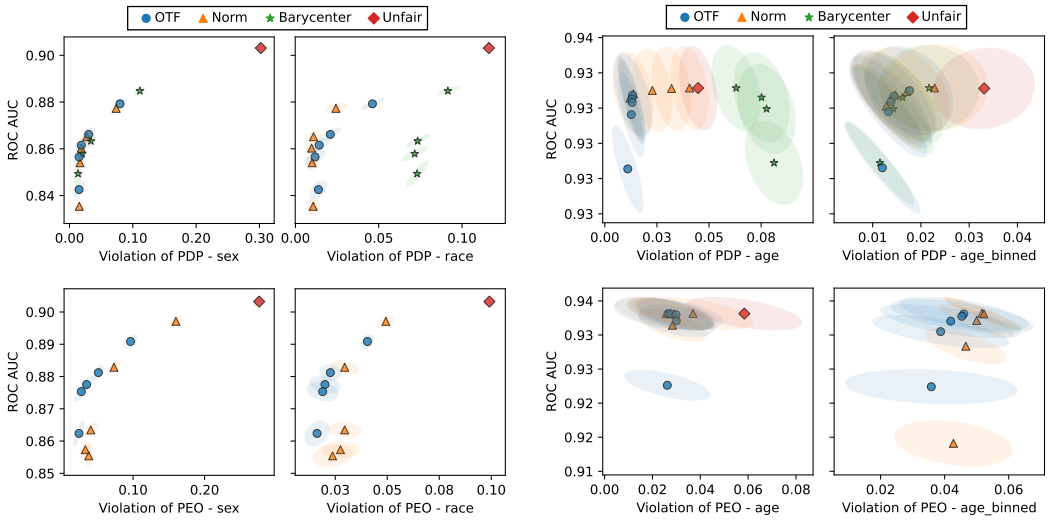

(a) Adult dataset

(b) Bank dataset

Figure 1: Test set results for the methods that were trained to reduce the evaluated fairness measure (PDP or PEO). Violation of PDP (and PEO) is computed as the maximal absolute Pearson correlation between the probability scores (conditioned on the output labels) and each sensitive attribute.

The test set results on the Bank dataset in Fig. 1b are clearly subject to a high amount of noise. Still, we see that OTF and Norm can be used to effectively minimize either PDP or PEO with respect to the continuous AGE attribute. Implementations of differentiable regularizers that can accommodate such continuous attributes are scarcely available. The Barycenter method can only minimize the OT cost between a limited number of score distributions, meaning it is naturally ill-equipped for continuous attributes. In separate runs, we therefore also trained OTF and Norm to minimize unfairness with respect to AGE_BINNED. We then see similar results as for the Adult dataset: our method is on par with the baselines for PDP, but clearly outperforms the Norm baseline for PEO.

5.4 Limitations and Impact

We nuance our results by pointing out that the OTF method has a higher computation complexity ($O(n(n + d_{\mathcal{F}}))$) than Norm, which has complexity $O(nd_{\mathcal{F}})$ with n the batch size and $d_{\mathcal{F}}$ the number of fairness constraints. Furthermore, as discussed in Sec. 2.1, we only evaluate *linear* fairness notions and leave an extension to non-linear notions of fairness [9] for future work.

The fairness of a system should always be judged based on a holistic consideration of the context of the system [3] and relevant ethical principles [4]. Without it, any method to improve fairness may reinforce the underlying injustice that led to the risks of discrimination that we aim to solve [22].

6 Conclusion

In this paper, we constructed the Optimal Transport to Fairness (OTF) method as a differentiable fairness regularization term, which combines the advantages of flexible linear fairness constraints with Optimal Transport theory, thereby taking non-protected similarities between individuals into account. In experiments, we show that our final method achieves a similar AUC-fairness trade-off as other cost terms for notions of fairness inspired by demographic parity, and significantly better for notions inspired by equalized odds. At the same time, OTF empirically displays why its increased flexibility, i.e. its ability to improve all linear fairness notions with respect to multiple sensitive attributes that may be categorical or continuous, provides an advantage over previous fairness applications of OT.

References

[1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudik, John Langford, and Hanna Wallach. A Reductions Approach to Fair Classification. In *Proceedings of the 35th International Conference on Machine Learning*, pages 60–69. PMLR, July 2018.

[2] Wael Alghamdi, Shahab Asoodeh, Hao Wang, Flavio P. Calmon, Dennis Wei, and Karthikeyan Natesan Ramamurthy. Model Projection: Theory and Applications to Fair Machine Learning. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2711–2716. IEEE, June 2020.

[3] Cynthia L. Bennett and Os Keyes. What is the point of fairness? disability, AI and the complexity of justice. *ACM SIGACCESS Accessibility and Computing 125*, page 5:1, March 2020.

[4] Reuben Binns. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 514–524, Barcelona Spain, January 2020. ACM.

[5] Maarten Buyl and Tijl De Bie. The KL-Divergence Between a Graph Model and its Fair I-Projection as a Fairness Regularizer. In *Machine Learning and Knowledge Discovery in Databases*, pages 351–366. Springer International Publishing, 2021.

[6] Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building Classifiers with Independency Constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pages 13–18, December 2009.

[7] Toon Calders and Sicco Verwer. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, September 2010.

[8] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. Optimized Pre-Processing for Discrimination Prevention. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[9] L. Elisa Celis, Lingxiao Huang, Vijay Keswani, and Nisheeth K. Vishnoi. Classification with Fairness Constraints: A Meta-Algorithm with Provable Guarantees. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 319–328, Atlanta GA USA, January 2019. ACM.

[10] Rumman Chowdhury. Sharing learnings about our image cropping algorithm. https://blog.twitter.com/engineering/en_us/topics/insights/2021/sharing-learnings-about-our-image-cropping-algorithm, May 2021.

[11] Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. Fair Regression with Wasserstein Barycenters. *arXiv:2006.07286 [cs, math, stat]*, June 2020.

[12] Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.

[13] Jeffrey Dastin. Amazon scraps secret AI recruiting tool that showed bias against women | Reuters. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G, October 2018.

[14] Dheeru Dua and Graff Casey. UCI Machine Learning Repository. http://archive.ics.uci.edu/ml, 2017.

[15] 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*, ITCS '12, pages 214–226, New York, NY, USA, January 2012. Association for Computing Machinery.

[16] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouve, and Gabriel Peyré. Interpolating between Optimal Transport and MMD using Sinkhorn Divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690. PMLR, April 2019.

[17] Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya-Polo, and Tomaso Poggio. Learning with a Wasserstein loss. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, pages 2053–2061, Cambridge, MA, USA, December 2015. MIT Press.

[18] Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning Generative Models with Sinkhorn Divergences. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, March 2018.

[19] Paula Gordaliza, Eustasio Del Barrio, Gamboa Fabrice, and Jean-Michel Loubes. Obtaining Fairness using Optimal Transport Theory. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2357–2365. PMLR, May 2019.

[20] Thibaut Le Gouic, Jean-Michel Loubes, and Philippe Rigollet. Projection to Fairness in Statistical Learning. *arXiv:2005.11720 [cs, math, stat]*, June 2020.

[21] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[22] Anna Lauren Hoffmann. Where fairness fails: Data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7):900–915, June 2019.

[23] Heinrich Jiang and Ofir Nachum. Identifying and Correcting Label Bias in Machine Learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 702–712. PMLR, June 2020.

[24] Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. Wasserstein Fair Classification. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, pages 862–872. PMLR, August 2020.

[25] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, October 2012.

[26] Toshihiro Kamishima, Shotaro Akaho, and Jun Sakuma. Fairness-aware Learning through Regularization Approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, pages 643–650, Vancouver, BC, Canada, December 2011. IEEE.

[27] Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How We Analyzed the COMPAS Recidivism Algorithm. https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm, May 2016.

[28] Michael Lohaus, Michael Perrot, and Ulrike Von Luxburg. Too Relaxed to Be Fair. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6360–6369. PMLR, November 2020.

[29] MehrabiNinareh, MorstatterFred, SaxenaNripsuta, LermanKristina, and GalstyanAram. A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys (CSUR)*, July 2021.

[30] Manisha Padala and Sujit Gujar. FNNC: Achieving fairness through neural networks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI'20, pages 2277–2283, Yokohama, Yokohama, Japan, January 2021.

[31] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, February 2019.

[32] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. *arXiv:1803.00567 [stat]*, March 2020.

[33] Nian Si, Karthyek Murthy, Jose Blanchet, and Viet Anh Nguyen. Testing Group Fairness via Optimal Transport Projections. *arXiv:2106.01070 [cs, math, stat]*, June 2021.

[34] Chiappa Silvia, Jiang Ray, Stepleton Tom, Pacchiano Aldo, Jiang Heinrich, and Aslanides John. A General Approach to Fairness with Optimal Transport. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3633–3640, April 2020.

[35] Bahar Taskesen, Jose Blanchet, Daniel Kuhn, and Viet Anh Nguyen. A Statistical Test for Probabilistic Fairness. *arXiv:2012.04800 [cs, stat]*, December 2020.

[36] Sahil Verma and Julia Rubin. Fairness definitions explained. In *Proceedings of the International Workshop on Software Fairness*, pages 1–7, Gothenburg Sweden, May 2018. ACM.

[37] Michael Wick, swetasudha panda, and Jean-Baptiste Tristan. Unlocking Fairness: A Trade-off Revisited. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[38] Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro. Learning Non-Discriminatory Predictors. In *Proceedings of the 2017 Conference on Learning Theory*, pages 1920–1953. PMLR, June 2017.

[39] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness Constraints: A Flexible Approach for Fair Classification. *Journal of Machine Learning Research*, 20(75):1–42, 2019.

[40] Meike Zehlike, Philipp Hacker, and Emil Wiedemann. Matching code and law: Achieving algorithmic fairness with optimal transport. *Data Mining and Knowledge Discovery*, 34(1):163–200, January 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] In Sec. 3, we proposed the OTF cost term and deployed tricks to minimize it effectively. In Sec. 5, we show that the results of using our proposed method are close to the baselines or surpassing it.
   (b) Did you describe the limitations of your work? [Yes] We describe limitations of our considered fairness notions in Sec. 2.1 and the complexity of our proposed OTF cost term in Sec. 3.2.3. We briefly discuss overall limitations of fairness methods in Sec. 5.4.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We briefly discuss the danger that inconsiderately applied fairness methods can pose in Sec. 5.4
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Our method is constructed in several steps. When required, we discuss and add additional assumptions.
   (b) Did you include complete proofs of all theoretical results? [Yes] Proofs of all propositions and derivations of other results are included in the Appendix.
3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] There are references to the datasets we used in the Appendix. Code and a pipeline to run it are included in the supplemental material.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Initial training details are given in Sec.5.2. More details of the hyperparameters are included in the Appendix.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Since we show the trade-off between two goals (accuracy and fairness) we show the confidence ellipses.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] This is mentioned at the end of Sec. 5.2.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] For the Barycenter code, a reference to the Github implementation is included. For datasets, we refer to the UCI repository.

(b) Did you mention the license of the assets? [Yes] For the Barycenter code, we include the license. We refer to the data according to the requested citation policy.

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] The datasets we use are popular, public datasets in machine learning fairness literature. Obtaining consent is infeasible. However, the data is highly anonymized and its popularity promotes reproducibility in the community.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] As briefly mentioned in the Appendix, the data is anonymous and free from any personal identifiers.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]