
Wasserstein Fair Classification

Ray Jiang*
DeepMind
rayjiang@google.com

Aldo Pacchiano*
UC Berkeley, DeepMind
pacchiano@berkeley.edu

Tom Stepleton
DeepMind
stepleton@google.com

Heinrich Jiang
Google Research
heinrichj@google.com

Silvia Chiappa
DeepMind
csilvia@google.com

Abstract

We propose an approach to fair classification that enforces independence between the classifier outputs and sensitive information by minimizing Wasserstein-1 distances. The approach has desirable theoretical properties and is robust to specific choices of the threshold used to obtain class predictions from model outputs. We introduce different methods that enable hiding sensitive information at test time or have a simple and fast implementation. We show empirical performance against different fairness baselines on several benchmark fairness datasets.

1 INTRODUCTION

The increasing use of machine learning in decision-making scenarios that have serious implications for individuals and society, such as health care, criminal risk assessment, social services, hiring, financial lending, and online advertising (De Fauw et al., 2018; Dieterich et al., 2016; Eubanks, 2018; Hoffman et al., 2018; Malekipirbazari and Aksakalli, 2015; Perlich et al., 2014), is raising concern that bias in the data and model inaccuracies can lead to decisions that are “unfair” towards unrepresented or historically discriminated groups.

This concern has motivated researches to investigate ways of ensuring that sensitive information (e.g. race and gender) does not unfairly influence the decisions. In the classification case considered in this paper, the most widely used approach is to enforce statistical independence between class predictions and sensitive attributes, a criterion called demographic parity (Feldman et al., 2015).

* Equal contribution.

In the common scenario in which the model outputs continuous values from which class predictions are obtained through thresholds, this approach would however ensure fairness only with respect to the particular choice of thresholds. Furthermore, as independence constraints on the class predictions are difficult to impose in practice, uncorrelation constraints on the model outputs are often imposed instead.

In this paper, we propose an approach that overcomes these limitations by imposing independence constraints directly on the model outputs. This is achieved through enforcing small Wasserstein distances between the distributions of the model outputs corresponding to groups of individuals with different sensitive attributes. We demonstrate that using Wasserstein-1 distances to the barycenter is optimal, in the sense that it achieves independence with minimal changes to the class predictions that would have been obtained without constraints. We introduce a Wasserstein-1 penalized logistic regression method that learns the optimal transport map in the logistic model parameters, with a variation that has the advantage of being demographically blind at test time. In addition, we provide a simpler and faster post-processing method. We show that the proposed methods outperform previous approaches in the literature on four benchmark fairness datasets.

2 STRONG DEMOGRAPHIC PARITY

Let $\mathcal{D} = \{(\mathbf{a}^n, \mathbf{x}^n, y^n)\}_{n=1}^N$ be a sequence of N *i.i.d.* samples drawn from an unknown probability distribution over $\mathcal{A} \times \mathcal{X} \times \mathcal{Y} = \mathbb{N}^k \times \mathbb{R}^d \times \{0, 1\}$. Each datapoint $(\mathbf{a}^n, \mathbf{x}^n, y^n)$ corresponds to information from an individual (or community): y^n indicates a binary class, each element a_i^n of \mathbf{a}^n corresponds to a different sensitive attribute, e.g. to the gender of the individual, and \mathbf{x}^n is a feature vector that, possibly together with \mathbf{a}^n , can be used to form a prediction $\hat{y}^n \in \{0, 1\}$ of the class y^n . We denote with $\mathcal{D}_{\mathbf{a}} = \{(\mathbf{a}^n, \mathbf{x}^n, y^n) \in \mathcal{D} \text{ s.t. } \mathbf{a}^n = \mathbf{a}\}$ the set

of N_a individuals belonging to group a . We indicate with $\mathbf{A}, \mathbf{X}, Y$ and \hat{Y} the random variables corresponding to \mathbf{a}^n, y^n and \hat{y}^n , and with $p(\cdot)$ or $p_X(\cdot)$ probability density functions (pdfs), where the latter is used to emphasize the associated random variable.

Many classifiers, rather than a binary class prediction \hat{y}^n , output a non-binary value s^n . In the logistic regression case considered in this paper, $s^n \in \Omega = [0, 1]$ indicates the *model belief* that individual n belongs to class 1, i.e. $s^n = \mathbb{P}(Y = 1 | \mathbf{A} = \mathbf{a}^n, \mathbf{X} = \mathbf{x}^n)$ ¹. From s^n , a class prediction $\hat{y}^n \in \{0, 1\}$ is obtained using a threshold $\tau \in \Omega$, i.e. $\hat{y}^n := \mathbb{1}_{s^n > \tau}$, where $\mathbb{1}_{s^n > \tau}$ equals to one if $s^n > \tau$ and zero otherwise. We call the random variable S corresponding to s^n the *belief variable*, and denote with S_a the belief variable for group a , i.e. with pdf $p(S_a) = p(S | \mathbf{A} = a)$.

We are interested in ensuring that sensitive information does not influence the decisions. This is often achieved by imposing that the model satisfies a fairness criterion called *demographic parity (DP)*, defined as

$$\mathbb{P}(\hat{Y} = 1 | \mathbf{A} = a) = \mathbb{P}(\hat{Y} = 1 | \mathbf{A} = \bar{a}), \quad \forall a, \bar{a} \in \mathcal{A}.$$

DP can equivalently be expressed as requiring statistical independence between \hat{Y} and \mathbf{A} , denoted as $\hat{Y} \perp \mathbf{A}$.

Enforcing demographic parity at a given threshold τ does not necessarily imply that the criterion is satisfied for other thresholds. Furthermore, to alleviate difficulties in optimizing on the class prediction \hat{Y} , relaxations are often considered, such as imposing the constraint $\mathbb{E}[S | \mathbf{A} = a] = \mathbb{E}[S | \mathbf{A} = \bar{a}] \quad \forall a, \bar{a} \in \mathcal{A}$, where $\mathbb{E}[\cdot]$ denotes expectation (Goh et al., 2016; Zafar et al., 2017).

To deal with these limitations, we propose an approach that enforces statistical independence between S and \mathbf{A} , $S \perp \mathbf{A}$. We call this fairness criterion *strong demographic parity (SDP)*, as it ensures that the decision does not depend on the sensitive attribute regardless of the threshold τ used, since $S \perp \mathbf{A}$ implies $\hat{Y} \perp \mathbf{A}$ for any value of τ . SDP can be defined as

$$p_{S_a} = p_{S_{\bar{a}}}, \quad \forall a, \bar{a} \in \mathcal{A}.$$

In Remark 1, we prove that this definition is equivalent to²

$$\mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_a > \tau) - \mathbb{P}(S_{\bar{a}} > \tau)| = 0, \quad \forall a, \bar{a} \in \mathcal{A},$$

¹Throughout the paper, we use $\mathbb{P}(\cdot)$ to indicate probability measures associated with the corresponding probability spaces $(\mathcal{O}, \mathcal{F}, \mathbb{P}(\cdot))$ where \mathcal{F} is a σ -algebra on the sample output space \mathcal{O} .

²We omit the brackets from the expectation to simplify the notation.

where $U(\Omega)$ denotes the uniform distribution over Ω . This result leads us to use

$$\sum_{\substack{a, \bar{a} \in \mathcal{A} \\ s.t. a \neq \bar{a}}} \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_a > \tau) - \mathbb{P}(S_{\bar{a}} > \tau)|,$$

as a measure of dependence of S on \mathbf{A} , the we call *strong pairwise demographic disparity (SPDD)*.

3 WASSERSTEIN FAIR CLASSIFICATION

We suggest to achieve SDP by enforcing the model output pdfs corresponding to groups of individuals with different sensitive attributes, $\{p_{S_a}\}_{a \in \mathcal{A}}$, to coincide with their Wasserstein-1 barycenter distribution $p_{\bar{S}}$. The use of the Wasserstein distance is motivated because this distance is defined and computable even between distributions with disjoint supports. This is critical because the empirical estimates $\{\hat{p}_{S_a}\}, \hat{p}_{\bar{S}}$ of $\{p_{S_a}\}$ and $p_{\bar{S}}$ used to implement the methods and their supports are typically disjoint.

3.1 OPTIMALITY OF WASSERSTEIN-1 DISTANCE

Preliminary. Given two pdfs p_X and p_Y on \mathcal{X} and \mathcal{Y} , a transportation map $T : \mathcal{X} \rightarrow \mathcal{Y}$ is defined by $\int_{\mathcal{B}} p_Y(y) dy = \int_{T^{-1}(\mathcal{B})} p_X(x) dx$ for any measurable subset $\mathcal{B} \subset \mathcal{Y}$ (indicating that the mass of the set \mathcal{B} with respect to the density p_Y equals the mass of the set $T^{-1}(\mathcal{B})$ with respect to the density p_X). Let \mathcal{T} be the set of transportation maps from \mathcal{X} to \mathcal{Y} , and $c : \mathcal{X} \times \mathcal{Y} \rightarrow [0, \infty]$ be a cost function such that $c(x, T(x))$ indicates the cost of transporting x to $T(x)$. In the original formulation (Monge, 1781), the optimal transport map T^* is the one that minimizes the total transportation cost, i.e.

$$T^* = \arg \min_{T \in \mathcal{T}} \int_{x \in \mathcal{X}} c(x, T(x)) p_X(x) dx.$$

To address limitations of this formulation, Kantorovich (1942) reformulated the optimal transport problem as finding an optimal pdf $p_{X \times Y}$ in the set $\Gamma(p_X, p_Y)$ of joint pdfs on $\mathcal{X} \times \mathcal{Y}$ with marginals over Y and X given by p_X and p_Y such that

$$\gamma^* = \arg \min_{\gamma \in \Gamma(f_X, f_Y)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) p_{X \times Y} dx dy.$$

The p -Wasserstein distance is defined as

$$\mathcal{W}_p(p_X, p_Y) = \min_{\gamma \in \Gamma(f_X, f_Y)} \left(\int_{\mathcal{X} \times \mathcal{Y}} d(x, y)^p p_{X \times Y} dx dy \right)^{\frac{1}{p}},$$

where $\mathcal{X} = \mathcal{Y}$, d is a distance on \mathcal{X} , and $p \geq 1$.

Fair Optimal Post-Processing. Let us first consider the problem of post-processing the beliefs of a model to achieve SDP while making minimal model class prediction changes.

Let S_1 and S_2 be two belief variables with values in $\Omega = [0, 1]$ and pdfs p_{S_1} and p_{S_2} , and let $T : \Omega \rightarrow \Omega$ be a transportation map satisfying $\int_{\mathcal{B}} p_{S_2}(y) dy = \int_{T^{-1}(\mathcal{B})} p_{S_1}(x) dx$ for any measurable subset $\mathcal{B} \subset \Omega$. Let \mathcal{T} be the set of all such transportation maps. A class prediction $\hat{y} = \mathbb{1}_{s_1 > \tau}$ changes due to transportation $T(s_1)$ if and only if $\tau \in (m_{s_1}^T, M_{s_1}^T)$ where $m_{s_1}^T = \min[s_1, T(s_1)]$ and $M_{s_1}^T = \max[s_1, T(s_1)]$. This observation leads to the following result.

Proposition 1. *Given two belief variables S_1 and S_2 in $\Omega = [0, 1]$ with pdfs p_{S_1} and p_{S_2} , the following three quantities are equal:*

- (i) $\mathcal{W}_1(p_{S_1}, p_{S_2}) = \min_{T \in \mathcal{T}} \int_{x \in \Omega} |x - T(x)| p_{S_1}(x) dx.$
- (ii) $\mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_1 > \tau) - \mathbb{P}(S_2 > \tau)|.$
- (iii) *Expected class prediction changes due to transporting p_{S_1} into p_{S_2} through the map T^**

$$\mathbb{E}_{\tau \sim U(\Omega), x \sim p_{S_1}} \mathbb{P}(\tau \in (m_x^{T^*}, M_x^{T^*})).$$

Proof. In the one-dimensional case of p_{S_1} and p_{S_2} , the total transportation cost $\mathcal{W}_1(p_{S_1}, p_{S_2})$ can be written as

$$\begin{aligned} \mathcal{W}_1(p_{S_1}, p_{S_2}) &= \int_{x=0}^1 |P_{S_1}^{-1}(x) - P_{S_2}^{-1}(x)| dx \\ &= \int_{\tau=0}^1 |P_{S_1}(\tau) - P_{S_2}(\tau)| d\tau \\ &\quad \text{(by Lemma 6 in Appendix C)} \\ &= \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_1 \leq \tau) - \mathbb{P}(S_2 \leq \tau)| \\ &= \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_1 > \tau) - \mathbb{P}(S_2 > \tau)|, \end{aligned}$$

where P_{S_1} and P_{S_2} are the cumulative distribution functions of S_1 and S_2 respectively. This prove that (i) equals (ii).

The expected class prediction changes due to applying the transportation map T is given by

$$\begin{aligned} \mathbb{E}_{\tau \sim U(\Omega)} \mathbb{P}_{x \sim p_{S_1}}(\tau \in (m_x^T, M_x^T)) \\ &= \int_{\tau=0}^1 \int_x |x - T(x)| p_{S_1}(x) dx d\tau \\ &= \int_x |x - T(x)| p_{S_1}(x) dx. \end{aligned}$$

³The proof of this equality can be found in Rachev and Rüschendorf (1998).

Thus,

$$\begin{aligned} \mathcal{W}_1(p_{S_1}, p_{S_2}) &= \min_{T \in \mathcal{T}} \int_x |x - T(x)| p_{S_1}(x) dx \\ &= \int_x |x - T^*(x)| p_{S_1}(x) dx \\ &= \mathbb{E}_{\tau \sim U(\Omega)} \mathbb{P}_{x \sim p_{S_1}}(\tau \in (m_x^{T^*}, M_x^{T^*})). \end{aligned}$$

This prove that (i) equals (iii). \square

Remark 1. *Notice that (ii) = $\mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_1 > \tau) - \mathbb{P}(S_2 > \tau)| = 0$ if and only if $p_{S_1} = p_{S_2}$. Indeed, by Proposition 1 and the property of the \mathcal{W}_1 metric, (ii) = 0 $\iff \mathcal{W}_1(p_{S_1}, p_{S_2}) = 0 \iff p_{S_1} = p_{S_2}$.*

To reach SDP, we need to achieve $p_{S_a} = p^* \forall a \in \mathcal{A}$, where $p^* \in \mathcal{P}(\Omega)$, the space of pdfs on Ω . We would like to choose transportation maps T and a target distribution p^* such that the transportation process from p_{S_a} to p^* incurs minimal total expected class prediction changes. Assume that the groups are all disjoint, so that the per-group transportation maps T are independent from each other. Let $\mathbb{T}(p^*)$ be the set of transportation maps with elements T such that, restricted to group a , T is a transportation map from p_{S_a} to p^* (i.e. $\mathbb{T}(p^*) = \{T \in \mathcal{T}(p_S, p^*) \mid T(S)|_{A=a} = T_a \in \mathcal{T}_a = \mathcal{T}(p_{S_a}, p^*)\}$ where $\mathcal{T}(p_S, p^*)$ denotes the space of transportation maps from p_S to p^*). We would like to obtain

$$\begin{aligned} \min_{\substack{T \in \mathbb{T}(p^*) \\ p^* \in \mathcal{P}(\Omega)}} \mathbb{E}_{\tau \sim U(\Omega)} \mathbb{P}_{x \sim p_S}(\tau \in (m_x^T, M_x^T)) \\ &= \min_{\substack{T \in \mathbb{T}(p^*) \\ p^* \in \mathcal{P}(\Omega)}} \sum_{\mathbf{a} \in \mathcal{A}} \underbrace{p(\mathbf{A} = \mathbf{a})}_{p_{\mathbf{a}}} \mathbb{E}_{\tau \sim U(\Omega)} \mathbb{P}_{x \sim p_{S_{\mathbf{a}}}}(\tau \in (m_x^T, M_x^T)) \\ &= \min_{p^* \in \mathcal{P}(\Omega)} \sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \min_{T \in \mathcal{T}_{\mathbf{a}}} \mathbb{E}_{\tau \sim U(\Omega)} \mathbb{P}_{x \sim p_{S_{\mathbf{a}}}}(\tau \in (m_x^T, M_x^T)) \\ &= \min_{p^* \in \mathcal{P}(\Omega)} \sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \min_{T \in \mathcal{T}_{\mathbf{a}}} \int_{x \in \Omega} |x - T(x)| p_{S_{\mathbf{a}}}(x) dx \\ &= \min_{p^* \in \mathcal{P}(\Omega)} \sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*). \end{aligned}$$

Therefore we are interested in

$$p_{\bar{S}} = \arg \min_{p^* \in \mathcal{P}(\Omega)} \sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*), \quad (1)$$

which coincides with the Wasserstein-1 barycenter with normalized subgroup size as weight to every group distribution p_{S_a} (Agueh and Carlier, 2011).

In summary, we have demonstrated that the optimal post-processing procedure that minimizes total expected model prediction changes is to use the Wasserstein-1 optimal transport map T^* to transport all group distributions p_{S_a} to their weighted barycenter distribution $p_{\bar{S}}$.

Optimal Trade-Offs. We have shown that post-processing the beliefs of a model through optimal transportation achieves SDP (and therefore SPDD = 0) whilst minimizing expected prediction changes. We now examine the case in which, after transportation, SDP is not attained, *i.e.* SPDD is positive. By triangle inequality

$$\begin{aligned} \text{SPDD} &\leq 2(|\mathcal{A}|-1) \sum_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_{\mathbf{a}} > \tau) - \mathbb{P}(\bar{S} > \tau)| \\ &= 2(|\mathcal{A}|-1) \sum_{\mathbf{a} \in \mathcal{A}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{\bar{S}}). \end{aligned}$$

We call this upper bound on SPDD *pseudo-SPDD*. Pseudo-SPDD is the tightest upper bound to SPDD among all possible target distributions by the definition of the barycenter $p_{\bar{S}}$ and Proposition 1. Indeed

$$\begin{aligned} &\sum_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_{\mathbf{a}} > \tau) - \mathbb{P}(\bar{S} > \tau)| \\ &= \sum_{\mathbf{a} \in \mathcal{A}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{\bar{S}}) \leq \sum_{\mathbf{a} \in \mathcal{A}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{S^0}) \\ &= \sum_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\tau \sim U(\Omega)} |\mathbb{P}(S_{\mathbf{a}} > \tau) - \mathbb{P}(S^0 > \tau)|, \end{aligned}$$

for any distribution $p_{S^0} \in \mathcal{P}(\Omega)$. Since SPDD is difficult to derive optimal trade-offs for, we do that with respect to the pseudo-SPDD as the measure of fairness instead.

We are interested in changing $p_{S_{\mathbf{a}}}$ to $p_{S_{\mathbf{a}}^*}$, $\forall \mathbf{a} \in \mathcal{A}$, to reach a fairness bound $\lambda \in \mathbb{R}_+$ for pseudo-SPDD such that the required model prediction changes are minimal in expectation. This is obtained by choosing the $p_{S_{\mathbf{a}}^*} \in \mathcal{P}(\Omega)$ that minimizes the total expected prediction changes, which equals $\sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{S_{\mathbf{a}}^*})$ by Proposition 1, while bounding the pseudo-SPDD by λ , *i.e.* $\sum_{\mathbf{a} \in \mathcal{A}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{\bar{S}}) \leq \lambda$. Assuming that the groups are disjoint, we can optimize each group transportation in turn independently assuming the other groups are fixed. This gives

$$\begin{aligned} p_{S_{\mathbf{a}}^*} &= \arg \min_{p^* \in \mathcal{P}(\Omega) \text{ s.t. } \mathcal{W}_1(p^*, p_{\bar{S}}) \leq \lambda - \gamma} p_{\mathbf{a}} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*) \\ &= \arg \min_{p^* \in \mathcal{P}(\Omega) \text{ s.t. } \mathcal{W}_1(p^*, p_{\bar{S}}) \leq \lambda - \gamma} \mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*), \end{aligned}$$

where $\gamma = \sum_{\bar{\mathbf{a}} \in \mathcal{A} \setminus \mathbf{a}} \mathcal{W}_1(p_{S_{\bar{\mathbf{a}}}}, p_{\bar{S}})$. By triangle inequality, $\mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*) \geq |\mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{\bar{S}}) - \mathcal{W}_1(p^*, p_{\bar{S}})|$. The distance $\mathcal{W}_1(p_{S_{\mathbf{a}}}, p^*)$ reaches its minimum if and only if p^* lies on a shortest path between $p_{S_{\mathbf{a}}}$ and $p_{\bar{S}}$. Thus it is optimal to transport $p_{S_{\mathbf{a}}}$ along any shortest path between itself and $p_{\bar{S}}$ in the Wasserstein-1 metric space. In the approach proposed in the next section, we approximate transporting group distributions along these shortest paths with hyperparameter tuning of a gradient descent method to minimize $\mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{\bar{S}})$ for every group.

Empirical Computation of the Barycenter. In practice, as building the barycenter from the population distributions $p_{S_{\mathbf{a}}}$ is impossible, we use the empirical distributions $\hat{p}_{S_{\mathbf{a}}}$ obtained from $\mathcal{D}_{\mathbf{a}}$. The choice is justified by the following result:

Lemma 1. *If the samples in \mathcal{D} are i.i.d., as $|\mathcal{D}| \rightarrow \infty$, if $\mathcal{W}_1(p_{S_{\mathbf{a}}}, p_{S_{\mathbf{a}}}) < \infty$ for all \mathbf{a} , the empirical barycenter distribution satisfies $\lim_{\sum_{\mathbf{a}} \hat{p}_{\mathbf{a}}} \mathcal{W}_1(\hat{p}_{\bar{S}}, \hat{p}_{S_{\mathbf{a}}}) \rightarrow \sum_{\mathbf{a}} p_{\mathbf{a}} \mathcal{W}_1(p_{\bar{S}}, p_{S_{\mathbf{a}}})$ almost surely⁴.*

The proof is given in Appendix A.

In the next two sections we introduce two different approaches to achieve SDP with Wasserstein-1 distances: A penalization approach to logistic regression and a simpler practical approach consisting in post-processing model beliefs.

3.2 WASSERSTEIN-1 PENALIZED LOGISTIC REGRESSION

The average logistic regression loss function over $\mathcal{D} = \{\mathbf{a}^n, \mathbf{x}^n, y^n\}_{n=1}^N$ is given by

$$J_{\mathcal{D}}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^N -y^n \log s^n - (1 - y^n) \log(1 - s^n),$$

where the model belief that individual n belongs to class 1, s^n , is obtained as $s^n = \sigma(\boldsymbol{\theta}^\top \mathbf{w}^n) = 1/(1 + e^{-\boldsymbol{\theta}^\top \mathbf{w}^n})$, with $\mathbf{w}^n = (\mathbf{x}^n, \mathbf{a}^n, 1)^\top$, and where $\boldsymbol{\theta} \in \mathbb{R}^{d+k+1}$ are the model parameters. We denote with $\{s_{\mathbf{a}}^i\}$ the model beliefs for group \mathbf{a} and with $\{\bar{s}^j\}_{j=1}^N$ the atoms of $p_{\bar{S}}$.

The gradient of $J_{\mathcal{D}}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is given by

$$\nabla_{\boldsymbol{\theta}} J_{\mathcal{D}}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^N \mathbf{w}^n (\sigma(\boldsymbol{\theta}^\top \mathbf{w}^n) - y^n).$$

We propose to find model parameters $\boldsymbol{\theta}^*$ that minimize the population level logistic loss $\mathbb{E}[J_{\mathcal{D}}(\boldsymbol{\theta})]$ under the constraint of small Wasserstein-1 distances $\mathcal{W}_1(\hat{p}_{S_{\mathbf{a}}}, \hat{p}_{\bar{S}})$ between $\hat{p}_{S_{\mathbf{a}}}$ and the empirical barycenter $\hat{p}_{\bar{S}}$, $\forall \mathbf{a} \in \mathcal{A}$.

The Wasserstein-1 distance between any two empirical distributions \hat{p}_b and \hat{p}_c underlying two datasets $\{b^i\}_{i=1}^{N_b}, \{c^j\}_{j=1}^{N_c} \subset \mathbb{R}$ is given by

$$\mathcal{W}_1(\hat{p}_b, \hat{p}_c) = \min_{T_{b,c} \in U(b,c)} \langle T_{b,c}, \mathbf{C} \rangle, \quad (2)$$

where $U(b,c) = \{T \in \mathbb{R}^{N_b \times N_c} \text{ s.t. } T_{b,c} \mathbf{1}_c = \frac{1}{N_b} \mathbf{1}_b \text{ and } T_{b,c}^\top \mathbf{1}_b = \frac{1}{N_c} \mathbf{1}_c\}$ with $\mathbf{1}_c$ denoting a vector of ones of size N_c . The brackets $\langle \cdot, \cdot \rangle$ denote the trace

⁴See Klenke (2013) for a formal definition of almost sure convergence of random variables.

dot product and C is the cost matrix associated with the Wasserstein-1 cost function c of elements $C_{i,j} = |b^i - c^j|$.

In particular, the Wasserstein-1 distance $\mathcal{W}_1(\hat{p}_{S_a}, \hat{p}_{\bar{S}})$ can be computed by solving the optimization problem of Eq. (2) with cost matrix $C_a^\theta \in \mathbb{R}^{N_a \times \bar{N}}$ satisfying

$$(C_a^\theta)_{i,j} = |s_a^i - \bar{s}^j|,$$

where the upper script θ in C_a^θ is maintained to remind the reader that model predictions are a function of the model parameter θ .

The Wasserstein-1 penalized logistic regression objective is given by

$$J_{\mathcal{W}_1}(\theta) = \alpha J_{\mathcal{D}}(\theta) + (1 - \alpha)\beta \sum_{\mathbf{a} \in \mathcal{A}} \mathcal{W}_1(\hat{p}_{S_a}, \hat{p}_{\bar{S}}), \quad (3)$$

where α and β are penalization coefficients.

Lemma 2. *If the datasets $\{b^i\}_{i=1}^{N_b}, \{c^j\}_{j=1}^{N_c} \subset \mathbb{R}$ have empirical distributions \hat{p}_b and \hat{p}_c , and C is the cost matrix of elements $C_{i,j} = |b^i - c^j|$:*

$$\nabla_C \mathcal{W}_1(\hat{p}_b, \hat{p}_c) = T_{b,c}^*,$$

where $T_{b,c}^* = \arg \min_{T_{b,c} \in U(b,c)} \langle T_{b,c}, C \rangle$ is the optimal coupling resulting from the optimization objective of Eq. (2).

Proof. The result follows immediately from the subgradient rule for a pointwise max function (see Boyd and Vandenberghe (2004)). \square

Lemma 3. *The gradient of $J_{\mathcal{W}_1}(\theta)$ equals:*

$$\alpha \nabla_{\theta} J_{\mathcal{D}}(\theta) + (1 - \alpha)\beta \left(\sum_{\mathbf{a} \in \mathcal{A}} \sum_{i,j} T_a^*(\theta) \nabla_{\theta} |s_a^i - \bar{s}^j| \right),$$

where T_a^* is the optimal coupling between \hat{p}_{S_a} and $\hat{p}_{\bar{S}}$ ⁵.

Proof. This formula is a consequence of the chain rule and Lemma 1. \square

Computation Method. We propose to optimize the Wasserstein penalized logistic loss objective (Eq. (3)) via gradient descent. The procedure is detailed in Algorithm 1. We start by describing how to perform Step 2. under the assumption that $\hat{p}_{\bar{S}}$ and $\{\bar{s}^i\}_{i=1}^{\bar{N}}$ have been computed. The computation of the optimal coupling family $\{T_a^*\}$ hinges on the following Lemma.

Lemma 4. *If $\{b^i\}_{i=1}^{N_b}, \{c^j\}_{j=1}^{N_c} \subset \mathbb{R}$, and $B_i = [i \times (N_c - 1) + 1, \dots, i \times N_c]$ for all i and $C_j = [j \times (N_b - 1) + 1, \dots, j \times N_b]$ for all j , then: $(T_{b,c}^*)_{i,j} = \frac{\#|B_i \cap C_j|}{N_b N_c}$.*

⁵Recall that T_a^* is a function of θ .

Algorithm 1 Wass-1 Penalized Logistic Regression

Input: Dataset $\mathcal{D} = \{(\mathbf{a}^n, \mathbf{x}^n, y^n)\}_{n=1}^N$, **penalization coefficients** α, β , **gradient step size** η , **number of optimization rounds** M , **frequency of barycenter computation** K .

Compute datasets $\{\mathcal{D}_a\}$.

Initialize model parameters θ_0 .

for $m = 1, \dots, M$ **do**

1. Compute the barycenter distribution $\hat{p}_{\bar{S}}$ (Flamary and Courty (2017)) once every K steps, and $\{\bar{s}^i\}_{i=1}^{\bar{N}}$.
2. Compute optimal couplings $\{T_a^*\}$ as defined in Lemma 3.
3. Update parameter $\theta_m = \theta_{m-1} - \eta \nabla_{\theta} J_{\mathcal{W}_1}(\theta_m)$.

end

Return: θ_M .

This lemma characterizes the coupling matrix between the empirical distributions of two datasets made of real numbers. When $N_b = N_c$ and the datasets are $\{b^i\}_b^N$ and $\{c^i\}_c^N$, with $b^1 < \dots < b^N$, and $a^1 < \dots < a^N$, then the optimal coupling equals $1/N \times I_N$ where I_N denotes the $N \times N$ identity matrix. Lemma 4 extends this simple case to the general case of datasets of arbitrary orderings and sizes, see Deshpande et al. (2018) for a proof. It is easy to see that the optimal coupling $T_{b,c}^*$ is sparse and has at most $\mathcal{O}(N_b + N_c)$ nonzero entries (see Cuturi (2013)). As a consequence, the computation of $\nabla_{\theta} J_{\mathcal{W}_1}(\theta)$ can be performed in linear time $\mathcal{O}(\sum_{\mathbf{a}} (N_a + \bar{N}))$ where $N_a = |\mathcal{D}_a|$. In the computation of $\nabla_{\theta} J_{\mathcal{W}_1}(\theta)$ only the nonzero entries of $T_{b,c}^*$ matter.

We compute the empirical barycenter $\hat{p}_{\bar{S}}$ and $\{\bar{s}^i\}_{i=1}^{\bar{N}}$, using the POT library by Flamary and Courty (2017). We fix the support of potential barycenters to bins of equal-width spanning the $[0, 1]$ interval, and use the iterative KL-projection method proposed by Benamou et al. (2015). We then generate a number of samples from the normalized probability distribution of the computed barycenter.

Demographically-Blind Wasserstein-1 Penalized Logistic Regression.

In real-world applications, the use of sensitive attributes might be prohibited when deploying a system. We therefore consider the variation where $\mathbf{w}^n = (\mathbf{x}^n, 1)^\top$. This variation still uses the sensitive attributes to calculate the Wasserstein-1 loss but, by not including them into the feature set, does not require knowledge of sensitive information at test time.

3.3 WASSERSTEIN-1 POST-PROCESSING

In this section, we propose a simple, fast quantile matching method to post-process the beliefs of a classifier

Algorithm 2 Wass-1 Post-Processing

Input: dataset $\mathcal{D} = \{(\mathbf{a}^n, \mathbf{x}^n, y^n)\}_{n=1}^N$, **set of quantile bins** \mathcal{B} , **model beliefs** $\{s^n\}$

Compute datasets $\{\mathcal{D}_a\}$ and their barycenter $\bar{\mathcal{D}}$.

Define the i -th quantile of dataset \mathcal{D}_a as

$$q_{\mathcal{D}_a}(i) := \sup \left\{ s : \frac{1}{N_a} \sum_{n \text{ s.t. } \mathbf{a}^n = \mathbf{a}} \mathbb{1}_{s^n \leq s} \leq \frac{i-1}{|\mathcal{B}|} \right\},$$

and its inverse as $q_{\mathcal{D}_a}^{-1}(s) := \sup\{i \in \mathcal{B} : q_{\mathcal{D}_a}(i) \leq s\}$.

Return: $\{q_{\bar{\mathcal{D}}}(q_{\mathcal{D}_a}^{-1}(s^n))\}$.

trained on \mathcal{D} . This method corresponds to an approximate Wasserstein-1 optimal transport map by the formulation of Rachev and Rüschendorf (1998):

$$\mathcal{W}_1(p_{S_a}, p_{\bar{S}}) = \int_{\tau=0}^1 |P_{S_a}^{-1}(\tau) - P_{\bar{S}}^{-1}(\tau)| d\tau.$$

The procedure is detailed in Algorithm 2. For each group \mathbf{a} , we compute quantiles of \hat{p}_{S_a} and map all group beliefs belonging in each quantile bin to the supremum of those belonging to the corresponding quantile bin of $\hat{p}_{\bar{S}}$.

3.4 GENERALIZATION

The following lemma addresses generalization of the Wasserstein-1 objective. Assume $\mathcal{W}_1(p_{S_a}, p_{\bar{S}}) \leq L$ for all $\mathbf{a} \in \mathcal{A}$. Let P_S, P_{S_a} and $P_{\bar{S}}$ be the cumulative density functions of S, S_a and \bar{S} . Assume these random variables all have domain $\Omega = [0, 1]$ and that all $P \in \{P_S, P_{\bar{S}}\} \cup \{P_{S_a}\}_{\mathbf{a} \in \mathcal{A}}$ are continuous, then:

Lemma 5. For any $\epsilon, \delta > 0$, if $\min[\bar{N}, \min_{\mathbf{a}} [N_{\mathbf{a}}]] \geq \frac{16 \log(2|\mathcal{A}|/\delta) |\mathcal{A}|^2 \max[1, L]^2}{\epsilon^2}$, with probability $1 - \delta$:

$$\sum_{\mathbf{a} \in \mathcal{A}} p_{\mathbf{a}} \mathcal{W}_1(p_{S_a}, p_{\bar{S}}) \leq \sum_{\mathbf{a} \in \mathcal{A}} \hat{p}_{\mathbf{a}} \mathcal{W}_1(\hat{p}_{S_a}, \hat{p}_{\bar{S}}) + \epsilon.$$

In other words, provided access to sufficient samples, a low value of $\sum_{\mathbf{a}} \hat{p}_{\mathbf{a}} \mathcal{W}_1(\hat{p}_{S_a}, \hat{p}_{\bar{S}})$ implies a low value for $\sum_{\mathbf{a}} p_{\mathbf{a}} \mathcal{W}_1(p_{S_a}, p_{\bar{S}})$ with high probability and therefore good performance at test time.

The proof is given in Appendix B.

Lemma 5 implies that under appropriate conditions, the value of the population objective of the Wasserstein cost is upper bounded by the empirical Wasserstein cost plus a small constant.

4 RELATED WORK

Broadly speaking, we can group current literature on fair classification and regression into three main approaches.

The first approach consists in pre-processing the data to remove bias, or in extracting representations that do not contain sensitive information during training (Beutel et al., 2017; Calders et al., 2009; Calmon et al., 2017; Edwards and Storkey, 2016; Feldman et al., 2015; Fish et al., 2015; Kamiran and Calders, 2009, 2012; Louizos et al., 2016; Zemel et al., 2013; Žliobaite et al., 2011). This approach includes current methods to fairness using Wasserstein distances consisting in achieving SDP through transportation of features (Del Barrio et al., 2019; Johndrow and Lum, 2019).

The second approach consists in performing a post-processing of the model outputs (Doherty et al., 2012; Feldman, 2015; Hardt et al., 2016).

The third approach consists in enforcing fairness notions by imposing constraints into the optimization, or by using an adversary. Some methods transform the constrained optimization problem via the method of Lagrange multipliers (Agarwal et al., 2018; Corbett-Davies et al., 2017; Cotter et al., 2018; Goh et al., 2016; Narasimhan, 2018; Wu et al., 2018; Zafar et al., 2017). Other work similar in spirit adds penalties to the objective (Donini et al., 2018; Komiyama et al., 2018). Adversarial methods maximize the system ability to predict Y while minimizing the ability to predict \mathbf{A} (Zhang et al., 2018).

5 EXPERIMENTS

In this section, we evaluate the methods introduced in Sections 3.2 and 3.3 on four datasets from the UCI repository (Lichman, 2013). For penalized logistic regression, we refer to the method in which sensitive information is included in the feature set, *i.e.* $\mathbf{w}^n = (\mathbf{x}^n, \mathbf{a}^n, 1)^\top$, as **Wass-1 Penalty**; and to the demographically-blind variant in which sensitive information is not included, *i.e.* $\mathbf{w}^n = (\mathbf{x}^n, 1)^\top$, as **Wass-1 Penalty DB**. We refer to the post-processing method as **Wass-1 Post-Process**. We also include a variant of this method using \hat{p}_S instead of the barycenter $\hat{p}_{\bar{S}}$ (**Wass-1 Post-Process \hat{p}_S**), which gives a simpler algorithm that only requires computing basic quantile functions. We compare these methods with the following baselines:

Unconstrained: Logistic regression with no fairness constraints.

Hardt’s Post-Process: Post-processing of the logistic regression beliefs s^n of all individuals in group \mathbf{a} by adding $0.5 - \tau_{\mathbf{a}}$, where the threshold $\tau_{\mathbf{a}}$ is found using the method of Hardt et al. (2016). This ensures that DP is satisfied at threshold $\tau = 0.5$.

Constrained Optimization: Lagrangian-based method (see *e.g.* Eban et al. (2017); Goh et al. (2016)) using a linear model as the underlying predictor and equal

Table 1: **Adult Dataset – German Credit Dataset**

	Adult					German				
	Err-.5	Err-Exp	DD-.5	SDD	SPDD	Err-.5	Err-Exp	DD-.5	SDD	SPDD
Unconstrained	.142	.198	.413	.426	.806	.248	.319	.124	.102	.103
Hardt’s Post-Process	.165	.289	.327	.551	1.058	.248	.333	.056	.045	.045
Constrained Opt.	.205	.198	.065	.087	.166	.318	.320	.173	.149	.149
Adv. Constr. Opt.	.219	.207	.0	.114	.203	.306	.307	.0	.021	.021
Wass-1 Penalty	.199	.208	.014	.022	.044	.306	.311	.0	.003	.003
Wass-1 Penalty DB	.230	.233	.010	.012	.023	.306	.309	.0	.010	.010
Wass-1 Post-Process	.174	.214	.013	.017	.042	.258	.327	.068	.023	.023
Wass-1 Post-Process \hat{p}_S	.165	.216	.032	.022	.059	.248	.320	.056	.025	.025

positive prediction rate between each group \mathcal{D}_a and \mathcal{D} as fairness constraints with threshold $\tau = 0$.

Adv. Constr. Opt.: The same as the previous method, but with more fairness constraints. Specifically, the fairness constraints are equal positive prediction rates for a set of thresholds from -2 to 2 in increments of 0.2 on the output of the linear model.

5.1 TRAINING DETAILS

In the approaches Unconstrained, Hardt’s Post-Process, Wass-1 Penalty, and Wass-1 Post-Process, we trained a logistic regression model using Scikit-Learn with default hyper-parameters (Pedregosa and et al., 2011).

For Wass-1 Penalty (Algorithm 1), as initial model parameters θ_0 we used the ones given by the trained logistic regression. We swept over penalization coefficients $\alpha = [0, 0.5]$, $\beta = [10^{-2}, 3 \cdot 10^{-2}, 10^{-1}, 3 \cdot 10^{-1}, 1, 3, 10, 30, 10^2]$, gradient step sizes $\eta = [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$, set the maximum number of training steps to $M = 80,000$, and computed the barycenter once every $K > M$ steps, effectively only once after the initialization of θ_0 . In the computation of the barycenter (using the POT library by Flarmy and Courty (2017)), we swept over numbers of bins $B = [50, 90]$, entropy penalty $\delta = [10^{-3}, 5 \cdot 10^{-3}, 10^{-2}]$, and used number of iterations $M = 1,000$. The time complexity of our implementation is $\mathcal{O}(N \log(N))$. Our gradient steps take on average ~ 0.02 seconds.

For Wass-1 Post-Process (Algorithm 2), we used a number of bins $|\mathcal{B}| = 100$.

For Constrained Optimization, we used the hinge loss as objective and the hinge relaxation for the fairness constraints. We trained by jointly optimizing the model parameters and Lagrange multipliers on the Lagrangian using ADAM with the default step-size of 0.001 and mini-batch size of 100 , and trained for 50 steps. We allowed an

additive slack of 0.05 on the constraints, as otherwise we found feasibility issues leading to degenerate classifiers.

5.2 DATASETS

The UCI Adult Dataset. The Adult dataset contains 14 attributes including age, working class, education level, marital status, occupation, relationship, race, gender, capital gain and loss, working hours, and nationality for 48,842 individuals; 32,561 and 16,281 for the training and test sets respectively. The goal is to predict whether the individual’s annual income is above or below \$50,000.

Pre-processing and Sensitive Attributes. We pre-processed the data in the same way as done in Goh et al. (2016); Zafar et al. (2017). The categorical features were encoded into binary features (one for each category), and the continuous features were transformed into binary encodings depending on five quantile values, obtaining a total of 122 features. As sensitive attributes, we considered race (Black and White) and gender (female and male), obtaining four groups corresponding to black females, white females, black males, and white males.

The UCI German Credit Dataset. This dataset contains 20 attributes for 1,000 individuals applying for loans. Each applicant is classified as a good or bad credit risk, *i.e.* as likely or not likely to repay the loan. We randomly divided the dataset into training and test sets of sizes 670 and 330 respectively.

Pre-processing and Sensitive Attributes. We did not do any pre-processing. As sensitive attributes, we considered age (≤ 30 and > 30 years old), obtaining two groups.

The UCI Bank Marketing Dataset. This dataset contains 20 attributes for 41,188 individuals. Each individual is classified as subscribed or not to a term deposit. We divided the dataset into train and test sets of sizes 32,950

Table 2: **Bank Marketing Dataset – Community & Crime Dataset**

	Bank Marketing					Community & Crime				
	Err-.5	Err-Exp	DD-.5	SDD	SPDD	Err-.5	Err-Exp	DD-.5	SDD	SPDD
Unconstrained	.094	.138	.135	.134	.61	.116	.195	.581	1.402	7.649
Hardt’s Post-Process	.097	.181	.018	.367	1.057	.321	.441	.226	.536	2.679
Constrained Opt.	.105	.110	.049	.026	.076	.289	.263	.193	.369	2.003
Adv. Constr. Opt.	.105	.105	.050	.064	.184	.303	.275	.022	.312	1.628
Wass-1 Penalty	.114	.151	.001	.015	.050	.313	.315	.0	.008	.039
Wass-1 Penalty DB	.114	.131	.001	.006	.018	.313	.315	.0	.011	.051
Wass-1 Post-Process	.100	.144	.016	.020	.062	.321	.363	.226	.133	.680
Wass-1 Post-Process \hat{p}_S	.097	.141	.014	.020	.063	.321	.335	.226	.159	.822

and 8,238 respectively.

Pre-processing and Sensitive Attributes. We pre-processed the data as for the Adult dataset. We transformed the categorical features into binary ones, and the continuous features into five binary features based on five quantile bins, obtaining a total of 60 features. We also subtracted the mean from cons.price.idx, cons.conf.idx, euribor3m, and nr.employed to make them zero-centered. As sensitive attributes, we considered age, which was discretized based on five quantiles leading to five groups.

The UCI Communities & Crime Dataset. This dataset contains 135 attributes for 1994 communities; 1495 and 499 for the training and test sets respectively. The goal is to predict whether a community has high (above the 70-th percentile) crime rate.

Pre-processing and Sensitive Attributes. We pre-processed the data as in Wu et al. (2018). As sensitive attributes, we considered race (Black, White, Asian and Hispanic), thresholded at the median to form height groups.

5.3 RESULTS

We compared the different methods using the following metrics:

Err-.5: Binary classification error using threshold $\tau = 0.5$, i.e. $\text{Err-.5} = \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{\hat{y}^n \neq y^n}$.

Err-Exp: As above, but averaging over 100 uniformly-spaced thresholds $\tau \in [0, 1]$.

DD-.5: Demographic disparity at $\tau = 0.5$, summed over all groups $\mathbf{a} \in \mathcal{A}$, i.e. $\text{DD-.5} = \sum_{\mathbf{a} \in \mathcal{A}} |\mathbb{P}(S_{\mathbf{a}} > 0.5) - \mathbb{P}(S > 0.5)|$, where e.g. $\mathbb{P}(S > \tau)$ is estimated as $\mathbb{P}(S > \tau) \approx \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{s^n > \tau}$.

SDD (strong demographic disparity): As above, but averaging over 100 uniformly-spaced thresholds $\tau \in [0, 1]$, i.e. $\text{SDD} = \sum_{\mathbf{a} \in \mathcal{A}} \mathbb{E}_{\tau \sim U([0,1])} |\mathbb{P}(S_{\mathbf{a}} > \tau) - \mathbb{P}(S > \tau)|$. We use this metric to compare

with other baselines that use the full-dataset belief distribution.

SPDD: $\text{SPDD} = \sum_{\mathbf{a}, \bar{\mathbf{a}} \in \mathcal{A}} \mathbb{E}_{\tau \sim U([0,1])} |\mathbb{P}(S_{\mathbf{a}} > \tau) - \mathbb{P}(S_{\bar{\mathbf{a}}} > \tau)|$. This metric is the most important, target-neutral, (un)fairness measurement as it does not depend on the target distribution, e.g. the full-dataset belief distribution or the barycenter.

Figure 1 shows overlaying model belief histograms for four demographic groups and their barycenter in the Adult dataset. Wasserstein-1 Penalty effectively matches all group histograms to the barycenter after training for 10,000 steps with $\beta = 100$.

The main experiment results are shown in Tables 1 and 2⁶. Focusing on the three more relevant metrics – namely Err-Exp as the robust error measure, SDD as the conventional fairness comparison metric, and SPDD as the target-neutral, preferred fairness metric (according to which we picked the best hyperparameter settings) – we can see that Wass-1 Penalty and Wass-1 Penalty DB have lowest SDD and SPDD (blue) on the German and Crime datasets and on the Adult and Bank datasets respectively. The fairness performance of these two methods are followed closely by the simpler Wass-1 Post-Process methods on all datasets. Hardt’s Post-Process method incurs largest errors (red) on all datasets. After the Unconstrained baseline, Constrained Optimization and Adv. Contr. Opt. give lowest error on the Adult, Bank and Crime datasets, whilst Constrained Optimization and Wass-1 Penalty (DB) give lowest error on the German dataset. Overall the Wasserstein-1 methods gave best fairness performance on all the datasets with similar or lower compromise on accuracy than the baselines.

Since Wass-1 Penalty is trained by gradient descent, early-stopping can be an effective way to control trade-off between accuracy and fairness. Figure 2 shows a typical

⁶Given the deterministic baseline logistic regression model, all standard deviations are on the order of 10^{-4} or below.




Figure 1: Histograms of model beliefs for groups of Black females, Black males, White females, and White males, and their barycenter on the Adult dataset using Wass-1 Penalty. Top: Initial state. Bottom: After 10,000 training steps with $\alpha = 0$, $\beta = 100$ each group histogram matches the barycenter.

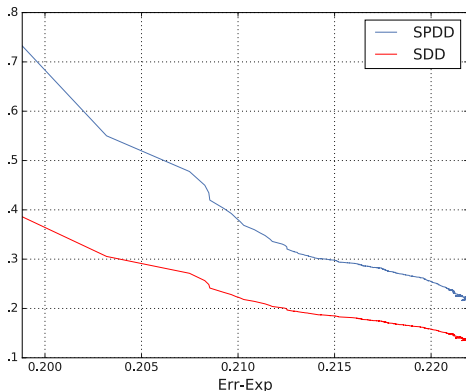


Figure 2: Err-Exp v.s. SDD, Err-Exp v.s. SPDD trade-off curves on Bank test set using Wass-1 Penalty DB, points plotted every 100 steps over 80,000 total training steps.

example of two trade-off curves between SDD/SPDD and Err-Exp. Though not always the case, often as the learning model moves towards the fairness goal of SDP, model accuracy decreases (Err-Exp increases).

6 CONCLUSIONS

We introduced an approach to ensure that the output of a classification system does not depend on sensitive information using the Wasserstein-1 distance. We demonstrated that using the Wasserstein-1 barycenter enables

us to reach independence with minimal modifications of the model decisions. We introduced two methods with different desirable properties, a Wasserstein-1 constrained method that does not necessarily require access to sensitive information at deployment time, and an alternative fast and practical approximation method that requires knowledge of sensitive information at test time. We showed that these methods outperform previous approaches in the literature.

Acknowledgements

The authors would like to thank Mark Rowland for useful feedback on the manuscript.

References

- A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, and H. Wallach. A reduction approach to fair classification. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924, 2011.
- J. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 2(37):A1111–A1138, 2015.

- A. Beutel, J. Chen, Z. Zhao, and E. H. Chi. Data decisions and theoretical implications when adversarially learning fair representations. *CoRR*, abs/1707.00075, 2017.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. In *Data mining workshops, 2009. ICDMW'09. IEEE international conference on*, pages 13–18, 2009.
- F. Calmon, D. Wei, B. Vinzamuri, K. N. Ramamurthy, and K. R. Varshney. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems 30*, pages 3995–4004, 2017.
- S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806, 2017.
- A. Cotter, H. Jiang, and K. Sridharan. Two-player games for efficient non-convex constrained optimization. *CoRR*, abs/1804.06500, 2018.
- M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26*, pages 2292–2300, 2013.
- J. De Fauw et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine*, 24(9):1342–1350, 2018.
- E. Del Barrio, F. Gamboa, P. Gordaliza, and J.-M. Loubes. Obtaining fairness using optimal transport theory. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2357–2365, 2019.
- I. Deshpande, Z. Zhang, and A. G. Schwing. Generative modeling using the sliced Wasserstein distance. In *The IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- W. Dieterich, C. Mendoza, and T. Brennan. Compas risk scales: Demonstrating accuracy equity and predictive parity, 2016.
- N. A. Doherty, A. V. Kartasheva, and R. D. Phillips. Information effect of entry into credit ratings market: The case of insurers’ ratings. *Journal of Financial Economics*, 106(2):308–330, 2012.
- M. Donini, L. Oneto, S. Ben-David, J. Shawe-Taylor, and M. Pontil. Empirical risk minimization under fairness constraints. In *Advances in Neural Information Processing Systems 31*, pages 2791–2801, 2018.
- E. Eban, M. Schain, A. Mackey, A. Gordon, R. Rifkin, and G. Elidan. Scalable learning of non-decomposable objectives. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 832–840, 2017.
- H. Edwards and A. Storkey. Censoring representations with an adversary. In *4th International Conference on Learning Representations*, 2016.
- V. Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin’s Press, 2018.
- M. Feldman. Computational fairness: Preventing machine-learned discrimination. 2015.
- M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268, 2015.
- B. Fish, J. Kun, and A. D. Lelkes. Fair boosting: a case study. In *FAT/ML Workshop*, 2015.
- R. Flamary and N. Courty. POT python optimal transport library, 2017. <https://github.com/rflamary/POT>.
- G. Goh, A. Cotter, M. Gupta, and M. P. Friedlander. Satisfying real-world goals with dataset constraints. In *Advances in Neural Information Processing Systems 29*, pages 2415–2423, 2016.
- M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems 29*, pages 3315–3323, 2016.
- M. Hoffman, L. B. Kahn, and D. Li. Discretion in hiring. *The Quarterly Journal of Economics*, 133(2):765–800, 2018.
- J. Johndrow and K. Lum. An algorithm for removing sensitive information: Application to race-independent recidivism prediction. *The Annals of Applied Statistics*, 13(1):189–220, 2019.
- F. Kamiran and T. Calders. Classifying without discriminating. In *Computer, Control and Communication, 2009. IC4 2009. 2nd International Conference on*, pages 1–6, 2009.
- F. Kamiran and T. Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- L. Kantorovich. On the transfer of masses (in Russian). *Doklady Akademii Nauk*, 37(2):227–229, 1942.
- A. Klenke. *Probability Theory: A Comprehensive Course*. Springer Science & Business Media, 2013.
- J. Komiyama, A. Takeda, J. Honda, and H. Shimao. Non-convex optimization for regression with fairness constraints. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2742–2751, 2018.

- C.-A. Laisant. Intégration des fonctions inverses. *Nouvelles annales de mathématiques, journal des candidats aux coles polytechnique et normale*, 5 (4):253–257, 1905.
- M. Lichman. UCI Machine Learning Repository, 2013. <http://archive.ics.uci.edu/ml>.
- C. Louizos, K. Swersky, Y. Li, M. Welling, and R. Zemel. The variational fair autoencoder. In *4th International Conference on Learning Representations*, 2016.
- M. Malekipirbazari and V. Aksakalli. Risk assessment in social lending via random forests. *Expert Systems with Applications*, 42(10):4621–4631, 2015.
- Pascal Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pages 1269–1283, 1990.
- G. Monge. Memoire sur la theorie des déblais et des remblais. *Histoire de l' Académie des Sciences de Paris*, 1781.
- H. Narasimhan. Learning with complex loss functions and constraints. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pages 1646–1654, 2018.
- F. Pedregosa and et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011.
- C. Perlich, B. Dalessandro, T. Raeder, O. Stitelman, and F. Provost. Machine learning for targeted display advertising: Transfer learning in action. *Machine Learning*, 95(1):103–127, 2014.
- S. Rachev and L. Rüschendorf. *Mass Transportation Problems*. Springer, 1998.
- J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *CoRR*, abs/1707.00087, 2017.
- S. Wu, M. Kearns, S. Neel, and A. Roth. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2564–2572, 2018.
- M. B. Zafar, I. Valera, M. Gomez Rodriguez, and K. P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 962–970, 2017.
- R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning*, pages 325–333, 2013.
- B. Zhang, H Lemoine, and M. Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340, 2018.
- I. Žliobaite, F. Kamiran, and T. Calders. Handling conditional discrimination. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining*, pages 992–1001, 2011.