

Certifying and Removing Disparate Impact*

Michael Feldman
Haverford College
mfeldman@haverford.edu

Sorelle A. Friedler
Haverford College
sorelle@cs.haverford.edu

John Moeller
University of Utah
moeller@cs.utah.edu

Carlos Scheidegger
University of Arizona
cscheid@cscheid.net

Suresh Venkatasubramanian[†]
University of Utah
suresh@cs.utah.edu

ABSTRACT

What does it mean for an algorithm to be biased? In U.S. law, unintentional bias is encoded via *disparate impact*, which occurs when a selection process has widely different outcomes for different groups, even as it appears to be neutral. This legal determination hinges on a definition of a protected class (ethnicity, gender) and an explicit description of the process.

When computers are involved, determining disparate impact (and hence bias) is harder. It might not be possible to disclose the process. In addition, even if the process is open, it might be hard to elucidate in a legal setting how the algorithm makes its decisions. Instead of requiring access to the process, we propose making inferences based on the *data* it uses.

We present four contributions. First, we link disparate impact to a measure of classification accuracy that while known, has received relatively little attention. Second, we propose a test for disparate impact based on how well the protected class can be predicted from the other attributes. Third, we describe methods by which data might be made unbiased. Finally, we present empirical evidence supporting the effectiveness of our test for disparate impact and our approach for both masking bias and preserving *relevant* information in the data. Interestingly, our approach resembles some actual selection practices that have recently received legal scrutiny.

Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications—*Data Mining*; J.4 [Computer Applications]: Social and Behavioral Sciences; K.5.m [Computing Milieux]: Legal Aspects of Computing—*Miscellaneous*

Keywords

Disparate impact, fairness, machine learning

* Author ordering on this paper is alphabetical.

[†]Corresponding author.

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

KDD'15, August 10-13, 2015, Sydney, NSW, Australia.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3664-2/15/08 ...\$15.00.

1. INTRODUCTION

In *Griggs v. Duke Power Co.* [19], the US Supreme Court ruled a business hiring decision illegal if it resulted in disparate impact by race even if the decision was not explicitly determined based on race. The Duke Power Co. was forced to stop using intelligence test scores and high school diplomas, qualifications largely correlated with race, to make hiring decisions. The Griggs decision gave birth to the legal doctrine of *disparate impact*, which today is the predominant legal theory used to determine *unintended* discrimination in the U.S. Note that disparate *impact* is different from disparate *treatment*, which refers to *intended* or direct discrimination. Ricci v. DeStefano [21] examined the relationship between the two notions, and disparate impact remains a topic of legal interest.

Today, algorithms are being used to make decisions both large and small in almost all aspects of our lives, whether they involve mundane tasks like recommendations for buying goods, predictions of credit rating prior to approving a housing loan, or even life-altering decisions like sentencing guidelines after conviction [6]. How do we know if these algorithms are biased, involve illegal discrimination, or are unfair?

These concerns have generated calls, by governments and NGOs alike, for research into these issues [17, 23]. In this paper, we introduce and address two such problems with the goals of quantifying and then removing disparate impact.

While the Supreme Court has resisted a “rigid mathematical formula” defining disparate impact [20], we will adopt a generalization of the 80 percent rule advocated by the US Equal Employment Opportunity Commission (EEOC) [24]. We note that disparate impact itself is not illegal; in hiring decisions, business necessity arguments can be made to excuse disparate impact.

DEFINITION 1.1 (DISPARATE IMPACT (“80% RULE”).) *Given data set $D = (X, Y, C)$, with protected attribute X (e.g., race, sex, religion, etc.), remaining attributes Y , and binary class to be predicted C (e.g., “will hire”), we will say that D has disparate impact if*

$$\frac{\Pr(C = \text{YES} | X = 0)}{\Pr(C = \text{YES} | X = 1)} \leq \tau = 0.8$$

for positive outcome class YES and majority protected attribute 1 where $\Pr(C = c | X = x)$ denotes the conditional probability (evaluated over D) that the class outcome is $c \in C$ given protected attribute $x \in X$.¹

¹Note that under this definition disparate impact is determined based on the given data set and decision outcomes.

The two problems we consider address identifying and removing disparate impact. The *disparate impact certification problem* is to guarantee that, given D , any classification algorithm aiming to predict some C' (which is potentially different from the given C) from Y would not have disparate impact. By certifying any outcomes C' , and not the process by which they were reached, we follow legal precedent in making no judgment on the algorithm itself, and additionally ensure that potentially sensitive algorithms remain proprietary. The *disparate impact removal problem* is to take some data set D and return a data set $\bar{D} = (X, \bar{Y}, C)$ that can be certified as not having disparate impact. The goal is to change only the remaining attributes Y , leaving C as in the original data set so that the ability to classify can be preserved as much as possible.

1.1 Results

We have four main contributions.

We first introduce these problems to the computer science community and develop its theoretical underpinnings. The study of the EEOC's 80% rule as a specific class of loss function does not appear to have received much attention in the literature. We link this measure of disparate impact to the balanced error rate (BER). We show that any decision exhibiting disparate impact can be converted into one where the protected attribute leaks, i.e. can be predicted with low BER.

Second, this theoretical result gives us a procedure for certifying the impossibility of disparate impact on a data set. This procedure involves a particular regression algorithm which minimizes BER. We connect BER to disparate impact in a variety of settings (point and interval estimates, and distributions). We discuss these two contributions in Sections 3 and 4.

In Section 5, we show how to transform the input dataset so that predictability of the protected attribute is impossible. We show that this transformation still preserves much of the signal in the unprotected attributes and has nice properties in terms of closeness to the original data distribution.

Finally, we present a detailed empirical study in Section 6. We show that our algorithm certifying lack of disparate impact on a data set is effective, such that with the three classifiers we used certified data sets don't show disparate impact. We demonstrate the fairness / utility tradeoff for our partial repair procedures. Comparing to related work, we find that for any desired fairness value we can achieve a higher accuracy than other fairness procedures. This is likely due to our emphasis on changing the data to achieve fairness, thus allowing any strong classifier to be used for prediction.

Our procedure for detecting disparate impact goes through an actual classification algorithm. As we show in our experiments, a better classifier provides a more sensitive detector. We believe this is notable. As algorithms get better at learning patterns, they become more able to introduce subtle biases into the decision-making process by finding subtle dependencies among features. But this very sophistication helps *detect* such biases as well via our procedure! Thus, data mining can be used to verify the fairness of such algorithms as well.

2. RELATED WORK

There is, of course, a long history of legal work on disparate impact. There is also related work under the name *statistical*

Notably, it does not use a broader sample universe, and does not take into account statistical significance as has been advocated by some legal scholars [16].

discrimination in Economics. We will not survey such work here. Instead, we direct the reader to the survey of Romei and Ruggieri [18] and to a discussion of the issues specific to data mining and disparate impact [1]. Here, we focus on data mining research relating to combating discrimination. This research can be broadly categorized in terms of methods that achieve fairness by modifying the classifiers and those that achieve fairness by modifying data.

Kamishima et al. [9,10] develop a regularizer for classifiers to penalize prejudicial outcomes and show that this can reduce indirect prejudice (their name for implicit discrimination like disparate impact) while still allowing for accurate classification. They note that as prejudicial outcomes are decreased, the classification accuracy is also decreased. Our work falls into the category of algorithms that change the input data. Previous work has focused on changing the class values of the original data in such a way so that the total number of class changes is small [2,8], while we will keep the class values the same for training purposes and change the data itself. Calders et al. [2] have also previously examined one method for changing the data in which different data items are given weights and the weights are adjusted to achieve fairness. In this category of work, as well, there is worry that the change to the data will decrease the classification accuracy, and Calders et al. have formalized this as a fairness/utility tradeoff [2]. We additionally note that lower classification accuracy may actually be the desired result, if that classification accuracy was due to discriminatory decision making in the past.

An important related work is the approach of "fairness through awareness" of Dwork et al. [3] and Zemel et al. [25]. Dwork et al. [3] focus on the problem of *individual* fairness; their approach posits the existence of a similarity measure between individual entities and seeks to find classifiers that ensure similar outcomes on individuals that are similar, via a Lipschitz condition. In the work of Zemel et al. [25], this idea of protecting individual fairness is combined with a statistical group-based fairness criterion that is similar to the approach we take in this work. A key contribution of their work is that they learn a modified representation of the data in which fairness is ensured while attempting to preserve fidelity with the original classification task. While this group fairness measure is similar to ours in spirit, it does not match the legal definition we base our work on. Another paper that also (implicitly) defines fairness on an individual basis is the work by Thanh et al. [11]. Their proposed repair mechanism changes class attributes of the data (rather than the data itself).

Pedreschi, Ruggieri and Turini [14,15] have examined the "80% rule" that we study in this paper as part of a larger class of measures based on a classifier's confusion matrix.

3. DISPARATE IMPACT AND ERROR RATES

We start by reinterpreting the "80% rule" in terms of more standard statistical measures of quality of a classifier. This presents notational challenges. The terminology of "right" and "wrong", "positive" and "negative" that is used in classification is an awkward fit when dealing with majority and minority classes, and selection decisions. For *notational convenience only*, we will use the convention that the protected class X takes on two values: $X = 0$ for the "minority" class and $X = 1$ for the "default" class. For example, in most gender-discrimination scenarios the value 0 would be assigned to "female" and 1 to "male". We will denote a successful binary classification outcome C (say, a hiring decision) by $C = \text{YES}$

and a failure by $C = \text{NO}$. Finally, we will map the majority class to “positive” examples and the minority class to “negative” examples with respect to the classification outcome, all the while reminding the reader that this is merely a convenience to do the mapping, and does not reflect any judgments about the classes. The advantage of this mapping is that it renders our results more intuitive: a classifier with high “error” will also be one that is least biased, because it is unable to distinguish the two classes.

Table 1 describes the *confusion matrix* for a classification with respect to the above attributes where each entry is the probability of that particular pair of outcomes for data sampled from the input distribution (we use the empirical distribution when referring to a specific data set).

Outcome	$X = 0$	$X = 1$
$C = \text{NO}$	a	b
$C = \text{YES}$	c	d

Table 1: A confusion matrix

The 80% rule can then be quantified as:

$$\frac{c/(a+c)}{d/(b+d)} \geq 0.8$$

Note that the traditional notion of “accuracy” includes terms in the numerator from both columns, and so cannot be directly compared to the 80% rule. Still, other class-sensitive error metrics are known, and more directly relate to the 80% rule:

DEFINITION 3.1 (CLASS-CONDITIONED ERROR METRICS). *The sensitivity of a test (informally, its true positive rate) is defined as the conditional probability of returning YES on “positive” examples (a.k.a. the majority class). In other words,*

$$\text{sensitivity} = \frac{d}{b+d}$$

The specificity of a test (its true negative rate) is defined as the conditional probability of returning NO on “negative” examples (a.k.a. the minority) class. I.e.,

$$\text{specificity} = \frac{a}{a+c}$$

DEFINITION 3.2 (LIKELIHOOD RATIO (POSITIVE)). *The likelihood ratio positive, denoted by LR_+ , is given by*

$$\text{LR}_+(C, X) = \frac{\text{sensitivity}}{1 - \text{specificity}} = \frac{d/(b+d)}{c/(a+c)}$$

We can now restate the 80% rule in terms of a data set.

DEFINITION 3.3 (DISPARATE IMPACT). *A data set has disparate impact if*

$$\text{LR}_+(C, X) > \frac{1}{\tau} = 1.25$$

It will be convenient to work with the reciprocal of LR_+ , which we denote by

$$\text{DI} = \frac{1}{\text{LR}_+(C, X)}.$$

This will allow us to discuss the value associated with disparate impact before the threshold is applied.

Multiple classes. Disparate impact is defined only for two classes. In general, one might imagine a multivalued class

attribute (for example, like ethnicity). In this paper, we will assume that a multivalued class attribute has one value designated as the “default” or majority class, and will compare each of the other values pairwise to this default class. While this ignores zero-sum effects between the different class values, it reflects the current binary nature of legal thought on discrimination. A more general treatment of joint discrimination among multiple classes is beyond the scope of this work.

4. COMPUTATIONAL FAIRNESS

Our notion of computational fairness starts with two players, Alice and Bob. Alice runs an algorithm \mathcal{A} that makes decisions based on some input. For example, Alice may be an employer using \mathcal{A} to decide who to hire. Specifically, \mathcal{A} takes a data set D with protected attribute X and unprotected attributes Y and makes a (binary) decision C . By law, Alice is not allowed to use X in making decisions, and claims to use only Y . It is Bob’s job to verify that on the data D , Alice’s algorithm \mathcal{A} is not liable for a claim of disparate impact.

Trust model. We assume that Bob does not have access to \mathcal{A} . Further, we assume that Alice has good intentions: specifically, that Alice is not secretly using X in \mathcal{A} while lying about it. While assuming Alice is lying about the use of X might be more plausible, it is much harder to detect. More importantly, from a functional perspective, it does not matter whether Alice uses X explicitly or uses proxy attributes Y that have the same effect: this is the core message from the Griggs case that introduced the doctrine of disparate impact. In other words, our *certification* process is indifferent to Alice’s intentions, but our *repair* process will assume good faith.

We summarize our main idea with the following intuition:

If Bob cannot predict X given the other attributes of D , then \mathcal{A} is fair with respect to Bob on D .

4.1 Predictability and Disparate Impact

We now present a formal definition of predictability and link it to the legal notion of disparate impact. Recall that $D = (X, Y, C)$ where X is the protected attribute, Y is the remaining attributes, and C is the class outcome to be predicted.

The basis for our formulation is a procedure that predicts X from Y . We would like a way to measure the quality of this predictor in a way that a) can be optimized using standard predictors in machine learning and b) can be related to LR_+ . The standard notions of accuracy of a classifier fail to do the second (as discussed earlier) and using LR_+ directly fails to satisfy the first constraint.

The error measure we seek turns out to be the *balanced error rate* BER.

DEFINITION 4.1 (BER). *Let $f : Y \rightarrow X$ be a predictor of X from Y . The balanced error rate BER of f on distribution \mathcal{D} over the pair (X, Y) is defined as the (unweighted) average class-conditioned error of f . In other words,*

$$\text{BER}(f(Y), X) = \frac{\Pr[f(Y) = 0 | X = 1] + \Pr[f(Y) = 1 | X = 0]}{2}$$

DEFINITION 4.2 (PREDICTABILITY). *X is said to be ϵ -predictable from Y if there exists a function $f : Y \rightarrow X$ such that*

$$\text{BER}(f(Y), X) \leq \epsilon.$$

This motivates our definition of ϵ -fairness, as a data set that is *not* predictable.

DEFINITION 4.3 (ϵ -FAIRNESS). A data set $D = (X, Y, C)$ is said to be ϵ -fair if for any classification algorithm $f: Y \rightarrow X$

$$\text{BER}(f(Y), X) > \epsilon$$

with (empirical) probabilities estimated from D .

Recall the definition of disparate impact from Section 3. We will be interested in examining the potential disparate impact of a classifier $g: Y \rightarrow C$ and will consider the value $\text{DI}(g) = 1/\text{LR}_+(g(Y), X)$ as it relates to the threshold τ . Where g is clear from context, we will refer to this as DI.

The justification of our definition of fairness comes from the following theorem:

THEOREM 4.1. A data set is $(1/2 - \beta/8)$ -predictable if and only if it admits disparate impact, where β is the fraction of elements in the minority class ($X = 0$) that are selected ($C = 1$).

PROOF. We will start with the direction showing that disparate impact implies predictability. Suppose that there exists some function $g: Y \rightarrow C$ such that $\text{LR}_+(g(y), c) \geq \frac{1}{\tau}$. We will create a function $\psi: C \rightarrow X$ such that $\text{BER}(\psi(g(y)), x) < \epsilon$ for $(x, y) \in D$. Thus the combined predictor $f = \psi \circ g$ satisfies the definition of predictability.

Consider the confusion matrix associated with g , depicted in Table 2. Set $\alpha \triangleq \frac{b}{b+d}$ and $\beta \triangleq \frac{c}{a+c}$. Then we can write

Prediction	$X = 0$	$X = 1$
$g(y) = \text{NO}$	a	b
$g(y) = \text{YES}$	c	d

Table 2: Confusion matrix for g

$$\text{LR}_+(g(y), X) = \frac{1-\alpha}{\beta} \text{ and } \text{DI}(g) = \frac{\beta}{1-\alpha}.$$

We define the *purely biased* mapping $\psi: C \rightarrow X$ as $\psi(\text{YES}) = 1$ and $\psi(\text{NO}) = 0$. Finally, let $\phi: Y \rightarrow X = \psi \circ g$. The confusion matrix for ϕ is depicted in Table 3. Note that the confusion matrix for ϕ is identical to the matrix for g .

Prediction	$X = 0$	$X = 1$
$\phi(Y) = 0$	a	b
$\phi(Y) = 1$	c	d

Table 3: Confusion matrix for ϕ

We can now express $\text{BER}(\phi)$ in terms of this matrix. Specifically, $\text{BER}(\phi) = \frac{\alpha+\beta}{2}$.

Representations. We can now express contours of the DI and BER functions as curves in the unit square $[0, 1]^2$. Reparametrizing $\pi_1 = 1 - \alpha$ and $\pi_0 = \beta$, we can express the error measures as $\text{DI}(g) = \frac{\pi_0}{\pi_1}$ and $\text{BER}(\phi) = \frac{1+\pi_0-\pi_1}{2}$

As a consequence, any classifier g with $\text{DI}(g) = \delta$ can be represented in the $[0, 1]^2$ unit square as the line $\pi_1 = \pi_0/\delta$. Any classifier ϕ with $\text{BER}(\phi) = \epsilon$ can be written as the function $\pi_1 = \pi_0 + 1 - 2\epsilon$.

Let us now fix the desired DI threshold τ , corresponding to the line $\pi_1 = \pi_0/\tau$. Notice that the region $\{(\pi_0, \pi_1) \mid \pi_1 \geq \pi_0/\tau\}$ is the region where one would make a finding of disparate impact (for $\tau = 0.8$).

Now given a classification that admits a finding of disparate impact, we can compute β . Consider the point $(\beta, \beta/\tau)$ at which the line $\pi_0 = \beta$ intersects the DI curve $\pi_1 = \pi_0/\tau$. This

point lies on the BER contour $(1 + \beta - \beta/\tau)/2 = \epsilon$, yielding $\epsilon = 1/2 - \beta(\frac{1}{\tau} - 1)/2$. In particular, for the DI threshold of $\tau = 0.8$, the desired BER threshold is

$$\epsilon = \frac{1}{2} - \frac{\beta}{8}$$

and so disparate impact implies predictability.

With this infrastructure in place, the other direction of the proof is now easy. To show that predictability implies disparate impact, we will use the same idea of a purely biased classifier. Suppose there is a function $f: Y \rightarrow X$ such that $\text{BER}(f(y), x) \leq \epsilon$. Let $\psi^{-1}: X \rightarrow C$ be the inverse purely biased mapping, i.e. $\psi^{-1}(1) = \text{YES}$ and $\psi^{-1}(0) = \text{NO}$. Let $g: Y \rightarrow C = \psi^{-1} \circ f$. Using the same representation as before, this gives us $\pi_1 \geq 1 + \pi_0 - 2\epsilon$ and therefore

$$\frac{\pi_0}{\pi_1} \leq \frac{\pi_0}{1 + \pi_0 - 2\epsilon} = 1 - \frac{1 - 2\epsilon}{\pi_0 + 1 - 2\epsilon}$$

Recalling that $\text{DI}(g) = \frac{\pi_0}{\pi_1}$ and that $\pi_0 = \beta$ yields $\text{DI}(g) \leq 1 - \frac{1-2\epsilon}{\beta+1-2\epsilon} = \tau$. For $\tau = 0.8$, this again gives us a desired BER threshold of $\epsilon = \frac{1}{2} - \frac{\beta}{8}$. \square

Note that as ϵ approaches $1/2$ the bound tends towards the trivial (since any binary classifier has BER at most $1/2$). In other words, as β tends to 0, the bound becomes vacuous.

This points to an interesting line of attack to evade a disparate impact finding. Note that β is the (class conditioned) rate at which members of the protected class are selected. Consider now a scenario where a company is being investigated for discriminatory hiring practices. One way in which the company might defeat such a finding is by interviewing (but not hiring) a large proportion of applicants from the protected class. This effectively drives β down, and the observation above says that in this setting their discriminatory practices will be harder to detect, because our result can not guarantee that a classifier will have error significantly less than 0.5.

Observe that in this analysis we use an extremely weak classifier to prove the existence of a relation between predictability and disparate impact. It is likely that using a better classifier (for example the Bayes optimal classifier or even a classifier that optimizes BER) might yield a stronger relationship between the two notions.

Dealing with uncertainty. In general, β might be hard to estimate from a fixed data set, and in practice we might only know that the true value of β lies in a range $[\beta_\ell, \beta_u]$. Since the BER threshold varies monotonically with β , we can merely use β_ℓ to obtain a conservative estimate.

Another source of uncertainty is in the BER estimate itself. Suppose that our classifier yields an error that lies in a range $[\gamma, \gamma']$. Again, because of monotonicity, we will obtain an interval of values $[\tau, \tau']$ for DI. Note that if (using a Bayesian approach) we are able to build a distribution over BER, this distribution will then transfer over to the DI estimate as well.

4.2 Certifying (lack of) DI with SVMs

The above argument gives us a way to determine whether a data set is potentially amenable to disparate impact (in other words, whether there is insufficient information to detect a protected attribute from the provided data).

Algorithm. We run a classifier that optimizes BER on the given data set, attempting to predict the protected attributes X from

the remaining attributes Y . Suppose the error in this prediction is ϵ . Then using the estimate of β from the data, we can substitute this into the equation above and obtain a threshold ϵ' . If $\epsilon' < \epsilon$, then we can declare the data set free from disparate impact.

Assume that we have an optimal classifier with respect to BER. Then we know that all classifiers will incur a BER of at least ϵ . By Theorem 4.1, this implies that no classifier on D will exhibit disparate impact, and so our certification is correct.

The only remaining question is what classifier is used by this algorithm. The usual way to incorporate class sensitivity into a classifier is to use different costs for misclassifying points in different classes. A number of class-sensitive cost measures fall into this framework, and there are algorithms for optimizing these measures (see [12] for a review), as well as a general (but expensive) method due to Joachims that does a clever grid search over a standard SVM to optimize a large family of class-sensitive measures [7]. Oddly, BER is not usually included among the measures studied.

Formally, as pointed out by Zhao et al [26], BER is not a cost-sensitive classification error measure because the weights assigned to class-specific misclassification depend on the relative class sizes (so they can be normalized). However, for any given data set we know the class sizes and can reweight accordingly. We adapt a standard hinge-loss SVM to incorporate class-sensitivity and optimize for (regularized) BER. This adaptation is standard, and yields a cost function that can be optimized using AdaBoost. For more details, see the extended version of this paper [5].

5. REMOVING DISPARATE IMPACT

Once Bob’s certification procedure has made a determination of (potential) disparate impact on D , Alice might request a *repaired* version \bar{D} of D , where any attributes in D that could be used to predict X have been changed so that \bar{D} would be certified as ϵ -fair. We now describe how to construct such a set $\bar{D} = (X, \bar{Y}, C)$ such that \bar{D} does not have disparate impact in terms of protected attribute X . While for notational simplicity we will assume that X is used directly in what follows, in practice the attribute used to stratify the data for repair need not directly be the protected attribute or even a single protected attribute. In the case of the Texas Top 10% Rule that admits the top ten percent of every high school class in Texas to the University of Texas [22], the attribute used to stratify is the high school attended, which is an attribute that correlates with race. If repair of multiple protected attributes is desired, the joint distribution can be used to stratify the data. (We will look into the effects of this experimentally in Section 6.2.)

Of course, it is important to change the data in such a way that predicting the class is still possible. Specifically, our goal will be to preserve the relative per-attribute ordering as follows. Given protected attribute X and a *single* numerical attribute Y , let $Y_x = \Pr(Y|X = x)$ denote the marginal distribution on Y conditioned on $X = x$. Let $F_x : Y_x \rightarrow [0, 1]$ be the cumulative distribution function for values $y \in Y_x$ and let $F_x^{-1} : [0, 1] \rightarrow Y_x$ be the associated quantile function (i.e. $F_x^{-1}(1/2)$ is the value of y such that $\Pr(Y \geq y|X = x) = 1/2$). We will say that F_x *ranks* the values of Y_x .

Let \bar{Y} be the repaired version of Y in \bar{D} . We will say that \bar{D} *strongly preserves rank* if for any $y \in Y_x$ and $x \in X$, its “repaired” counterpart $\bar{y} \in \bar{Y}_x$ has $F_x(y) = F_x(\bar{y})$. Strongly preserving rank in this way, despite changing the true values




Figure 1: Consider the fake probability density functions shown here where the blue curve shows the distribution of SAT scores (Y) for $X = \text{female}$, with $\mu = 550, \sigma = 100$, while the red curve shows the distribution of SAT scores for $X = \text{male}$, with $\mu = 400, \sigma = 50$. The resulting fully repaired data is the distribution in black, with $\mu = 475, \sigma = 75$. Male students who originally had scores in the 95th percentile, i.e., had scores of 500, are given scores of 625 in the 95th percentile of the new distribution in \bar{Y} , while women with scores of 625 in \bar{Y} originally had scores of 750.

of Y , appears to allow Alice’s algorithm to continue choosing stronger (higher ranked) applicants over weaker ones. We present experimental evidence for this in Section 6.

With this motivation, we now give a repair algorithm that strongly preserves rank and ensures that $\bar{D} = (X, \bar{Y}, C)$ is fair (i.e., is ϵ -fair for $\epsilon = 1/2$). In the discussion that follows, for the sake of clarity we will treat Y as a single attribute over a totally-ordered domain. To handle multiple totally-ordered attributes Y_1, \dots, Y_k we will repair each attribute individually.

We define a “median” distribution A in terms of its quantile function F_A^{-1} : $F_A^{-1}(u) = \text{median}_{x \in X} F_x^{-1}(u)$. The choice of the term “median” is not accidental.

LEMMA 5.1. *Let A be a distribution such that $F_A^{-1}(u) = \text{median}_{x \in X} F_x^{-1}(u)$. Then A is also the distribution minimizing $\sum_{x \in X} d(Y_x, C)$ over all distributions C , where $d(\cdot, \cdot)$ is the earthmover distance on \mathbb{R} .*

PROOF. For any two distributions P and Q on the line, the earthmover distance (using the underlying Euclidean distance $d(x, y) = |x - y|$ as the metric) can be written as

$$d(P, Q) = \int_0^1 |F_P^{-1}(u) - F_Q^{-1}(u)| du$$

In other words, the map $P \mapsto F_P^{-1}$ is an isometric embedding of the earthmover distance into ℓ_1 .

Consider now a set of points $p_1, \dots, p_n \in \ell_1$. Their 1-median – the point minimizing $\sum_i \|p_i - c\|_1$ – is the point whose j^{th} coordinate is the median of the j^{th} coordinates of the p_i . This is precisely the definition of the distribution A (in terms of F_A^{-1}). \square

Algorithm. Our repair algorithm creates \bar{Y} , such that for all $y \in Y_x$, the corresponding $\bar{y} = F_A^{-1}(F_x(y))$. The resulting $\bar{D} = (X, \bar{Y}, C)$ changes only Y while the protected attribute and class remain the same as in the original data, thus preserving the ability to predict the class. See Figure 1 for an example.

Notes. We note that this definition is reminiscent of the method by which partial rankings are combined to form a total ranking. The rankings are “Kemeny”-ized by finding a ranking that minimizes the sum of distances to the original rankings. However, there is a crucial difference in our procedure. Rather than merely reorganizing the data into a total ranking, we are modifying the data to construct this consensus distribution.

THEOREM 5.1. \bar{D} is fair and strongly preserves rank.

PROOF. In order to show that \bar{D} strongly preserves rank, recall that we would like to show that $F_x(y) = F_x(\bar{y})$ for all $x \in X$, $\bar{y} \in \bar{Y}_x$, and $y \in Y_x$. Since, by definition of our algorithm, $\bar{y} = F_A^{-1}(F_x(y))$, we know that $F_x(\bar{y}) = F_x(F_A^{-1}(F_x(y)))$, so we would like to show that $F_x(F_A^{-1}(z)) = z$ for all $z \in [0, 1]$ and for all x . Recall that $F_A^{-1}(z) = \text{median}_{x \in X} F_x^{-1}(z)$.

Suppose the above claim is not true. Then there are two values $z_1 < z_2$ and some value x such that $F_x(F_A^{-1}(z_1)) > F_x(F_A^{-1}(z_2))$. That is, there is some x and two elements $y_1 = F_A^{-1}(z_1)$, $y_2 = F_A^{-1}(z_2)$ such that $y_1 > y_2$. Now we know that $y_1 = \text{median}_{x \in X} F_x^{-1}(z_1)$. Therefore, if $y_1 > y_2$ it must be that there are strictly less than $|X|/2$ elements of the set $\{F_x^{-1}(z_1) | x \in X\}$ below y_2 . But by the assumption that $z_1 < z_2$, we know that each element of $\{F_x^{-1}(z_1) | x \in X\}$ is above the corresponding element of $\{F_x^{-1}(z_2) | x \in X\}$ and there are $|X|/2$ elements of this latter set below y_2 by definition. Hence we have a contradiction and so a flip cannot occur, which means that the claim is true.

Note that the resulting \bar{Y}_x distributions are the same for all $x \in X$, so there is no way for Bob to differentiate between the protected attributes. Hence the algorithm is 1-fair. \square

This repair has the effect that if you consider the \bar{Y} values at some rank z , the probability of the occurrence of a data item with attribute $x \in X$ is the same as the probability of the occurrence of x in the full population. This informal observation gives the intuitive backing for the lack of predictability of X from \bar{Y} and, hence, the lack of disparate impact in the repaired version of the data.

5.1 Partial Repair

Since the repair process outlined above is likely to degrade Alice’s ability to classify accurately, she might want a *partially repaired* data set instead. This in effect creates a tradeoff between the ability to classify accurately and the fairness of the resulting data. This tradeoff can be achieved by simply moving each inverse quantile distribution only part way towards the median distribution. Let $\lambda \in [0, 1]$ be the amount of repair desired, where $\lambda = 0$ yields the unmodified data set and $\lambda = 1$ is the fully repaired version described above. Recall that $F_x : Y_x \rightarrow [0, 1]$ is the function giving the rank of y . The repair algorithm for $\lambda = 1$ creates \bar{Y} such that $\bar{y} = F_A^{-1}(F_x(y))$ where A is the median distribution.

In the partial repair setting we will be creating a differential distribution A_x for each protected value $x \in X$ and setting $\bar{y} = F_{A_x}^{-1}(F_x(y))$. Consider the ordered set of all y at rank u in their respective conditional distributions i.e the set $U(u) = \{F_x^{-1}(u) | x \in X\}$. We can associate with U the cumulative function $UF(u, y) = |\{y' \geq y | y \in U(u)\}| / |U(u)|$ and define the associated quantile function $UF^{-1}(u, \alpha) = y$ where

$UF(u, y) = \alpha$. We can restate the full repair algorithm in this formulation as follows: for any (x, y) , $\bar{y} = UF^{-1}(F_x(y), 1/2)$.

We now describe two different approaches to performing a partial repair, each with their own advantages and disadvantages. Intuitively, these repair methods differ in which space they operate in: the *combinatorial* space of ranks or the *geometric* space of values.

5.1.1 A Combinatorial Repair

The intuition behind this repair strategy is that each item, rather than being moved to the median of its associated distribution, is only moved part of the way there, with the amount moved being proportional (in rank) to its distance from the median.

DEFINITION 5.1 (COMBINATORIAL REPAIR). Fix an x and consider any pair (x, y) . Let $r = F_x(y)$ be the rank of y conditioned on $X = x$. Suppose that in the set $U(r)$ (the collection of all $y' \in Y$ with rank r in their respective conditional distributions) the rank of y is ρ . Then we replace y by $\bar{y} \in U(r)$ whose rank in $U(r)$ is $\rho' = \lfloor (1 - \lambda)\rho + \lambda/2 \rfloor$. Formally, $\bar{y} = UF^{-1}(r, \rho')$. We call the resulting data set \bar{D}_λ .

While this repair is intuitive and easy to implement, it does not satisfy the property of strong rank preservation. In other words, it is possible that two pairs (x, y) and (x, y') with $y > y'$ to be repaired in a way that $\bar{y} < \bar{y}'$. While this could potentially affect the quality of the resulting data (we discuss this in Section 6.2), it does not affect the fairness properties of the repair. Indeed, we formulate the fairness properties of this repair as a formal conjecture.

CONJECTURE 5.1. \bar{D}_λ is $f(\lambda)$ -fair for a monotone function f .

5.1.2 A Geometric Repair

The algorithm above has an easy-to-describe *operational* form. It does not however admit a *functional* interpretation as an optimization of a certain distance function, like the full repair. For example, it is not true that for $\lambda = 1/2$ the modified distributions \bar{Y} are equidistant (under the earthmover distance) between the original unrepaired distributions and the full repair. The algorithm we propose now does have this property, as well as possessing a simple operational form. The intuition is that rather than doing a linear interpolation in *rank space* between the original item and the fully repaired value, it does a linear interpolation in the original *data space*.

DEFINITION 5.2 (GEOMETRIC REPAIR). Let F_A be the cumulative distribution associated with A , the result performing a full repair on the conditional cumulative distributions as described in Section 5. Given a conditional distribution $F_x(y)$, its λ -partial repair is given by

$$\bar{F}_x^{-1}(\alpha) = (1 - \lambda)F_x^{-1}(\alpha) + \lambda(F_A)^{-1}(\alpha)$$

Linear interpolation allows us to connect this repair to the underlying earthmover distance between repaired and unrepaired distributions. In particular,

THEOREM 5.2. For any x , $d(Y_x, \bar{Y}_x) = \lambda d(Y_x, Y_A)$ where Y_A is the distribution on Y in the full repair, and \bar{Y}_x is the λ -partial repair. Moreover, the repair strongly preserves rank.

PROOF. The earthmover distance bound follows from the proof of Lemma 5.1 and the isometric mapping between the

earthmover distance between Y_x and \bar{Y}_x and the ℓ_1 distance between F_x and \bar{F}_x . Rank preservation follows by observing that the repair is a *linear* interpolation between the original data and the full repair (which preserves rank by Lemma 5.1). \square

5.2 Fairness / Utility Tradeoff

The reason partial repair may be desired is that increasing fairness may result in a loss of utility. Here, we make this intuition precise. Let $\bar{D}_\lambda = (X, \bar{Y}, C)$ be the partially repaired data set for some value of $\lambda \in [0, 1]$ as described above (where $\bar{D}_{\lambda=0} = D$). Let $\bar{g}_\lambda: \bar{Y} \rightarrow C$ be the classifier with the utility we are trying to measure.

DEFINITION 5.3 (UTILITY). *The utility of a classifier $\bar{g}_\lambda: \bar{Y} \rightarrow C$ with respect to some partially repaired data set \bar{D}_λ is*

$$\gamma(\bar{g}_\lambda, \bar{D}_\lambda) = 1 - \text{BER}(\bar{g}_\lambda(\bar{y}), c).$$

If the classifier $\bar{g}_{\lambda=0}: Y \rightarrow C$ has an error of zero on the unrepaired data, then the utility is 1. More commonly, $\gamma(\bar{g}_{\lambda=0}, \bar{D}_{\lambda=0}) < 1$. In our experiments, we will investigate how γ decreases as λ increases.

6. EXPERIMENTS

We will now consider the certification algorithm and repair algorithm’s fairness/utility tradeoff experimentally on three data sets. The first is the *Ricci data set* at the heart of the Ricci v. DeStefano case [21]. It consists of 118 test taker entries, each including information about the firefighter promotion exam taken (Lieutenant or Captain), the score on the oral section of the exam, the written score, the combined score, and the race of the test taker (black, white, or Hispanic). In our examination of the protected race attribute, we will group the black and Hispanic test takers into a single non-white category. The classifier originally used to determine which test takers to promote was the simple threshold classifier that allowed anyone with a combined score of at least 70% to be eligible for promotion [13]. Although the true number of people promoted was chosen from the eligible pool according to their ranked ordering and the number of slots available, for simplicity in these experiments we will describe all eligible candidates as having been promoted. We use a random two-thirds / one-third split for the training / test data.

The other two data sets we will use are from the UCI Machine Learning Repository². So that we can compare our results to those of Zemel et al. [25], we will use the same data sets and the same decisions about what constitutes a sensitive attribute as they do. First, we will look at the *German credit* data set, also considered by Kamiran and Calders [8]. It contains 1000 instances, each of which consists of 20 attributes and a categorization of that instance as GOOD or BAD. The protected attribute is Age. In the examination of this data set with respect to their discriminatory measure, Kamiran and Calders found that the most discrimination was possible when splitting the instances into YOUNG and OLD at age 25 [8]. We will discretize the data accordingly to examine this potential worst case. We use a random two-thirds / one-third split for the training / test data.

We also look at the *Adult income* data set, also considered by Kamishima et al. [10]. It contains 48,842 instances, each of which consists of 14 attributes and a categorization of that person as making more or less than \$50,000 per year. The

²<http://archive.ics.uci.edu/ml>

protected attribute we will examine is Gender. Race is also an attribute in the data, and it will be excluded for classification purposes, except for when examining the effects of having multiple protected attributes - in this case, race will be categorized as white and non-white. The training / test split given in the original data is also used for our experiments.

For each of these data sets, we look at a total of 21 versions of the data - the original data set plus 10 partially or fully repaired attribute sets for each of the combinatorial and geometric partial repairs. These are the repaired attributes for $\lambda \in [0, 1]$ at increments of 0.1. Data preprocessing was applied before the partial repair algorithm was run.

Preprocessing. Datasets were preprocessed as follows:

1. Remove all protected attributes from Y . This ensures that we are not trying to learn a classifier that depends on other protected attributes that might correlate with the target protected attribute. (The repair process does still get to know X .)
2. Remove all unordered categorical features since our repair procedure assumes that the space of values is ordered. Ordered categories are converted to integers.³
3. Scale each feature so that the minimum is zero and the maximum is one.

Classifiers. Three different classifiers were used as oracles for measuring discrimination (under the disparate impact measure and a measure by Zemel et al. [25]), and to test the accuracy of a classification after repair. The classifiers used for our experimental tasks were provided by the Scikit-learn⁴ python package.

LR: Logistic Regression: Liblinear’s [4] logistic regression algorithm for L2 regularization and logistic loss. The classifier was configured to weight the examples automatically so that classes were weighted equally.

SVM: Support Vector Machine: Liblinear’s [4] linear SVM algorithm for L2 regularization and L2 loss. The classifier was configured to weight the examples automatically so that classes were weighted equally.

GNB: Gaussian Naïve Bayes: Scikit-Learn’s naïve Bayes algorithm with a balanced class prior.

Parameter selection and cross-validation. LR and SVM classifiers were cross-validated using three-fold cross validation and the best parameter based on BER was chosen. We cross-validated the parameter controlling the tradeoff between regularization and loss, and 13 parameters between 10^{-3} and 10^3 , with logarithms uniformly spaced, were searched.

Repair details. The repair procedure requires a ranking of each attribute. The numeric values and ordered categorical attributes were ordered in the natural way and then quantiles were used as the ranks. Since the repair assumes that there is a point at each quantile value in each protected class, the

³On the Adult Income data, it happens that all missing values were of these unordered categorical columns, so no data sets had missing values after this step.

⁴<http://scikit-learn.org>.




Figure 2: Lack of predictability (BER) of the protected attributes on the German Credit Adult Income, and Ricci data sets as compared to the disparate impact found in the test set when the class is predicted from the non-protected attributes. The certification algorithm guarantees that points to the right of the BER threshold are also above $\tau = 0.8$, the threshold for legal disparate impact. For clarity, we only show results using the combinatorial repair, but the geometric repair results follow the same pattern.

quantiles were determined in the following way. For each attribute, the protected class with the smallest number of members was determined. This size determined how many quantile buckets to create. The other protected classes were then appropriately divided into the same number of quantile buckets, with the median value in each bucket chosen as a representative value for that quantile. Each quantile value in the fully repaired version is the median of the representative values for that quantile. The combinatorial partial repair determines all valid values for an attribute and moves the original data part way to the fully repaired data within this space. The geometric repair assumes all numeric values are allowed for the partial repair.

6.1 Certification

The goal in this section is to experimentally validate our certification algorithm, described in subsection 4.2. On each of the data sets described above, we attempt to predict the protected attribute from the remaining attributes. The resulting BER is compared to $DI(g)$ where $g : Y \rightarrow C$, i.e., the disparate impact value as measured when some classifier attempts to predict the class given the non-protected attributes. From the underlying data, we can calculate the BER threshold $\epsilon = 1/2 - \beta/8$. Above this threshold, any classifier applied to the data will have disparate impact. The threshold is chosen conservatively so as to preclude false positives (times when we falsely declare the data to be safe from disparate impact).

In Figure 2 we can see that there are no data points greater than the BER threshold and also much below $\tau = 0.8$, the threshold for legal disparate impact. The only false positives are a few points very close to the line. This is likely because the β value, as measured from the data, has some error. We can also see, from the points close to the BER threshold line

on its left but below τ that while we chose the threshold conservatively, we were not overly conservative. Still, using a classifier other than the purely biased one in the certification algorithm analysis might allow this threshold to be tightened.

The points in the upper left quadrant of these charts represent false negatives of our certification algorithm on a specific data set and a specific classifier. However, our certification algorithm guarantees lack of disparate impact over *any* classifier, so these are not false negatives in the traditional sense. In fact, when a single data set is considered over all classifiers, we see that all such data sets below the BER threshold have some classifier that has DI close to or below $\tau = 0.8$.

One seemingly surprising artifact in the charts is the vertical line in the Adult Income data chart at $BER = 0.5$ for the GNB repair. Recall that the chart is based off of two different confusion matrices - the BER comes from predicting gender while the disparate impact is calculated when predicting the class. In a two class system, the BER cannot be any higher than 0.5, so while the ability to predict the gender cannot get any worse, the resulting fairness of the class predictions can still improve, thus causing the vertical line in the chart.

6.2 Fairness / Utility Tradeoff

The goal in this section is to determine how much the partial repair procedure degrades utility. Using the same data sets as described above, we will examine how the utility (see Definition 5.3) changes DI (measuring fairness) increases. Utility will be defined with respect to the data labels. Note that this may itself be faulty data, in that the labels may not themselves provide the best possible utility based on the underlying, but perhaps unobservable, desired outcomes. For example, the results on the test from the Ricci data may not perfectly measure a firefighter’s ability and so outcomes based on that test may not correctly predict who should be promoted. Still, in the absence of knowledge of more precise data, we will use these labels to measure utility. For the Ricci data, which is unlabeled, we will assume that the true labels are those provided by the simple threshold classifier used on the non-repaired version of the Ricci data, i.e. that anyone with a score of at least 70% should pass the exam. Disparate impact (DI) for all data sets is measured with respect to the predicted outcomes on the test set as differentiated by protected attribute. The SVM described above is used to classify on the Adult Income and German Credit data sets while the Ricci data uses the simple threshold classifier. The utility $(1 - BER)$ shown is based on the confusion matrix of the original labels versus the labels predicted by these classifiers.

The results, shown in Figure 3, demonstrate the expected decay over utility as fairness increases. Each unrepaired data set begins with $DI < 0.8$, i.e., it would fail the 80% rule, and we are able to repair it to a legal value. For the Adult Income data set, repairing the data fully only results in a utility loss from about 74% to 72%, while for the German Credit data, repairing the data fully reduces the utility from about 72% to 50% - essentially random. We suspect that this difference in decay is inherent to the class decisions in the data set (and the next section will show that other existing fairness repairs face this same decay). We suspect that the lack of linearity in the utility decay in the German Credit data after it has fairness greater than $DI = 0.8$ is due to this low utility.

Looking more closely at the charts, we notice that some of the partially repaired data points have $DI > 1$. Since DI is calculated with respect to fixed majority and minority classes,




Figure 3: Disparate impact (DI) vs. utility (1-BER) from our combinatorial and geometric partial repair processes using the SVM to classify on the Adult Income and German Credit data sets and the simple threshold classifier on the Ricci data set. Recall that only points with $DI \geq \tau = 0.8$ are legal. $DI = 1.0$ represents full fairness.

this happens when the classifier has given a good outcome to proportionally more minority than majority class members. These points should be considered unfair to the majority class.

Figure 3 also shows that combinatorial and geometric repairs have similar DI and utility values for all partial repair data sets. This means that either repair can be used.

Multiple Protected Attributes. Our repair procedure can operate over the joint distribution of multiple protected attributes. To examine how this affects utility, we considered the Adult Income data set repaired by gender only, race only, and over both gender and race. For the repairs with respect to race, a binary racial categorization of white and non-white is used. Repairs with respect to both race and gender are taken over the joint distribution. In the joint distribution case, the DI calculated is the average of the DI of each of the three protected sets (white women, non-white men, and non-white women) with respect to the advantaged group (white men). The classifier used to predict the class from the non-protected attributes is the SVM described earlier.

The results, shown in Figure 4, show that the utility loss over the joint distribution is close to the maximum of the utility loss over each protected attribute considered on its own. In other words, the loss does not compound. These good results are likely due in part to the size of the data set allowing each subgroup to still be large enough. On such data sets, allowing all protected attributes to be repaired appears reasonable.

6.3 Comparison to previous work

Here, we compare our results to related work on the German credit data and Adult income data sets. Logistic regression is used as a baseline comparison, fair naive Bayes is the solution from Kamiran and Calders [8], regularized logistic regression is the repair method from Kamishima et al. [10], and learned fair representations is Zemel et al.’s solution [25]. All comparison data is taken from Zemel et al.’s implementations [25]. Zemel et al. define discrimination as $(1 - \alpha) - \beta$. So that increasing Zemel scores mean that fairness has increased, as is the case with DI, we will look at the *Zemel fairness* score

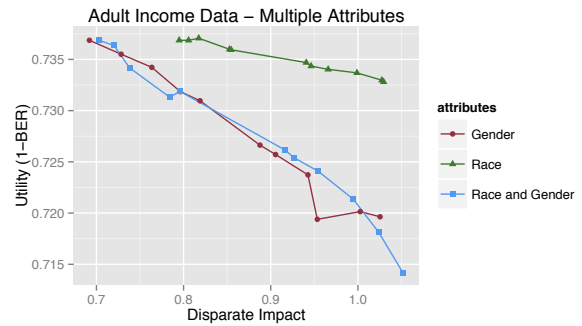


Figure 4: Disparate impact (DI) vs. utility (1-BER) from our combinatorial and geometric partial repair processes using the SVM as the classifier. For clarity in the figure, only the combinatorial repairs are shown, though the geometric repairs follow the same pattern.

which we define as $1 - ((1 - \alpha) - \beta) = 2 \cdot \text{BER}$. Accuracy is the usual rate of successful classification. Unlike the compared works, we do not choose a single partial repair point. Figure 5 shows our fairness and accuracy results for both combinatorial and geometric partial repairs for values of $\lambda \in [0, 1]$ at increments of 0.1 using all three classifiers described above.

Figure 5 shows that our method can be flexible with respect to the chosen classifier. Since the repair is done over the data, we can choose a classification algorithm appropriate to the data set. For example, on the Adult Income data set the repairs based on Naïve Bayes have better accuracy at high values of fairness than the repairs based on Logistic Regression. On the German and Adult data sets our results show that for any fairness value a partially repaired data set at that value can be chosen and a classifier applied to achieve accuracy that is better than competing methods.

Since the charts in Figure 5 include unrepaired data, we can also separate the effects of our classifier choices from the effects of the repair. In each classifier repair series, the data point with the lowest Zemel fairness (furthest to the left) is the original data. Comparing the original data point when the LR classifier was used to the LR classifier used by Zemel et al. as a comparison baseline, we see a large jump in both fairness and accuracy. Configuring the classifier to weight classes equally may have accounted for this improvement.

7. LIMITATIONS AND FUTURE WORK

Our experiments show a substantial difference in the performance of our repair algorithm depending on the specific algorithms we chose. Given the myriad classification algorithms used in practice, there is a clear need for a future systematic study of the relationship between dataset features, algorithms, and repair performance.

In addition, our discussion of disparate impact is necessarily tied to the legal framework as defined in United States law. It would be valuable in future work to collect the legal frameworks of different jurisdictions, and investigate whether a single unifying formulation is possible.

Finally, we note that the algorithm we present operates only on numerical attributes. Although we are satisfied with its performance, we chose this setting mostly for its relative theoretical simplicity. A natural avenue for future work is to investigate generalizations of our repair procedures for

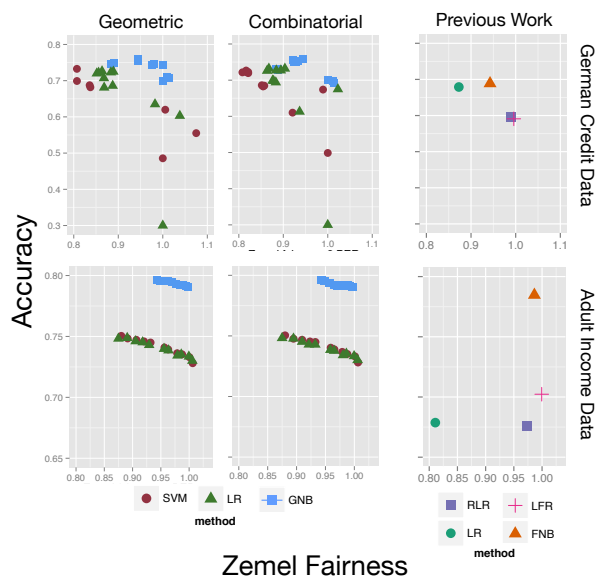


Figure 5: Zemel fairness vs. accuracy from our combinatorial and geometric partial repairs as compared to previous work. Legend: RLR, Regularized Logistic Regression [10]; LFR, Learned Fair Representations [25]; FNB, Fair Naïve Bayes [8]; GNB, Gaussian Naïve Bayes with balanced prior; LR, Logistic Regression; SVM, Support Vector Machine.

datasets with different attribute types, such as categorical data, vector-valued attributes, etc.

8. ACKNOWLEDGMENTS

This research was funded in part by the NSF under grant BIGDATA-1251049. Thanks to Deborah Karpatkin, David Robinson, and Natalie Shapero for helping us understand the legal interpretation of disparate impact. Any misunderstandings about these issues in this paper are our own. Thanks also to Mark Gould for pointing us to the Griggs v. Duke decision, which helped to set us down this path in the first place.

9. REFERENCES

- [1] S. Barocas and A. D. Selbst. Big data’s disparate impact. Technical report, available at SSRN: <http://ssrn.com/abstract=2477899>, 2014.
- [2] T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. In *ICDM Workshop Domain Driven Data Mining*, pages 13–18, 2009.
- [3] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proc. of Innovations in Theoretical Computer Science*, 2012.
- [4] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *J. of Machine Learning Research*, 9:1871–1874, 2008.
- [5] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. *CoRR*, abs/1412.3756, 2014.
- [6] H. Hodson. No one in control: The algorithms that run our lives. *New Scientist*, Feb. 04, 2015.
- [7] T. Joachims. A support vector method for multivariate performance measures. In *Proc. of Intl. Conf. on Machine Learning*, pages 377–384. ACM, 2005.

- [8] F. Kamiran and T. Calders. Classifying without discriminating. In *Proc. of the IEEE International Conference on Computer, Control and Communication*, 2009.
- [9] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware classifier with prejudice remover regularizer. *Machine Learning and Knowledge Discovery in Databases*, pages 35–50, 2012.
- [10] T. Kamishima, S. Akaho, and J. Sakuma. Fairness aware learning through regularization approach. In *Proc. of Intl. Conf. on Data Mining*, pages 643–650, 2011.
- [11] B. T. Luong, S. Ruggieri, and F. Turini. k-nn as an implementation of situation testing for discrimination discovery and prevention. In *Proc. of Intl. Conf. on Knowledge Discovery and Data Mining, KDD ’11*, pages 502–510, 2011.
- [12] A. Menon, H. Narasimhan, S. Agarwal, and S. Chawla. On the statistical consistency of algorithms for binary classification under class imbalance. In *Proc. 30th. ICM*, pages 603–611, 2013.
- [13] W. Miao. Did the results of promotion exams have a disparate impact on minorities? Using statistical evidence in Ricci v. DeStefano. *J. of Stat. Ed.*, 19(1), 2011.
- [14] D. Pedreschi, S. Ruggieri, and F. Turini. Integrating induction and deduction for finding evidence of discrimination. In *Proc. of Intl. Conf. on Artificial Intelligence and Law, ICAIL ’09*, pages 157–166, 2009.
- [15] D. Pedreschi, S. Ruggieri, and F. Turini. A study of top-k measures for discrimination discovery. In *Proc. of Symposium on Applied Computing, SAC ’12*, pages 126–131, 2012.
- [16] J. L. Peresie. Toward a coherent test for disparate impact discrimination. *Indiana Law Journal*, 84(3):Article 1, 2009.
- [17] J. Podesta, P. Pritzker, E. J. Moniz, J. Holdren, and J. Zients. Big data: seizing opportunities, preserving values. Executive Office of the President, May 2014.
- [18] A. Romei and S. Ruggieri. A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review*, pages 1–57, April 3 2013.
- [19] Supreme Court of the United States. Griggs v. Duke Power Co. 401 U.S. 424, March 8, 1971.
- [20] Supreme Court of the United States. Watson v. Fort Worth Bank & Trust. 487 U.S. 977, 995, 1988.
- [21] Supreme Court of the United States. Ricci v. DeStefano. 557 U.S. 557, 174, 2009.
- [22] Texas House of Representatives. House bill 588. *75th Legislature*, 1997.
- [23] The Leadership Conference. Civil rights principles for the era of big data. <http://www.civilrights.org/press/2014/civil-rights-principles-big-data.html>, Feb. 27, 2014.
- [24] The U.S. EEOC. Uniform guidelines on employee selection procedures, March 2, 1979.
- [25] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *Proc. of Intl. Conf. on Machine Learning*, pages 325–333, 2013.
- [26] M.-J. Zhao, N. Edakunni, A. Pocock, and G. Brown. Beyond Fano’s inequality: bounds on the optimal F-score, BER, and cost-sensitive risk and their implications. *J. of Machine Learning Research*, 14(1):1033–1090, 2013.