Scalable and Stable Surrogates for Flexible Classifiers with Fairness Constraints

Anonymous Author(s)
Affiliation
Address
email

Abstract

We investigate how fairness relaxations scale to flexible classifiers like deep neural networks for images and text. We propose an easy-to-use and robust way of imposing fairness constraints when training, and through this framework prove that some prior fairness surrogates exhibit degeneracies for non-convex models. We resolve these problems via three new surrogates: an adaptive data re-weighting, and two smooth upper-bounds that are provably more robust than some previous methods. Our surrogates perform comparably to the state-of-the-art on low-dimensional fairness benchmarks, while achieving superior accuracy and stability for more complex computer vision and natural language processing tasks.

1 Introduction

A growing body of research has explored the ethical and societal implications of increasingly automated decision-making. One critical concern is whether particular machine learning algorithms perform differently on average for members of marginalized groups. In addition to being an ethical question, this is a legal one: Attributes such as race and gender are legally-protected categories, and an employer whose decision-making process disproportionately hurts one group can be litigated for committing adverse or disparate impact. One quantitative notion of fairness is called *demographic parity*, which states that a decision maker should accept (or classify as positive) a roughly equal proportion of each group [8]. An employer who hires many male applicants, but few female applicants, violates demographic parity by disproportionately allocating resources to one sensitive group.

Machine learning is also increasingly used by law enforcement, where ethical questions become civil rights questions. Facial recognition and risk assessment algorithms have been shown to have disparities in accuracy between black and white defendants [2, 7], which may lead to the arrest of an innocent person [26]. The *equalized odds* [24] fairness criterion requires a decision maker to have the same false negative rates, as well as the same false positive rates, for each sensitive group. *Equality of opportunity* requires only one of those two equalities to hold, whichever is considered more ethically important (e.g., in law enforcement, we may ensure that innocent individuals have similar treatment). Other notions of fairness include *predictive parity* [14], which matches classifier precision across groups, or the minimization of *mutual information* [31] between predictions and sensitive attributes. Except in degenerate cases, some pairs of criteria cannot be achieved simultaneously [3], meaning a domain expert must make a moral decision about which fairness constraint they aim to satisfy.

We formalize the problem of fair classification as learning the highest-accuracy model that satisfies a constraint on some fairness criterion (either demographic parity or equality of opportunity). This constrained optimization problem is mathematically difficult to solve, which has lead to the creation of a variety of fair classification algorithms. Prior work has mostly limited experiments to small tabular data sets like COMPAS [2], obscuring the fact that many previous methods are either heuristic, or have guarantees that are limited to linear classifiers with convex losses.

Our first major contribution is an optimization framework for comparing fairness relaxations and their fairness-accuracy tradeoffs. We are the first to show (theoretically and empirically) what conditions

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

may lead relaxation methods to fail or converge to degenerate solutions. We further show that this problem grows worse for non-convex classifiers like deep neural networks. To address this weakness, we promote three new fairness surrogates with promising theoretical properties: the *logistic difference*, *logistic sum*, and *sigmoid* relaxations. These surrogates are more efficient to optimize and robust to outliers than prior work. The logistic sum and sigmoid bounds have provable fairness guarantees; the logistic difference has weaker guarantees, but often has excellent empirical performance.

We also provide some of the first empirical results on fair relaxations for large-scale image and text classification. To enable these experiments, we show how our fairness surrogates may be optimized via stochastic gradient descent, and applied to data where the sensitive attribute is not always observed. Our methods are comparable to the state-of-the-art on small tabular datasets, and show clear improvements for deep learning of flexible (but fair) image and text classifiers.

2 Background: Fair Classification

Let $\mathcal{D} = \{(\mathbf{x}, s, y)\}^N$ be a set of N training examples with predictor attributes (features) $\mathbf{x} \in \mathbb{R}^p$, target class label $y \in \{0, 1\}$, and sensitive attribute $s \in \{a, b\}$ encoding membership in protected groups. Let $f_\theta(\mathbf{x})$ be a parameterized classifier, such as logistic regression or a neural network, that predicts $y = 1$ when $f_\theta(\mathbf{x}) > 0$. We define the *difference in demographic parity* (DDP) to equal

$$\text{DDP}(\theta) = P(f_\theta(X) > 0 \mid S = a) - P(f_\theta(X) > 0 \mid S = b). \tag{1}$$

If $\text{DDP}(\theta) \gg 0$, group a receives an unfair share of positive outcomes; if $\text{DDP}(\theta) \ll 0$, group b is favored. We seek models where $\text{DDP}(\theta) \approx 0$. The *difference in equality of opportunity* (DEO) is

$$\text{DEO}(\theta) = P(f_\theta(X) > 0 \mid S = a, Y = 0) - P(f_\theta(X) > 0 \mid S = b, Y = 0), \tag{2}$$

where fair classifiers have $\text{DEO}(\theta) \approx 0$. If false negative rates are of more concern than false positives, $\text{DEO}(\theta)$ may be modified to condition on $Y = 1$ instead of $Y = 0$. We seek the best classifier whose empirical fairness disparity $\Delta(\theta)$ has magnitude smaller than some tolerance $\epsilon > 0$:

$$\min_\theta \quad \frac{1}{N} \sum_{(\mathbf{x},s,y)\in\mathcal{D}} \mathcal{L}(f_\theta(\mathbf{x}), y), \text{ such that } |\Delta(\theta)| \leq \epsilon. \tag{3}$$

Our experiments focus on the case where $\mathcal{L}(.)$ is the negative log-likelihood (or cross-entropy) loss,

$$\mathcal{L}(f_\theta(\mathbf{x}), y) = -y \log(\sigma(f_\theta(\mathbf{x}))) - (1 - y) \log(\sigma(-f_\theta(\mathbf{x}))),$$

with $\sigma(z) = 1/(1 + e^{-z})$. This model family includes logistic regression and many neural networks. We set $\Delta(\theta) = \widehat{\text{DDP}}(\theta)$ for demographic parity, $\Delta(\theta) = \widehat{\text{DEO}}(\theta)$ for equality of opportunity, where

$$\widehat{\text{DDP}}(\theta) = \frac{1}{N_a} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=a}} 1_{\{f_\theta(\mathbf{x})>0\}} - \frac{1}{N_b} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=b}} 1_{\{f_\theta(\mathbf{x})>0\}}, \tag{4}$$

$$\widehat{\text{DEO}}(\theta) = \frac{1}{N_{a0}} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=a, y=0}} 1_{\{f_\theta(\mathbf{x})>0\}} - \frac{1}{N_{b0}} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=b, y=0}} 1_{\{f_\theta(\mathbf{x})>0\}}. \tag{5}$$

There are N_s data points in group s, and N_{s0} points in group s with label $y = 0$. To achieve equalized odds, we enforce two constraints as in Eq. (5), one for false positives and one for false negatives.

In addition to selecting the fairness metric, a user must specify an unfairness tolerance $\epsilon > 0$. Although we might wish to bring unfairness to exactly zero, this may not be possible except in trivial cases [20, 46], or might come at a huge accuracy cost [24] with no guarantee of generalized fairness on test data [29]. Moreover, positive tolerances often have legal justifications: adverse impact hiring laws dictate that the ratio of hiring rates for two sensitive groups should be at least 4/5 [19].

2.1 Surrogates for Tractable Optimization

The constrained optimization problem of Eq. (3) is typically intractable due to the discontinuous constraints $\Delta(\theta)$, which make gradient-based algorithms for constrained optimization inapplicable or ineffective. We overcome this intractability by replacing constraints with smooth surrogates.

Linear relaxations. The widely-used family of linear relaxations [49] seeks to improve demographic parity by replacing the intractable difference in acceptance rates with the covariance between the sensitive attribute and model output. Letting $z = 1$ if $s = a$ and $z = 0$ if $s = b$, we then have

$$\text{Cov}(z, f_\theta(\mathbf{x})) \approx \frac{1}{N} \sum_{i=1}^{N} (z_i - \bar{z}) f_\theta(\mathbf{x}_i) \propto \mathbb{E}[f_\theta(\mathbf{x}) | Z = 1] - \mathbb{E}[f_\theta(\mathbf{x}) | Z = 0]. \tag{6}$$

Because the sensitive attribute is binary, this covariance equals a difference in expectations. Bringing this *first moment relaxation* [46] to zero is equivalent to forcing the output distributions for the two sensitive groups to have equal means. By modifying the expectations in Eq. (6) to also condition on $Y = 0$, others have applied equivalent relaxations (up to constants) to equality of opportunity [4, 15]. However, matching the mean response across groups is insufficient to ensure fairness: Classifiers often still differ in higher moments, causing disparities on how groups are treated [37].

Non-linear relaxations. Zafar et al. [50, 48] also propose a relaxation for equality of opportunity, the covariance between the sensitive attribute and the model response *for only misclassified points*:

$$\mathrm{Cov}(z, g(f_\theta(\mathbf{x})) \mid Y = 0) \propto \mathbb{E}[g(f_\theta(\mathbf{x})) \mid Z = 1, Y = 0] - \mathbb{E}[g(f_\theta(\mathbf{x})) \mid Z = 0, Y = 0], \quad (7)$$

where $g(f_\theta(\mathbf{x})) = \max(f_\theta(\mathbf{x}), 0)$. This constraint is non-convex, but for classifiers like SVMs with convex loss functions, convex-concave programming [42] may be used to find a local optimum. However, the gap between the relaxation and the true target persists: even when this covariance equals 0, we have no guarantees about the difference in false positive rates between the two groups.

2.2 Other Related Work

Post-Processing of Unfair Classifiers. Post-processing methods learn a model via an unconstrained loss with no regard to fairness, and then enforce fairness as a post-processing step [24] by setting group-specific decision boundaries to equalize false positive rates (or acceptance rates). However, this requires the sensitive attribute at test time. Besides being impractical in some applications, such classifiers may make different decisions for individuals with identical features $\mathbf{x}$. Such *disparate treatment* constitutes a different form of unfairness and can be illegal in some domains [17].

Wei et al. [45] avoid requiring sensitive attributes by training a separate classifier to estimate $p(S \mid X)$. Morally, this is not much different than disparate treatment; instead of treating individuals differently based on their sensitive attributes, the model treats them differently based on its estimate of their sensitive attribute. The authors also acknowledge that while their method is effective at small reductions in unfairness, it often fails to achieve large reductions. Without knowledge of true sensitive attributes, post-processing algorithms may lack the maneuverability to dramatically reduce unfairness.

Fairness by Density Estimation. Celis et al. [9] optimize classifiers with a constraint on the true fairness discrepancy of the data-generating distribution. This requires an explicit density estimate for the training data, for which they use a Gaussian Naive Bayes model that would poorly capture high-dimensional images or text. Generative modeling of such data is a challenging research problem.

Data Set Stability. Huang & Vishnoi [28] examine the stability of fair classifiers under data perturbations, and propose a regularizer that can be added to other fairness algorithms, including ours. We instead analyze and improve fairness surrogates for complex, non-convex classifiers.

Fair Mixup. Instead of penalizing a surrogate to the fairness discrepancy, Chuang & Mroueh [10] penalize the gradient of $f_\theta(X)$ along interpolated paths between data with opposite attributes. This approach is less interpretable, and does not guarantee demographic parity nor equality of opportunity.

Fairness in Computer Vision. Large-scale image data sets are often unannotated and thus lack guarantees that they are balanced along sensitive attributes such as race or gender [22]. Dataset balance has become a focus of conversation in fair computer vision [26], but it is insufficient to guarantee fairness [44]. Other work in fair computer vision has explored issues of bias amplification [25, 52], counterfactual fairness [12], and transfer learning [16, 32]. Whereas relaxation methods explore fairness-accuracy trade-offs, many vision papers use fairness as a vehicle to boost accuracy [41, 35, 43]. By improving classifier performance on under-represented sub-populations, overall accuracy can be increased, but groups may still receive disparate impact from the model's decisions.

Fairness in Natural Language Processing. Much of the previous work on fair language models has focused on domain-specific examples of bias, such as word embeddings [5, 54, 21, 6, 55] or coreference resolution [40, 53]. In contrast, we are not concerned with latent bias in the language model's word representation; we want to ensure that predictive models do not discriminate against authors based on their identity. It has been shown that gender expression can be predicted from writing [34], so models could discriminate against authors of one gender if fairness is not mandated.

3 A Lagrangian Framework for Relaxations

The absolute value in Eq. (3) induces two linear fairness constraints, a lower bound and an upper bound, at most one of which can be active. For the optimal unconstrained classifier parameters $\theta^* = \arg\min \frac{1}{N} \sum_{(\mathbf{x},s,y)\in\mathcal{D}} \mathcal{L}(f_\theta(\mathbf{x}), y)$, we assume (without loss of generality) that $\Delta(\theta^*) > \epsilon$. If

$\Delta(\theta^*) < -\epsilon$, we swap the sensitive groups a, b to negate the discrepancy and obtain $\Delta(\theta^*) > \epsilon$. If $|\Delta(\theta^*)| \leq \epsilon$, the unconstrained classifier is already sufficiently fair and optimally solves Eq. (3).

3.1 Fair Classification via Smooth Penalties

Letting $g(r)$ be a continuous and monotonically increasing surrogate for the indicator functions arising in exact fairness constraints, we define the following tractable optimization problem:

$$\min_{\theta} \frac{1}{N} \sum_{(\mathbf{x},s,y)\in\mathcal{D}} \mathcal{L}(f_\theta(\mathbf{x}), y), \text{ s.t. } \Delta_g(\theta) \leq \tau. \tag{8}$$

For demographic parity, we set $\Delta_g(\theta) = \widehat{\mathrm{DDP}}_g(\theta)$, where

$$\widehat{\mathrm{DDP}}_g(\theta) = \frac{1}{N_a} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=a}} g(f_\theta(\mathbf{x})) - \frac{1}{N_b} \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D} \\ s=b}} g(f_\theta(\mathbf{x})). \tag{9}$$

For equality of opportunity, we similarly define $\widehat{\mathrm{DEO}}_g(\theta)$ by replacing the indicators in Eq. (5) with surrogates $g(r)$. When $g(r) = r$, for DDP we recover [50], and for DEO we recover [4, 15]. In this case, both $g(r)$ and $-g(r)$ are linear and convex, so when the loss is also convex we have a convex optimization problem. When $g(r) = \max(0, r)$, for DEO we recover the relaxation in [48, 50]. For non-linear $g(r)$, Eq. (8) is non-convex even for convex losses due to the $-g(r)$ term in $\Delta_g(\theta)$.

Unlike some prior work, we employ the KKT conditions [39] to optimize Eq. (8) via its Lagrangian:

$$\min_{\theta} \quad \mathcal{L}_{g,\lambda}(\theta), \quad \text{where } \mathcal{L}_{g,\lambda}(\theta) = \frac{1}{N} \sum_{(\mathbf{x},s,y)\in\mathcal{D}} \mathcal{L}(f_\theta(\mathbf{x}), y) + \lambda \Delta_g(\theta). \tag{10}$$

For every relaxed-fairness tolerance τ, there exists a $\lambda \geq 0$ such that the constrained optimum of (8) is a minimizer of (10). Furthermore, second-order sufficient conditions [39, Chap. 12.5] guarantee that every local minima $\hat{\theta}$ of (10), with fairness discrepancy $\Delta_g(\hat{\theta})$, is a feasible local minima of the constrained objective (8) with $\tau = \Delta_g(\hat{\theta})$. For non-convex losses $\mathcal{L}$, gradient descent may reach a fixed point with sub-optimal likelihood. But this is a concern for any optimizer, and in experiments we find that deep neural networks are flexible enough to achieve effective fairness-accuracy tradeoffs.

We can rewrite the fairness-penalized objective of Eq. (10) as a sum over data examples:

$$\mathcal{L}_{g,\lambda}(\theta) = \frac{1}{N} \sum_{(\mathbf{x},s,y)\in\mathcal{D}} \Big(\mathcal{L}(f_\theta(\mathbf{x}), y) + \lambda R_g(\mathbf{x}, y, s) \Big). \tag{11}$$

For demographic parity and equality of opportunity, the fairness regularizers are given by

$$R_g(\mathbf{x}, y, s) = \begin{cases} \frac{g(f_\theta(\mathbf{x}))}{p_a} & \text{if } s = a, \\ -\frac{g(f_\theta(\mathbf{x}))}{p_b} & \text{if } s = b. \end{cases} \ (DP). \quad R_g(\mathbf{x}, y, s) = \begin{cases} \frac{g(f_\theta(\mathbf{x}))}{p_{a0}} & \text{if } s = a, y = 0, \\ -\frac{g(f_\theta(\mathbf{x}))}{p_{b0}} & \text{if } s = b, y = 0, \\ 0 & \text{otherwise.} \end{cases} \ (EO)$$

Here, $p_s = N_s/N$ is the proportion of data in sensitive group s, and $p_{s0} = N_{s0}/N$.

Our framework differs from prior work that introduced these surrogates in that we do not seek $\Delta_g = 0$. We know Δ_g will decrease as we increase λ, but it might be necessary to bring $\Delta_g < 0$ to achieve $\Delta \leq \epsilon$ for relaxations (like the linear) where Δ_g does not rigorously bound Δ. In practice, we need not observe Δ_g: we can increase λ and track changes in Δ until we find a λ^* with $\Delta \leq \epsilon$, if it exists.

3.2 Degeneracy of Fairness Relaxations

When we choose an unbounded surrogate, we may observe pathological behavior where Δ_g becomes negative and large enough in magnitude for the optimizer to ignore predictive loss. For example in Fig. 1, the linear relaxation performs reasonably well on an initial synthetic dataset. We then add an outlier with $(x_1, x_2) = (10^3, 10^3)$ and $y = 1$. The unconstrained model sees no change, because the new point is already classified correctly. If the outlier has $s = b$, then the relaxed difference becomes much more negative. This outlier causes the classifier to degenerate for even small λ values, focusing on improving the fairness objective at great detriment to the predictive objective.

The optimization framework of Sec. 3.1 may fail in two ways: the surrogate problem (8) could be infeasible if the model is not expressive enough, or a poor choice of surrogate $g(r)$ may cause there to be no value of τ for which the solution to (8) solves the true objective (3). We now provide theoretical conditions under which fairness surrogates fail to achieve fairness. Proofs are in the supplement.

Consider a model with (possibly learned) data representation $\Phi(\mathbf{x})$, output $f_\theta(\mathbf{x}) = \mathbf{w}^T\Phi(x)$, and Bernoulli negative log-likelihood loss $\mathcal{L}(\mathbf{w}^L\Phi(\mathbf{x}), y)$. This family includes logistic regression, where $\Phi(\mathbf{x}) = \mathbf{x}$ is fixed and the model parameters $\theta = \mathbf{w}$. It also includes most neural network architectures for binary classification, where $\Phi(\mathbf{x})$ is parameterized and learned at training time so that $\theta = \mathbf{w} \cup \theta_\Phi$.

We minimize the Lagrangian $\mathcal{L}_{g,\lambda}$ of Eq. (11) with predictors $\Phi(\mathbf{x})$ and a fairness surrogate $g(r)$ that is continuous and monotonically increasing, where limits $\delta_g^+ = \lim_{r\to\infty} g'(r)$ and $\delta_g^- = \lim_{r\to-\infty} g'(r)$ exist. Let μ_{jsy} be the mean of $\Phi(\mathbf{x})_j$, the jth entry of $\Phi(\mathbf{x})$, for the N_{sy} data points with sensitive group s and label y. Also let $\mu_{js.}$ be the mean of $\Phi(\mathbf{x})_j$ for data with group s and any label, and $\mu_{j.y}$ be the mean of $\Phi(\mathbf{x})_j$ for data with label y and any group.

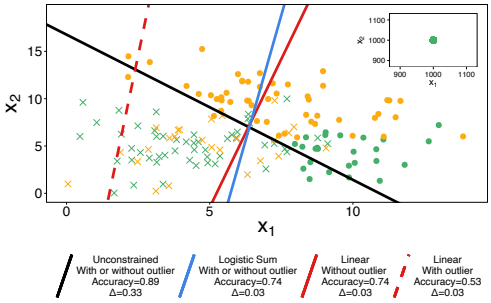

Figure 1: Synthetic dataset with large latent fairness discrepancy. Circles have true label $y = 1$ and crosses have true label $y = 0$. Color encodes sensitive attribute, with green indicating group b. For both linear and logistic sum (see Sec. 4.2) fairness formulations, we trained a logistic regression model, searching across a dense grid of λ values to find one for which $\widehat{\text{DDP}} \leq 0.03$. Adding a single outlier at $(10^3, 10^3)$ with $y = 1, s = b$ does not change the unconstrained or logistic sum decision boundaries, but linear deteriorates dramatically.

Theorem 1 (Degeneracy for DP)**.** *Consider a feature j such that $\Phi(\boldsymbol{x})_j \geq 0$ for all $\boldsymbol{x} \in \mathcal{D}$.*
(1) If $\delta_g^- \neq 0$ and $\mu_{ja.} > \mu_{jb.}$, $\lambda > \lambda_j^* = \frac{\mu_{j.1}p_1}{\delta_g^-(\mu_{ja.}-\mu_{jb.})}$, then $\lim_{\boldsymbol{w}_j\to-\infty} \mathcal{L}_{g,\lambda}(\theta) = -\infty$.
(2) If $\delta_g^+ \neq 0$ and $\mu_{ja.} < \mu_{jb.}$, $\lambda > \lambda_j^* = \frac{\mu_{j.0}p_0}{\delta_g^+(\mu_{jb.}-\mu_{ja.})}$, then $\lim_{\boldsymbol{w}_j\to+\infty} \mathcal{L}_{g,\lambda}(\theta) = -\infty$.
If either (1) or (2) hold and $g(r)$ is linear, then for this fixed $\Phi(\boldsymbol{x})$, $\mathcal{L}_{g,\lambda}$ has no stationary points.

Theorem 2 (Degeneracy for EOO)**.** *Consider a feature j such that $\Phi(\boldsymbol{x})_j \geq 0$ for all $\boldsymbol{x} \in \mathcal{D}$.*
(1) If $\delta_g^- \neq 0$ and $\mu_{ja0} > \mu_{jb0}$, $\lambda > \lambda_j^* = \frac{\mu_{j.1}p_1}{\delta_g^-(\mu_{ja0}-\mu_{jb0})}$, then $\lim_{\boldsymbol{w}_j\to-\infty} \mathcal{L}_{g,\lambda}(\theta) = -\infty$.
(2) If $\delta_g^+ \neq 0$ and $\mu_{ja0} < \mu_{jb0}$, $\lambda > \lambda_j^* = \frac{\mu_{j.0}p_0}{\delta_g^+(\mu_{jb0}-\mu_{ja0})}$, then $\lim_{\boldsymbol{w}_j\to+\infty} \mathcal{L}_{g,\lambda}(\theta) = -\infty$.
If either (1) or (2) hold and $g(r)$ is linear, then for this fixed $\Phi(\boldsymbol{x})$, $\mathcal{L}_{g,\lambda}$ has no stationary points.

For constraint multipliers $\lambda > \lambda_j^*$, the Lagrangian is unbounded below and can be reduced arbitrarily by changing a single parameter $\mathbf{w}_j$. Each attribute j satisfying (1) or (2) induces a critical λ_j^*, so the critical value λ^* for the full model is the minimal λ_j^*. Because the predictive loss is bounded below, divergence must be caused by Δ_g approaching $-\infty$. Note this divergence is *not* a failure of the KKT conditions: for every finite τ, the corresponding constraint multiplier λ is bounded in $0 \leq \lambda \leq \lambda^*$.

While we may solve the surrogate problem (11) for any finite τ via some $\lambda \leq \lambda^*$, for general $g(r)$ there is no guarantee we satisfy our true fairness constraints (3). This is illustrated in Fig. 2, which highlights gaps between the surrogate and true rate differences. Degeneracy occurs for the linear relaxation $g(r) = r$, where $\delta_g^- = \delta_g^+ = 1$, of DDP on the Adult data set. For the ReLU relaxation of DEO, which takes $g(r) = \max(0, r)$ so that $\delta_g^- = 0$ but $\delta_g^+ = 1$, degeneracy also occurs on the COMPAS data set; this is possible because $-g(r)$ and Δ_g are unbounded below. Even in less extreme cases like Fig. 1 where we recover a solution to the true target, if we need to set $\tau \ll 0$ and thus $\lambda \approx \lambda^*$, the surrogate's magnitude may cause the learned model to mostly ignore predictive loss.

While L_2 regularization will prevent $\mathbf{w}_j$ from fully diverging, it may still grow large enough in magnitude to substantially hurt performance. For example, when constraining DDP on the Adult dataset, the feature `occupation_Protective-serv` induces $\lambda^* \approx 0.066$. Regularizing above that threshold results in huge losses to mean log-likelihood as $\mathbf{w}_j$ diverges (see Fig. 3). This degeneracy is not clearly visible from accuracy plots because only 4% of the data has `occupation_Protective-serv`=1.

Degeneracy becomes more prevalent for deep neural networks, which can learn one feature $\Phi(\mathbf{x})_j$ such that λ^* is minimized and very negative training loss is achievable, while doing little to improve the true fairness. The assumption that $\Phi(\mathbf{x})_j \geq 0$ is satisfied by many common attributes including discrete indicators and ReLU non-linearities in neural networks. As we prove in the supplement, notationally complicated generalizations of Theorems 1-2 exist that characterize arbitrary $\Phi(\mathbf{x})$.

3.3 Related Work: SearchFair and Fair Reductions

The SearchFair algorithm [37] was the first to reduce relaxations $\Delta_g < 0$ to enforce fairness. They prove that when the fairness-penalized loss is *convex*, $\text{DDP}(\theta(\lambda))$ for the (unknown) true data

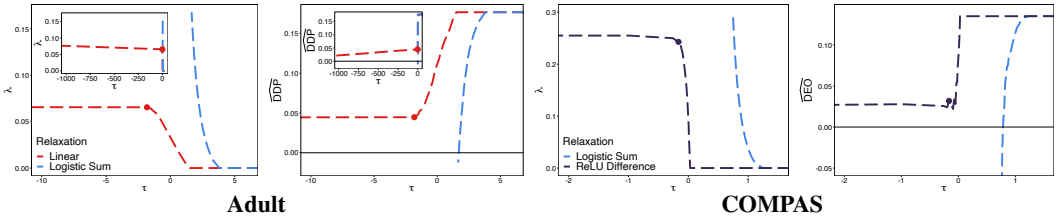

Adult **COMPAS**

Figure 2: Relationship between surrogate threshold τ and multiplier λ (first & third plots) or non-relaxed unfairness (second & fourth plots). When τ is large, the constraint is inactive. Highlighted points show models with critical $\lambda = \lambda^*$ derived from Theorem 1 (left, for DDP) or Theorem 2 (right, for DEO). Our logistic sum surrogate provably never degenerates (see Sec. 4.2). We avoid numerical instabilities via very weak L_2-regularization of weights $\mathbf{w}$, with coefficient 2^{-20}, so that models with $\lambda > \lambda^*$ have bounded loss.

distribution is continuous with respect to the Lagrange multiplier λ. Assuming there exists λ^- such that $\text{DDP}(\theta(\lambda^-)) < 0$, and λ^+ such that $\text{DDP}(\theta(\lambda^+)) > 0$, there exists a root $\lambda^- < \lambda^\circ < \lambda^+$ with $\text{DDP}(\theta(\lambda^\circ)) = 0$. SearchFair then uses binary search to find a λ° such that $\widehat{\text{DDP}}(\theta(\lambda^\circ)) \approx 0$.

A limiting assumption of SearchFair is that λ^- and λ^+ exist for any fairness surrogate. One of our key contributions is to demonstrate theoretically (via Theorems 1-2) and empirically (via Figs. 2-3) that this is not true: binary search cannot be used to find the root of a function with no root. Note that Lohaus et al. [37] explicitly consider both relaxations whose degeneracies are demonstrated in Fig. 2.

Continuity of $\text{DDP}(\theta(\lambda))$, as required to guarantee the existence of the λ° sought by SearchFair, is *not* guaranteed for non-convex losses like (deep) neural networks. We do not use SearchFair's binary search for this reason, and because (as motivated in Sec. 2) we instead seek $|\text{DDP}(\theta)| < \epsilon$ for $\epsilon > 0$.

The Lagrangian dual [1, 38] provides an iterative approach to choosing Lagrange multipliers for multiple fairness constraints. Agarwal et al. [1] "reduce" each iteration to a cost-sensitive classification problem. While classification with heterogeneous costs is widely studied, for some settings of the Lagrange multipliers, the induced "costs" will become negative. Our results show that this may cause typically convex relaxations to become non-convex, and thus cost-sensitive learning may fail.

4 Logistic Fairness Relaxations

We analyze relaxations based on the *log-sigmoid* $g(r) = -\log \sigma(-r)$ and *sigmoid* $g(r) = \sigma(r)$.

4.1 Logistic Differences for Improved Fairness

The log-sigmoid surrogate $g(r) = -\log \sigma(-r)$ has the same form as the log-likelihood: $g(f_\theta(\mathbf{x})) = \mathcal{L}(f_\theta(\mathbf{x}), 0)$. When using that loss function and the DEO criterion, the two terms in Eq. (11) combine nicely, and the relaxed objective can be interpreted as a simple reweighting of points: positive examples have weight 1, negative examples of group a have weight $1 + \frac{\lambda}{p_{a0}}$, and negative examples of group b have weight $1 - \frac{\lambda}{p_{b0}}$. Thus for logistic regression with $\lambda \leq p_{b0}$, the Lagrangian formulation in Eq. (10) is *convex* even though the corresponding constraints in Eq. (8) are non-convex.

This analysis aligns with our earlier theorems. Because the log-sigmoid surrogate has $\delta_g^- = 0$, condition (1) in Theorem 2 will never be met. Condition (2) may only be satisfied when λ is large.

Theorem 3. *For EOO and any $\Phi(\boldsymbol{x})$ such that $\Phi(\boldsymbol{x})_j \geq 0$ for all $\boldsymbol{x} \in \mathcal{D}$, $\lambda_j^* = \frac{\mu_{j.0} p_0}{\delta_g^+(\mu_{jb0} - \mu_{ja0})} \geq p_{b0}$.*

Thus for small constraint multipliers $\lambda \leq p_{b0}$, the objective function is convex and bounded. For $\lambda > \lambda^*$, degeneracy as in Fig. 2 is possible, but this is not observed in any of our experiments. A generalization Theorem 3, that allows arbitrary $\Phi(\mathbf{x})$, is proven in the supplement. Note that similar simplifications of the DP criterion are not possible, and no comparable bound exists in that case.

Related Work: Data Re-weighting. Re-weighting of data points is a classic way to correct deficiencies of generative models [27, 23]. For fair classification, Kamiran & Calders [30] pre-process data to encourage demographic parity by assigning weight $w(s, y) = \frac{p(s)p(y)}{p(s \wedge y)}$ to points from group s with label y. Probabilities are estimated from the training data so that y and s are independent under the weighted empirical distribution. However, this lack of dependence does not guarantee demographic parity; relative to these weights, classifiers may either under- or over-balance sensitive groups.

Our logistic difference relaxation instead implies that *some* amount of weight increase for the lower-accuracy group, and weight decrease for the higher-accuracy group, will lead to approximate fairness. The search for λ must be integrated with training; it cannot be determined from raw data statistics.

4.2 Logistic Sums with Fairness Guarantees

The primary weakness of Eq. (8) is that we lack meaningful guarantees about the relationship between Δ_g and Δ. To address this problem, we rewrite the empirical DDP as

$$\widehat{\mathrm{DDP}}(\theta) = \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D}\\ s=a}} \frac{1_{\{f_\theta(\mathbf{x})>0\}}}{N_a} + \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D}\\ s=b}} \frac{1_{\{-f_\theta(\mathbf{x})>0\}}}{N_b} - 1. \tag{12}$$

Replacing indicators with convex upper bounds $g(r)$, the corresponding Lagrangian becomes:

$$\min_\theta \quad \frac{1}{N}\sum_{(\mathbf{x},s,y)\in\mathcal{D}} \mathcal{L}(f_\theta(\mathbf{x}), y) \text{ such that } \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D}\\ s=a}} \frac{g(f_\theta(\mathbf{x}))}{N_a} + \sum_{\substack{(\mathbf{x},s,y)\in\mathcal{D}\\ s=b}} \frac{g(-f_\theta(\mathbf{x}))}{N_b} \le \tau + 1. \tag{13}$$

We need $g(r) \ge 1_{\{r>0\}}$ for all $r \in \mathbb{R}$, as satisfied by the hinge function $g(r) = \max(0, 1+r)$ and the smooth log-sigmoid $g(r) = -\log\sigma(-r)$. Because the constraint in Eq. (13) is a rigorous upper bound on $\widehat{\mathrm{DDP}}(\theta)$, if we set $\tau \le \epsilon$, we are guaranteed that $\widehat{\mathrm{DDP}}(\theta) \le \Delta_g(\theta) \le \tau \le \epsilon$. This bound prevents the pathological behavior we demonstrated for other relaxations, where Δ_g can be brought arbitrarily negative without sufficiently reducing $\widehat{\mathrm{DDP}}$. An identical derivation can be applied to the equality of opportunity criterion, resulting in a surrogate involving only the points with $y = 0$.

In practice, the Lagrangian lets us skip the need to select τ entirely. For expressive models where (13) is feasible, there exists some surrogate threshold τ, and therefore λ, which reduces fairness discrepancy below ϵ. A practitioner must simply search λ to find a value satisfying their tolerance ϵ.

We can rewrite Eq. (13) as in Eq. (11), but with a new fairness regularizer R_g^+:

$$R_g^+(\mathbf{x}, y, s) = \frac{g(f_\theta(\mathbf{x}))}{p_a} \text{ if } s = a, \qquad R_g^+(\mathbf{x}, y, s) = \frac{g(-f_\theta(\mathbf{x}))}{p_b} \text{ if } s = b. \tag{14}$$

In addition to providing an upper bound on the true fairness discrepancy, this logistic sum relaxation has an attractive interpretation. $R_g^+(\mathbf{x}, y, s)$ is the negative log-likelihood loss of predicting s with $f_\theta(\mathbf{x})$, with added weights p_a, p_b. Thus we can interpret our Lagrangian objective as trying to predict *both* y and s with our model output, and summing the negative log-likelihood loss in both predictions.

Related Work: Lower-Upper Bounds. Wu et al. [47] used a similar transform to replace the difference in rates with a sum. They then add a second constraint to Eq. (13), using a concave surrogate $h(r) \le 1_{\{r>0\}}$ to lower bound $\widehat{\mathrm{DDP}}(\theta)$. Their experiments take $g(r) = \max(r+1, 0)$ and $h(r) = -g(-r) + 1$. If this optimization problem is solved, we know the resulting model has $-\gamma \le \widehat{\mathrm{DDP}} \le \tau$. However, choosing γ and τ is difficult: hinge functions are loose bounds on the 0-1 loss, so typically $-\gamma \ll \widehat{\mathrm{DDP}} \ll \tau$. Lohaus et al. [37] argue that for SVMs and meaningful values of γ, τ, the constraints induce disjoint sets and no solution exists. Compared to this method, we identify the active fairness constraint and only require one upper bound. We also avoid the need to explicitly specify a threshold τ (let alone γ) by directly varying the constraint multiplier λ. This simplification is critical in enabling tractable application to more complex classifiers like deep neural networks.

4.3 Sigmoid Sums for Tight Bounds

We examine a third and final relaxation, the non-convex sigmoid function $g(r) = \sigma(r)$. Because $-\sigma(r) = \sigma(-r) - 1$, the difference and sum constructions lead to the same constraint. When viewed as a difference, we see $\delta_g^- = \delta_g^+ = 0$, preventing degeneracy. When viewed as a sum, we see it as a scaled upper bound on the indicator function, $2\sigma(r) \ge 1_{\{r>0\}}$. Thus like logistic sum, there provably exists some τ such that the solution to the surrogate achieves the true target fairness.

While sigmoids are widely used to relax indicator functions, there seems to be little work using them as surrogates for DDP or DEO. Kamishima et al. [31] use sigmoids to parameterize a different, mutual information fairness criterion. The cost of this approach is that $g(r)$ and $R_g(\mathbf{x}, y, s) = R_g^+(\mathbf{x}, y, s)$ are non-convex, introducing local minima even for convex predictive losses. For deep neural networks, the vanishing gradients of the sigmoid function may lead to greater optimization difficulties.

5 Experiments and Data Sets

We compare relaxations by training fairness-regularized logistic regression on tabular data, and fairness-regularized deep neural networks on image and text data. To ensure equal treatment, for all models the sensitive attribute is not used as a predictive attribute. For tabular data, we use the standard

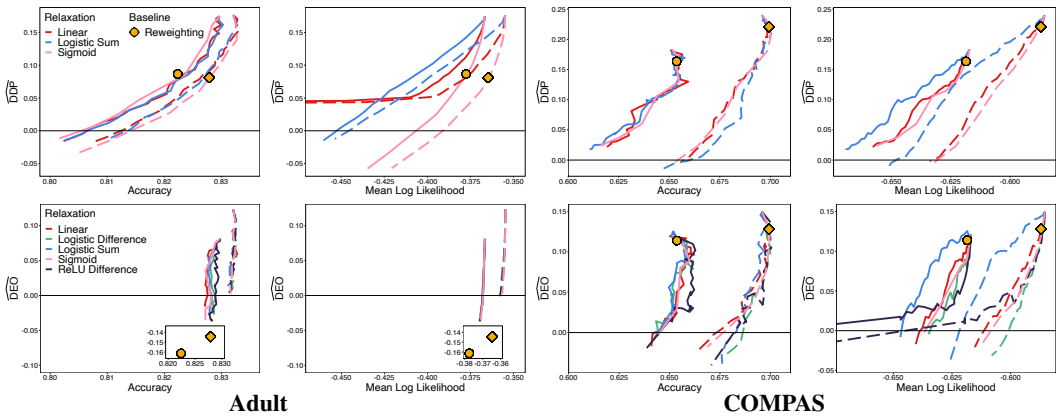

Figure 3: Fair logistic regression results on the Adult and COMPAS data sets. Each relaxation is run over a dense grid of λ values until $\Delta(\theta) < 0$ is achieved on training data (or for demographic parity on Adult, until degeneracy occurs as in Fig. 2). Dashed lines show results on training data, solid lines on test data. For the reweighting baseline [30], the diamond shows training results and the circle test results.

fairness benchmarks of the Adult census [33] and COMPAS recidivism data sets [2]. Methods for hyperparameter selection, data pre-processing, and optimization are detailed in the supplement.

Relaxations. For demographic parity, we compare the *linear* relaxation $g(r) = r$, our *logistic sum* upper bound $g(r) = -\log \sigma(-r)$ with regularizer R_g^+, and our *sigmoid* upper bound $g(r) = \sigma(r)$. For equality of opportunity, we compare these three relaxations to the *logistic difference* formulation with $g(r) = -\log \sigma(-r)$ and difference regularizer R_g. We also consider the non-linear *ReLU difference* relaxation of [50, 48], where $g(r) = \max(0, r)$ with difference regularizer R_g. As a final baseline, we test the reweighting technique of Kamiran & Calders [30] as outlined in Sec. 4.1.

5.1 Fair Deep Learning with Image and Text Data

To our knowledge, we are the first to explore the performance of general-purpose fairness relaxations on deep neural networks for text and image classification. For the computer vision tasks, we use a randomly-initialized wide residual network architecture (WRN-50-2), which in 2016 broke established records on many visual recognition benchmarks including CIFAR-100 [51]. For the text classification task, we use Google's transformer-based BERT, a state-of-the-art language model [13].

CelebA. The CelebFaces Attributes data set [36] contains images of celebrities' faces, with 40 annotated binary attributes. As in other work, we use *smiling* as the target attribute and *sex* as the sensitive attribute. We discard the remaining binary attributes and use the image data for prediction.

Faces of the World. The Faces of the World (FOTW) data set [18] comes from the 2016 Looking at People CVPR Challenge. We use *smiling* as the target attribute and *gender* as the sensitive attribute. To demonstrate the performance cost of Lohaus et al. [37]'s convexity requirement, we compare our fairness-regularized deep neural network to their kernel-based approach.

Yelp Text Data. We use a subset of the Yelp review data set [11], which includes the text of reviews along with metadata including the reviewer's first name and the review's star rating. We perform the sentiment analysis task of predicting if a reviewer gives ≥ 4 or < 4 stars from the text of the review. We use *gender* as the sensitive attribute, which is inferred from reviewers' first names.

Sensitive Attribute Unknown. The Yelp and FOTW data sets include examples for which the sensitive attribute is unknown or cannot be inferred. These data therefore do not contribute towards the fairness constraint, but do affect the predictive loss and improve classifier accuracy.

Minibatch loss. Because we have written the objective of Eq. (11) as a sum over data points, we can create an unbiased estimate of it from mini-batches $\mathcal{D}^{\{1\}}, \ldots, \mathcal{D}^{\{T\}}$, where $\mathcal{D}^{\{t\}} = \{(\mathbf{x}, s, y)\}^M$ has M points drawn randomly from the full data set. By estimating via mini-batch the point-wise sum in Eq. (11), we are not evaluating the batch's surrogate discrepancy Δ_g. Instead, by using the N_{sy} and p_{sy} statistics for the entire data set (not just the batch), we evaluate the batch's contribution to the full data unfairness Δ_g. This is possible because demographic parity and equality of opportunity decompose as sums over data examples. This scalable stochastic approximation cannot be performed for other fairness criteria like predictive parity [14] or mutual information [31].

5.2 Analysis of Results

Linear Degeneracy. On CelebA, the linear relaxation degenerates almost immediately, achieving mean log-likelihoods 1000 times worse than competitors. This is predicted by our claim that neural

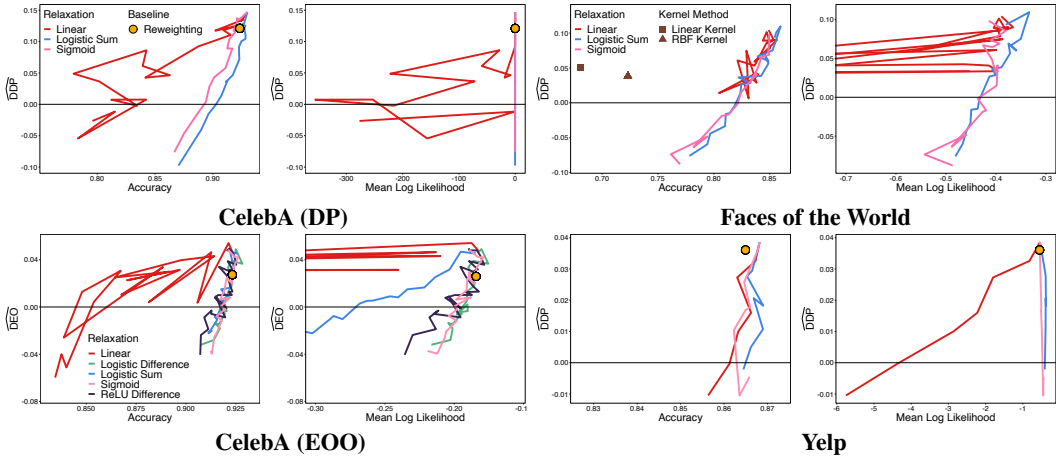

Figure 4: Fairness-regularized deep neural network results on computer vision and natural language processing data sets. Each simulation uses different random seeds for drawing minibatches, revealing instability in some relaxations. All results are shown on test data. The kernel-based model optimizes a convex loss function, using the SearchFair method of binary search to find a model with $\Delta(\theta) = 0$ on training data.

networks may design $\Phi(\mathbf{x})$ to induce cause degeneracy as in Theorems 1-2. Degeneracy is less severe for FOTW and Yelp, which are initialized from pre-trained models. We conjecture that pre-training safeguards against the model learning to optimize fairness without regard for prediction.

Smoothness Matters. The logistic difference and ReLU difference surrogates have the same asymptotic limits, and therefore the same potential for degeneracy with large constraint multipliers. However, we only observe degeneracy for the ReLU difference formulation. Due to a mismatch between the smooth predictive loss and the ReLU surrogate, the Lagrangian may be non-convex for any $\lambda > 0$, in contrast with Theorem 3. In the supplement, additional experiments show that optimization convergence for non-smooth surrogates is about 10 times slower than for smooth log-logistics.

The Cost of Guarantees. The logistic upper bound formulation is guaranteed to not degenerate, but this comes at the cost of poorer likelihood. We conjecture that this is due to the surrogate punishing high-confidence model outputs. The sigmoid formulation has the same fairness guarantees, but has slightly worse accuracy for large-scale data; we conjecture this is due to additional local optima.

The Lagrangian View. The re-weighting baseline often only reduces the fairness discrepancy marginally: data set balance is insufficient to ensure fairness. This baseline can be thought of as a way to set λ before learning based on data statistics. Its poor fairness gains argue in favor of tuning multipliers via iterative training, despite the computational demands of training several models.

The Cost of Convexity. The binary search method of SearchFair is only guaranteed to find a root when the loss function is convex. In the case of image data, forcing a linear model by using a kernel method comes at a large accuracy cost. Though SearchFair finds a point with $\Delta(\theta) = 0$ on the training split, due to sampling noise this does not guarantee perfect fairness on test data.

Limitations. Our strongest theoretical guarantees are about global minima, which non-convex optimization may fail to find. Training may converge to a model with arbitrarily worse likelihood than the true minimizer of Eq. (8). However, this is a problem for all deep learning methods. Because our surrogates induce losses which have been widely studied in deep learning, we conjecture that the same architectures that enable standard deep learning will be effective for our fair extensions.

We also acknowledge that fairness constraints are not a universal solution to problems of bias in artificial intelligence. Some tasks might be inherently unethical to automate, or the very existence of the task itself might be unethical and reinforce existing power structures. Demonstrating equitable group-level performance does not solve these problems, and might obfuscate deeper problems.

6 Discussion

We analyze three new fairness surrogates: the logistic difference, logistic sum, and sigmoid relaxations. They are easier to optimize than non-smooth alternatives in the literature, and provably more stable (especially for non-convex classifiers) than linear or ReLU relaxations. Unlike some prior work, our Lagrangian framework allows easy integration with non-convex deep neural networks. Results on complex vision and NLP tasks confirm our approach is robust, scalable, and broadly applicable.

References

[1] Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. A reductions approach to fair classification. *arXiv preprint arXiv:1803.02453*, 2018.

[2] Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias. propublica, may 23, 2016, 2016.

[3] Barocas, S., Hardt, M., and Narayanan, A. *Fairness and Machine Learning*. fairmlbook.org, 2019. `http://www.fairmlbook.org`.

[4] Bechavod, Y. and Ligett, K. Penalizing unfairness in binary classification. *arXiv preprint arXiv:1707.00044*, 2017.

[5] Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pp. 4349–4357, 2016.

[6] Brunet, M.-E., Alkalay-Houlihan, C., Anderson, A., and Zemel, R. Understanding the origins of bias in word embeddings. In *International Conference on Machine Learning*, pp. 803–811. PMLR, 2019.

[7] Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pp. 77–91. PMLR, 2018.

[8] Calders, T. and Verwer, S. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.

[9] Celis, L. E., Huang, L., Keswani, V., and Vishnoi, N. K. Classification with fairness constraints: A meta-algorithm with provable guarantees. In *Proceedings of the conference on fairness, accountability, and transparency*, pp. 319–328, 2019.

[10] Chuang, C.-Y. and Mroueh, Y. Fair mixup: Fairness via interpolation. In *International Conference on Learning Representations*, 2021. URL `https://openreview.net/forum?id=DNl5s5BXeBn`.

[11] Dataset, Y., 2014. URL `https://www.yelp.com/dataset`.

[12] Denton, E., Hutchinson, B., Mitchell, M., and Gebru, T. Detecting bias with generative counterfactual face attribute augmentation. *arXiv preprint arXiv:1906.06439*, 2019.

[13] Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/v1/n19-1423. URL `https://doi.org/10.18653/v1/n19-1423`.

[14] Dieterich, W., Mendoza, C., and Brennan, T. Compas risk scales: Demonstrating accuracy equity and predictive parity. *Northpoint Inc*, 7(7.4):1, 2016.

[15] Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S., and Pontil, M. Empirical risk minimization under fairness constraints. In *Advances in Neural Information Processing Systems*, pp. 2791–2801, 2018.

[16] Dwork, C., Immorlica, N., Kalai, A. T., and Leiserson, M. Decoupled classifiers for group-fair and efficient machine learning. In *Conference on Fairness, Accountability and Transparency*, pp. 119–133, 2018.

[17] EEOC, T. Uniform guidelines on employee selection procedures, 1979.

[18] Escalera, S., Baró, X., Escalante, H. J., and Guyon, I. Chalearn looking at people: A review of events and resources. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 1594–1601. IEEE, 2017.

[19] Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 259–268, 2015.

[20] Friedler, S. A., Scheidegger, C., and Venkatasubramanian, S. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236*, 2016.

[21] Gonen, H. and Goldberg, Y. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In *Proceedings of NAACL-HLT*, pp. 609–614, 2019.

[22] Grover, A., Choi, K., Shu, R., and Ermon, S. Fair generative modeling via weak supervision. *arXiv preprint arXiv:1910.12008*, 2019.

[23] Grover, A., Song, J., Kapoor, A., Tran, K., Agarwal, A., Horvitz, E. J., and Ermon, S. Bias correction of learned generative models using likelihood-free importance weighting. In *Advances in Neural Information Processing Systems*, pp. 11058–11070, 2019.

[24] Hardt, M., Price, E., Srebro, N., et al. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pp. 3315–3323, 2016.

[25] Hendricks, L. A., Burns, K., Saenko, K., Darrell, T., and Rohrbach, A. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*, pp. 793–811. Springer, 2018.

[26] Hill, K. Wrongfully accused by an algorithm. *The New York Times, June*, 24, 2020.

[27] Horvitz, D. G. and Thompson, D. J. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.

[28] Huang, L. and Vishnoi, N. Stable and fair classification. In *International Conference on Machine Learning*, pp. 2879–2890. PMLR, 2019.

[29] Ji, D., Smyth, P., and Steyvers, M. Can i trust my fairness metric? assessing fairness with unlabeled data and bayesian inference. *arXiv preprint arXiv:2010.09851*, 2020.

[30] Kamiran, F. and Calders, T. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.

[31] Kamishima, T., Akaho, S., Asoh, H., and Sakuma, J. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 35–50. Springer, 2012.

[32] Kim, B., Kim, H., Kim, K., Kim, S., and Kim, J. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9012–9020, 2019.

[33] Kohavi, R. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *Kdd*, volume 96, pp. 202–207, 1996.

[34] Koppel, M., Argamon, S., and Shimoni, A. R. Automatically categorizing written texts by author gender. *Literary and linguistic computing*, 17(4):401–412, 2002.

[35] Li, Y. and Vasconcelos, N. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9572–9581, 2019.

[36] Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

[37] Lohaus, M., Perrot, M., and von Luxburg, U. Too relaxed to be fair. In *International Conference on Machine Learning*, 2020.

[38] Lokhande, V. S., Akash, A. K., Ravi, S. N., and Singh, V. Fairalm: Augmented lagrangian method for training fair models with little regret. In *European Conference on Computer Vision*, pp. 365–381. Springer, 2020.

[39] Nocedal, J. and Wright, S. *Numerical optimization*. Springer Science & Business Media, 2006.

[40] Rudinger, R., Naradowsky, J., Leonard, B., and Van Durme, B. Gender bias in coreference resolution. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pp. 8–14, 2018.

[41] Ryu, H. J., Mitchell, M., and Adam, H. Improving smiling detection with race and gender diversity. *arXiv preprint arXiv:1712.00193*, 1(2):7, 2017.

[42] Shen, X., Diamond, S., Gu, Y., and Boyd, S. Disciplined convex-concave programming. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 1009–1014. IEEE, 2016.

[43] Wang, M., Deng, W., Hu, J., Tao, X., and Huang, Y. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 692–702, 2019.

[44] Wang, T., Zhao, J., Yatskar, M., Chang, K.-W., and Ordonez, V. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5310–5319, 2019.

[45] Wei, D., Ramamurthy, K. N., and Calmon, F. Optimized score transformation for fair classification. In *International Conference on Artificial Intelligence and Statistics*, pp. 1673–1683. PMLR, 2020.

[46] Woodworth, B., Gunasekar, S., Ohannessian, M. I., and Srebro, N. Learning non-discriminatory predictors. *arXiv preprint arXiv:1702.06081*, 2017.

[47] Wu, Y., Zhang, L., and Wu, X. On convexity and bounds of fairness-aware classification. In *The World Wide Web Conference*, pp. 3356–3362, 2019.

[48] Zafar, M. B., Valera, I., Gomez Rodriguez, M., and Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, pp. 1171–1180, 2017.

[49] Zafar, M. B., Valera, I., Rogriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pp. 962–970. PMLR, 2017.

[50] Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gummadi, K. P. Fairness constraints: A flexible approach for fair classification. *Journal of Machine Learning Research*, 20(75):1–42, 2019. URL `http://jmlr.org/papers/v20/18-262.html`.

[51] Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

[52] Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457*, 2017.

[53] Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 2, 2018.

[54] Zhao, J., Zhou, Y., Li, Z., Wang, W., and Chang, K.-W. Learning gender-neutral word embeddings. *arXiv preprint arXiv:1809.01496*, 2018.

[55] Zhao, J., Wang, T., Yatskar, M., Cotterell, R., Ordonez, V., and Chang, K.-W. Gender bias in contextualized word embeddings. *arXiv preprint arXiv:1904.03310*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sec 5.2.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Sec 5.2.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Each theorem is immediately preceeded by its assumptions.
   (b) Did you include complete proofs of all theoretical results? [Yes] In the supplement.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Yes, as an anonymous repo in the supplement.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In the supplement.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Due to computational cost of training deep neural networks, we use a single train-test split. However, because the optimization objective changes only slightly as the Lagrange multiplier is varied, sensitivity of optimization methods to random seed can be seen in the smoothness of the curves in result figures.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In the supplement.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We used existing code for kernel-based methods, details provided in supplement
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] We only use benchmark dataset that have been widely studied. Further, all data sets contain only publicly available information: census data for Adult, Freedom of Information Act (FOIA) results for COMPAS, public comments for Yelp and publicly-available pictures for CelebA and Faces of the World. This eliminates the concern of revealing private personal information.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]