

MIT Open Access Articles

Why is my classifier discriminatory?

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Sontag, David and Johansson, Fredrik D. 2018. "Why is my classifier discriminatory?." Advances in Neural Information Processing Systems, 2018-December.

As Published: <https://papers.nips.cc/paper/2018/hash/1f1baa5b8edac74eb4eaa329f14a0361-Abstract.html>

Persistent URL: <https://hdl.handle.net/1721.1/137319>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Why Is My Classifier Discriminatory?

Irene Y. Chen
MIT
iychen@mit.edu

Fredrik D. Johansson
MIT
fredrikj@mit.edu

David Sontag
MIT
dsontag@csail.mit.edu

Abstract

Recent attempts to achieve fairness in predictive models focus on the balance between fairness and accuracy. In sensitive applications such as healthcare or criminal justice, this trade-off is often undesirable as any increase in prediction error could have devastating consequences. In this work, we argue that the fairness of predictions should be evaluated in context of the data, and that unfairness induced by inadequate samples sizes or unmeasured predictive variables should be addressed through data collection, rather than by constraining the model. We decompose cost-based metrics of discrimination into bias, variance, and noise, and propose actions aimed at estimating and reducing each term. Finally, we perform case-studies on prediction of income, mortality, and review ratings, confirming the value of this analysis. We find that data collection is often a means to reduce discrimination without sacrificing accuracy.

1 Introduction

As machine learning algorithms increasingly affect decision making in society, many have raised concerns about the fairness and biases of these algorithms, especially in applications to healthcare or criminal justice, where human lives are at stake (Angwin et al., 2016; Barocas & Selbst, 2016). It is often hoped that the use of automatic decision support systems trained on observational data will remove human bias and improve accuracy. However, factors such as data quality and model choice may encode unintentional discrimination, resulting in systematic disparate impact.

We study fairness in prediction of outcomes such as recidivism, annual income, or patient mortality. Fairness is evaluated with respect to *protected groups* of individuals defined by attributes such as gender or ethnicity (Ruggieri et al., 2010). Following previous work, we measure discrimination in terms of differences in prediction cost across protected groups (Calders & Verwer, 2010; Dwork et al., 2012; Feldman et al., 2015). Correcting for issues of data provenance and historical bias in labels is outside of the scope of this work. Much research has been devoted to constraining models to satisfy cost-based fairness in prediction, as we expand on below. *The impact of data collection on discrimination has received comparatively little attention.*

Fairness in prediction has been encouraged by adjusting models through regularization (Bechavod & Ligett, 2017; Kamishima et al., 2011), constraints (Kamiran et al., 2010; Zafar et al., 2017), and representation learning (Zemel et al., 2013). These attempts can be broadly categorized as model-based approaches to fairness. Others have applied data preprocessing to reduce discrimination (Hajian & Domingo-Ferrer, 2013; Feldman et al., 2015; Calmon et al., 2017). For an empirical comparison, see for example Friedler et al. (2018). Inevitably, however, restricting the model class or perturbing training data to improve fairness may harm predictive accuracy (Corbett-Davies et al., 2017).

A *tradeoff* of predictive accuracy for fairness is sometimes difficult to motivate when predictions influence high-stakes decisions. In particular, post-hoc correction methods based on randomizing predictions (Hardt et al., 2016; Pleiss et al., 2017) are unjustifiable for ethical reasons in clinical tasks

such as severity scoring. Moreover, as pointed out by Woodworth et al. (2017), post-hoc correction may lead to suboptimal predictive accuracy compared to other equally fair classifiers.

Disparate predictive accuracy can often be explained by insufficient or skewed sample sizes or inherent unpredictability of the outcome given the available set of variables. With this in mind, we propose that fairness of predictive models should be analyzed in terms of model bias, model variance, and outcome noise *before* they are constrained to satisfy fairness criteria. This exposes and separates the adverse impact of inadequate data collection and the choice of the model on fairness. The cost of fairness need not always be one of predictive accuracy, but one of investment in data collection and model development. In high-stakes applications, the benefits often outweigh the costs.

In this work, we use the term “discrimination” to refer to specific kinds of differences in the predictive power of models when applied to different protected groups. In some domains, such differences may not be considered discriminatory, and it is critical that decisions made based on this information are sensitive to this fact. For example, in prior work, researchers showed that causal inference may help uncover which sources of differences in predictive accuracy introduce unfairness (Kusner et al., 2017). In this work, we assume that observed differences are considered discriminatory and discuss various means of explaining and reducing them.

Main contributions We give a procedure for analyzing discrimination in predictive models with respect to cost-based definitions of group fairness, emphasizing the impact of data collection. First, we propose the use of bias-variance-noise decompositions for separating sources of discrimination. Second, we suggest procedures for estimating the value of collecting additional training samples. Finally, we propose the use of clustering for identifying subpopulations that are discriminated against to guide additional variable collection. We use these tools to analyze the fairness of common learning algorithms in three tasks: predicting income based on census data, predicting mortality of patients in critical care, and predicting book review ratings from text. We find that the accuracy in predictions of the mortality of cancer patients vary by as much as 20% between protected groups. In addition, our experiments confirm that discrimination level is sensitive to the quality of the training data.

2 Background

We study fairness in prediction of an outcome $Y \in \mathcal{Y}$. Predictions are based on a set of covariates $X \in \mathcal{X} \subseteq \mathbb{R}^k$ and a *protected attribute* $A \in \mathcal{A}$. In mortality prediction, X represents the medical history of a patient in critical care, A the self-reported ethnicity, and Y mortality. A model is considered fair if its errors are distributed similarly across protected groups, as measured by a cost function γ . Predictions learned from a training set d are denoted $\hat{Y}_d := h(X, A)$ for some $h : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{Y}$ from a class \mathcal{H} . The protected attribute is assumed to be binary, $\mathcal{A} = \{0, 1\}$, but our results generalize to the non-binary case. A dataset $d = \{(x_i, a_i, y_i)\}_{i=1}^n$ consists of n samples distributed according to $p(X, A, Y)$. When clear from context, we drop the subscript from \hat{Y}_d .

A popular cost-based definition of fairness is the *equalized odds* criterion, which states that a binary classifier \hat{Y} is fair if its false negative rates (FNR) and false positive rates (FPR) are equal across groups (Hardt et al., 2016). We define FPR and FNR with respect to protected group $a \in \mathcal{A}$ by


$$\text{FPR}_a(\hat{Y}) := \mathbb{E}_X[\hat{Y} \mid Y = 0, A = a], \quad \text{FNR}_a(\hat{Y}) := \mathbb{E}_X[1 - \hat{Y} \mid Y = 1, A = a].$$

Exact equality, $\text{FPR}_0(\hat{Y}) = \text{FPR}_1(\hat{Y})$, is often hard to verify or enforce in practice. Instead, we study the *degree* to which such constraints are violated. More generally, we use differences in *cost functions* γ_a between protected groups $a \in \mathcal{A}$ to define the *level of discrimination* Γ ,

$$\Gamma^\gamma(\hat{Y}) := \left| \gamma_0(\hat{Y}) - \gamma_1(\hat{Y}) \right|. \quad (1)$$

In this work we study cost functions $\gamma_a \in \{\text{FPR}_a, \text{FNR}_a, \text{ZO}_a\}$ in binary classification tasks, with $\text{ZO}_a(\hat{Y}) := \mathbb{E}_X[\mathbb{1}[\hat{Y} \neq Y] \mid A = a]$ the *zero-one loss*. In regression problems, we use the group-specific *mean-squared error* $\text{MSE}_a := \mathbb{E}_X[(\hat{Y} - Y)^2 \mid A = a]$. According to (1), predictions \hat{Y} satisfy equalized odds on d if $\Gamma^{\text{FPR}}(\hat{Y}) = 0$ and $\Gamma^{\text{FNR}}(\hat{Y}) = 0$.

Calibration and impossibility A score-based classifier is *calibrated* if the prediction score assigned to a unit equals the fraction of positive outcomes for all units assigned similar scores. It



(a) For identically distributed protected groups and unaware outcome (see below), bias and noise are equal in expectation. Perceived discrimination is only due to variance.

(b) Heteroskedastic noise, i.e. $\exists x, x' : N(x) \neq N(x')$, may contribute to discrimination even for an optimal model if protected groups are not identically distributed.

(c) One choice of model may be more suited for one protected group, even under negligible noise and variance, resulting in a difference in expected bias, $\bar{B}_0 \neq \bar{B}_1$.

Figure 1: Scenarios illustrating how properties of the training set and model choice affect perceived discrimination in a binary classification task, under the assumption that outcomes are *unaware*, i.e. $p(Y | X, A) = p(Y | X)$ and $p(\hat{Y} | X, A) = p(\hat{Y} | X)$. Through bias-variance-noise decompositions (see Section 3.1), we can identify which of these dominate in their effect on fairness. We propose procedures for addressing each component in Section 4, and use them in experiments (see Section 5) to mitigate discrimination in income prediction and prediction of ICU mortality.

is impossible for a classifier to be calibrated in every protected group and satisfy multiple cost-based fairness criteria at once, unless accuracy is perfect or base rates of outcomes are equal across groups (Chouldechova, 2017). A relaxed version of this result (Kleinberg et al., 2016) applies to the discrimination level Γ . Inevitably, both constraint-based methods and our approach are faced with a choice between which fairness criteria to satisfy, and at what cost.

3 Sources of perceived discrimination

There are many potential sources of discrimination in predictive models. In particular, the choice of hypothesis class \mathcal{H} and learning objective has received a lot of attention (Calders & Verwer, 2010; Zemel et al., 2013; Fish et al., 2016). However, data collection—the chosen set of predictive variables X , the sampling distribution $p(X, A, Y)$, and the training set size n —is an equally integral part of deploying fair machine learning systems in practice, and it should be guided to promote fairness. Below, we tease apart sources of discrimination through bias-variance-noise decompositions of cost-based fairness criteria. In general, we may think of noise in the outcome as the effect of a set of unobserved variables U , potentially interacting with X . Even the optimal achievable error for predictions based on X may be reduced further by observing parts of U . In Figure 1, we illustrate three common learning scenarios and study their fairness properties through bias, variance, and noise.

To account for randomness in the sampling of training sets, we redefine discrimination level (1) in terms of the *expected cost* $\bar{\gamma}_a(\hat{Y}) := \mathbb{E}_D[\gamma_a(\hat{Y}_D)]$ over draws of a random training set D .

Definition 1. The *expected discrimination level* $\bar{\Gamma}(\hat{Y})$ of a predictive model \hat{Y} learned from a random training set D , is

$$\bar{\Gamma}(\hat{Y}) := \left| \mathbb{E}_D \left[\gamma_0(\hat{Y}_D) - \gamma_1(\hat{Y}_D) \right] \right| = \left| \bar{\gamma}_0(\hat{Y}) - \bar{\gamma}_1(\hat{Y}) \right| .$$

$\bar{\Gamma}(\hat{Y})$ is not observed in practice when only a single training set d is available. If n is small, it is recommended to estimate $\bar{\Gamma}$ through re-sampling methods such as bootstrapping (Efron, 1992).

3.1 Bias-variance-noise decompositions of discrimination level

An algorithm that learns models \hat{Y}_D from datasets D is given, and the covariates X and size of the training data n are fixed. We assume that \hat{Y}_D is a deterministic function $\hat{y}_D(x, a)$ given the training set D , e.g. a thresholded scoring function. Following Domingos (2000), we base our analysis on decompositions of loss functions L evaluated at points (x, a) . For decompositions of costs $\gamma_a \in \{\text{ZO}, \text{FPR}, \text{FNR}\}$ we let this be the zero-one loss, $L(y, y') = \mathbb{1}[y \neq y']$, and for

$\gamma_a = \text{MSE}$, the squared loss, $L(y, y') = (y - y')^2$. We define the *main prediction* $\tilde{y}(x, a) = \arg \min_{y'} \mathbb{E}_D[L(\hat{Y}_D, y') \mid X = x, A = a]$ as the average prediction over draws of training sets for the squared loss, and the majority vote for the zero-one loss. The (*Bayes*) *optimal prediction* $y^*(x, a) = \arg \min_{y'} \mathbb{E}_Y[L(Y, y') \mid X = x, A = a]$ achieves the smallest expected error with respect to the random outcome Y .

Definition 2 (Bias, variance and noise). Following Domingos (2000), we define bias B , variance V and noise N at a point (x, a) below.

$$\begin{aligned} B(\hat{Y}, x, a) &= L(y^*(x, a), \tilde{y}(x, a)) & N(x, a) &= \mathbb{E}_Y[L(y^*(x, a), Y) \mid X = x, A = a] \\ V(\hat{Y}, x, a) &= \mathbb{E}_D[L(\tilde{y}(x, a), \hat{y}_D(x, a))] \end{aligned} \quad (2)$$

Here, y^* , \tilde{y} and \hat{y} are all deterministic functions of (x, a) , while Y is a random variable.

In words, the bias B is the loss incurred by the main prediction relative to the optimal prediction. The variance V is the average loss incurred by the predictions learned from different datasets relative to the main prediction. The noise N is the remaining loss independent of the learning algorithm, often known as the Bayes error. We use these definitions to decompose $\bar{\Gamma}$ under various definitions of γ_a .

Theorem 1. With $\bar{\gamma}_a$ the group-specific zero-one loss or class-conditional versions (e.g. FNR, FPR), or the mean squared error, $\bar{\gamma}_a$ and the discrimination level $\bar{\Gamma}$ admit decompositions of the form

$$\bar{\gamma}_a(\hat{Y}) = \underbrace{\bar{N}_a}_{\text{Noise}} + \underbrace{\bar{B}_a(\hat{Y})}_{\text{Bias}} + \underbrace{\bar{V}_a(\hat{Y})}_{\text{Variance}} \quad \text{and} \quad \bar{\Gamma} = |(\bar{N}_0 - \bar{N}_1) + (\bar{B}_0 - \bar{B}_1) + (\bar{V}_0 - \bar{V}_1)|$$

where we leave out \hat{Y} in the decomposition of $\bar{\Gamma}$ for brevity. With B, V defined as in (2), we have

$$\bar{B}_a(\hat{Y}) = \mathbb{E}_X[B(\tilde{y}, X, a) \mid A = a] \quad \text{and} \quad \bar{V}_a(\hat{Y}) = \mathbb{E}_{X, D}[c_v(X)V(\hat{Y}_D, X, a) \mid A = a].$$

For the zero-one loss, $c_v(x, a) = 1$ if $\hat{y}_m(x, a) = y^*(x, a)$, otherwise $c_v(x, a) = -1$. For the squared loss $c_v(x, a) = 1$. The noise term for population losses is

$$\bar{N}_a := \mathbb{E}_X[c_n(X, a)L(y^*(X, a), Y) \mid A = a]$$

and for class-conditional losses w.r.t class $y \in \{0, 1\}$,

$$\bar{N}_a(y) := \mathbb{E}_X[c_n(X, a)L(y^*(X, a), y) \mid A = a, Y = y].$$

For the zero-one loss, and class-conditional variants, $c_n(x, a) = 2\mathbb{E}_D[\mathbb{1}[\hat{y}_D(x, a) = y^*(x, a)]] - 1$ and for the squared loss, $c_n(x, a) = 1$.

Proof sketch. Conditioning and exchanging order of expectation, the cases of mean squared error and zero-one losses follow from Domingos (2000). Class-conditional losses follow from a case-by-case analysis of possible errors. See the supplementary material for a full proof. \square

Theorem 1 points to distinct sources of perceived discrimination. Significant differences in bias $\bar{B}_0 - \bar{B}_1$ indicate that the chosen model class is not flexible enough to fit both protected groups well (see Figure 1c). This is typical of (misspecified) linear models which approximate non-linear functions well only in small regions of the input space. Regularization or post-hoc correction of models effectively increase the bias of one of the groups, and should be considered only if there is reason to believe that the original bias is already minimal.

Differences in variance, $\bar{V}_0 - \bar{V}_1$, could be caused by differences in sample sizes n_0, n_1 or group-conditional feature variance $\text{Var}(X \mid A)$, combined with a high capacity model. Targeted collection of training samples may help resolve this issue. Our decomposition does not apply to post-hoc randomization methods (Hardt et al., 2016) but we may treat these in the same way as we do random training sets and interpret them as increasing the variance \bar{V}_a of one group to improve fairness.

When noise is significantly different between protected groups, discrimination is partially unrelated to model choice and training set size and may only be reduced by measuring additional variables.

Proposition 1. If $\bar{N}_0 \neq \bar{N}_1$, no model can be 0-discriminatory in expectation without access to additional information or increasing bias or variance w.r.t. to the Bayes optimal classifier.

Proof. By definition, $\bar{\Gamma} = 0 \implies (\bar{N}_1 - \bar{N}_0) = (\bar{B}_0 - \bar{B}_1) + (\bar{V}_0 - \bar{V}_1)$. As the Bayes optimal classifier has neither bias nor variance, the result follows immediately. \square

In line with Proposition 1, most methods for ensuring algorithmic fairness reduce discrimination by trading off a difference in noise for one in bias or variance. However, this trade-off is only motivated if the considered predictive model is close to Bayes optimal *and* no additional predictive variables may be measured. Moreover, if noise is homoskedastic in regression settings, post-hoc randomization is ill-advised, as the difference in Bayes error $\bar{N}_0 - \bar{N}_1$ is zero, and discrimination is caused only by model bias or variance (see the supplementary material for a proof).

Estimating bias, variance and noise Group-specific variance \bar{V}_a may be estimated through sample splitting or bootstrapping (Efron, 1992). In contrast, the noise \bar{N}_a and bias \bar{B}_a are difficult to estimate when X is high-dimensional or continuous. In fact, no convergence results of noise estimates may be obtained without further assumptions on the data distribution (Antos et al., 1999). Under some such assumptions, noise may be approximately estimated using distance-based methods (Devijver & Kittler, 1982), nearest-neighbor methods (Fukunaga & Hummels, 1987; Cover & Hart, 1967), or classifier ensembles (Tumer & Ghosh, 1996). When comparing the discrimination level of two different models, noise terms cancel, as they are independent of the model. As a result, *differences* in bias may be estimated even when the noise is not known (see the supplementary material).

Testing for significant discrimination When sample sizes are small, perceived discrimination may not be statistically significant. In the supplementary material, we give statistical tests both for the discrimination level $\Gamma(\hat{Y})$ and the difference in discrimination level between two models \hat{Y}, \hat{Y}' .

4 Reducing discrimination through data collection

In light of the decomposition of Theorem 1, we explore avenues for reducing group differences in bias, variance, and noise without sacrificing predictive accuracy. In practice, predictive accuracy is often artificially limited when data is expensive or impractical to collect. With an investment in training samples or measurement of predictive variables, both accuracy and fairness may be improved.

4.1 Increasing training set size

Standard regularization used to avoid overfitting is not guaranteed to improve or preserve fairness. An alternative route is to collect more training samples and reduce the impact of the bias-variance trade-off. When supplementary data is collected from the same distribution as the existing set, covariate shift may be avoided (Quionero-Candela et al., 2009). This is often achievable; labeled data may be expensive, such as when paying experts to label observations, but given the means to acquire additional labels, they would be drawn from the original distribution. To estimate the value of increasing sample size, we predict the discrimination level $\bar{\Gamma}(\hat{Y}_D)$ as D increases in size.

The curve measuring generalization performance of predictive models as a function of training set size n is called a Type II *learning curve* (Domhan et al., 2015). We call $\bar{\gamma}_a(\hat{Y}, n) := \mathbb{E}[\gamma_a(\hat{Y}_{D_n})]$, as a function of n , the learning curve with respect to protected group a . We define the discrimination learning curve $\bar{\Gamma}(\hat{Y}, n) := |\bar{\gamma}_0(\hat{Y}, n) - \bar{\gamma}_1(\hat{Y}, n)|$ (see Figure 2a for an example). Empirically, learning curves behave asymptotically as *inverse power-law* curves for diverse algorithms such as deep neural networks, support vector machines, and nearest-neighbor classifiers, even when model capacity is allowed to grow with n (Hestness et al., 2017; Mukherjee et al., 2003). This observation is also supported by theoretical results (Amari, 1993).

Assumption 1 (Learning curves). *The population prediction loss $\bar{\gamma}(\hat{Y}, n)$, and group-specific losses $\bar{\gamma}_0(\hat{Y}, n), \bar{\gamma}_1(\hat{Y}, n)$, for a fixed learning algorithm \hat{Y} , behave asymptotically as inverse power-law curves with parameters (α, β, δ) . That is, $\exists M, M_0, M_1$ such that for $n \geq M, n_a \geq M_a$,*

$$\bar{\gamma}(\hat{Y}, n) = \alpha n^{-\beta} + \delta \quad \text{and} \quad \forall a \in \mathcal{A} : \bar{\gamma}_a(\hat{Y}, n_a) = \alpha_a n_a^{-\beta_a} + \delta_a \quad (3)$$

Intercepts, δ, δ_a in (3) represent the asymptotic bias $\bar{B}(\hat{Y}_{D_\infty})$ and the Bayes error \bar{N} , with the former vanishing for consistent estimators. Accurately estimating δ from finite samples is often challenging as the first term tends to dominate the learning curve for practical sample sizes.

In experiments, we find that the inverse power-laws model fit group conditional (γ_a) and class-conditional (FPR, FNR) errors well, and use these to extrapolate $\bar{\Gamma}(\hat{Y}, n)$ based on estimates from subsampled data.

4.2 Measuring additional variables

When discrimination $\bar{\Gamma}$ is dominated by a difference in noise, $\bar{N}_0 - \bar{N}_1$, fairness may not be improved through model selection alone without sacrificing accuracy (see Proposition 1). Such a scenario is likely when available covariates are not equally predictive of the outcome in both groups. We propose identification of clusters of individuals in which discrimination is high as a means to guide further variable collection—if the variance in outcomes within a cluster is not explained by the available feature set, additional variables may be used to further distinguish its members.

Let a random variable C represent a (possibly stochastic) clustering such that $C = c$ indicates membership in cluster c . Then let $\rho_a(c)$ denote the expected prediction cost for units in cluster c with protected attribute a . As an example, for the zero-one loss we let

$$\rho_a^{\text{ZO}}(c) := \mathbb{E}_X[\mathbb{1}[\hat{Y} \neq Y] \mid A = a, C = c],$$

and define ρ analogously for false positives or false negatives. Clusters c for which $|\rho_0(c) - \rho_1(c)|$ is large identify groups of individuals for which discrimination is worse than average, and can guide targeted collection of additional variables or samples. In our experiments on income prediction, we consider particularly simple clusterings of data defined by subjects with measurements above or below the average value of a single feature $x(c)$ with $c \in \{1, \dots, k\}$. In mortality prediction, we cluster patients using topic modeling. As measuring additional variables is expensive, the utility of a candidate set should be estimated before collecting a large sample (Koepeke & Bilenko, 2012).

5 Experiments

We analyze the fairness properties of standard machine learning algorithms in three tasks: prediction of income based on national census data, prediction of patient mortality based on clinical notes, and prediction of book review ratings based on review text.¹ We disentangle sources of discrimination by assessing the level of discrimination for the full data, estimating the value of increasing training set size by fitting Type II learning curves, and using clustering to identify subgroups where discrimination is high. In addition, we estimate the Bayes error through non-parametric techniques.

In our experiments, we omit the sensitive attribute A from our classifiers to allow for closer comparison to previous works, e.g. Hardt et al. (2016); Zafar et al. (2017). In preliminary results, we found that fitting separate classifiers for each group increased the error rates of both groups due to the resulting smaller sample size, as classifiers could not learn from other groups. As our model objective is to maximize accuracy over all data points, our analysis uses a single classifier trained on the entire population.


5.1 Income prediction

Predictions of a person’s salary may be used to help determine an individual’s market worth, but systematic underestimation of the salary of protected groups could harm their competitiveness on the job market. The Adult dataset in the UCI Machine Learning Repository (Lichman, 2013) contains 32,561 observations of yearly income (represented as a binary outcome: over or under \$50,000) and twelve categorical or continuous features including education, age, and marital status. Categorical attributes are dichotomized, resulting in a total of 105 features.

We follow Pleiss et al. (2017) and strive to ensure fairness across genders, which is excluded as a feature from the predictive models. Using an 80/20 train-test split, we learn a random forest predictor, which is well-calibrated for both groups (Brier (1950) scores of 0.13 and 0.06 for men and women). We find the difference in zero-one loss $\Gamma^{\text{ZO}}(\hat{Y})$ has a 95%-confidence interval² $.085 \pm .069$ with decision thresholds at 0.5. At this threshold, the false negative rates are 0.388 ± 0.026 and 0.448 ± 0.064 for men and women respectively, and the false positive rates 0.111 ± 0.011 and

¹A synthetic experiment validating group-specific learning curves is left to the supplementary material.

²Details for computing statistically significant discrimination can be found in the supplementary material.



(a) Group differences in false positive rates and false negative rates for a random forest classifier decrease with increasing training set size.

Method	E_{low}	E_{up}	group
Mahalanobis	–	0.29	men
(Mahalanobis, 1936)	–	0.13	women
Bhattacharyya	0.001	0.040	men
(Bhattacharyya, 1943)	0.001	0.027	women
Nearest Neighbors	0.10	0.19	men
(Cover & Hart, 1967)	0.04	0.07	women

(b) Estimation of Bayes error lower and upper bounds (E_{low} and E_{up}) for zero-one loss of men and women. Intervals for men and women are non-overlapping for Nearest Neighbors.

Figure 2: Discrimination level and noise estimation in income prediction with the Adult dataset.

0.033 ± 0.008 . We focus on random forest classifiers, although we found similar results for logistic regression and decision trees.

We examine the effect of varying training set size n on discrimination. We fit inverse power-law curves to estimates of $FPR(\hat{Y}, n)$ and $FNR(\hat{Y}, n)$ using repeated sample splitting where at least 20% of the full data is held out for evaluating generalization error at every value of n . We tune hyperparameters for each training set size for decision tree classifiers and logistic regression but tuned over the entire dataset for random forest. We include full training details in the supplementary material. Metrics are averaged over 50 trials. See Figure 2a for the results for random forests. Both FPR and FNR decrease with additional training samples. The discrimination level Γ^{FNR} for false negatives decreases by a striking 40% when increasing the training set size from 1000 to 10,000. This suggests that trading off accuracy for fairness at small sample sizes may be ill-advised. Based on fitted power-law curves, we estimate that for unlimited training data drawn from the same distribution, we would have $\Gamma^{FNR}(\hat{Y}) \approx 0.04$ and $\Gamma^{FPR}(\hat{Y}) \approx 0.08$.

In Figure 2b, we compare estimated upper and lower bounds on noise (E_{low} and E_{up}) for men and women using the Mahalanobis and Bhattacharyya distances (Devijver & Kittler, 1982), and a k -nearest neighbor method (Cover & Hart, 1967) with $k = 5$ and 5-fold cross validation. Men have consistently higher noise estimates than women, which is consistent with the differences in zero-one loss found using all models. For nearest neighbors estimates, intervals for men and women are non-overlapping, which suggests that noise may contribute substantially to discrimination.

To guide attempts at reducing discrimination further, we identify clusters of individuals for whom false negative predictions are made at different rates between protected groups, with the method described in Section 4.2. We find that for individuals in executive or managerial occupations (12% of the sample), false negatives are more than twice as frequent for women (0.412) as for men (0.157). For individuals in all other occupations, the difference is significantly smaller, 0.543 for women and 0.461 for men, despite the fact that the disparity in outcome base rates in this cluster is large (0.26 for men versus 0.09 for women). A possible reason is that in managerial occupations the available variable set explains a larger portion of the variance in salary for men than for women. If so, further sub-categorization of managerial occupations could help reduce discrimination in prediction.

5.2 Intensive care unit mortality prediction

Unstructured medical data such as clinical notes can reveal insights for questions like mortality prediction; however, disparities in predictive accuracy may result in discrimination of protected groups. Using the MIMIC-III dataset of all clinical notes from 25,879 adult patients from Beth Israel Deaconess Medical Center (Johnson et al., 2016), we predict hospital mortality of patients in critical care. Fairness is studied with respect to five self-reported ethnic groups of the following proportions: Asian (2.2%), Black (8.8%), Hispanic (3.4%), White (70.8%), and Other (14.8%). Notes were collected in the first 48 hours of an intensive care unit (ICU) stay; discharge notes were excluded. We only included patients that stayed in the ICU for more than 48 hours. We use the tf-idf statistics of the 10,000 most frequent words as features. Training a model on 50% of the data, selecting

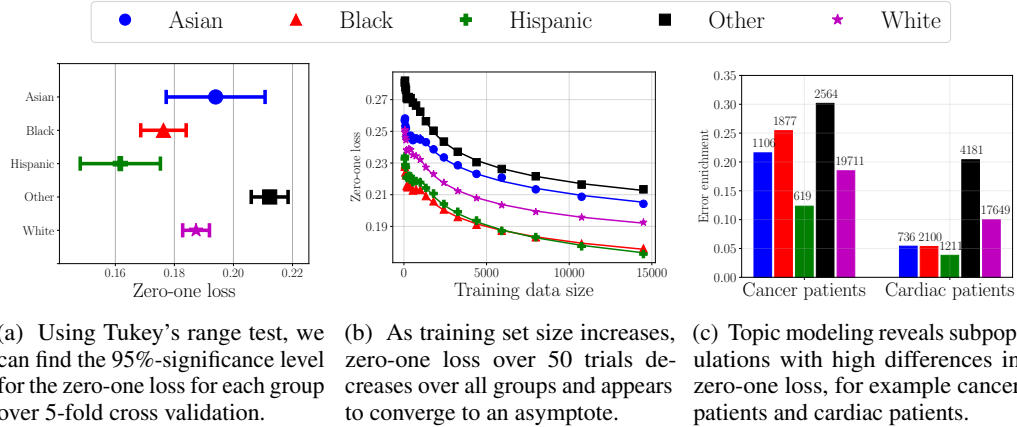


Figure 3: Mortality prediction from clinical notes using logistic regression. Best viewed in color.

hyper-parameters on 25%, and testing on 25%, we find that logistic regression with L1-regularization achieves an AUC of 0.81. The logistic regression is well-calibrated with Brier scores ranging from 0.06-0.11 across the five groups; we note better calibration is correlated with lower prediction error.

We report cost and discrimination level in terms of generalized zero-one loss (Pleiss et al., 2017). Using an ANOVA test (Fisher, 1925) with $p < 0.001$, we reject the null hypothesis that loss is the same among all five groups. To map the 95% confidence intervals, we perform pairwise comparisons of means using Tukey’s range test (Tukey, 1949) across 5-fold cross-validation. As seen in Figure 3a, patients in the Other and Hispanic groups have the highest and lowest generalized zero-one loss, respectively, with relatively few overlapping intervals. Notably, the largest ethnic group (White) does not have the best accuracy, whereas smaller ethnic groups tend towards extremes. While racial groups differ in hospital mortality base rates (Table 1 in the Supplementary material), Hispanic (10.3%) and Black (10.9%) patients have very different error rates despite similar base rates.

To better understand the discrimination induced by our model, we explore the effect of changing training set size. To this end, we repeatedly subsample and split the data, holding out at least 20% of the full data for testing. In Figure 3b, we show loss averaged over 50 trials of training a logistic regression on increasingly larger training sets; estimated inverse power-law curves show good fits. We see that some pairwise differences in loss decrease with additional training data.

Next, we identify clusters for which the difference in prediction errors between protected groups is large. We learn a topic model with $k = 50$ topics generated using Latent Dirichlet Allocation (Blei et al., 2003). Topics are concatenated into an $n \times k$ matrix Q where q_{ic} designates the proportion of topic $c \in [k]$ in note $i \in [n]$. Following prior work on enrichment of topics in clinical notes (Marlin et al., 2012; Ghassemi et al., 2014), we estimate the probability of patient mortality Y given a topic c as $\hat{p}(Y|C = c) := (\sum_{i=1}^n y_i q_{ic}) / (\sum_{i=1}^n q_{ic})$ where y_i is the hospital mortality of patient i . We compare relative error rates given protected group and topic using binary predicted mortality \hat{y}_i , actual mortality y_i , and group a_i for patient i through

$$\hat{p}(\hat{Y} \neq Y | A = a', C = c) = \frac{\sum_{i=1}^n \mathbb{1}(y_i \neq \hat{y}_i) \mathbb{1}(a_i = a') q_{ic}}{\sum_{i=1}^n \mathbb{1}(a_i = a') q_{ic}}$$

which follows using substitution and conditioning on A . These error rates were computed using a logistic regression with L1 regularization using an 80/20 train-test split over 50 trials. While many topics have consistent error rates across groups, some topics (e.g. cardiac patients or cancer patients as shown in Figure 3c) have large differences in error rates across groups. We include more detailed topic descriptions in the supplementary material. Once we have identified a subpopulation with particularly high error, for example cancer patients, we can consider collecting more features or collecting more data from the same data distribution. We find that error rates differ between 0.12 and 0.30 across protected groups of cancer patients, and between 0.05 and 0.20 for cardiac patients.

5.3 Book review ratings

In the supplementary material, we study prediction of book review ratings from review texts (Gnanesh, 2017). The protected attribute was chosen to be the gender of the author as determined from Wikipedia. In the dataset, the difference in mean-squared error $\Gamma^{\text{MSE}}(\hat{Y})$ has 95%-confidence interval 0.136 ± 0.048 with $\text{MSE}_M = 0.224$ for reviews for male authors and $\text{MSE}_F = 0.358$. Strikingly, our findings suggest that $\Gamma^{\text{MSE}}(\hat{Y})$ may be completely eliminated by additional targeted sampling of the less represented gender.

6 Discussion

We identify that existing approaches for reducing discrimination induced by prediction errors may be unethical or impractical to apply in settings where predictive accuracy is critical, such as in healthcare or criminal justice. As an alternative, we propose a procedure for analyzing the different sources contributing to discrimination. Decomposing well-known definitions of cost-based fairness criteria in terms of differences in bias, variance, and noise, we suggest methods for reducing each term through model choice or additional training data collection. Case studies on three real-world datasets confirm that collection of additional samples is often sufficient to improve fairness, and that existing post-hoc methods for reducing discrimination may unnecessarily sacrifice predictive accuracy when other solutions are available.

Looking forward, we can see several avenues for future research. In this work, we argue that identifying clusters or subpopulations with high predictive disparity would allow for more targeted ways to reduce discrimination. We encourage future research to dig deeper into the question of local or context-specific unfairness in general, and into algorithms for addressing it. Additionally, extending our analysis to intersectional fairness (Buolamwini & Gebru, 2018; Hébert-Johnson et al., 2017), e.g. looking at both gender and race or all subdivisions, would provide more nuanced grappling with unfairness. Finally, additional data collection to improve the model may cause unexpected delayed impacts (Liu et al., 2018) and negative feedback loops (Ensign et al., 2017) as a result of distributional shifts in the data. More broadly, we believe that the study of fairness in non-stationary populations is an interesting direction to pursue.

Acknowledgements

The authors would like to thank Yoni Halpern and Hunter Lang for helpful comments, and Zeshan Hussain for clinical guidance. This work was partially supported by Office of Naval Research Award No. N00014-17-1-2791 and NSF CAREER award #1350965.

References

- Amari, Shun-Ichi. A universal theorem on learning curves. *Neural networks*, 6(2):161–166, 1993.
- Angwin, Julia, Larson, Jeff, Mattu, Surya, and Kirchner, Lauren. Machine bias. *ProPublica*, May, 23, 2016.
- Antos, András, Devroye, Luc, and Györfi, László. Lower bounds for bayes error estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(7):643–645, 1999.
- Barocas, Solon and Selbst, Andrew D. Big data’s disparate impact. *Cal. L. Rev.*, 104:671, 2016.
- Bechavod, Yahav and Ligett, Katrina. Learning fair classifiers: A regularization-inspired approach. *arXiv preprint arXiv:1707.00044*, 2017.
- Bhattacharyya, Anil. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.*, 35:99–109, 1943.
- Blei, David M, Ng, Andrew Y, and Jordan, Michael I. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- Brier, Glenn W. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.

- Buolamwini, Joy and Gebru, Timnit. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, pp. 77–91, 2018.
- Calders, Toon and Verwer, Sicco. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.
- Calmon, Flavio, Wei, Dennis, Vinzamuri, Bhanukiran, Ramamurthy, Karthikeyan Natesan, and Varshney, Kush R. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pp. 3995–4004, 2017.
- Chouldechova, Alexandra. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv preprint arXiv:1703.00056*, 2017.
- Corbett-Davies, Sam, Pierson, Emma, Feller, Avi, Goel, Sharad, and Huq, Aziz. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 797–806. ACM, 2017.
- Cover, Thomas and Hart, Peter. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- Devijver, Pierre A. and Kittler, Josef. *Pattern recognition: a statistical approach*. Sung Kang, 1982.
- Domhan, Tobias, Springenberg, Jost Tobias, and Hutter, Frank. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- Domingos, Pedro. A unified bias-variance decomposition. In *Proceedings of 17th International Conference on Machine Learning*, pp. 231–238, 2000.
- Dwork, Cynthia, Hardt, Moritz, Pitassi, Toniann, Reingold, Omer, and Zemel, Richard. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 214–226. ACM, 2012.
- Efron, Bradley. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pp. 569–593. Springer, 1992.
- Ensign, Danielle, Friedler, Sorelle A., Neville, Scott, Scheidegger, Carlos Eduardo, and Venkatasubramanian, Suresh. Runaway feedback loops in predictive policing. *CoRR*, abs/1706.09847, 2017. URL <http://arxiv.org/abs/1706.09847>.
- Feldman, Michael, Friedler, Sorelle A, Moeller, John, Scheidegger, Carlos, and Venkatasubramanian, Suresh. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 259–268. ACM, 2015.
- Fish, Benjamin, Kun, Jeremy, and Lelkes, Ádám D. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pp. 144–152. SIAM, 2016.
- Fisher, R.A. *Statistical methods for research workers*. Edinburgh Oliver & Boyd, 1925.
- Friedler, Sorelle A, Scheidegger, Carlos, Venkatasubramanian, Suresh, Choudhary, Sonam, Hamilton, Evan P, and Roth, Derek. A comparative study of fairness-enhancing interventions in machine learning. *arXiv preprint arXiv:1802.04422*, 2018.
- Fukunaga, Keinosuke and Hummels, Donald M. Bayes error estimation using parzen and k-nn procedures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):634–643, 1987.
- Ghassemi, Marzyeh, Naumann, Tristan, Doshi-Velez, Finale, Brimmer, Nicole, Joshi, Rohit, Rumshisky, Anna, and Szolovits, Peter. Unfolding physiological state: Mortality modelling in intensive care units. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 75–84. ACM, 2014.
- Gnanesh. Goodreads book reviews, 2017. URL <https://www.kaggle.com/gnanesh/goodreads-book-reviews>.

- Hajian, Sara and Domingo-Ferrer, Josep. A methodology for direct and indirect discrimination prevention in data mining. *IEEE transactions on knowledge and data engineering*, 25(7):1445–1459, 2013.
- Hardt, Moritz, Price, Eric, Srebro, Nati, et al. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pp. 3315–3323, 2016.
- Hébert-Johnson, Ursula, Kim, Michael P, Reingold, Omer, and Rothblum, Guy N. Calibration for the (computationally-identifiable) masses. *arXiv preprint arXiv:1711.08513*, 2017.
- Hestness, Joel, Narang, Sharan, Ardalani, Newsha, Damos, Gregory, Jun, Heewoo, Kianinejad, Hassan, Patwary, Md, Ali, Mostofa, Yang, Yang, and Zhou, Yanqi. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.
- Johnson, Alistair EW, Pollard, Tom J, Shen, Lu, Lehman, Li-wei H, Feng, Mengling, Ghassemi, Mohammad, Moody, Benjamin, Szolovits, Peter, Celi, Leo Anthony, and Mark, Roger G. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3, 2016.
- Kamiran, Faisal, Calders, Toon, and Pechenizkiy, Mykola. Discrimination aware decision tree learning. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pp. 869–874. IEEE, 2010.
- Kamishima, Toshihiro, Akaho, Shotaro, and Sakuma, Jun. Fairness-aware learning through regularization approach. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pp. 643–650. IEEE, 2011.
- Kleinberg, Jon, Mullainathan, Sendhil, and Raghavan, Manish. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- Koepke, Hoyt and Bilenko, Mikhail. Fast prediction of new feature utility. *arXiv preprint arXiv:1206.4680*, 2012.
- Kusner, Matt J, Loftus, Joshua, Russell, Chris, and Silva, Ricardo. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, pp. 4069–4079, 2017.
- Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Liu, Lydia T, Dean, Sarah, Rolf, Esther, Simchowitz, Max, and Hardt, Moritz. Delayed impact of fair machine learning. *arXiv preprint arXiv:1803.04383*, 2018.
- Mahalanobis, Prasanta Chandra. On the generalized distance in statistics. National Institute of Science of India, 1936.
- Marlin, Benjamin M, Kale, David C, Khemani, Robinder G, and Wetzel, Randall C. Unsupervised pattern discovery in electronic health care data using probabilistic clustering models. In *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium*, pp. 389–398. ACM, 2012.
- Mukherjee, Sayan, Tamayo, Pablo, Rogers, Simon, Rifkin, Ryan, Engle, Anna, Campbell, Colin, Golub, Todd R, and Mesirov, Jill P. Estimating dataset size requirements for classifying dna microarray data. *Journal of computational biology*, 10(2):119–142, 2003.
- Pleiss, Geoff, Raghavan, Manish, Wu, Felix, Kleinberg, Jon, and Weinberger, Kilian Q. On fairness and calibration. In *Advances in Neural Information Processing Systems*, pp. 5684–5693, 2017.
- Quionero-Candela, Joaquin, Sugiyama, Masashi, Schwaighofer, Anton, and Lawrence, Neil D. *Dataset shift in machine learning*. The MIT Press, 2009.
- Ruggieri, Salvatore, Pedreschi, Dino, and Turini, Franco. Data mining for discrimination discovery. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 4(2):9, 2010.
- Tukey, John W. Comparing individual means in the analysis of variance. *Biometrics*, pp. 99–114, 1949.

- Tumer, Kagan and Ghosh, Joydeep. Estimating the bayes error rate through classifier combining. In *Pattern Recognition, 1996., Proceedings of the 13th International Conference on*, volume 2, pp. 695–699. IEEE, 1996.
- Woodworth, Blake, Gunasekar, Suriya, Ohannessian, Mesrob I, and Srebro, Nathan. Learning non-discriminatory predictors. *Conference On Learning Theory*, 2017.
- Zafar, Muhammad Bilal, Valera, Isabel, Gomez Rodriguez, Manuel, and Gummadi, Krishna P. Fairness constraints: Mechanisms for fair classification. *arXiv preprint arXiv:1507.05259*, 2017.
- Zemel, Richard S, Wu, Yu, Swersky, Kevin, Pitassi, Toniann, and Dwork, Cynthia. Learning fair representations. *ICML (3)*, 28:325–333, 2013.