

On the Power of Randomization in Fair Classification and Representation

Sushant Agarwal
sushant.agarwal@uwaterloo.ca
University of Waterloo
Canada

Amit Deshpande
amitdesh@microsoft.com
Microsoft Research
India

ABSTRACT

Fair classification and fair representation learning are two important problems in supervised and unsupervised fair machine learning, respectively. Fair classification asks for a classifier that maximizes accuracy on a given data distribution subject to fairness constraints. Fair representation maps a given data distribution over the original feature space to a distribution over a new representation space such that all classifiers over the representation satisfy fairness. In this paper, we examine the power of randomization in both these problems to minimize the loss of accuracy that results when we impose fairness constraints. Previous work on fair classification has characterized the optimal fair classifiers on a given data distribution that maximize accuracy subject to fairness constraints, e.g., Demographic Parity (DP), Equal Opportunity (EO), and Predictive Equality (PE). We refine these characterizations to demonstrate when the optimal randomized fair classifiers can surpass their deterministic counterparts in accuracy. We also show how the optimal randomized fair classifier that we characterize can be obtained as a solution to a convex optimization problem. Recent work has provided techniques to construct fair representations for a given data distribution such that any classifier over this representation satisfies DP. However, the classifiers on these fair representations either come with no or weak accuracy guarantees when compared to the optimal fair classifier on the original data distribution. Extending our ideas for randomized fair classification, we improve on these works, and construct DP-fair, EO-fair, and PE-fair representations that have provably optimal accuracy and suffer no accuracy loss compared to the optimal DP-fair, EO-fair, and PE-fair classifiers respectively on the original data distribution.

CCS CONCEPTS

• **Theory of computation** → **Unsupervised learning and clustering**; • **Computing methodologies** → **Supervised learning by classification**; *Learning latent representations*.

KEYWORDS

fairness; demographic parity; equal opportunity; randomization; machine learning; classification; representation

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

FACCT '22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9352-2/22/06...\$15.00
<https://doi.org/10.1145/3531146.3533209>

ACM Reference Format:

Sushant Agarwal and Amit Deshpande. 2022. On the Power of Randomization in Fair Classification and Representation. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FACCT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3531146.3533209>

1 INTRODUCTION

With the proliferation of machine learning models in several domains such as banking, education, healthcare, law enforcement, customer service etc. with direct social and economic impact on individuals, it has become imperative to build these models ethically and responsibly to avoid amplification of their biases. Assessing and mitigating bias to sensitive or underprivileged demographic groups is an important aspect of building ethical and responsible machine learning models, and understandably, a large part of recent literature on this has focused on the fairness of machine learning models [4, 5, 9].

Fairness in classification has been an important topic of study for machine learning models because it corresponds to automated decision-making scenarios in sensitive application domains. Fairness of classification in literature has been largely divided into two subtopics, focusing on either individual fairness [6, 17, 20, 31], or group fairness [21, 35]. Group-fair classification requires a classifier to achieve certain outcome metrics to be equal or near-equal across sensitive demographic groups (e.g., race, gender). Legal precedents on disparate impact and the four-fifths rule [5] have given rise to statistical definitions and quantitative metrics of group-fairness such as Demographic Parity (DP), Equal Opportunity (EO), and Predictive Equality (PE) [21, 35]. Literature on group-fair classification has looked at various trade-offs of different metrics of group-fairness with each other [25] as well as their trade-offs with desiderata such as accuracy, interpretability, and privacy [2, 3]. Given the inevitable trade-off between accuracy and fairness, it is natural to ask the theoretical question of finding the *optimal fair classifier* on a given data distribution that maximizes accuracy subject to certain group-fairness constraints. Previous work on this problem has tried to characterize optimal group-fair classifiers and proposed methods to train such group-fair classifiers [7, 10, 12, 21, 29], and proposed methods to train group-fair classifiers [1, 23, 36]. The characterization has considered both group-aware and group-blind classifiers, based on whether the use of sensitive attributes or group membership information at inference time is allowed or prohibited.

Classification by machine learning models is a part of supervised learning. Machine learning models that perform unsupervised learning (e.g., representation learning, generative models, clustering) do not directly correspond to decision-making scenarios but they are equally important in modern machine learning suites and equally

vulnerable to the hidden risks of amplifying social and economic harms [8, 30]. The outcome-based definitions of group-fairness such as DP and EO do not immediately apply to unsupervised learning models for fair representation and fair generation. Previous work on building fair representations has tried to incorporate certain differentiable proxies for fairness and accuracy objectives to train fair representations [15, 19, 34]. This approach does not immediately guarantee group-fairness of models trained on the representations. To circumvent this, it has been natural to define the objective of fair representation as finding an embedding or a representation map from the original data to a new representation, so that any downstream classifier on the new representation must satisfy group-fairness constraints. Subsequently, it reduces our objective to only finding the classifier of maximum accuracy on the above fair representation. Another practical motivation for the above definition of fair representation is that it supports the scenario where a data regulator (e.g., a government body) releases the data to data users (e.g., a recruitment agency) in such a way that any model or classifier built by the data users on this data cannot violate a desirable fairness guarantee.

The fairness-accuracy trade-off in fair classification implies that there is an accuracy loss (or price of fairness) when we maximize accuracy subject to fairness constraints, as opposed to unconstrained accuracy maximization. In fair representation, accuracy-maximizing classifiers built on fair representations suffer an accuracy loss compared to the optimal fair classifiers on the original data distribution because there is a potential information loss when we go from the original data distribution to its fair representation. The focus of our paper is to investigate how randomization can help minimize the accuracy loss in the fair classification and fair representation problems stated above. We study these problems for three of the most popular group-fairness constraints (DP, EO, PE), in the group-aware setting where the classifier and the representation map are allowed to use the sensitive group membership.

Roadmap of the paper. The paper is organized as follows. Section 1.1 briefly summarizes our main results. Then, Section 1.2 discusses related works in the literature on fair classification and fair representation. We introduce the setup in Section 2, and formally state the problems we study in this paper. Section 3 characterizes the optimal randomized DP-fair, EO-fair, and PE-fair classifiers, and shows how they can be obtained as the solution to a convex optimization problem. Section 4 constructively characterizes the optimal randomized DP-fair, EO-fair, and PE-fair representations, that have provably optimal accuracy, and suffer no accuracy loss compared to the optimal fair randomized classifier on the original data distribution. We conclude and talk about potential directions for future work in Section 5.

1.1 Our Results

Our main contributions are summarized below.

- We mathematically characterize the *optimal randomized fair classifier* that maximizes accuracy on a given data distribution subject to group-fairness constraints (DP, EO, PE). We do this in the group-aware setting and prove that the optimal randomized group-fair classifier is a randomized threshold classifier that can be obtained as the solution to a convex

optimization problem. We give an example of a data distribution to show that the accuracy of the optimal randomized group-fair classifier can surpass the accuracy of the optimal deterministic group-fair classifier. This is in contrast with the deterministic Bayes Optimal classifier that maximizes accuracy on a given data distribution without any fairness constraints, and is optimal among all deterministic as well as randomized classifiers on the distribution.

- Extending our ideas for randomized fair classification, we construct randomized fair representations for DP, EO, and PE constraints such that the optimal classifier on this representation incurs no *cost of fair representation*, or in other words, no accuracy loss compared to the optimal fair randomized classifier on the original data distribution. This improves upon previous work on fair representations that primarily focuses on DP and gives no or weak provable guarantees for accuracy.

1.2 Related Work

Given a joint data distribution on features and binary class labels, it is well-known that the optimal classifier is given by the Bayes Optimal classifier that applies a threshold of $1/2$ on the class-probability function, i.e., the probability of the class label being positive given the feature vector. A natural extension of this under fairness constraints is to characterize the optimal fair classifier on a given joint data distribution over features, class labels and group membership, that maximizes accuracy subject to group-fairness constraints. This question can be asked for both group-aware and group-blind classifiers, based on whether using a sensitive attribute is allowed or prohibited for inference. Previous work has studied this extensively: Given a predictive score function, Hardt et al. [21] characterize the optimal group-aware fair classifier for EO as a group-dependent threshold classifier on the predictive score value. Under the assumption that the class-probability function has a strictly positive density on $[0, 1]$, Corbett-Davies et al. [13] characterize the optimal group-aware fair classifiers for DP and PE as group-dependent threshold classifiers on the class-probability function. Menon and Williamson [29] characterize the optimal group-blind (and group-aware) fair classifiers for DP and EO as instance-dependent (and group-dependent) threshold classifiers on the class-probability function. Celis et al. [7] extend this characterization to optimal fair classifiers under multiple simultaneous fairness constraints. Chzhen et al. (Proposition 2.3 in [10]) give an explicit expression for the optimal deterministic group-aware fair classifier for EO, describing it as a recalibrated Bayes classifier. Zhao et al. [39] characterize the optimal group-aware fair classifier for DP using oracle access to the Bayes optimal classifiers on the underlying sensitive groups. Agarwal et al. [1] reduce fair classification to a sequence of cost-sensitive classifications, whose solutions give a randomized classifier of least empirical error subject to the desired fairness constraints. Their randomized classifier samples a random classifier from a given hypothesis class and then uses it to make the prediction. Our work refines the above results on the characterization of optimal fair classifiers to demonstrate when randomized fair classifiers have an advantage over deterministic fair classifiers. We further extend these ideas for optimal randomized

fair classification to construct randomized fair representations that achieve optimal accuracy and suffer no accuracy loss compared to the optimal fair classifiers on the original data distribution. In recent concurrent work, Zeng et al. [38] characterize the optimal fair randomized classifiers for DP, EO, and PE, however, they do not highlight the advantages of randomization, and do not consider the problem of fair representations. We note that our characterization results assume full access to the data distribution. In practice, this does not hold, and one only has access to a finite sample drawn from the distribution. To (approximately) realise our classifiers in practice, we can make use of techniques similar to the ones suggested in many of the above works, but we do not focus on this.

The outcome-based notions of fairness for classification need to be modified to apply to unsupervised learning problems such as representation learning. There is a long line of work on defining and constructing fair representations [11, 16, 18, 22, 26–28, 32, 33, 37]. Feldman et al. [18] construct DP-fair representations by introducing the paradigm of total repair, which involves transforming the original training dataset to a new dataset, such that the protected attribute is now independent of the class label. In particular, total repair involves mapping the group-wise conditional distributions from the original feature space to a common distribution in the new representation space, so that it is impossible to predict the protected attribute from the new sanitised dataset. This transformation is done in a manner that preserves as much of the information in the original dataset as possible. Many follow up works also give techniques for DP-fair representation that fall under the paradigm of total repair [16, 22]. However, most of these works provide weak or no accuracy guarantees for the optimal classifier on the fair representation. Our construction for DP-fair representation also falls under the paradigm of total repair, but in contrast, we construct a fair representation such that the optimal classifier on this representation incurs no accuracy loss compared to the optimal fair classifier on the original data distribution. In addition, we also construct fair representations for notions of fairness apart from DP, such as EO and PE. Our representation technique works for any distribution over an arbitrary feature space, which does not always hold for previous techniques that use geometric properties of continuous feature spaces. McNamara et al. [28] and Zhao et al. [39] also study the cost of using a fair representation in terms of accuracy loss. McNamara et al. [28] term this loss as the *cost of mistrust*, which is similar to what we call the *cost of fair representation*. We define the cost of fair representation as the difference between the accuracy of the optimal fair classifier on the original data distribution and the accuracy of the optimal classifier on the fair representation. McNamara et al. [28] provide a closed form expression for the cost of mistrust, and also derive an upper bound on its value. Their result works for a fairness-regularized accuracy maximization objective that does not imply a bound on the cost of fair representation defined in our paper. Apart from the above results that are closest to our paper, there is also plenty of work on fair representation learning in practice [15, 26, 27, 37]. Zemel et al. [37] develop a heuristic technique for learning DP-fair representations, where they formulate fairness as an optimization problem of finding an intermediate data representation that best encodes the data, while removing any information about the sensitive group memberships. Their optimization problem has terms to encourage

both these goals simultaneously. However, Zemel et al. [37] and its follow-up works [15, 26, 27] provide heuristics with no provable accuracy or fairness guarantees.

While we focus on the benefits of randomization, it is important to note that randomized decisions for fair machine learning can have practical limitations in real-world deployment, e.g., inconsistent decisions for the same individual when repeated, as discussed in Cotter et al. [14].

2 PROBLEM FORMULATION

2.1 Setup

Consider a learning problem, where we are given a distribution \mathcal{P} over $\mathcal{X} \times \mathcal{Z} \times \mathcal{Y}$, where $\mathcal{Z} = \{A, D\}$ represents the protected group membership, \mathcal{X} represents all the other features, and $\mathcal{Y} = \{0, 1\}$ represents the label set (we adopt the standard convention of associating the label 1 with success or acceptance). Our results also hold when there are multiple groups, but for ease of exposition, we restrict our analysis to the case of 2 groups (advantaged group A , and disadvantaged group D). A deterministic group-aware classification rule f is a map

$$f : \mathcal{X} \times \mathcal{Z} \rightarrow \{0, 1\},$$

that assigns a binary label to each point (also known as feature vector) in the feature space $\mathcal{X} \times \mathcal{Z}$. We denote the hypothesis class of all such deterministic functions f by \mathcal{H} .

The aim of the learning problem is to find a classifier in the hypothesis class that minimises a loss function. We consider the standard 0-1 loss function ℓ , whose expected value is given by

$$\mathcal{L}(f) = \mathbb{E}[\ell(f)] = \Pr[f(X, Z) \neq Y], \quad (1)$$

where the probability is over $(X, Z, Y) \sim \mathcal{P}$.

Remark 2.1. Henceforth, all probabilities will be over $(X, Z, Y) \sim \mathcal{P}$, unless explicitly stated.

Our results also work for the more general loss function ℓ_α , known in literature as cost-sensitive risk [29], that assigns a weight α to False Positive errors, and a weight $(1 - \alpha)$ to False Negative errors. The expected loss is a weighted linear combination of the False Positive rate (FPR) and False Negative rate (FNR), where

$$\begin{aligned} \text{FPR}(f) &= \Pr[f(X, Z) = 1 \mid Y = 0] \quad \text{and} \\ \text{FNR}(f) &= \Pr[f(X, Z) = 0 \mid Y = 1]. \end{aligned}$$

We want to minimize the expected loss, given by

$$\mathcal{L}_\alpha(f) = \mathbb{E}[\ell_\alpha(f)] = \alpha(1 - \pi)\text{FPR}(f) + (1 - \alpha)\pi\text{FNR}(f),$$

where $\pi = \Pr[Y = 1]$. By assigning equal weight to False Positive errors and False Negative errors by setting $\alpha = \frac{1}{2}$, we recover the standard 0-1 loss function, given by Equation 1.

Remark 2.2. If α is not mentioned, it is assumed to be $\frac{1}{2}$. Our results hold for any $0 < \alpha < 1$, but for ease of exposition, we often restrict our analysis to $\alpha = \frac{1}{2}$.

In the standard 0-1 loss setting, the optimal classifier f^* (the Bayes Optimal Classifier) is given by

$$f^*(x, z) = \mathcal{T}_{\frac{1}{2}}(\Pr[Y = 1 \mid X = x, Z = z]),$$

and $\mathcal{T}_\gamma(\beta)$ is the threshold function that outputs 1 if $\beta \geq \gamma$, and 0 otherwise. The optimal classifier for cost-sensitive risk ℓ_α has a very similar form [12], and is given by

$$f_\alpha^*(x, z) = \mathcal{T}_\alpha(\Pr[Y = 1 \mid X = x, Z = z]).$$

The only difference is that, instead of thresholding the success probability for a point at $\frac{1}{2}$, we threshold at α .

In many situations we may want to consider a fairness-aware learning problem, that aims to find the most accurate hypothesis in \mathcal{H} that is also group-fair, and does not discriminate based on the protected group feature \mathcal{Z} . We consider a few popular notions of fairness, that we define below.

Definition 2.3 (Demographic Parity). f satisfies Demographic Parity with selection rate r , if

$$\Pr[f(X, Z) = 1 \mid Z = A] = \Pr[f(X, Z) = 1 \mid Z = D] = r.$$

Definition 2.4 (Predictive Equality, or Equal FPR). f satisfies Predictive Equality if the false positive rates are equal for both groups, i.e.,

$$\Pr[f(X, Z) = 1 \mid Y = 0, Z = A] = \Pr[f(X, Z) = 1 \mid Y = 0, Z = D].$$

Definition 2.5 (Equal Opportunity, or Equal FNR [21]). f satisfies Equal Opportunity if the false negative rates are equal for both groups, i.e.,

$$\Pr[f(X, Z) = 0 \mid Y = 1, Z = A] = \Pr[f(X, Z) = 0 \mid Y = 1, Z = D].$$

Remark 2.6. Classifiers satisfying these notions of fairness will be referred to as DP-fair, PE-fair, and EO-fair respectively. Our results hold for all three fairness notions. However, we can obtain analogous results for EO from the same proof techniques as that of PE. Hence, we only discuss the results for PE. In addition, some previous works also consider equal True Positive rate (TPR), and equal True Negative rate (TNR), as notions of fairness. However, obtaining equal TPR is equivalent to obtaining equal FNR, and obtaining equal TNR is equivalent to obtaining equal FPR, and hence we do not separately discuss these two notions.

2.2 Randomized Fair Classification Problem

We now consider a relaxation of deterministic classifiers, and allow a classifier to be randomized. A randomized classification rule f is a function

$$f : \mathcal{X} \times \mathcal{Z} \rightarrow [0, 1],$$

where $f(x, z)$ denotes the probability of $(x, z) \in \mathcal{X} \times \mathcal{Z}$ being mapped to 1. We denote the hypothesis class of all such randomized functions f by $\tilde{\mathcal{H}}$. The notions of loss and fairness previously defined for the deterministic case can be easily modified for the randomized case. We investigate whether the ability to use randomization could lead to more powerful classifiers, and consequently study the following problem.

PROBLEM 2.7. *What is the most accurate classifier in $\tilde{\mathcal{H}}$ satisfying a particular notion of fairness?*

Note that ability to use randomization can only lead to more powerful classifiers, because $\mathcal{H} \subseteq \tilde{\mathcal{H}}$. In particular, a randomized classification rule can map a feature vector (x, z) to 1 with probability $f(x, z)$, and to 0 with probability $(1 - f(x, z))$. In contrast, a

deterministic classifier can either choose to map (x, z) strictly to 1, or strictly to 0, which is captured in the randomized setting by setting $f(x, z) = 1$, or $f(x, z) = 0$, respectively. To see how randomization can improve the accuracy of fair classifiers, let us look at an example.

Example 2.8. Consider the following distribution \mathcal{P} over $\mathcal{X} \times \mathcal{Z} \times \mathcal{Y}$, where $\mathcal{X} = \{x_1, x_2\}$, $\mathcal{Z} = \{A, D\}$, and $\mathcal{Y} = \{0, 1\}$.

$$\mathcal{P}(x_1, A, 1) = \frac{3}{8}, \mathcal{P}(x_1, A, 0) = \frac{1}{8}, \mathcal{P}(x_1, D, 1) = \frac{1}{8}, \mathcal{P}(x_1, D, 0) = \frac{1}{8}$$

$$\mathcal{P}(x_2, A, 1) = 0, \mathcal{P}(x_2, A, 0) = 0, \mathcal{P}(x_2, D, 1) = 0, \mathcal{P}(x_2, D, 0) = \frac{1}{4}.$$

There are only 2 deterministic classifiers satisfying DP, either the constant 1 classifier f_1 , or the constant 0 classifier f_0 , with $\mathcal{L}(f_1) = \mathcal{L}(f_0) = \frac{1}{2}$. On the other hand, consider the following randomized classifier f , where

$$f(x_1, A) = \frac{1}{2}, f(x_1, D) = 1, f(x_2, A) = f(x_2, D) = 0.$$

It is easy to see that f satisfies DP, and $\mathcal{L}(f) = \frac{3}{8}$, hence improving the accuracy of the optimal fair deterministic classifiers f_0 and f_1 .

2.3 Randomized Fair Representation Problem

A common way used to obtain fair classifiers is to pre-process the data by mapping the data from the original domain space to a new representation space. Essentially, a data regulator modifies the data before the data user is allowed to train a model on it. The goal of the data regulator is to sanitise the data such that models trained on it are fair, while also maintaining predictive accuracy by preserving as much information of the original data set as possible. A data representation is a randomized mapping

$$R : \mathcal{X} \times \mathcal{Z} \rightarrow \tilde{\mathcal{X}},$$

for some representation space $\tilde{\mathcal{X}}$. A deterministic classification rule g over this representation is a map $g : \tilde{\mathcal{X}} \rightarrow \{0, 1\}$, that assigns a binary label to each point in the new representation space $\tilde{\mathcal{X}}$. For a fixed R , we define the following hypothesis class $\tilde{\mathcal{H}}_R = \{g \circ R \mid g \text{ is a deterministic classifier over } \tilde{\mathcal{X}}\}$.

A data regulator wants to ensure that the classifier trained by the data user on the sanitised data is fair. However, the data user could be adversarial, and actively try and be unfair. Hence, the data regulator needs to ensure that every classifier over the representation is fair, and this motivates the following definition of a fair representation.

Definition 2.9 (Fair Representation). A representation R is fair if every $h \in \tilde{\mathcal{H}}_R$ is fair.

Note that for any R , $\tilde{\mathcal{H}}_R \subseteq \tilde{\mathcal{H}}$, and hence the accuracy of the optimal classifier in $\tilde{\mathcal{H}}$ cannot be lesser than the accuracy of the optimal classifier in $\tilde{\mathcal{H}}_R$. Similarly, if the representation R is fair, then the accuracy of the optimal fair classifier in $\tilde{\mathcal{H}}$ cannot be lesser than the accuracy of the optimal classifier in $\tilde{\mathcal{H}}_R$. Hence, the fairness of the representation may come at a cost, and we define the following quantity below, that measures the difference in the accuracy of the optimal fair classifier in $\tilde{\mathcal{H}}$, and the optimal classifier in $\tilde{\mathcal{H}}_R$.

Definition 2.10 (Cost of Fair Representation R (denoted by $\text{CFR}(R)$)).

$$\text{CFR}(R) = \min_{h \in \mathcal{H}_R} \mathcal{L}(h) - \min_{f \in \mathcal{H}} \mathcal{L}(f).$$

Consequently, we focus on the following problem.

PROBLEM 2.11. *Let R_{fair} be the set of fair representations. What is the representation $R' \in R_{\text{fair}}$ that yields the minimal cost of fair representation?*

Remark 2.12. For simplicity of analysis in later proofs, we only consider deterministic classifiers over the randomized representations. It is easy to see that allowing for randomized classifiers over randomized representations provides no additional power, and does not affect the minimum possible cost of fair representation, as the randomness of the optimal classifier on a given representation could be incorporated in the representation itself.

3 CHARACTERIZING THE RANDOMIZED FAIR CLASSIFIER OF MAXIMUM ACCURACY

We now introduce some notation and terminology that we will need later on.

Definition 3.1 (Cell [24]). Consider a randomized partition of the feature space $\mathcal{X} \times \mathcal{Z}$ into multiple disjoint components. We call these components cells, and denote a cell by C .

A natural partition we will encounter is the following, where each feature vector $(x, z) \in \mathcal{X} \times \mathcal{Z}$ is a separate cell, denoted by $C_{x,z}$. We consider randomized partitions, i.e., a feature vector (x, z) does not have to strictly lie within one cell, but could be divided into multiple cells. For example, feature vector (x, z) could be mapped to a cell C_1 with probability $\gamma(x, z)$, and to another cell C_2 with probability $(1 - \gamma(x, z))$.

Definition 3.2 (Score [24]). The score S of a cell C is the probability that a random instance drawn from \mathcal{P} , given that it lies in C , has label 1, i.e.,

$$S(C) = \Pr[Y = 1 \mid (X, Z) \in C].$$

Previously, we encountered classifiers that threshold based on success probabilities (i.e., scores), such as the Bayes Optimal classifier. We formally define them below.

Definition 3.3 (Score Threshold Classifiers). Given $t \in [0, 1]$, a score-threshold classifier \mathcal{T}_t maps $C_{x,z}$ to 1 if $S(C_{x,z}) \geq t$, and to 0 otherwise.

Randomized classifiers give us the ability to threshold by probability mass, instead of just thresholding by scores. To explain this better, we introduce the notion of group-wise sorted cells.

Definition 3.4 (Group-wise Sorted Cells). Define $C_z = \bigcup_{x \in \mathcal{X}} C_{x,z}$, where the component cells of C_A and C_D are arranged in descending order of scores S . If two or more cells from the same group have the same score, merge them. This gives us cells sorted in strictly decreasing order of scores.

By $C_z(t)$, denote the topmost cells of C_z comprising of t fraction of the total probability mass of C_z , i.e., $\frac{\mathcal{P}(C_z(t))}{\mathcal{P}(C_z)} = t$. Note that this may involve splitting a cell into 2 parts randomly. For example, in Example 2.8, $C_A(\frac{1}{2})$ would involve splitting $C_{x_1,A}$ into two equal parts randomly. However, in the deterministic setting, only $C_A(0)$ and $C_A(1)$ are defined, and $C_A(\frac{1}{2})$ does not make sense. Hence, in the deterministic setting, we cannot in general threshold by probability mass for an arbitrary threshold value.

Definition 3.5 (Mass Threshold Classifiers). Let the instances in the feature space accepted and rejected by f be denoted by $\mathcal{A}(f)$ and $\mathcal{R}(f)$ respectively. By $\tilde{\mathcal{T}}_t$, we denote the mass-threshold classifier such that $\mathcal{A}(\tilde{\mathcal{T}}_t) = \bigcup_{z \in \mathcal{Z}} C_z(t)$, and $\mathcal{R}(\tilde{\mathcal{T}}_t) = \bigcup_{z \in \mathcal{Z}} (C_z - C_z(t))$.

For example, in Example 2.8, the randomised classifier f we constructed is actually the randomized mass-threshold classifier $\tilde{\mathcal{T}}_{\frac{1}{2}}$. Note that a mass-threshold classifier needs to use randomization on (at most) only one cell in each group, and is deterministic on all other cells.

3.1 Characterizing the Optimal Randomized DP-Fair Classifier

Here, we answer Problem 2.7 for DP, with the following characterization result.

Theorem 1. There exists $r' \in [0, 1]$, such that the most accurate DP-fair classifier in \mathcal{H} is given by the mass-threshold classifier $\tilde{\mathcal{T}}_{r'}$.

PROOF. We start by proving Proposition 3.6, and use it to prove Lemma 3.7, which in turn implies Theorem 1. \square

PROPOSITION 3.6. *For any $r \in [0, 1]$, $\tilde{\mathcal{T}}_r$ is DP-fair.*

PROOF. Let the elements of a cell C accepted and rejected by f be denoted by $\mathcal{A}(f(C))$ and $\mathcal{R}(f(C))$ respectively. For both $z = A$, and $z = D$, $\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) = C_z(r)$, and $\mathcal{R}(\tilde{\mathcal{T}}_r(C_z)) = C_z - C_z(r)$.

Selection rate of $\tilde{\mathcal{T}}_r$ on group z is given by

$$\frac{\mathcal{P}(\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)))}{\mathcal{P}(C_z)} = \frac{\mathcal{P}(C_z(r))}{\mathcal{P}(C_z)}.$$

And by definition of $C_z(r)$,

$$\frac{\mathcal{P}(C_z(r))}{\mathcal{P}(C_z)} = r.$$

Hence $\tilde{\mathcal{T}}_r$ has the same selection rate (i.e., r) for both groups, and satisfies Demographic Parity. \square

We then proceed to prove Lemma 3.7.

LEMMA 3.7. *The most accurate DP-fair classifier in $\tilde{\mathcal{H}}$ with a fixed selection rate $r \in [0, 1]$ is given by $\tilde{\mathcal{T}}_r$.*

PROOF. Consider an arbitrary classifier with selection rate r , denoted by f_r , such that $f_r \neq \tilde{\mathcal{T}}_r$. We will now argue that $\mathcal{L}(\tilde{\mathcal{T}}_r) < \mathcal{L}(f_r)$.

We know that the (expected) loss is the (weighted) sum of false positive and false negative errors. Hence,

$$\mathcal{L}(f) = \mathcal{P}(\mathcal{A}(f))(1 - S(\mathcal{A}(f))) + \mathcal{P}(\mathcal{R}(f))S(\mathcal{R}(f)).$$

Since we know that

$$\mathcal{P}(\mathcal{A}(f_r)) = \mathcal{P}(\mathcal{A}(\tilde{\mathcal{T}}_r)) = r, \text{ and } \mathcal{P}(\mathcal{R}(f_r)) = \mathcal{P}(\mathcal{R}(\tilde{\mathcal{T}}_r)) = 1 - r,$$

it suffices to show that

$$\mathcal{S}(\mathcal{A}(\tilde{\mathcal{T}}_r)) > \mathcal{S}(\mathcal{A}(f_r)), \text{ and } \mathcal{S}(\mathcal{R}(\tilde{\mathcal{T}}_r)) < \mathcal{S}(\mathcal{R}(f_r)).$$

We prove the former first. We can prove the latter with exactly the same reasoning, and hence omit the details. It suffices to show that

$$\mathcal{S}(\mathcal{A}(\tilde{\mathcal{T}}_r(C_z))) > \mathcal{S}(\mathcal{A}(f_r(C_z))), \text{ if } \mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) \neq \mathcal{A}(f_r(C_z)).$$

To show the above, it is enough to observe that

$$\mathcal{S}(\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) - (\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) \cap \mathcal{A}(f_r(C_z)))) >$$

$$\mathcal{S}(\mathcal{A}(f_r(C_z)) - (\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) \cap \mathcal{A}(f_r(C_z)))).$$

This is easy to see, because

$$(\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) - (\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) \cap \mathcal{A}(f_r(C_z)))) \subseteq C_z(r),$$

and

$$(\mathcal{A}(f_r(C_z)) - (\mathcal{A}(\tilde{\mathcal{T}}_r(C_z)) \cap \mathcal{A}(f_r(C_z)))) \subseteq (C_z - C_z(r)),$$

and given $P \subseteq C_z(r)$, $Q \subseteq (C_z - C_z(r))$, such that $P \cap Q = \emptyset$, then $\mathcal{S}(P) > \mathcal{S}(Q)$, because the component cells of C_z are sorted in strictly descending order of scores.

This completes the proof of the lemma. \square

Armed with this lemma, it is easy to see that the most accurate classifier in \mathcal{H} satisfying Demographic Parity is given by $\tilde{\mathcal{T}}_{r'}$, where r' denotes the optimum selection rate, i.e.,

$$r' \in \operatorname{argmin}_{r \in [0,1]} \mathcal{L}(\tilde{\mathcal{T}}_r),$$

which completes the proof.

3.1.1 Finding the Optimal Randomized DP-Fair Classifier by Convex Optimization. Having characterized the optimal classifier, we now discuss how to find it as the solution to a convex optimization problem. The main result is the following theorem.

Theorem 2. $\mathcal{L}(\tilde{\mathcal{T}}_r)$ is a convex, piecewise linear, continuous function over $[0, 1]$. The optimal classifier $\tilde{\mathcal{T}}_{r'}$ as in Theorem 1, where

$$r' \in \operatorname{argmin}_{r \in [0,1]} \mathcal{L}(\tilde{\mathcal{T}}_r),$$

is therefore the solution to a convex optimisation problem.

We first introduce some terminology that will be needed in the upcoming proofs.

Definition 3.8 (Score boundaries). Consider the component cells of groupwise sorted C_A and C_D . Then,

$$\mathcal{I} = \mathcal{I}_A \cup \mathcal{I}_D, \text{ where } \mathcal{I}_z = \{0\} \cup \{r \mid \exists C_z^i \in C_z \mid C_z^{\leq i} = C_z(r)\},$$

and $C_z^{\leq i}$ denotes the topmost cells of C_z comprising of all cells until and including C_z^i . \mathcal{I}_z essentially consists of all the boundary points in C_z , where the component cells change score.

We now proceed with the proof.

PROOF. We start by showing that $\mathcal{L}(\tilde{\mathcal{T}}_r)$ is piecewise linear and continuous. We do this by proving that the loss function is linear between any two points in the set \mathcal{I} .

Consider any $r_i \in \mathcal{I} = \{0, r_a, r_b, \dots, 1\}$ in sorted order, and define a corresponding cell C_i , where

$$C_i = \mathcal{A}(\tilde{\mathcal{T}}_{r_i}) - \mathcal{A}(\tilde{\mathcal{T}}_{r_{i-}}),$$

and r_{i-} denotes the element in \mathcal{I} preceding r_i .

Consider any $k \geq 0$, such that, $r_{i-} + k \leq r_i$. We know that

$$\mathcal{S}(\mathcal{A}(\tilde{\mathcal{T}}_{r_{i-}+k}) - \mathcal{A}(\tilde{\mathcal{T}}_{r_{i-}})) = \mathcal{S}(C_i),$$

because the cell C_i is split into two parts at random, and hence both parts have the same score as C_i .

We know that loss is the (weighted) sum of false positive and false negative errors. Hence,

$$\begin{aligned} \mathcal{L}(\tilde{\mathcal{T}}_{r_{i-}+k}) &= \mathcal{L}(\tilde{\mathcal{T}}_{r_{i-}}) + k(1 - \mathcal{S}(C_i)) - k\mathcal{S}(C_i) \\ &= \mathcal{L}(\tilde{\mathcal{T}}_{r_{i-}}) + k(1 - 2\mathcal{S}(C_i)). \end{aligned}$$

Looking at the form of the expression of the loss as above, it is easy to see that the loss function is linear in the interval $[r_{i-}, r_i]$, for any r_i , and is therefore piecewise linear and continuous over $[0, 1]$.

We now proceed to prove that the loss function is convex. We begin by observing the following. \square

Observation 3.9. Consider the interval between $[r_{i-}, r_i]$, for any $r_i \in \mathcal{I}$. The slope of the loss function $\mathcal{L}(\tilde{\mathcal{T}}_r)$ at any point in this interval is $(1 - 2\mathcal{S}(C_i))$, and hence

- (1) $\mathcal{L}(\tilde{\mathcal{T}}_r)$ is strictly decreasing if $\mathcal{S}(C_i) > \frac{1}{2}$, and the decrease is steeper if the score is higher.
- (2) $\mathcal{L}(\tilde{\mathcal{T}}_r)$ is strictly increasing if $\mathcal{S}(C_i) < \frac{1}{2}$, and the increase is steeper if the score is lower.
- (3) $\mathcal{L}(\tilde{\mathcal{T}}_r)$ is constant if $\mathcal{S}(C_i) = \frac{1}{2}$.

Because the component cells of C_z are sorted in strictly descending order of scores, it is easy to see that, for every r_i , $\mathcal{S}(C_i) < \mathcal{S}(C_{i-})$. Hence, there exists a cell $C_{i'}$, such that for every $r_j, r_k \in \mathcal{I}$ satisfying $r_j < r_{i'} < r_k$, one of the following cases hold.

- (1) $\mathcal{S}(C_j) > \frac{1}{2} > \mathcal{S}(C_{i'}) > \mathcal{S}(C_k)$.
- (2) $\mathcal{S}(C_j) > \mathcal{S}(C_{i'}) > \frac{1}{2} > \mathcal{S}(C_k)$.
- (3) $\mathcal{S}(C_j) > \mathcal{S}(C_{i'}) = \frac{1}{2} > \mathcal{S}(C_k)$.

In case (1), the loss function is strictly decreasing with shrinking rate of decrease upto $r_{i'}$, and strictly increasing with growing rate of increase after $r_{i'}$. In case (2), the loss function is strictly decreasing with shrinking rate of decrease upto $r_{i'}$, and strictly increasing with growing rate of increase after $r_{i'}$. In case (3), the loss function is strictly decreasing with shrinking rate of decrease upto $r_{i'}$, constant between $r_{i'}$ and $r_{i'}$, and strictly increasing with growing rate of increase after $r_{i'}$. In each case, it is easy to see that the loss function is convex.

This concludes the proof.

Remark 3.10. We note that the optima is unique if there exists no cell C_i , such that $\mathcal{S}(C_i) = \frac{1}{2}$.

In addition, to find the optimum selection rate r' , we just need to evaluate $\mathcal{L}(\tilde{\mathcal{T}}_r)$ for all $r \in \mathcal{I}$, instead of for all $r \in [0, 1]$, as a corollary of Observation 3.11. This can drastically reduce search time, for example in the case where the distribution \mathcal{P} is discrete.

Observation 3.11.

$$\min_{r \in \mathcal{I}} \mathcal{L}(\tilde{\mathcal{T}}_r) = \min_{r \in [0, 1]} \mathcal{L}(\tilde{\mathcal{T}}_r).$$

PROOF. Consider any $r_i \in \mathcal{I}$. Enough to show that for any r such that $r_{i-} < r < r_i$, either $\mathcal{L}(\tilde{\mathcal{T}}_{r_i}) \leq \mathcal{L}(\tilde{\mathcal{T}}_r)$, or $\mathcal{L}(\tilde{\mathcal{T}}_{r_{i-}}) \leq \mathcal{L}(\tilde{\mathcal{T}}_r)$, which follows directly from Observation 3.9. \square

3.2 Characterizing the Optimal Randomized PE-Fair Classifier

In the previous section on Demographic Parity, we considered mass-threshold classifiers $\tilde{\mathcal{T}}_t$, that select $C_z(t)$, and reject $C_z - C_z(t)$, for both $z = A$, and $z = D$. $\tilde{\mathcal{T}}_t$ applies the same threshold t to both groups A and D . In this section, we consider groupwise mass-threshold classifiers $\tilde{\mathcal{T}}_{t_A, t_D}$ that apply different thresholds t_A and t_D to groups A and D respectively.

Definition 3.12 (Groupwise Mass-Threshold Classifier). A groupwise mass-threshold classifier $\tilde{\mathcal{T}}_{t_A, t_D}$ is defined such that $\mathcal{A}(\tilde{\mathcal{T}}_{t_A, t_D}(C_z)) = C_z(t_z)$, for both $z = A$, and $z = D$.

Denote the False Positive rate of a classifier f restricted to a cell C by $\text{FPR}(f(C))$. Given $r \in (0, 1]$, there is a unique classifier $\tilde{\mathcal{T}}_{t_A, t_D}$, such that $\text{FPR}(\tilde{\mathcal{T}}_{t_A, t_D}(C_A)) = \text{FPR}(\tilde{\mathcal{T}}_{t_A, t_D}(C_D)) = r$. Denote this classifier by f_r . Given $r = 0$, $\tilde{\mathcal{T}}_{t_A, t_D}$ need not be unique as there could exist cells with score 1. In that case, we define f_0 to be the unique groupwise mass-threshold classifier accepting exactly the cells with score 1.

Now, we answer Problem 2.7 for PE, with the following characterisation result.

Theorem 3. There exists $r' \in [0, 1]$, such that the optimal PE-fair classifier in $\tilde{\mathcal{H}}$ is given by the groupwise mass-threshold classifier $f_{r'}$.

PROOF. We start by first proving Lemma 3.13, which implies Theorem 3. \square

LEMMA 3.13. *The most accurate PE-fair classifier in $\tilde{\mathcal{H}}$ with a false positive rate $r \in [0, 1]$ is given by f_r .*

PROOF. Consider an arbitrary classifier $f \neq f_r$ with false positive rate r for both groups. We will now argue that $\mathcal{L}(f_r) < \mathcal{L}(f)$.

Let the loss of a classifier f restricted to a cell C be denoted by $\mathcal{L}(f(C))$. Enough to show that

$$\mathcal{L}(f(C_z)) > \mathcal{L}(f_r(C_z)), \text{ if } \mathcal{A}(f(C_z)) \neq \mathcal{A}(f_r(C_z)).$$

Denote the FP's and FN's of a classifier f on cell C by $\text{FP}(f(C))$ and $\text{FN}(f(C))$ respectively. We know that loss is the sum of false positive and false negative errors. Hence, we need to show that $\mathcal{P}(\text{FP}(f(C_z))) + \mathcal{P}(\text{FN}(f(C_z))) > \mathcal{P}(\text{FP}(f_r(C_z))) + \mathcal{P}(\text{FN}(f_r(C_z)))$.

Because

$$\mathcal{P}(\text{FP}(f(C_z))) = \mathcal{P}(\text{FP}(f_r(C_z))),$$

it is enough to show that

$$\mathcal{P}(\text{FN}(f(C_z))) > \mathcal{P}(\text{FN}(f_r(C_z))).$$

Using the fact that the cells of C_z are sorted in strictly decreasing order of scores, it is easy to see that

$$\mathcal{S}(\mathcal{A}(f(C_z))) < \mathcal{S}(\mathcal{A}(f_r(C_z))).$$

We know that,

$$\begin{aligned} \mathcal{P}(\text{FP}(f(C_z))) &= \mathcal{P}(\text{FP}(f(\mathcal{A}(f(C_z))))) \\ &= \mathcal{P}(\mathcal{A}(f(C_z)))(1 - \mathcal{S}(\mathcal{A}(f(C_z)))). \end{aligned}$$

Hence,

$$\begin{aligned} \mathcal{S}(\mathcal{A}(f(C_z))) &< \mathcal{S}(\mathcal{A}(f_r(C_z))) \\ \implies \mathcal{P}(\mathcal{A}(f(C_z))) &< \mathcal{P}(\mathcal{A}(f_r(C_z))) \\ \implies \mathcal{P}(\mathcal{R}(f(C_z))) &> \mathcal{P}(\mathcal{R}(f_r(C_z))). \end{aligned}$$

Because the cells of C_z are sorted in strictly decreasing order of scores,

$$\mathcal{S}(\mathcal{R}(f(C_z))) > \mathcal{S}(\mathcal{R}(f_r(C_z))).$$

Combining the above, we get that

$$\mathcal{P}(\text{FN}(f(C_z))) > \mathcal{P}(\text{FN}(f_r(C_z))).$$

This completes the proof of the lemma. \square

Armed with this lemma, it is easy to see that the most accurate PE-fair classifier in $\tilde{\mathcal{H}}$ is given by $f_{r'}$, where r' denotes the optimum false positive rate, i.e.,

$$r' \in \underset{r \in [0, 1]}{\text{argmin}} \mathcal{L}(f_r).$$

This completes the proof of the theorem.

3.2.1 Finding the Optimal Randomized PE-Fair Classifier by Convex Optimization. After characterizing the optimal classifier, we discuss how to find it as the solution to a convex optimization problem. The main result is the below theorem.

Theorem 4. $\mathcal{L}(f_r)$ is a convex, piecewise linear, continuous function over $[0, 1]$. Hence, the optimal classifier $f_{r'}$ in Theorem 3, where

$$r' \in \underset{r \in [0, 1]}{\text{argmin}} \mathcal{L}(f_r),$$

is given by the solution to a convex optimisation problem.

We first introduce some terminology that will be needed in the upcoming proofs. Denote the groupwise thresholds of f_r by r^A and r^D respectively, i.e., $f_r = \tilde{\mathcal{T}}_{r^A, r^D}$.

Definition 3.14 (FP-Boundaries). Recall the set of score-boundaries \mathcal{I} . We then define the set of FP-boundaries \mathcal{I}_{FP} as

$$\mathcal{I}_{FP} = \{r \mid r^A \in \mathcal{I}, \text{ or } r^D \in \mathcal{I}\}.$$

\mathcal{I}_{FP} essentially consists of all the false positive rates r , such that the corresponding groupwise threshold classifier $f_r = \tilde{\mathcal{T}}_{r^A, r^D}$ has a threshold at a point in the set of score boundaries \mathcal{I} .

PROOF. We first prove that $\mathcal{L}(f_r)$ is piecewise linear and continuous, by proving that the loss function is linear between any two points in the set \mathcal{I}_{FP} .

Consider $r_i \in \mathcal{I}_{FP} = \{0, r_a, r_b, \dots, 1\}$, and define a corresponding cell C_i , where

$$C_i = \mathcal{A}(f_{r_i}) - \mathcal{A}(f_{r_{i-}}).$$

Consider any $k \geq 0$, such that, $r_{i-} + k \leq r_i$. We know that

$$\mathcal{S}(\mathcal{A}(f_{r_{i-}+k}) - \mathcal{A}(f_{r_{i-}})) = \mathcal{S}(C_i),$$

because the cell C_i is split into two parts at random, and hence both parts have the same score as C_i .

We know that loss is the (weighted) sum of false positive and false negative errors. Hence,

$$\begin{aligned} \mathcal{L}(f_{r_{i-}+k}) &= \\ \mathcal{L}(f_{r_{i-}}) + \frac{k}{r_i - r_{i-}} \mathcal{P}(C_i)(1 - \mathcal{S}(C_i)) - \frac{k}{r_i - r_{i-}} \mathcal{P}(C_i)\mathcal{S}(C_i) &= \\ \mathcal{L}(f_{r_{i-}}) + \frac{k}{r_i - r_{i-}} \mathcal{P}(C_i)(1 - 2\mathcal{S}(C_i)) = \mathcal{L}(f_{r_{i-}}) + kc(1 - 2\mathcal{S}(C_i)), \end{aligned}$$

for some $c > 0$. Looking at the form of the expression of the loss as above, it is easy to see that the loss function is linear in the interval $[r_{i-}, r_i]$, for any r_i , and is therefore piecewise linear and continuous over $[0, 1]$.

We now proceed to prove that the loss function is convex. We begin by observing the following. \square

Observation 3.15. Consider the interval $[r_{i-}, r_i]$, for any $r_i \in \mathcal{I}_{FP}$. For any point in this interval, the slope of the loss function $\mathcal{L}(f_r)$ is $c(1 - 2\mathcal{S}(C_i))$ for some $c > 0$, and hence

- (1) $\mathcal{L}(f_r)$ is strictly decreasing if $\mathcal{S}(C_i) > \frac{1}{2}$, and the decrease is steeper if the score is higher.
- (2) $\mathcal{L}(f_r)$ is strictly increasing if $\mathcal{S}(C_i) < \frac{1}{2}$, and the increase is steeper if the score is lower.
- (3) $\mathcal{L}(f_r)$ is constant if $\mathcal{S}(C_i) = \frac{1}{2}$.

Because the component cells of C_2 are sorted in strictly descending order of scores, it is easy to see that, for every r_i , $\mathcal{S}(C_i) < \mathcal{S}(C_{i-})$. Hence, there exists a cell $C_{i'}$, such that for every $r_j, r_k \in \mathcal{I}_{FP}$ such that $r_j < r_{i'} < r_k$, one of the following cases hold.

- (1) $\mathcal{S}(C_j) > \frac{1}{2} > \mathcal{S}(C_{i'}) > \mathcal{S}(C_k)$.
- (2) $\mathcal{S}(C_j) > \mathcal{S}(C_{i'}) > \frac{1}{2} > \mathcal{S}(C_k)$.
- (3) $\mathcal{S}(C_j) > \mathcal{S}(C_{i'}) = \frac{1}{2} > \mathcal{S}(C_k)$.

These are the same 3 cases as in the proof of Theorem 2. Hence, in each case, we apply the same reasoning as in the proof of Theorem 2 to see that the loss function is convex.

This concludes the proof. \square

Remark 3.16. We note that the optima is unique if there exists no cell C_i , such that $\mathcal{S}(C_i) = \frac{1}{2}$.

In addition, to find the optimum false positive rate r' , we just need to evaluate $\mathcal{L}(f_r)$ for all $r \in \mathcal{I}_{FP}$, instead of for all $r \in [0, 1]$, as a corollary of Observation 3.17. This could drastically reduce search time, for example in the case where the distribution \mathcal{P} is discrete.

Observation 3.17.

$$\min_{r \in \mathcal{I}_{FP}} \mathcal{L}(f_r) = \min_{r \in [0, 1]} \mathcal{L}(f_r).$$

PROOF. Consider $r_i \in \mathcal{I}_{FP}$. Enough to show that for any r such that $r_{i-} < r < r_i$, either $\mathcal{L}(f_{r_i}) \leq \mathcal{L}(f_r)$, or $\mathcal{L}(f_{r_{i-}}) \leq \mathcal{L}(f_r)$, which follows directly from Observation 3.15. \square

4 OPTIMAL RANDOMIZED FAIR REPRESENTATION

Many previous works have given heuristics to construct fair representations, but in most cases they do not provide accuracy guarantees for the representation. Hence, the cost of fairness for the representations they give may be high, resulting in a loss in utility.

Here we construct a fair representation that has 0 cost of fairness, implying optimal accuracy. The representation we construct follows naturally from our previous analysis of optimal fair randomized classifier.

4.1 Characterizing the Optimal Randomized DP-Fair Representation

Now, we answer Problem 2.11 for DP, with the following result.

Theorem 5. There exists a DP-fair representation R_{DP} , such that cost of fairness for R_{DP} is 0.

PROOF. We prove this by explicitly constructing such a representation. Recall that we had previously defined the set of score boundaries \mathcal{I} . As in the proof of Theorem 2, consider any $r_i \in \mathcal{I} = \{0, r_a, r_b, \dots, 1\}$, and define a corresponding cell C_i , where

$$C_i = \mathcal{A}(\tilde{\mathcal{T}}_{r_i}) - \mathcal{A}(\tilde{\mathcal{T}}_{r_{i-}}),$$

and r_{i-} denotes the element in \mathcal{I} preceding r_i .

Consider a representation R_{DP} , which maps each C_i to a distinct $\tilde{x}_i \in \tilde{\mathcal{X}}$, such that given distinct $r_i, r_j \in \mathcal{I}$, $\tilde{x}_i \neq \tilde{x}_j$. We now proceed to prove that R_{DP} is DP-fair and has 0 cost of fairness.

Each classifier f over R_{DP} either accepts or rejects \tilde{x}_i . Hence, each f either accepts or rejects C_i . Denote the set of all $i \in \mathcal{I}$, such that f accepts C_i , by \mathcal{I}_f . It is easy to see that for both groups A , and D , selection rate is given by

$$\sum_{i \in \mathcal{I}_f} (r_i - r_{i-}).$$

Hence, f satisfies demographic parity, and R_{DP} is DP-fair.

It is easy to see that for every $r \in \mathcal{I}$, the threshold classifier $\tilde{\mathcal{T}}_r$ either accepts or rejects C_i . Hence, for every $r \in \mathcal{I}$, the threshold classifier $\tilde{\mathcal{T}}_r$ is expressible over R_{DP} , i.e., $\tilde{\mathcal{T}}_r \in \tilde{\mathcal{H}}_{R_{DP}}$. Now, let us say we are given an optimal fair classifier $\tilde{\mathcal{T}}_{r'}$, as in Theorem 1. Then, there exists $r \in \mathcal{I}$, such $\mathcal{L}(\tilde{\mathcal{T}}_r) = \mathcal{L}(\tilde{\mathcal{T}}_{r'})$, which follows from Observation 3.11. Since $r \in \mathcal{I}$, $\tilde{\mathcal{T}}_r \in \tilde{\mathcal{H}}_{R_{DP}}$, and the cost of fairness for R_{DP} is 0.

This concludes the proof. \square

4.2 Characterizing the Optimal Randomized PE-Fair Representation

Now, we answer Problem 2.11 for EO, with the following result.

Theorem 6. There exists a PE-fair representation R_{PE} , such that cost of fairness for R_{PE} is 0.

PROOF. We prove this by explicitly constructing such a representation. Consider $r_i \in \mathcal{I}_{FP} = \{0, r_a, r_b, \dots, 1\}$. As in the proof of Theorem 4, define a corresponding cell C_i , where

$$C_i = \mathcal{A}(f_{r_i}) - \mathcal{A}(f_{r_{i-}}).$$

Now, define representation R_{PE} , which maps each C_i to a distinct $\tilde{x}_i \in \tilde{\mathcal{X}}$, such that given distinct $r_i, r_j \in \mathcal{I}_{FP}$, $\tilde{x}_i \neq \tilde{x}_j$. We now proceed to prove that R_{PE} is PE-fair and has 0 cost of fairness.

Each classifier f over R_{PE} either accepts or rejects \tilde{x}_i . Hence, each f either accepts or rejects C_i . Denote the set of all $i \in \mathcal{I}_{FP}$,

such that f accepts C_i , by \mathcal{I}_f^{FP} . It is easy to see that for both groups A, D , false positive rate is given by

$$\sum_{i \in \mathcal{I}_f^{FP}} (r_i - r_{i-}).$$

Hence, f satisfies equal false positive rates, and R_{PE} is PE-fair.

It is easy to see that for every $r \in \mathcal{I}_{FP}$, groupwise threshold classifier $f_r = \tilde{\mathcal{H}}_{r,A,r,D}$ either accepts or rejects C_i . Hence, for every $r \in \mathcal{I}_{FP}$, threshold classifier f_r is expressible over R_{PE} , i.e., $f_r \in \tilde{\mathcal{H}}_{R_{PE}}$. Now, let us say we are given an optimal fair classifier f_r^* , as in Theorem 3. Then, there exists $r \in \mathcal{I}_{FP}$, such that $\mathcal{L}(f_r) = \mathcal{L}(f_r^*)$, which follows from Observation 3.17. Since $r \in \mathcal{I}_{FP}$, $f_r \in \tilde{\mathcal{H}}_{R_{PE}}$, and the cost of fairness for R_{PE} is 0.

This concludes the proof. \square

5 CONCLUSION AND FUTURE WORK

In this paper, we examined the power of randomization in the problems of fair classification and representation. In particular, we mathematically characterized the optimal randomized DP-fair, PE-fair, and EO-fair classifiers in the group aware setting, and showed that they are in general more accurate than their deterministic counterparts. In addition, we also showed how the classifiers that we characterize can be obtained as a solution to a convex optimization problem. Extending our ideas for randomized fair classification, we construct DP-fair, PE-fair, and EO-fair representations such that the optimal classifier on this representation incurs no cost of fair representation, i.e., no accuracy loss compared to the optimal randomized fair classifier on the original data distribution. This improves upon previous work on fair representations that primarily focuses on DP-fairness and gives no or weak provable guarantees for the accuracy loss.

Some directions for further work include extending our results for binary classification to multi-class classification and regression. Another direction could be to look at relaxed or approximate versions of the fairness notions we considered, or to look at other popular notions of fairness. It would also be valuable to experimentally validate our theoretical claims.

ACKNOWLEDGMENTS

Funding/Support: The authors were employed by Microsoft Research India during the course of this work.

REFERENCES

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. In *International Conference on Machine Learning*. PMLR, 60–69.
- [2] Sushant Agarwal. 2021. Trade-Offs between Fairness and Interpretability in Machine Learning. In *IJCAI 2021 Workshop on AI for Social Good*.
- [3] Sushant Agarwal. 2021. Trade-Offs between Fairness and Privacy in Machine Learning. In *IJCAI 2021 Workshop on AI for Social Good*.
- [4] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [5] Solon Barocas and Andrew D Selbst. 2016. Big Data’s Disparate Impact. *California Law Review* 104 (2016), 671. <http://lawcat.berkeley.edu/record/1127463>
- [6] Yahav Bechavod, Christopher Jung, and Zhiwei Steven Wu. 2020. Metric-Free Individual Fairness in Online Learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*. <https://proceedings.neurips.cc/paper/2020/hash/80b618ebcac7aa97a6dac2ba65cb7e36-Abstract.html>
- [7] L. Elisa Celis, Lingxiao Huang, Vijay Keswani, and Nisheeth K. Vishnoi. 2020. Classification with Fairness Constraints: A Meta-Algorithm with Provable Guarantees. arXiv:1806.06055 [cs.LG]
- [8] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. 2020. Fair Generative Modeling via Weak Supervision. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 1887–1898. <http://proceedings.mlr.press/v119/choi20a.html>
- [9] Alexandra Chouldechova and Aaron Roth. 2020. A Snapshot of the Frontiers of Fairness in Machine Learning. *Commun. ACM* 63, 5 (2020), 82–89. <https://doi.org/10.1145/3376898>
- [10] Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. 2019. Leveraging Labeled and Unlabeled Data for Consistent Fair Binary Classification. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.), 12739–12750. <https://proceedings.neurips.cc/paper/2019/hash/ba51e6158bcacf80fd0d834950251e693-Abstract.html>
- [11] Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. 2020. Fair regression with Wasserstein barycenters. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/51c8bd2611e844ece5d80878eb770436-Abstract.html>
- [12] Sam Corbett-Davies and Sharad Goel. 2018. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. arXiv:1808.00023 [cs.CY]
- [13] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Halifax, NS, Canada) (KDD ’17)*. Association for Computing Machinery, New York, NY, USA, 797–806. <https://doi.org/10.1145/3097983.3098095>
- [14] Andrew Cotter, Maya Gupta, and Harikrishna Narasimhan. 2019. On Making Stochastic Classifiers Deterministic. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/file/5fc34ed307aac159a30d81181c99847e-Paper.pdf>
- [15] Elliot Creager, David Madras, Joern-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. 2019. Flexibly Fair Representation Learning by Disentanglement. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 1436–1445. <https://proceedings.mlr.press/v97/creager19a.html>
- [16] Eustasio del Barrio, Fabrice Gamboa, Paula Gordaliza, and Jean-Michel Loubes. 2018. Obtaining fairness using optimal transport theory. arXiv:1806.03195 [math.ST]
- [17] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science (ITCS) Conference*. Association for Computing Machinery, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [18] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Sydney, NSW, Australia) (KDD ’15)*. Association for Computing Machinery, 259–268. <https://doi.org/10.1145/2783258.2783311>
- [19] Aaron Ferber, Umang Gupta, Greg Ver Steeg, and Bistra Dilkina. 2021. Differentiable Optimal Adversaries for Learning Fair Representations. In *IJCAI 2021 Workshop on AI for Social Good*.
- [20] Will Fleisher. 2021. *What’s Fair about Individual Fairness?* Association for Computing Machinery, 480–490.
- [21] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, Vol. 29. <https://proceedings.neurips.cc/paper/2016/file/9d2682367c3935defcb1f9e247a97c0d-Paper.pdf>
- [22] James E. Johndrow and Kristian Lum. 2017. An algorithm for removing sensitive information: application to race-independent recidivism prediction. arXiv:1703.04957 [stat.AP]
- [23] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* 33, 1 (2012), 1–33.
- [24] Jon Kleinberg and Sendhil Mullainathan. 2019. Simplicity Creates Inequity: Implications for Fairness, Stereotypes, and Interpretability. arXiv:1809.04578 [cs.LG]
- [25] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. arXiv preprint arXiv:1609.05807 (2016).
- [26] Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard Zemel. 2017. The Variational Fair Autoencoder. arXiv:1511.00830 [stat.ML]

- [27] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2018. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*. PMLR, 3384–3393.
- [28] Daniel McNamara, Cheng Soon Ong, and Robert C. Williamson. 2019. Costs and Benefits of Fair Representation Learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*. Association for Computing Machinery, 263–270. <https://doi.org/10.1145/3306618.3317964>
- [29] Aditya Krishna Menon and Robert C Williamson. 2018. The cost of fairness in binary classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 107–118. <https://proceedings.mlr.press/v81/menon18a.html>
- [30] Orestis Papakyriakopoulos, Simon Hegelich, Juan Carlos Medina Serrano, and Fabienne Marco. 2020. Bias in Word Embeddings. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 446–457. <https://doi.org/10.1145/3351095.3372843>
- [31] Felix Petersen, Debarghya Mukherjee, Yuekai Sun, and Mikhail Yurochkin. 2021. Post-processing for Individual Fairness. *Advances in Neural Information Processing Systems* 34 (2021).
- [32] Anian Ruoss, Mislav Balunović, Marc Fischer, and Martin Vechev. 2020. Learning certified individually fair representations. *arXiv preprint arXiv:2002.10312* (2020).
- [33] Chiappa Silvia, Jiang Ray, Stepleton Tom, Pacchiano Aldo, Jiang Heinrich, and Aslanides John. 2020. A General Approach to Fairness with Optimal Transport. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 04 (Apr. 2020), 3633–3640. <https://doi.org/10.1609/aaai.v34i04.5771>
- [34] Zilong Tan, Samuel Yeom, Matt Fredrikson, and Ameet Talwalkar. 2020. Learning Fair Representations for Kernel Models. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 108)*. PMLR, 155–166. <https://proceedings.mlr.press/v108/tan20a.html>
- [35] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummedi. 2017. Fairness Beyond Disparate Treatment and Disparate Impact: Learning Classification without Disparate Mistreatment. In *Proceedings of the 26th International Conference on World Wide Web (WWW) (WWW'17)*. 1171–1180.
- [36] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Roriguez, and Krishna P Gummedi. 2017. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*. PMLR, 962–970.
- [37] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 28)*. PMLR, 325–333. <https://proceedings.mlr.press/v28/zemel13.html>
- [38] Xianli Zeng, Edgar Dobriban, and Guang Cheng. 2022. Bayes-Optimal Classifiers under Group Fairness. <https://doi.org/10.48550/ARXIV.2202.09724>
- [39] Han Zhao and Geoffrey J. Gordon. 2019. Inherent Tradeoffs in Learning Fair Representations. In *NeurIPS*. 15649–15659.