# Certifying Some Distributional Fairness with Subpopulation Decomposition

**Mintong Kang** *
Zhejiang University
kangmintong@zju.edu.cn
mintong2@illinois.edu

**Linyi Li** *
UIUC
linyi2@illinois.edu

**Maurice Weber**
ETH Zurich
webermau@inf.ethz.ch

**Yang Liu**
UC Santa Cruz
yangliu@ucsc.edu

**Ce Zhang**
ETH Zurich
ce.zhang@inf.ethz.ch

**Bo Li**
UIUC
lbo@illinois.edu

**Abstract**

Extensive efforts have been made to understand and improve the fairness of machine learning models based on different fairness measurement metrics, especially in high-stakes domains such as medical insurance, education, and hiring decisions. However, there is a lack of *certified fairness* on the end-to-end performance of an ML model. In this paper, we first formulate the certified fairness of an ML model trained on a given data distribution as an optimization problem based on the model performance loss bound on a fairness constrained distribution, which is within bounded distributional distance with the training distribution. We then propose a general fairness certification framework and instantiate it for both sensitive shifting and general shifting scenarios. In particular, we propose to solve the optimization problem by decomposing the original data distribution into analytical subpopulations and proving the convexity of the sub-problems to solve them. We evaluate our certified fairness on six real-world datasets and show that our certification is tight in the sensitive shifting scenario and provides non-trivial certification under general shifting. Our framework is flexible to integrate additional non-skewness constraints and we show that it provides even tighter certification under different real-world scenarios. We also compare our certified fairness bound with adapted existing distributional robustness bounds on Gaussian data and demonstrate that our method is significantly tighter.

## 1 Introduction

As machine learning (ML) has become ubiquitous [23, 17, 5, 10, 8, 12], fairness of ML have attracted a lot of attention from different perspectives. For instance, some automated hiring systems are biased towards males due to gender imbalanced training data [3]. Different approaches have been proposed to improve ML fairness, such as regularized training [15, 21, 25, 29], disentanglement [11, 27, 39], duality [43], low-rank matrix factorization [33], and distribution alignment [4, 28, 52].

In addition to existing approaches that *evaluate* fairness, it is important and challenging to provide *certification* for ML fairness. Recent studies have explored the certified fair *representation* of ML [38, 4, 35]. However, there lacks certified fairness on the *predictions* of an end-to-end ML model trained on an arbitrary data distribution. In addition, current fairness literature mainly focuses on training an ML model on a potentially (im)balanced distribution and evaluate its performance in a target domain measured by existing statistical fairness definitions [16, 19]. Since in practice these selected target domains can encode certain forms of unfairness of their own (e.g., sampling bias), the evaluation would be more informative if we can evaluate and certify fairness of an ML model on an *objective* distribution. Taking these factors into account, in this work, we aim to provide the first definition of *certified fairness* given an ML model and a training distribution by bounding its end-to-end performance on an objective, *fairness constrained distribution*. In particular, we define *certified fairness* as the worst-case upper bound of the ML prediction loss on a fairness constrained test distribution $\mathcal{Q}$, which is within a bounded distance to the training distribution $\mathcal{P}$. We mainly

*The first two authors contributed equally. Work done while Mintong Kang was a remote intern at UIUC.

focus on the base rate condition as the fairness constraint for $\mathcal{Q}$. We prove that our certified fairness based on a base rate constrained distribution will imply other other fairness metrics, such as demographic parity (DP) and equalized odds (EO) (proposition 1). Moreover, our framework is flexible to integrate other fairness constraints into $\mathcal{Q}$. We consider two scenarios: (1) *sensitive shifting* where only the joint distribution of sensitive attribute and label can be changed when optimizing $\mathcal{Q}$; and (2) *general shifting* where everything including the conditioned distribution of non-sensitive attributes can be changed. We then propose an effective *fairness certification framework* to compute the certificate.

In our fairness certification framework, we first formulate the problem as constrained optimization, where the fairness constrained distribution is encoded by base rate constraints. Our key technique is to decompose both training and the fairness constrained test distributions to several subpopulations based on sensitive attributes and target labels, which can be used to encode the base rate constraints. With such a decomposition, in sensitive shifting, we can decompose the distance constraint to subpopulation ratio constraints and prove the transformed low-dimensional optimization problem is convex and thus efficiently solvable. In general shifting case, we propose to solve it based on divide and conquer: we first partition the feasible space into different subpopulations, then optimize the density (ratio) of each subpopulation, apply relaxation on each subpopulation as a sub-problem, and finally prove the convexity of the sub-problems with respect to other low-dimensional variables. Our framework is applicable for any black-box ML models and any distributional shifts bounded by the Hellinger distance, which is a type of $f$-divergence studied in the literature [46, 13, 7, 24, 14].

To demonstrate the effectiveness and tightness of our framework, we evaluate our fairness bounds on *six* real-world fairness related datasets [3, 2, 18, 47]. We show that our certificate is tight under different scenarios. In addition, we verify that our framework is flexible to integrate additional constraints on $\mathcal{Q}$ and evaluate the certified fairness with additional non-skewness constraints, with which our fairness certificate is tighter. Finally, as the first work on certifying fairness of an end-to-end ML model, we adapt existing distributional robustness bound [42] for comparison to provide more intuition. Note that directly integrating the fairness constraint to the existing distributional robustness bound is challenging, which is one of the main contributions for our framework. We show that with the fairness constraints and our effective solution, our bound is strictly tighter.

**Technical Contributions**. In this work, we take the first attempt towards formulating and computing the *certified fairness* on an end-to-end ML model, which is trained on a given distribution. We make contributions on both theoretical and empirical fronts.

1. We formulate the *certified fairness* of an end-to-end ML model trained on a given distribution $\mathcal{P}$ as the worst-case upper bound of its prediction loss on a fairness constrained distribution $\mathcal{Q}$, which is within bounded distributional distance with $\mathcal{P}$.

2. We propose an effective fairness certification framework that simulates the problem as constrained optimization and solve it by decomposing the training and fairness constrained test distributions into subpopulations and proving the convexity of each sub-problem to solve it.

3. We evaluate our certified fairness on six real-world datasets to show its tightness and scalability. We also show that with additional distribution constraints on $\mathcal{Q}$, our certification would be tighter.

4. We show that our bound is strictly tighter than adapted distributional robustness bound on Gaussian dataset due to the added fairness constraints and our effective optimization approach.

**Related Work** Fairness in ML can be generally categorized into individual fairness and group fairness. Individual fairness guarantees that similar inputs should lead to similar outputs for a model and it is analyzed with optimization approaches [48, 32] and different types of relaxations [20]. Group fairness indicates to measure the *independence* between the sensitive features and model prediction, the *separation* which means that the sensitive features are statistically independent of model prediction given the target label, and the *sufficiency* which means that the sensitive features are statistically independent of the target label given the model prediction [26]. Different approaches are proposed to analyze group fairness via static analysis [45], interactive computation [40], and probabilistic approaches [1, 9, 6]. In addition, there is a line of work trying to certify the *fair representation* [38, 4, 35]. Our certified fairness differs from existing work from three perspectives: 1) we provide fairness certification considering the end-to-end model performance instead of the representation level, 2) we define and certify fairness based on a fairness constrained distribution which implies other fairness notions, and 3) our certified fairness can be computed for *any* black-box models trained on an arbitrary given data distribution.

# 2 Certified Fairness Based on Fairness Constrained Distribution

In this section, we first introduce preliminaries, and then propose the definition of *certified fairness* based on a bounded fairness constrained distribution, which to the best of our knowledge is the first formal fairness certification on end-to-end model prediction. We also show that our proposed certified fairness relates to established fairness definitions in the literature.

**Notations.** We consider the general classification setting: we denote by $\mathcal{X}$ and $\mathcal{Y} = [C]$ the feature space and labels, $[C] := \{1, 2, \cdots, C\}$. $h_\theta \colon \mathcal{X} \to \Delta^{|\mathcal{Y}|}$ represents a mapping function parameterized with $\theta \in \Theta$, and $\ell \colon \Delta^{|\mathcal{Y}|} \times \mathcal{Y} \to \mathbb{R}_+$ is a non-negative loss function such as cross-entropy loss. Within feature space $\mathcal{X}$, we identify a *sensitive* or *protected attribute* $\mathcal{X}_s$ that takes a finite number of values: $\mathcal{X}_s := [S]$, i.e., for any $X \in \mathcal{X}$, $X_s \in [S]$.

**Definition 1** (Base Rate)**.** Given a distribution $\mathcal{P}$ supported over $\mathcal{X} \times \mathcal{Y}$, the base rate for sensitive attribute value $s \in [S]$ with respect to label $y \in [C]$ is $b^{\mathcal{P}}_{s,y} = \Pr_{(X,Y)\sim\mathcal{P}}[Y = y \mid X_s = s]$.

Given the definition of base rate, we define a *fair base rate distribution* (in short as *fair distribution*).

**Definition 2** (Fair Base Rate Distribution)**.** A distribution $\mathcal{P}$ supported over $\mathcal{X}\times\mathcal{Y}$ is a fair base rate distribution if and only if for any label $y \in [C]$, the base rate $b^{\mathcal{P}}_{s,y}$ is equal across all $s \in [S]$, i.e., $\forall i \in [S], \forall j \in [S], b^{\mathcal{P}}_{i,y} = b^{\mathcal{P}}_{j,y}$.

*Remark.* In the literature, the concepts of fairness are usually directly defined at the model prediction level, where the criterion is whether the model prediction is fair against individual attribute changes [38, 35, 49] or fair at population level [53]. In this work, to certify the fairness of model prediction, we define a fairness constrained distribution on which we will certify the model prediction (e.g., bound the prediction error), rather than relying on the empirical fairness evaluation. In particular, we first define the fairness constrained distribution through the lens of base rate parity, i.e., the probability of being any class should be independent of sensitive attribute values, and then define the certified fairness of a given model based on its performance on the fairness constrained distribution as we will show next.

The choice of focusing on fair base rate may look restrictive but its definition aligns very well with the celebrated fairness definition Demographic Parity [50], which promotes that $\Pr[h_\theta(X) = 1|X_s = i] = \Pr[h_\theta(X) = 1|X_s = j]$. In this case, the prediction performance of $h_\theta$ on $\mathcal{Q}$ with fair base rate will relate directly to $\Pr[h_\theta(X) = 1|X_s = i]$. Secondly, under certain popular data generation process, the base rate sufficiently encodes the differences in distributions and a fair base rate will imply a homogeneous (therefore equal or "fair") distribution over $X, Y$: consider when $\Pr(X|Y = y, X_s = i)$ is the same across different group $X_s$. Then $\Pr(X, Y|X_s = i)$ is simply a linear combination of basis distributions $\Pr(X|Y = y, X_s = i)$, and the difference between different groups' joint distribution of $X, Y$ is fully characterized by the difference in base rate $\Pr(Y = y|X_s)$. This assumption will greatly enable trackable analysis and is not an uncommon modeling choice in the recent discussion of fairness when distribution shifts [51, 36].

## 2.1 Certified Fairness

Now we are ready to define the fairness certification based on the optimized fairness constrained distribution. We define the certification under two data generation scenarios: *general shifting* and *sensitive shifting*. In particular, consider the data generative model $\Pr(X_o, X_s, Y) = \Pr(Y) \Pr(X_s|Y) \Pr(X_o|Y, X_s)$, where $X_o$ and $X_s$ represent the non-sensitive and sensitive features, respectively. If all three random variables on the RHS are allowed to change, we call it *general shifting*; if both $\Pr(Y)$ and $\Pr(X_s|Y)$ are allowed to change to ensure the fair base rate (Def. 2) while $\Pr(X_o|Y, X_s)$ is the same across different groups, we call it *sensitive shifting*. In Section 3 we will introduce our certification framework for both scenarios.

**Problem 1** (Certified Fairness with General Shifting). *Given a training distribution $\mathcal{P}$ supported on $\mathcal{X} \times \mathcal{Y}$, a model $h_\theta(\cdot)$ trained on $\mathcal{P}$, and distribution distance bound $\rho > 0$, we call $\bar{\ell} \in \mathbb{R}$ a* fairness certificate *with general shifting, if $\bar{\ell}$ upper bounds*

$$\max_{\mathcal{Q}} \mathbb{E}_{(X,Y)\sim\mathcal{Q}}[\ell(h_\theta(X), Y)] \quad \text{s.t.} \quad \mathrm{dist}(\mathcal{P}, \mathcal{Q}) \le \rho, \quad \mathcal{Q} \textit{ is a fair distribution},$$

*where* $\mathrm{dist}(\cdot, \cdot)$ *is a predetermined distribution distance metric.*

In the above definition, we define the fairness certificate as the upper bound of the model's loss among all fair base rate distributions $\mathcal{Q}$ within a bounded distance from $\mathcal{P}$. Besides the bounded distance constraint $\mathrm{dist}(\mathcal{P}, \mathcal{Q}) \le \rho$, there is no other constraint between $\mathcal{P}$ and $\mathcal{Q}$ so this satisfies "*general shifting*". This bounded distance constraint, parameterized by a tunable parameter $\rho$, ensures that the test distribution should not be too far away from the training. In practice, the model $h_\theta$ may represent a DNN whose complex analytical forms would pose challenges for solving Problem 1. As a result, as we will show in Equation (2) we can query some statistics of $h_\theta$ trained on $\mathcal{P}$ as constraints to characterize $h_\theta$, and thus compute the upper bound certificate.

The feasible region of optimization problem 1 might be empty if the distance bound $\rho$ is too small, and thus we cannot provide fairness certification in this scenario, indicating that there is no nearby fair distribution and thus the fairness of the model trained on the highly "unfaired" distribution is generally low. In other words, if the training distribution $\mathcal{P}$ is unfair (typical case) and there is no feasible fairness constrained distribution $\mathcal{Q}$ within a small distance to $\mathcal{P}$, fairness cannot be certified.

This definition follows the intuition of typical real-world scenarios: The real-world training dataset is usually biased due to the limitation in data curation and collection processes, which causes the model to be unfair. Thus, when the trained models are evaluated on the real-world fairness constrained test distribution or ideal fair distribution, we hope that the model does not encode the training bias which would lead to low test performance. That is to say, the model performance on fairness constrained distribution is indeed a witness of the model's intrinsic fairness.

We can further constrain that the subpopulation of $\mathcal{P}$ and $\mathcal{Q}$ parameterized by $X_s$ and $Y$ does not change, which results in the following "*sensitive shifting*" fairness certification.

**Problem 2** (Certified Fairness with Sensitive Shifting). *Under the same setting as Problem 1, we call $\bar{\ell}$ a fairness certificate against sensitive shifting, if $\bar{\ell}$ upper bounds*

$$\max_{\mathcal{Q}} \mathbb{E}_{(X,Y)\sim\mathcal{Q}}[\ell(h_\theta(X), Y)]$$
$$\text{s.t.} \quad \mathrm{dist}(\mathcal{P}, \mathcal{Q}) \le \rho, \quad \mathcal{P}_{s,y} = \mathcal{Q}_{s,y} \, \forall s \in [S], y \in [C], \quad \mathcal{Q} \textit{ is a fair distribution},$$

*where $\mathcal{P}_{s,y}$ and $\mathcal{Q}_{s,y}$ are the subpopulations of $\mathcal{P}$ and $\mathcal{Q}$ on the support $\{(X,Y) : X \in \mathcal{X}, X_s = s, Y = y\}$ respectively, and* $\mathrm{dist}(\cdot, \cdot)$ *is a predetermined distribution distance metric.*

The definition adds an additional constraint between $\mathcal{P}$ and $\mathcal{Q}$ that each subpopulation, partitioned by the sensitive attribute $X_s$ and label $Y$, does not change. This constraint corresponds to the scenario where the distribution shifting between training and test distributions only happens on the proportions of different sensitive attributes and labels, and within each subpopulation the shifting is negligible.

In addition, to model the real-world test distribution, we may further request that the test distribution $\mathcal{Q}$ is not too skewed regarding the sensitive attribute $X_s$ by adding constraint (1). We will show that this constraint can also be integrated into our fairness certification framework flexibly in Section 4.3.

$$\forall i \in [S], \forall j \in [S], \left| \Pr_{(X,Y)\sim\mathcal{Q}}[X_s = i] - \Pr_{(X,Y)\sim\mathcal{Q}}[X_s = j] \right| \le \Delta_S. \tag{1}$$

**Connections to Other Fairness Measurements.** Though not explicitly stated, our goal of certifying the performance on a fair distribution $\mathcal{Q}$ relates to certifying established fairness definitions in the literature. Consider the following example: Suppose Problem 2 is feasible and returns a classifier $h_\theta$ that achieves certified fairness per group and per label class $\bar{l} := \Pr_{(X,Y)\sim\mathcal{Q}}[h_\theta(X) \ne Y | Y = y, X_s = i] \le \epsilon$ on $\mathcal{Q}$. We will then have the following proposition:

**Proposition 1.** *$h_\theta$ achieves $\epsilon$-Demographic Parity (DP) [50] and $\epsilon$-Equalized Odds (EO) [17]:*

- *$\epsilon$-DP:* $|\Pr_{\mathcal{Q}}[h_\theta(X)=1|X_s=i]-\Pr_{\mathcal{Q}}[h_\theta(X)=1|X_s=j]|\le\epsilon,\ \forall i,j.$
- *$\epsilon$-EO:* $|\Pr_{\mathcal{Q}}[h_\theta(X)=1|Y=y,X_s=i]-\Pr_{\mathcal{Q}}[h_\theta(X)=1|Y=y,X_s=j]|\le\epsilon,\forall y,i,j.$

The detailed proof is omitted to appendix D.1.

# 3 Fairness Certification Framework

We will introduce our fairness certification framework which efficiently computes the fairness certificate defined in Section 2.1. We first introduce our framework for *sensitive shifting* (Problem 2) which is less complex and shows our core methodology, then *general shifting* case (Problem 1).

Our framework focuses on using the Hellinger distance to bound the distributional distance in Problems 1 and 2. The Hellinger distance $H(\mathcal{P},\mathcal{Q})$ is defined in Def. 3 (in Appendix C.1). The Hellinger distance has some nice properties, e.g., $H(\mathcal{P},\mathcal{Q})\in[0,1]$, and $H(\mathcal{P},\mathcal{Q})=0$ if and only if $\mathcal{P}=\mathcal{Q}$ and the maximum value of 1 is attained when $\mathcal{P}$ and $\mathcal{Q}$ have disjoint support. The Hellinger distance is a type of $f$-divergences which are widely studied in ML distributional robustness literature [46, 13] and in the context of distributionally robust optimization [7, 24, 14]. Also, using Hellinger distance enables our certification framework to generalize to *total variation distance (or statistic distance)* $\delta(\mathcal{P},\mathcal{Q})$[1] directly with the connection, $H^2(\mathcal{P},\mathcal{Q})\le\delta(\mathcal{P},\mathcal{Q})\le\sqrt{2}H(\mathcal{P},\mathcal{Q})$ ([44], Equation 1). We leave the extension of our framework to other distance metrics as future work.

## 3.1 Core Idea: Subpopulation Decomposition

The core idea in our framework is (finite) subpopulation decomposition. Consider a generic optimization problem for computing the loss upper bound on a constrained test distribution $\mathcal{Q}$, given training distribution $\mathcal{P}$ and trained model $h_\theta(\cdot)$, we first characterize model $h_\theta(\cdot)$ based on some statistics, e.g., mean and variance for loss of the model: $h_\theta(\cdot)$ satisfies $e_j(\mathcal{P},h_\theta)\le v_j$, $1\le j\le L$. Then we characterize the properties (e.g., fair base rate) of the test distribution $\mathcal{Q}$: $g_j(\mathcal{Q})\le u_j$, $1\le j\le M$. As a result, we can upper bound the loss of $h_\theta(\cdot)$ on $\mathcal{Q}$ as the following optimization:

$$\max_{\mathcal{Q},\theta}\mathbb{E}_{(X,Y)\sim\mathcal{Q}}[\ell(h_\theta(X),Y)]\quad\text{s.t.}\quad H(\mathcal{P},\mathcal{Q})\le\rho,\quad e_j(\mathcal{P},h_\theta)\le v_j\,\forall j\in[L],\quad g_j(\mathcal{Q})\le u_j\,\forall j\in[M]. \tag{2}$$

Now we decompose the space $\mathcal{Z}:=\mathcal{X}\times\mathcal{Y}$ to $N$ partitions: $\mathcal{Z}:=\biguplus\mathcal{Z}_i$, where $\mathcal{Z}$ is the support of both $\mathcal{P}$ and $\mathcal{Q}$. Then, we denote $\mathcal{P}$ conditioned on $\mathcal{Z}_i$ by $\mathcal{P}_i$ and similarly $\mathcal{Q}$ conditioned on $\mathcal{Z}_i$ by $\mathcal{Q}_i$. As a result, we can write $\mathcal{P}=\sum_{i\in[N]}p_i\mathcal{P}_i$ and $\mathcal{Q}=\sum_{i\in[N]}q_i\mathcal{Q}_i$. Since $\mathcal{P}$ is known, $p_i$'s are known. In contrast, both $\mathcal{Q}_i$ and $q_i$'s are optimizable. Our key observation is that

$$H(\mathcal{P},\mathcal{Q})\le\rho\iff 1-\rho^2-\sum_{i=1}^N\sqrt{p_iq_i}(1-H(\mathcal{P}_i,\mathcal{Q}_i)^2)\le 0 \tag{3}$$

which leads to the following theorem.

**Theorem 1.** *The following constrained optimization upper bounds Equation (2):*

$$\max_{\mathcal{Q}_i,q_i,\rho_i,\theta}\quad\sum_{i=1}^N q_i\mathbb{E}_{(X,Y)\sim\mathcal{Q}_i}[\ell(h_\theta(X),Y)] \tag{4a}$$

$$\text{s.t.}\quad 1-\rho^2-\sum_{i=1}^N\sqrt{p_iq_i}(1-\rho_i^2)\le 0, \tag{4b}$$

$$H(\mathcal{P}_i,\mathcal{Q}_i)\le\rho_i\quad\forall i\in[N],\quad\sum_{i=1}^N q_i=1,\quad q_i\ge 0\quad\forall i\in[N],\quad\rho_i\ge 0\quad\forall i\in[N], \tag{4c}$$

$$e'_j(\{\mathcal{P}_i\}_{i\in[N]},\{p_i\}_{i\in[N]},h_\theta)\le v'_j\,\forall j\in[L],\quad g'_j(\{\mathcal{Q}_i\}_{i\in[N]},\{q_i\}_{i\in[N]})\le u'_j\,\forall j\in[M], \tag{4d}$$

*if $e_j(\mathcal{P},h_\theta)\le v_j$ implies $e'_j(\{\mathcal{P}_i\}_{i\in[N]},\{p_i\}_{i\in[N]},h_\theta)\le v'_j$ for any $j\in[L]$, and $g_j(\mathcal{Q})\le u_j$ implies $g'_j(\{\mathcal{Q}_i\}_{i\in[N]},\{q_i\}_{i\in[N]})\le u'_j$ for any $j\in[M]$.*

[1] $\delta(\mathcal{P},\mathcal{Q})=\sup_{A\in\mathcal{F}}|\mathcal{P}(A)-\mathcal{Q}(A)|$ where $\mathcal{F}$ is a $\sigma$-algebra of subsets of the sample space $\Omega$.

We provide a detailed proof in Appendix D.2. The constrained optimization (4) breaks down the whole distribution $\mathcal{Q}$ to subpopulations $\{\mathcal{Q}_i\}_{i=1}^N$ by introducing the ratio variables $q_i$ and distance variables $\rho_i$. As a result, we have an equivalent form of bounded distance constraint but defined on subpopulations: Equation (4b). Although the optimization problem (Equation (4)) may look more complicated then the original Equation (2), this optimization allows us to upper bound each subpopulation loss $\mathbb{E}_{(X,Y)\sim\mathcal{Q}_i}[\ell(h_\theta(X),Y)]$ individually and makes the whole optimization tractable.

## 3.2 Certified Fairness with Sensitive Shifting

For the sensitive shifting case, we instantiate Thm. 1 and obtain the following fairness certificate.

**Theorem 2.** *Given a distance bound $\rho > 0$, the following constrained optimization, which is **convex**, when feasible, provides a **tight** fairness certificate for Problem 2:*

$$\max_{k_s,r_y} \sum_{s=1}^{S}\sum_{y=1}^{C} k_s r_y E_{s,y}, \quad \text{s.t.} \quad \sum_{s=1}^{S} k_s = 1, \quad \sum_{y=1}^{C} r_y = 1, \quad k_s \ge 0 \quad \forall s \in [S], \quad r_y \ge 0 \quad \forall y \in [C],$$
$$1 - \rho^2 - \sum_{s=1}^{S}\sum_{y=1}^{C} \sqrt{p_{s,y} k_s r_y} \le 0,$$

*where $E_{s,y} := \mathbb{E}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X),Y)]$ and $p_{s,y} := \Pr_{(X,Y)\in\mathcal{P}}[X_s = s, Y = y]$ are constants.*

*Proof sketch.* We decompose distribution $\mathcal{P}$ and $\mathcal{Q}$ to $\mathcal{P}_{s,y}$'s and $\mathcal{Q}_{s,y}$'s according to their sensitive attribute and label values. In sensitive shifting, $\Pr(X_o|Y,X_s)$ is fixed, i.e., $\mathcal{P}_{s,y} = \mathcal{Q}_{s,y}$, which means $\mathbb{E}_{(X,Y)\sim\mathcal{Q}_{s,y}}[\ell(h_\theta(X),Y)] = E_{s,y}$ and $\rho_{s,y} = H(\mathcal{P}_{s,y},\mathcal{Q}_{s,y}) = 0$. We plug these properties into Thm. 1. Then, denoting $q_{s,y}$ to $\Pr_{(X,Y)\sim\mathcal{Q}}[X_s = s, Y = y]$, we can represent the fairness constraint in Def. 2 as $q_{s_0,y_0} = \left(\sum_{s=1}^{S} q_{s,y_0}\right)\left(\sum_{y=1}^{C} q_{s_0,y}\right)$ for any $s_0 \in [S]$ and $y_0 \in [C]$. Next, we parameterize $q_{s,y}$ with $k_s r_y$. Such parameterization simplifies the fairness constraint and allow us to prove the convexity of the resulting optimization. Since all the constraints are encoded equivalently, the problem formulation provides a tight certification. Detailed proof in Appendix D.3. □

As Thm. 2 suggests, we can exploit the expectation information $E_{s,y} = \mathbb{E}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X),Y)]$ and density information $p_{s,y} = \Pr_{(X,Y)\sim\mathcal{P}}[X_s = s, Y = y]$ of each $\mathcal{P}$'s subpopulation to provide a tight fairness certificate in sensitive shifting. The convex optimization problem with $(S + C)$ variables can be efficiently solved by existing packages.

## 3.3 Certified Fairness with General Shifting

For the general shifting case, we leverage Thm. 1 and the parameterization trick $q_{s,y} := k_s r_y$ used in Thm. 2 to reduce Problem 1 to the following constrained optimization.

**Lemma 3.1.** *Given a distance bound $\rho > 0$, the following constrained optimization, when feasible, provides a **tight** fairness certificate for Problem 1:*

$$\max_{k_s,r_y,\mathcal{Q},\rho_{s,y}} \quad \sum_{s=1}^{S}\sum_{y=1}^{C} k_s r_y \mathbb{E}_{(X,Y)\sim\mathcal{Q}_{s,y}}[\ell(h_\theta(X),Y)] \tag{6a}$$
$$\text{s.t.} \quad \sum_{s=1}^{S} k_s = 1, \quad \sum_{y=1}^{C} r_y = 1, \quad k_s \ge 0 \quad \forall s \in [S], \quad r_y \ge 0 \quad \forall y \in [C], \tag{6b}$$
$$\sum_{s=1}^{S}\sum_{y=1}^{C} \sqrt{p_{s,y} k_s r_y}(1 - \rho_{s,y}^2) \ge 1 - \rho^2 \tag{6c}$$
$$H(\mathcal{P}_{s,y},\mathcal{Q}_{s,y}) \le \rho_{s,y} \quad \forall s \in [S], y \in [C], \tag{6d}$$

*where $p_{s,y} := \Pr_{(X,Y)\in\mathcal{P}}[X_s = s, Y = y]$ is a fixed constant. The $\mathcal{P}_{s,y}$ and $\mathcal{Q}_{s,y}$ are the subpopulations of $\mathcal{P}$ and $\mathcal{Q}$ on the support $\{(X,Y) : X \in \mathcal{X}, X_s = s, Y = y\}$ respectively.*

*Proof sketch.* We show that Equation (6b) ensures a parameterization of $q_{s,y} = \Pr_{(X,Y)\in\mathcal{Q}}[X_s = s, Y = y]$ that satisfies fairness constraints on $\mathcal{Q}$. Then, leveraging Thm. 1 we prove that the constrained optimization provides a fairness certificate. Since all the constraints are either kept or equivalently encoded, this resulting certification is *tight*. Detailed proof in Appendix D.4. $\square$

To solve the constrained optimization, we first leverage Thm. 4 in Appendix C.2 [46] to upper bound the loss of $h_\theta(\cdot)$ within each shifted subpopulation $\mathcal{Q}_{s,y}$, i.e., an upper bound of $\mathbb{E}_{(X,Y)\sim\mathcal{Q}_{s,y}}[\ell(h_\theta(X), Y)]$ in Equation (6a), and then propose a novel technique to prove the convexity of sub-problems regarding our transformed decision variables, which is one of our main contributions. Here, we construct the sub-problems by partitioning the feasible region with different density ratios within each subpopulation. Concretely, the following theorem states our computable certificate for Problem 1, with detailed proof and explanation in Appendix D.5.

**Theorem 3.** *If for any $s \in [S]$ and $y \in [Y]$, $H(\mathcal{P}_{s,y}, \mathcal{Q}_{s,y}) \le \bar{\gamma}_{s,y}$ and $0 \le \sup_{(X,Y)\in\mathcal{X}\times\mathcal{Y}} \ell(h_\theta(X), Y) \le M$, given a distance bound $\rho > 0$, for any region granularity $T \in \mathbb{N}_+$, the following expression provides a fairness certificate for Problem 1:*

$$\bar{\ell} = \max_{\{i_s\in[T]:s\in[S]\},\{j_y\in[T]:y\in[C]\}} \mathsf{C}\left(\left\{\left[\frac{i_s-1}{T}, \frac{i_s}{T}\right]\right\}_{s=1}^{S}, \left\{\left[\frac{j_y-1}{T}, \frac{j_y}{T}\right]\right\}_{y=1}^{C}\right), \text{ where} \tag{7}$$

$$\mathsf{C}\left(\{[\underline{k_s}, \overline{k_s}]\}_{s=1}^S, \{[\underline{r_y}, \overline{r_y}]\}_{y=1}^C\right) = \max_{x_{s,y}} \sum_{s=1}^{S}\sum_{y=1}^{C} \Big(\overline{k_s}\overline{r_y}\left(E_{s,y} + C_{s,y}\right)_+ + \underline{k_s}\underline{r_y}\left(E_{s,y} + C_{s,y}\right)_- +$$
$$2\overline{k_s}\overline{r_y}\sqrt{x_{s,y}(1-x_{s,y})}\sqrt{V_{s,y}} - \underline{k_s}\underline{r_y}x_{s,y}(C_{s,y})_+ - \overline{k_s}\overline{r_y}x_{s,y}(C_{s,y})_-\Big) \tag{8a}$$

$$\text{s.t.} \quad \sum_{s=1}^{S}\underline{k_s} \le 1, \quad \sum_{s=1}^{S}\overline{k_s} \ge 1, \quad \sum_{y=1}^{C}\underline{r_y} \le 1, \quad \sum_{y=1}^{C}\overline{r_y} \ge 1, \tag{8b}$$

$$\sum_{s=1}^{S}\sum_{y=1}^{C}\sqrt{p_{s,y}\overline{k_s}\overline{r_y}x_{s,y}} \ge 1-\rho^2, \quad (1-\bar{\gamma}_{s,y}^2)^2 \le x_{s,y} \le 1 \quad \forall s \in [S], y \in [C], \tag{8c}$$

*where $(\cdot)_+ = \max\{\cdot, 0\}$, $(\cdot)_- = \min\{\cdot, 0\}$; $E_{s,y} = \mathbb{E}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)]$, $V_{s,y} = \mathbb{V}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)]$, $p_{s,y} = \Pr_{(X,Y)\sim\mathcal{P}}[X_s = s, Y = y]$, $C_{s,y} = M - E_{s,y} - \frac{V_{s,y}}{M - E_{s,y}}$, and $\bar{\gamma}_{s,y}^2 = 1 - (1 + (M - E_{s,y})^2/V_{s,y})^{-\frac{1}{2}}$. Equation (7) only takes $\mathsf{C}$'s value when it is feasible, and each $\mathsf{C}$ queried by Equation (7) is a **convex optimization**.*

**Implications.** Thm. 3 provides a fairness certificate for Problem 1 under two assumptions: (1) The loss function is bounded (by $M$). This assumption holds for several typical losses such as 0-1 loss and JSD loss. (2) The distribution shift between training and test distribution within each subpopulation is bounded by $\bar{\gamma}_{s,y}$, where $\bar{\gamma}_{s,y}$ is determined by the model's statistics on $\mathcal{P}$. In practice, this additional distance bound assumption generally holds, since $\bar{\gamma}_{s,y} \gg \rho$ for common choices of $\rho$.

In Thm. 3, we exploit three types of statistics of $h_\theta(\cdot)$ on $\mathcal{P}$ to compute the fairness certificates: the expectation $E_{s,y} = \mathbb{E}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)]$, the variance $V_{s,y} = \mathbb{V}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)]$, and the density $p_{s,y} = \Pr_{(X,Y)\sim\mathcal{P}}[X_s = s, Y = y]$, all of which are at the subpopulation level and a high-confidence estimation of them based on finite samples are tractable (Section 3.4).

Using Thm. 3, after determining the region granularity $T$, we can provide a fairness certificate for Problem 1 by solving $T^{SC}$ convex optimization problems, each of which has $SC$ decision variables. Note that the computation cost is independent of $h_\theta$, and therefore we can numerically compute the certificate for large DNN models used in practice. Specifically, when $S = 2$ (binary sensitive attribute) or $C = 2$ (binary classification) which is common in the fairness evaluation setting, we can construct the region for only one dimension $k_1$ or $r_1$, and use $1 - k_1$ or $1 - r_1$ for the other dimension. Thus, for the typical setting $S = 2, C = 2$, we only need to solve $T^2$ convex optimization problems.

Note that for Problem 2, our certificate in Thm. 2 is tight, whereas for Problem 1, our certificate in Thm. 3 is not. This is because in Problem 1, extra distribution shift exists within each subpopulation, i.e., $\Pr(X_o|Y, X_s)$ changes from $\mathcal{P}$ to $\mathcal{Q}$, and to bound such shift, we need to leverage Thm. 2.2 in [46] which has no tightness guarantee. Future work providing tighter bounds than [46] can be seamlessly incorporated into our framework to tighten our fairness certificate for Problem 1.

### 3.4 Dealing with Finite Sampling Error

In Section 3.2 and Section 3.3, we present Thm. 2 and Thm. 3 that provide computable fairness certificates for sensitive shifting and general shifting scenarios respectively. In these theorems, we need to know the quantities related to the training distribution and trained $\mathcal{P}$ and model $h_\theta(\cdot)$:

$$E_{s,y} = \mathop{\mathbb{E}}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)], V_{s,y} = \mathop{\mathbb{V}}_{(X,Y)\sim\mathcal{P}_{s,y}}[\ell(h_\theta(X), Y)], p_{s,y} = \Pr_{(X,Y)\sim\mathcal{P}}[X_s = s, Y = y]. \quad (9)$$

Section 3.3 further requires $C_{s,y}$ and $\bar{\gamma}_{s,y}$ which are functions of $E_{s,y}$ and $V_{s,y}$. However, a practical challenge is that common training distributions do not have an analytical expression that allows us to precisely compute these quantities. Indeed, we only have access to a finite number of individually drawn samples, i.e., the training dataset, from $\mathcal{P}$. Thus, we will provide high-confidence bounds for $E_{s,y}$, $V_{s,y}$, and $p_{s,y}$ in Lemma E.1 (stated in Appendix E.1).

For Thm. 2, we can replace $E_{s,y}$ in the objective by the upper bounds of $E_{s,y}$ and replace the concrete quantities of $p_{s,y}$ by interval constraints and the unit constraint $\sum_s \sum_y p_{s,y} = 1$, which again yields a convex optimization that can be effectively solved. For Thm. 3, we compute the confidence intervals of $C_{s,y}$ and $\rho_{s,y}$, then plug in either the lower bounds or the upper bounds to the objective (8a) based on the coefficient, and finally replace the concrete quantities of $p_{s,y}$ by interval constraints and the unit constraint $\sum_s \sum_y p_{s,y} = 1$. The resulting optimization is proved to be convex and provides an upper bound for any possible values of $E_{s,y}$, $V_{s,y}$, and $p_{s,y}$ within the confidence intervals. We defer the statement of Thm. 2 and Thm. 3 considering finite sampling error to Appendix E.2. To this point, we have presented our framework for computing high-confidence fairness certificates given access to model $h_\theta(\cdot)$ and a finite number of samples drawn from $\mathcal{P}$.

## 4 Experiments

In this section, we evaluate the certified fairness under both *sensitive shifting* and *general shifting* scenarios on six real-world datasets. We observe that under the sensitive shifting, our certified fairness bound is *tight* (Section 4.1); while the bound is less tight under general shifting (Section 4.2) which depends on the tightness of generalization within each subpopulation (details in Section 3.3). In addition, we show that our certification framework can flexibly integrate more constraints on $\mathcal{Q}$, leading to a tighter fairness certification (Section 4.3). Finally, we compare our certified fairness bound with existing distributional robustness bound [42] (section 4.4), since both consider a shifted distribution while our bound is optimized with an additional fairness constraint which is challenging to be directly integrated to the existing distributional robustness optimization. We show that with the fairness constraint and our optimization approach, our bound is much tighter.

**Dataset & Model.** We validate our certified fairness on *six* real-world datasets: Adult [3], Compas [2], Health [18], Lawschool [47], Crime [3], and German [3]. Details on the datasets and data processing steps are provided in Appendix F.1. Following the standard setup of fairness evaluation in the literature [38, 37, 30, 41], we consider the scenario that the sensitive attributes and labels take binary values. The ReLU network composed of 2 hidden layers of size 20 is used for all datasets.

**Fairness Certification**. We perform vanilla model training and then leverage our fairness certification framework to calculate the fairness certificate. Concretely, we input the trained model information on $\mathcal{P}$ and the framework would output the fairness certification for both sensitive shifting and general shifting scenarios following Thm. 2 and Thm. 3, respectively.

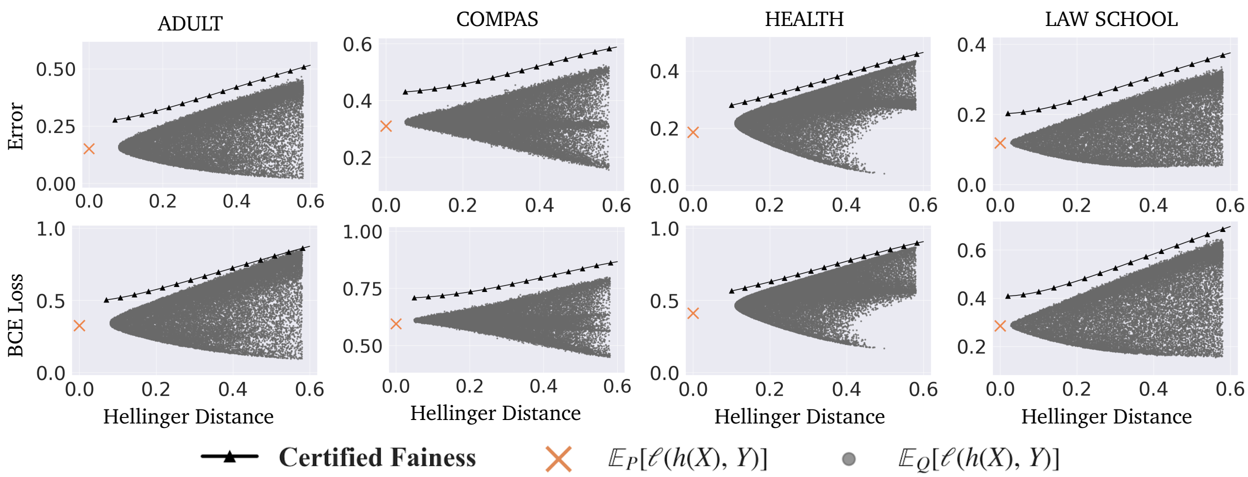


Figure 1: Certified fairness with sensitive shifting.

### 4.1 Certified Fairness with Sensitive Shifting

To evaluate how well our certificates capture the fairness risk in practice, we compare our certification bound with the empirical loss evaluated on randomly generated $30,000$ fairness constrained distributions $\mathcal{Q}$ shifted from $\mathcal{P}$. Under sensitive shifting, since each subpopulation divided by the sensitive attribute and label does not change (Section 2.1), we randomly generate distribution $\mathcal{Q}$ based on different combinations of subpopulation proportion $q_{s,y}$ sastisfying fair base rate, and then sample from each subpopulation of $\mathcal{P}$ individually according to the proportion $q_{s,y}$. The detailed generation steps are provided in Appendix F.2. We report the classification error (Error) and BCE loss as the evaluation metric. Figure 1 illustrates the certified fairness on Adult, Compas, Health, and Lawschool under sensitive shifting. More results on two relatively small datasets (Crime, German) are shown in Appendix F.5. From the results, we see that our certified fairness is tight in practice.

### 4.2 Certified Fairness with General Shifting

In the general shifting scenario, we similarly randomly generate $30,000$ fair distributions $\mathcal{Q}$ shifted from $P$. Different from sensitive shifting, the distribution conditioned on sensitive attribute $X_s$ and label $Y$ can also change in this scenario. Therefore, we construct another distribution $\mathcal{Q}'$ disjoint with $\mathcal{P}$ on non-sensitive attributes and mix $\mathcal{P}$ and $\mathcal{Q}'$ in each subpopulation individually guided by mixing parameters satisfying fair base rate constraint. Detailed generation steps are given in Appendix F.2. Since the fairness certification for general shifting requires bounded loss, we select classification error (Error) and Jensen-Shannon loss (JSD Loss) as the evaluation metric. Figure 2 illustrates the certified fairness with classification error metric under general shifting. Results of JSD loss and more results on two relatively small datasets (Crime, German) are in Appendix F.5.

### 4.3 Certified Fairness with Additional Non-Skewness Constraints

In Section 2.1, we discussed that to represent different real-world scenarios we can add more constraints such as Equation (1) to prevent the skewness of $\mathcal{Q}$, which can be flexibly incorporated into our certificate framework. Concretely, for sensitive shifting, we only need to add one more box constraint[2] $0.5 - \Delta_s/2 \le k_s \le 0.5 + \Delta_s/2$ where $\Delta_s$ is a parameter controlling the skewness of $\mathcal{Q}$, which still guarantees convexity. For general shifting, we only need to modify the region partition step[2], where we split $[0.5 - \Delta_s/2, 0.5 + \Delta_s/2]$ instead of $[0, 1]$. The certification results with additional constraints are in Figures 3(a) and 3(b), which suggests that if the added constraints are strict (i.e., smaller $\Delta_s$), the bound is tighter. More constraints w.r.t. labels can also be handled by our framework and the corresponding results as well as results on more datasets are in Appendix F.6.

[2]Note that such modification is only viable when sentive attributes take binary values, which is the typical scenario in the literature of fairness evaluation [38, 37, 30, 41].

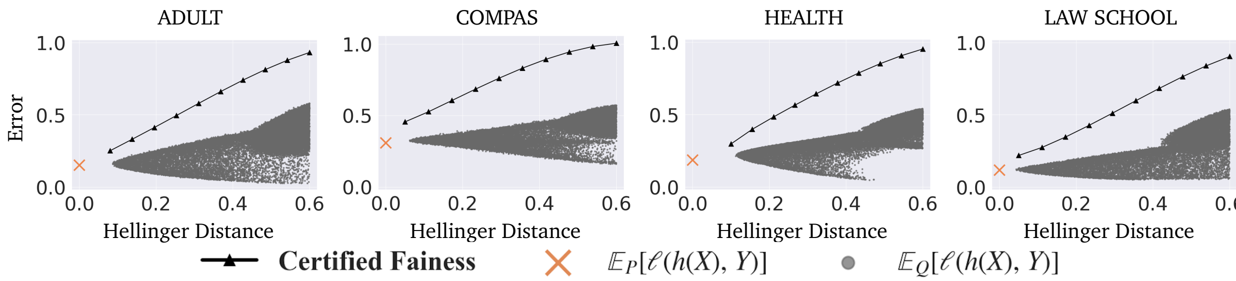


Figure 2: Certified fairness with general shifting.

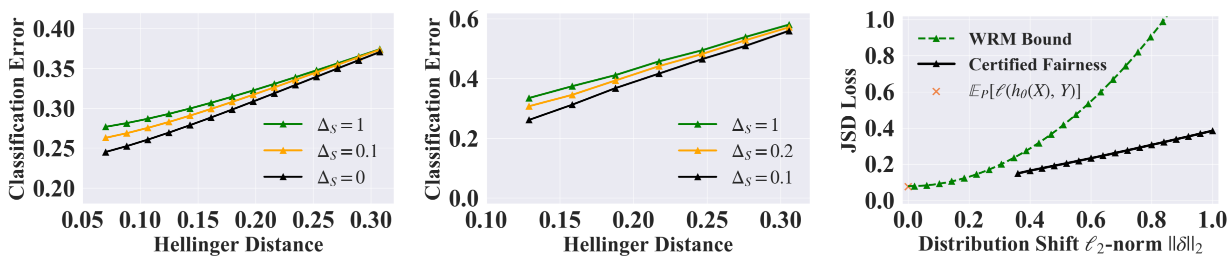


(a) Sentitive Shifting (b) General Shifting (c) Comparison with WRM

Figure 3: Certified fairness with additional non-skewness constraints on Adult dataset is shown in (a) (b). $\Delta_s$ controls the skewness of $\mathcal{Q}$ ($|\Pr_{(X,Y)\sim\mathcal{Q}}[X_s = 0] - \Pr_{(X,Y)\sim\mathcal{Q}}[X_s = 1]| \leq \Delta_s$). In (c), we compare our certified fairness bound with the distributional robustness bound [42].

### 4.4 Comparison with Distributional Robustness Bound

To the best of our knowledge, there is no existing work providing *certified fairness* on the end-to-end model performance. Thus, we try to compare our bound with the distributional robustness bound since both consider certain distribution shifts. However, it is challenging to directly integrate the fairness constraints into existing bounds. Therefore, we compare with the state-of-the-art distributional robustness certification WRM [42], which solves the similar optimization problem as ours except for the fairness constraint. For fair comparison, we construct a synthetic dataset following [42], on which there is a one-to-one correspondence between the Hellinger and Wasserstein distance used by WRM. We randomly select one dimensio as the sensitive attribute. Since WRM has additional assumptions on smoothness of models and losses, we use JSD loss and a small ELU network with 2 hidden layers of size 4 and 2 following their setting. More implementation details are in Appendix F.4. Results in Figure 3(c) suggest that 1) our certified fairness bound is much tighter than WRM given the additional fairness distribution constraint and our optimization framework; 2) with additional fairness constraint, our certificate problem could be infeasible under very small distribution distances since the fairness constrained distribution $\mathcal{Q}$ does not exist near the skewed original distribution $\mathcal{P}$; 3) with the fairness constraint, we provide non-trivial fairness certification even when the distribution shift is large.

## 5 Conclusion

In this paper, we provide the first *fairness certification* on end-to-end model performance, based on a fairness constrained distribution which has bounded distribution distance from the training distribution. We show that our fairness certification has strong connections with existing fairness notions such as group parity, and we provide an effective framework to calculate the certification under different scenarios. We provide both theoretical and empirical analysis of our fairness certification.

**Acknowledgements.** This work is partially supported by NSF grant No.1910100, NSF CNS No.2046726, C3 AI, and the Alfred P. Sloan Foundation.

## References


[1] Aws Albarghouthi, Loris D'Antoni, Samuel Drews, and Aditya V Nori. Fairsquare: probabilistic verification of program fairness. *Proceedings of the ACM on Programming Languages*, 1(OOPSLA):1–30, 2017.

[2] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. In *Ethics of Data and Analytics*, pages 254–264. Auerbach Publications, 2016.

[3] Arthur Asuncion and David Newman. Uci machine learning repository, 2007.

[4] Mislav Balunović, Anian Ruoss, and Martin Vechev. Fair normalizing flows. *arXiv preprint arXiv:2106.05937*, 2021.

[5] Solon Barocas and Andrew D Selbst. Big data's disparate impact. *Calif. L. Rev.*, 104:671, 2016.

[6] Osbert Bastani, Xin Zhang, and Armando Solar-Lezama. Probabilistic verification of fairness properties via concentration. *Proceedings of the ACM on Programming Languages*, 3(OOPSLA):1–27, 2019.

[7] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

[8] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1):3–44, 2021.

[9] YooJung Choi, Meihua Dang, and Guy Van den Broeck. Group fairness by probabilistic modeling with latent fair decisions. *arXiv preprint arXiv:2009.09031*, 2020.

[10] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.

[11] Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *International conference on machine learning*, pages 1436–1445. PMLR, 2019.

[12] Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad privacy settings. *Proceedings on privacy enhancing technologies*, 2015(1):92–112, 2015.

[13] John Duchi and Hongseok Namkoong. Variance-based regularization with convex objectives. *The Journal of Machine Learning Research*, 20(1):2450–2504, 2019.

[14] John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.

[15] Harrison Edwards and Amos Storkey. Censoring representations with an adversary. *arXiv preprint arXiv:1511.05897*, 2015.

[16] Jeffrey A Gliner. Reviewing qualitative research: Proposed criteria for fairness and rigor. *The Occupational Therapy Journal of Research*, 14(2):78–92, 1994.

[17] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.

[18] Kaggle Inc. Heritage health prize kaggle. `https://www.kaggle.com/c/hhp`, 2022.

[19] Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves, and Xiang Ren. On transferability of bias mitigation effects in language model fine-tuning. *arXiv preprint arXiv:2010.12864*, 2020.

[20] Philips George John, Deepak Vijaykeerthy, and Diptikalyan Saha. Verifying individual fairness in machine learning models. In *Conference on Uncertainty in Artificial Intelligence*, pages 749–758. PMLR, 2020.

[21] Thomas Kehrenberg, Myles Bartlett, Oliver Thomas, and Novi Quadrianto. Null-sampling for interpretable and fair representations. In *European Conference on Computer Vision*, pages 565–580. Springer, 2020.

[22] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.

[23] Himabindu Lakkaraju, Jon Kleinberg, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 275–284, 2017.

[24] Henry Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.

[25] Jiachun Liao, Chong Huang, Peter Kairouz, and Lalitha Sankar. Learning generative adversarial representations (gap) under fairness and censoring constraints. *arXiv preprint arXiv:1910.00411*, 1, 2019.

[26] Lydia T Liu, Max Simchowitz, and Moritz Hardt. The implicit fairness criterion of unconstrained learning. In *International Conference on Machine Learning*, pages 4051–4060. PMLR, 2019.

[27] Francesco Locatello, Gabriele Abbati, Thomas Rainforth, Stefan Bauer, Bernhard Schölkopf, and Olivier Bachem. On the fairness of disentangled representations. *Advances in Neural Information Processing Systems*, 32, 2019.

[28] Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard Zemel. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*, 2015.

[29] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*, pages 3384–3393. PMLR, 2018.

[30] Subha Maity, Debarghya Mukherjee, Mikhail Yurochkin, and Yuekai Sun. Does enforcing fairness mitigate biases caused by subpopulation shift? In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 25773–25784. Curran Associates, Inc., 2021.

[31] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.

[32] Daniel McNamara, Cheng Soon Ong, and Robert C. Williamson. Costs and benefits of fair representation learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, page 263–270, New York, NY, USA, 2019. Association for Computing Machinery.

[33] Luca Oneto, Michele Donini, Massimiliano Pontil, and Andreas Maurer. Learning fair and transferable representations with theoretical guarantees. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 30–39, 2020.

[34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[35] Momchil Peychev, Anian Ruoss, Mislav Balunović, Maximilian Baader, and Martin Vechev. Latent space smoothing for individually fair representations. *arXiv preprint arXiv:2111.13650*, 2021.

[36] Reilly Raab and Yang Liu. Unintended selection: Persistent qualification rate disparities and interventions. *Advances in Neural Information Processing Systems*, 34, 2021.

[37] Yuji Roh, Kangwook Lee, Steven Whang, and Changho Suh. Sample selection for fair and robust training. *Advances in Neural Information Processing Systems*, 34, 2021.

[38] Anian Ruoss, Mislav Balunovic, Marc Fischer, and Martin Vechev. Learning certified individually fair representations. *Advances in Neural Information Processing Systems*, 33:7584–7596, 2020.

[39] Mhd Hasan Sarhan, Nassir Navab, Abouzar Eslami, and Shadi Albarqouni. Fairness by learning orthogonal disentangled representations. In *European Conference on Computer Vision*, pages 746–761. Springer, 2020.

[40] Shahar Segal, Yossi Adi, Benny Pinkas, Carsten Baum, Chaya Ganesh, and Joseph Keshet. Fairness in the eyes of the data: Certifying machine-learning models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 926–935, 2021.

[41] Shubhanshu Shekhar, Greg Fields, Mohammad Ghavamzadeh, and Tara Javidi. Adaptive sampling for minimax fair classification. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 24535–24544. Curran Associates, Inc., 2021.

[42] Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.

[43] Jiaming Song, Pratyusha Kalluri, Aditya Grover, Shengjia Zhao, and Stefano Ermon. Learning controllable fair representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2164–2173. PMLR, 2019.

[44] Ton Steerneman. On the total variation and hellinger distance between signed measures; an application to product measures. *Proceedings of the American Mathematical Society*, 88(4):684–688, 1983.

[45] Caterina Urban, Maria Christakis, Valentin Wüstholz, and Fuyuan Zhang. Perfectly parallel fairness certification of neural networks. *Proceedings of the ACM on Programming Languages*, 4(OOPSLA):1–30, 2020.

[46] Maurice Weber, Linyi Li, Boxin Wang, Zhikuan Zhao, Bo Li, and Ce Zhang. Certifying out-of-domain generalization for blackbox functions. In *International Conference on Machine Learning*. PMLR, 2022.

[47] Linda F Wightman. Lsac national longitudinal bar passage study. lsac research report series. *Law School Admission Council*, 1998.

[48] Robin Winter, Floriane Montanari, Frank Noé, and Djork-Arné Clevert. Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chemical science*, 10(6):1692–1701, 2019.

[49] Samuel Yeom and Matt Fredrikson. Individual fairness revisited: Transferring techniques from adversarial robustness. *arXiv preprint arXiv:2002.07738*, 2020.

[50] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International conference on machine learning*, pages 325–333. PMLR, 2013.

[51] Xueru Zhang, Ruibo Tu, Yang Liu, Mingyan Liu, Hedvig Kjellstrom, Kun Zhang, and Cheng Zhang. How do fair decisions fare in long-term qualification? *Advances in Neural Information Processing Systems*, 33:18457–18469, 2020.

[52] Han Zhao, Amanda Coston, Tameem Adel, and Geoffrey J Gordon. Conditional learning of fair representations. *arXiv preprint arXiv:1910.07162*, 2019.

[53] Han Zhao and Geoff Gordon. Inherent tradeoffs in learning fair representations. *Advances in neural information processing systems*, 32, 2019.