

Understanding and Mitigating Accuracy Disparity in Regression

Jianfeng Chi¹ Yuan Tian¹ Geoffrey J. Gordon² Han Zhao³

Abstract

With the widespread deployment of large-scale prediction systems in high-stakes domains, *e.g.*, face recognition, criminal justice, etc., disparity in prediction accuracy between different demographic subgroups has called for fundamental understanding on the source of such disparity and algorithmic intervention to mitigate it. In this paper, we study the accuracy disparity problem in regression. To begin with, we first propose an *error decomposition theorem*, which decomposes the accuracy disparity into the distance between marginal label distributions and the distance between conditional representations, to help explain why such accuracy disparity appears in practice. Motivated by this error decomposition and the general idea of distribution alignment with statistical distances, we then propose an algorithm to reduce this disparity, and analyze its game-theoretic optima of the proposed objective functions. To corroborate our theoretical findings, we also conduct experiments on five benchmark datasets. The experimental results suggest that our proposed algorithms can effectively mitigate accuracy disparity while maintaining the predictive power of the regression models.

1. Introduction

Recent progress in machine learning has led to its widespread use in many high-stakes domains, such as criminal justice, healthcare, student loan approval, and hiring. Meanwhile, it has also been widely observed that accuracy disparity could occur inadvertently under various scenarios in practice (Barocas and Selbst, 2016). For example, errors are inclined to occur for individuals of certain underrepresented demographic groups (Kim, 2016). In other

¹Department of Computer Science, University of Virginia
²Machine Learning Department, Carnegie Mellon University
³Department of Computer Science, University of Illinois at Urbana-Champaign. Correspondence to: Jianfeng Chi <jc6ub@virginia.com>, Han Zhao <hanzhao@illinois.edu>.

Proceedings of the 38th International Conference on Machine Learning, PMLR 139, 2021. Copyright 2021 by the author(s).

cases, Buolamwini and Gebru (2018) showed that notable accuracy disparity exists across different racial and gender demographic subgroups on several real-world image classification systems. Moreover, Bagdasaryan et al. (2019) found out that a differentially private model even exacerbates such accuracy disparity. Such accuracy disparity across demographic subgroups not only raises concerns in high-stake applications but also can be utilized by malicious parties to cause information leakage (Yaghini et al., 2019; Zhao et al., 2020).

Despite the ample needs of accuracy parity, most prior work limits its scope to studying the problem in binary classification settings (Hardt et al., 2016; Zafar et al., 2017b; Zhao and Gordon, 2019; Jiang et al., 2019). Compared to the accuracy disparity problem in classification settings, accuracy disparity in regression is a more challenging but less studied problem, due to the fact that many existing algorithmic techniques designed for classification cannot be extended in a straightforward way when the target variable is continuous (Zhao et al., 2019). In a seminal work, Chen et al. (2018) analyzed the impact of data collection on accuracy disparity in general learning models. They provided a descriptive analysis of such parity gaps and advocated for collecting more training examples and introducing more predictive variables. While such a suggestion is feasible in applications where data collection and labeling is cheap, it is not applicable in domains where it is time-consuming, expensive, or even infeasible to collect more data, *e.g.*, in autonomous driving, education, etc.

Our Contributions In this paper, we provide a prescriptive analysis of accuracy disparity and aim at providing algorithmic interventions to reduce the disparity gap between different demographic subgroups in the regression setting. To start with, we first formally characterize why accuracy disparity appears in regression problems by depicting the feasible region of the underlying group-wise errors. Next, we derive an error decomposition theorem that decomposes the accuracy disparity into the distance between marginal label distributions and the distance between conditional rep-

¹Technically, accuracy disparity refers to (squared) error difference in our paper. We would like to use accuracy disparity throughout our paper since it is a more commonly used term in fairness problems.

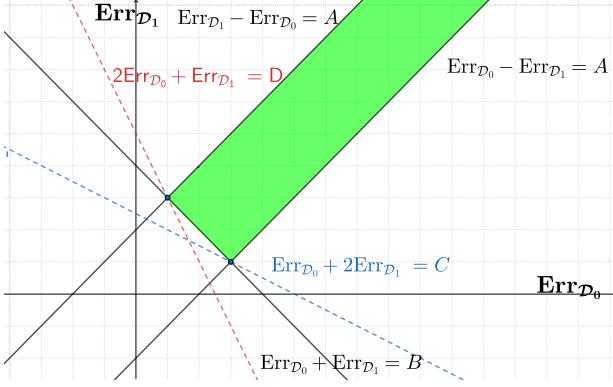


Figure 1: Geometric interpretation of accuracy disparity in regression. The green area corresponds to the feasible region of $\text{Err}_{\mathcal{D}_0}$ and $\text{Err}_{\mathcal{D}_1}$ under the hypothesis class \mathcal{H} . For any optimal hypothesis h which is solely designed to minimize the overall error, the best the hypothesis h can do is to intersect with one of the two bottom vertices of the green area, leading to accuracy disparity if the width of the feasible region is nonzero. See section 3.1 for more details.

resentations. We also provide a lower bound on the joint error across groups. Based on these results, we illustrate why regression models aiming to minimize the global loss will inevitably lead to accuracy disparity if the marginal label distributions or conditional representations differ across groups. See Figure 1 for illustration.

Motivated by the error decomposition theorem, we propose two algorithms to reduce accuracy disparity via joint distribution alignment with the total variation distance and the Wasserstein distance, respectively. Furthermore, we analyze the game-theoretic optima of the objective functions and illustrate the principle of our algorithms from a game-theoretic perspective. To corroborate the effectiveness of our proposed algorithms in reducing accuracy disparity, we conduct experiments on five benchmark datasets. Experimental results suggest that our proposed algorithms help to mitigate accuracy disparity while maintaining the predictive power of the regression models. We believe our theoretical results contribute to the understanding of why accuracy disparity occurs in machine learning models, and the proposed algorithms provides an alternative for intervention in real-world scenarios where accuracy parity is desired but collecting more data/features is time-consuming or infeasible.

2. Preliminaries

Notation We use $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{Y} \subseteq \mathbb{R}$ to denote the input and output space. We use X and Y to denote random variables which take values in \mathcal{X} and \mathcal{Y} , respectively. Lower case letters \mathbf{x} and y denote the instantiation of X and Y . We use $H(X)$ to denote the Shannon entropy of random

variable X , $H(X | Y)$ to denote the conditional entropy of X given Y , and $I(X; Y)$ to denote the mutual information between X and Y . To simplify the presentation, we use $A \in \{0, 1\}$ as the sensitive attribute, *e.g.*, gender, race, etc. Let \mathcal{H} be the hypothesis class of regression models. In other words, for $h \in \mathcal{H}$, $h : \mathcal{X} \rightarrow \mathcal{Y}$ is a predictor. Note that even if the predictor does not explicitly take the sensitive attribute A as an input variable, the prediction can still be biased due to the correlations with other input variables. In this work we study the stochastic setting where there is a joint distribution \mathcal{D} over X, Y and A from which the data are sampled. For $a \in \{0, 1\}$ and $y \in \mathbb{R}$, we use \mathcal{D}_a to denote the conditional distribution of \mathcal{D} given $A = a$ and \mathcal{D}^y to denote the conditional distribution of \mathcal{D} given $Y = y$. For an event E , $\mathcal{D}(E)$ denotes the probability of E under \mathcal{D} . Given a feature transformation function $g : \mathcal{X} \rightarrow \mathcal{Z}$ that maps instances from the input space \mathcal{X} to feature space \mathcal{Z} , we define $g_{\sharp}\mathcal{D} := \mathcal{D} \circ g^{-1}$ to be the induced (pushforward) distribution of \mathcal{D} under g , *i.e.*, for any event $E' \subseteq \mathcal{Z}$, $g_{\sharp}\mathcal{D}(E') := \mathcal{D}(\{x \in \mathcal{X} \mid g(x) \in E'\})$. We define $(\cdot)_+$ to be $\max\{\cdot, 0\}$.

For regression problems, given a joint distribution \mathcal{D} , the error of a predictor h under \mathcal{D} is defined as $\text{Err}_{\mathcal{D}}(h) := \mathbb{E}_{\mathcal{D}}[(Y - h(X))^2]$. To make the notation more compact, we may drop the subscript \mathcal{D} when it is clear from the context. Furthermore, we also use $\text{MSE}_{\mathcal{D}}(\hat{Y}, Y)$ to denote the mean squared loss between the predicted variable $\hat{Y} = h(X)$ and the true label Y over the joint distribution \mathcal{D} . Similarly, we also use $\text{CE}_{\mathcal{D}}(A \parallel \hat{A})$ to denote the cross-entropy loss between the predicted variable \hat{A} and the true label A over the joint distribution \mathcal{D} . Throughout the paper, we make the following standard boundedness assumption:

Assumption 2.1. There exists $M > 0$, such that for any hypothesis $\mathcal{H} \ni h : \mathcal{X} \rightarrow \mathcal{Y}$, $\|h\|_{\infty} \leq M$ and $|Y| \leq M$.

Problem Setup Our goal is to learn a regression model that is fair in the sense that the errors of the regressor are approximately equal across the groups given by the sensitive attribute A . We assume that the sensitive attribute A is only available to the learner during the training phase and is not visible during the inference phase. We would like to point out that there are many other different and important definitions of fairness (Narayanan, 2018) even in the sub-category of group fairness, and our discussion is by no means comprehensive. For example, two frequently used definitions of fairness in the literature are the so-called statistical parity (Dwork et al., 2012) and equalized odds (Hardt et al., 2016). Nevertheless, throughout this paper we mainly focus accuracy parity as our fairness notion, due to the fact that machine learning systems have been shown to exhibit substantial accuracy disparities between different demographic subgroups (Barocas and Selbst, 2016; Kim, 2016; Buolamwini and Gebru, 2018). This observation has already

brought huge public attention (*e.g.*, see New York Times, The Verge, and Insurance Journal) and calls for machine learning systems that (at least approximately) satisfy accuracy parity. For example, in a healthcare spending prediction system, stakeholders do not want the prediction error gaps to be too large among different demographic subgroups. Formally, accuracy parity is defined as follows:

Definition 2.1. Given a joint distribution \mathcal{D} , a predictor h satisfies accuracy parity if $\text{Err}_{\mathcal{D}_0}(h) = \text{Err}_{\mathcal{D}_1}(h)$.

In practice the exact equality of accuracy between two groups is often hard to ensure, so we define *error gap* to measure how well the model satisfies accuracy parity:

Definition 2.2. Given a joint distribution \mathcal{D} , the error gap of a hypothesis h is $\Delta_{\text{Err}}(h) := |\text{Err}_{\mathcal{D}_0}(h) - \text{Err}_{\mathcal{D}_1}(h)|$.

By definition, if a model satisfies accuracy parity, $\Delta_{\text{Err}}(h)$ will be zero. Next we introduce two distance metrics that will be used in our theoretical analysis and algorithm design:

- Total variation distance: it measures the largest possible difference between the probabilities that the two probability distributions can assign to the same event E . We use $d_{\text{TV}}(\mathcal{P}, \mathcal{Q})$ to denote the total variation:

$$d_{\text{TV}}(\mathcal{P}, \mathcal{Q}) := \sup_E |\mathcal{P}(E) - \mathcal{Q}(E)|.$$

- Wasserstein distance: the Wasserstein distance between two probability distributions is

$$W_1(\mathcal{P}, \mathcal{Q}) = \sup_{f \in \{f: \|f\|_L \leq 1\}} \left| \int_{\Omega} f d\mathcal{P} - \int_{\Omega} f d\mathcal{Q} \right|,$$

where $\|f\|_L$ is the Lipschitz semi-norm of a real-valued function of f and Ω is the sample space over which two probability distributions \mathcal{P} and \mathcal{Q} are defined. By the Kantorovich-Rubinstein duality theorem (Villani 2008), we recover the primal form of the Wasserstein distance, defined as

$$W_1(\mathcal{P}, \mathcal{Q}) := \inf_{\gamma \in \Gamma(\mathcal{P}, \mathcal{Q})} \int d(X, Y) d\gamma(X, Y),$$

where $\Gamma(\mathcal{P}, \mathcal{Q})$ denotes the collection of all couplings of \mathcal{P} and \mathcal{Q} , and X and Y denote the random variables with law \mathcal{P} and \mathcal{Q} respectively. Throughout this paper we use L_1 distance for $d(\cdot, \cdot)$, but extensions to other distances, *e.g.*, L_2 distance, is straightforward.

3. Main Results

In this section, we first characterize why accuracy disparity arises in regression models. More specifically, given a hypothesis $h \in \mathcal{H}$, we first prove a lower bound of joint errors. Then, we provide an error decomposition theorem

which upper bounds the accuracy disparity and decompose it into the distance between marginal label distributions and the distance between conditional representations. Based on these results, we give a geometric interpretation to visualize the feasible region of $\text{Err}_{\mathcal{D}_0}$ and $\text{Err}_{\mathcal{D}_1}$, and illustrate how error gap arises when learning a hypothesis h that minimizes the global square error. Motivated by the error decomposition theorem, we propose two algorithms to reduce accuracy disparity, connect the game-theoretic optima of the objective functions in our algorithms with our theorems, and describe the practical implementations of the algorithms. Due to the space limit, we defer all the detailed proofs to the appendix.

3.1. Bounds on Conditional Errors and Accuracy Disparity Gap

Before we provide the prescriptive analysis of the accuracy disparity problem in regression, it is natural to ask whether accuracy parity is achievable in the first place. Hence, we first provide a sufficient condition to achieve accuracy parity in regression.

Proposition 3.1. Assume both $\mathbb{E}_{\mathcal{D}_a}[Y]$ and $\mathbb{E}_{\mathcal{D}_a}[Y^2]$ are equivalent for any $A = a$, then using a constant predictor ensures accuracy parity in regression.

Proposition 3.1 states if the first two order moments of marginal label distributions are equal across different groups, then using a constant predictor leads to accuracy parity in regression. Proposition 3.1 is a relaxation of our proposed error decomposition theorem (Theorem 3.2) which requires the total variation distance between group-wise marginal label distributions to be zero. However, the condition rarely holds in real-world scenarios and it does not provide any insights to algorithm design. Next we provide more in-depth analysis to understand why accuracy disparity appears in regression models and provide algorithm interventions to mitigate the problem.

When we learn a predictor, the prediction function induces $X \xrightarrow{h} \hat{Y}$, where \hat{Y} is the predicted target variable given by hypothesis h . Hence for any distribution \mathcal{D}_0 (\mathcal{D}_1) of X , the predictor also induces a distribution $h_{\sharp}\mathcal{D}_0$ ($h_{\sharp}\mathcal{D}_1$) of \hat{Y} . Recall that the Wasserstein distance is metric, hence the following chain of triangle inequalities holds:

$$W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) \leq W_1(\mathcal{D}_0(Y), h_{\sharp}\mathcal{D}_0) + W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1) + W_1(h_{\sharp}\mathcal{D}_1, \mathcal{D}_1(Y))$$

Intuitively, $W_1(\mathcal{D}_a(Y), h_{\sharp}\mathcal{D}_a)$ measures the distance between the true marginal label distribution and the predicted one when $A = a$. This distance is related to the prediction error of function h conditioned on $A = a$:

Lemma 3.1. Let $\hat{Y} = h(X)$, then for $a \in \{0, 1\}$, $W_1(\mathcal{D}_a(Y), h_{\sharp}\mathcal{D}_a) \leq \sqrt{\text{Err}_{\mathcal{D}_a}(h)}$.

Now we can get the following theorem that characterizes the lower bound of joint error on different groups:

Theorem 3.1. Let $\hat{Y} = h(X)$ be the predicted variable, then $\text{Err}_{\mathcal{D}_0}(h) + \text{Err}_{\mathcal{D}_1}(h) \geq \frac{1}{2}[(W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1))_+]^2$.

In Theorem 3.1, we see that if the difference between marginal label distributions across groups is large, then statistical parity could potentially lead to a large joint error. Moreover, Theorem 3.1 could be extended to give a lower bound on the joint error incurred by h as well:

Corollary 3.1. Let $\hat{Y} = h(X)$ and $\alpha = \mathcal{D}(A = 0) \in [0, 1]$, we have $\text{Err}_{\mathcal{D}}(h) \geq \frac{1}{2} \min\{\alpha, 1 - \alpha\} \cdot [(W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1))_+]^2$.

Now we upper bound the error gap. We first relate the error gap to marginal label distributions and the predicted distributions conditioned on $Y = y$:

Theorem 3.2. If Assumption 2.1 holds, then for $\forall h \in \mathcal{H}$, let $\hat{Y} = h(X)$, the following inequality holds:

$$\begin{aligned} \Delta_{\text{Err}}(h) &\leq 8M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) \\ &\quad + 3M \min\{\mathbb{E}_{\mathcal{D}_0}[\|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]\|], \\ &\quad \mathbb{E}_{\mathcal{D}_1}[\|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]\|]\}. \end{aligned}$$

Remark We see that the error gap is upper bounded by two terms: the distance between marginal label distributions and the discrepancy between conditional predicted distributions across groups. Given a dataset, the distance between marginal label distributions is a constant since the marginal label distributions are fixed. For the second term, if we can minimize the discrepancy of the conditional predicted distribution across groups, we then have a model that is free of accuracy disparity when the marginal label distributions are well aligned.

Geometric Interpretation By Theorem 3.1 and Theorem 3.2, we can visually illustrate how accuracy disparity arises given data distribution and the learned hypothesis that aims to minimize the global square error. In Figure 1, given the hypothesis class \mathcal{H} , we use the line $\text{Err}_{\mathcal{D}_0} + \text{Err}_{\mathcal{D}_1} = B$ to denote the lower bound in Theorem 3.1 and the two lines $|\text{Err}_{\mathcal{D}_0} - \text{Err}_{\mathcal{D}_1}| = A$ to denote the upper bound in Theorem 3.2. These three lines form a feasible region (the green area) of $\text{Err}_{\mathcal{D}_0}$ and $\text{Err}_{\mathcal{D}_1}$ under the hypothesis class \mathcal{H} . For any optimal hypothesis h which is solely designed to minimize the overall error, the best the hypothesis h can do is to intersect with one of the two bottom vertices. For example, the hypotheses (the red dotted line and the blue dotted line) trying to minimize overall error intersect with the two vertices of the region to achieve the smallest $\text{Err}_{\mathcal{D}_0}$ -intercept ($\text{Err}_{\mathcal{D}_1}$ -intercept), due to the imbalance between these two

groups. However, since these two vertices are not on the diagonal of the feasible region, there is no guarantee that the hypothesis can satisfy accuracy parity ($\text{Err}_{\mathcal{D}_0} = \text{Err}_{\mathcal{D}_1}$), unless we can shrink the width of green area to zero.

3.2. Algorithm Design

Inspired by Theorem 3.2, we can mitigate the error gap by aligning the group distributions via minimizing the distance of the conditional distributions across groups. However, it is intractable to do so explicitly in regression problems since Y can take infinite values on \mathbb{R} . Next we will present two algorithms to approximately solve the problem through adversarial representation learning.

Given a Markov chain $X \xrightarrow{g} Z \xrightarrow{h} \hat{Y}$, we are interested in learning group-invariant conditional representations so that the discrepancy between the induced conditional distributions $\mathcal{D}_0^Y(Z = g(X))$ and $\mathcal{D}_1^Y(Z = g(X))$ is minimized. In this case, the second term of the upper bound in Theorem 3.2 is minimized. However, it is in general not feasible since Y is a continuous random variable. Instead, we propose to learn the representations of Z to minimize the discrepancy between the joint distributions $\mathcal{D}_0(Z = g(X), Y)$ and $\mathcal{D}_1(Z = g(X), Y)$. Next, we will show the distances between conditional predicted distributions $\mathcal{D}_0^Y(Z = g(X))$ and $\mathcal{D}_1^Y(Z = g(X))$ are minimized when we minimize the joint distributions $\mathcal{D}_0(Z = g(X), Y)$ and $\mathcal{D}_1(Z = g(X), Y)$ in Theorem 3.3 and Theorem 3.4.

To proceed, we first consider using the total variation distance to measure the distance between two distributions. In particular, we can choose to learn a binary discriminator $f : Z \times Y \rightarrow \hat{A}$ that achieves minimum binary classification error on discriminating between points sampled from two distributions. In practice, we use the cross-entropy loss as a convex surrogate loss. Formally, we are going to consider the following minimax game between g and f :

$$\min_{f \in \mathcal{F}} \max_g \text{CE}_{\mathcal{D}}(A \parallel f(g(X), Y)) \quad (1)$$

Interestingly, for the above equation, the optimal feature transformation g corresponds to the one that induces invariant conditional feature distributions.

Theorem 3.3. Consider the minimax game in (1). The equilibrium (g^*, f^*) of the game is attained when 1). $Z = g^*(X)$ is independent of A conditioned on Y ; 2). $f^*(Z, Y) = \mathcal{D}(A = 1 \mid Y, Z)$.

Since in the equilibrium of the game Z is independent of A conditioned on Y , the optimal $f^*(Z, Y)$ could also be equivalently written as $f^*(Z, Y) = \mathcal{D}(A = 1 \mid Y)$, i.e., the only useful information for the discriminator in the equilibrium is through the external information Y . In Theorem 3.3, the minimum cross-entropy loss that the discriminator (the

equilibrium of the game) can achieve is $H(A | Z, Y)$ (see Proposition A.1 in Appendix A). For any feature transform g , by the basic property of conditional entropy, we have:

$$\begin{aligned} \min_{f \in \mathcal{F}} \text{CE}_{\mathcal{D}}(A \| f(g(X), Y)) &= H(A | Z, Y) \\ &= H(A | Y) - I(A; Z | Y). \end{aligned}$$

We know that $H(A | Y)$ is a constant given the data distribution. The maximization of g in (11) is equivalent to the minimization of $\min_{Z=g(X)} I(A; Z | Y)$, and it follows that the optimal strategy for the transformation g is the one that induces conditionally invariant features, *e.g.*, $I(A; Z | Y) = 0$. Formally, we arrive at the following minimax problem:

$$\min_{h, g} \max_{f \in \mathcal{F}} \text{MSE}_{\mathcal{D}}(h(g(X)), Y) - \lambda \cdot \text{CE}_{\mathcal{D}}(A \| f(g(X), Y))$$

In the above formulation, the first term corresponds to the minimization of prediction loss of the target task and the second term is the loss incurred by the adversary f . As a whole, the minimax optimization problem expresses a trade-off (controlled by the hyper-parameter $\lambda > 0$) between accuracy and accuracy disparity through the representation learning function g .

Wasserstein Variant Similarly, if we choose to align joint distributions via minimizing Wasserstein distance, the following theorem holds.

Theorem 3.4. Let the optimal feature transformation $g^* := \arg \min_g W_1(\mathcal{D}_0(g(X), Y), \mathcal{D}_1(g(X), Y))$, then $\mathcal{D}_0^Y(Z = g^*(X)) = \mathcal{D}_1^Y(Z = g^*(X))$ almost surely.

One notable advantage of using the Wasserstein distance instead of the TV distance is that, the Wasserstein distance is a continuous functional of both the feature map g as well as the discriminator f (Arjovsky et al., 2017). Furthermore, if both g and f are continuous functions of their corresponding model parameters, which is the case for models we are going to use in experiments, the objective function will be continuous in both model parameters. This property of the Wasserstein distance makes it more favorable from an optimization perspective. Using the dual formulation, equivalently, we can learn a Lipschitz function $f : Z \times Y \rightarrow \mathbb{R}$ as a witness function:

$$\begin{aligned} \min_{h, g, Z_0 \sim g^* \mathcal{D}_0, Z_1 \sim g^* \mathcal{D}_1} \max_{f: \|f\|_L \leq 1} \text{MSE}_{\mathcal{D}}(h(g(X)), Y) \\ + \lambda \cdot |f(Z_0, Y) - f(Z_1, Y)|. \end{aligned}$$

Game-Theoretic Interpretation We provide a game-theoretic interpretation of our algorithms in Figure 2 to make our algorithms easier to follow.

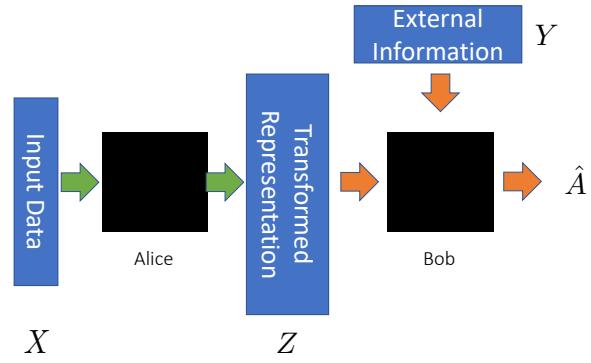


Figure 2: The game-theoretic illustration of our algorithms. Bob's goal is to guess the group membership A of each feature Z sent by Alice with the corresponding labels Y as the external information, while Alice's goal is to find a transformation from X to Z to confuse Bob.

As illustrated in Figure 2, consider Alice (encoder) and Bob (discriminator) participate a two-player game: upon receiving a set of inputs X , Alice applies a transformation to the inputs to generate the corresponding features Z and then sends them to Bob. Besides the features sent by Alice, Bob also has access to the external information Y , which corresponds to the corresponding labels for the set of features sent by Alice. Once having both the features Z and the corresponding labels Y from external resources, Bob's goal is to guess the group membership A of each feature sent by Alice, and to maximize his correctness as much as possible. On the other hand, Alice's goal is to compete with Bob, *i.e.*, to find a transformation to confuse Bob as much as she can. Different from the traditional game without external information, here due to the external information Y Bob has access to, Alice cannot hope to fully fool Bob, since Bob can gain some insights about the group membership A of features from the external label information anyway. Nevertheless, Theorem 3.3 and Theorem 3.4 both state that when Bob uses a binary discriminator or a Wasserstein discriminator to learn A , the best Alice could do is to learn a transformation g so that the transformed representation Z is insensitive to the values of A conditioned on any values of $Y = y$.

4. Experiments

Inspired by our theoretical results that decompose accuracy disparity into the distance between marginal label distributions and the distance between conditional representations, we propose two algorithms to mitigate it. In this section, we conduct experiments to evaluate the effectiveness of our

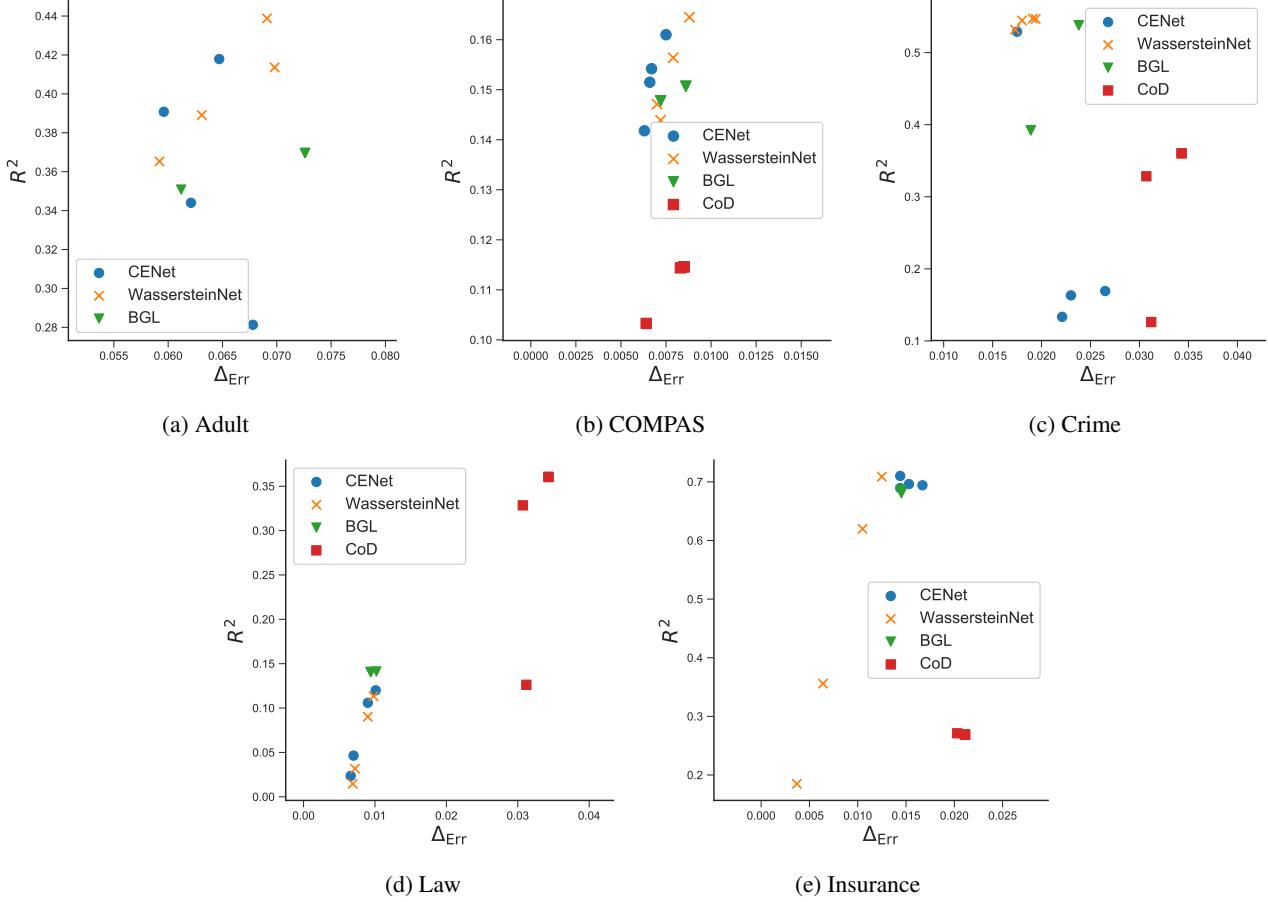


Figure 3: Overall results: R^2 regression scores and error gaps of different methods in five datasets. Our goal is to achieve high R^2 scores with small error gap values (*i.e.*, the most desirable points are located in the upper-left corner). Our proposed methods achieve the best trade-offs in Adult, COMPAS, Crime and Insurance datasets.

proposed algorithms in reducing the accuracy disparity.

4.1. Experimental Setup

Datasets We conduct experiments on five benchmark datasets: the Adult dataset (Dua and Graff, 2017), COMPAS dataset (Dieterich et al., 2016), Communities and Crime dataset (Dua and Graff, 2017), Law School dataset (Wightman and Ramsey, 1998) and Medical Insurance Cost dataset (Lantz, 2013). All datasets contain binary sensitive attributes (*e.g.*, male/female, white/non-white). We refer readers to Appendix B for detailed descriptions of the datasets and the data pre-processing pipelines. Note that although the Adult and COMPAS datasets are for binary classification tasks, recent evidences (Que and Belkin, 2016; Muthukumar et al., 2020; Hui and Belkin, 2021) suggest that square loss achieves comparable performance with cross-entropy loss and hinge loss. In this regard, we take them as regression tasks with two distinctive ordinal values.

Methods We term the proposed algorithms CENET and WASSERSTEINNET for our two proposed algorithms respectively and implement them using Pytorch (Paszke et al., 2019).² To the best of our knowledge, no previous study aims to minimize accuracy disparity in regression using representation learning. However, there are other similar fairness notions and mitigation techniques proposed for regression and we add them as our baselines: (1) Bounded group loss (BGL) (Agarwal et al., 2019), which asks for the prediction errors for any groups to remain below a predefined level ϵ ; (2) Coefficient of determination (CoD) (Komiyama et al., 2018), which asks for the coefficient of determination between the sensitive attributes and the predictions to remain below a predefined level ϵ .

For each dataset, we perform controlled experiments by

²Our code is publicly available at: <https://github.com/JFChi/Understanding-and-Mitigating-Accuracy-Disparity-in-Regression>

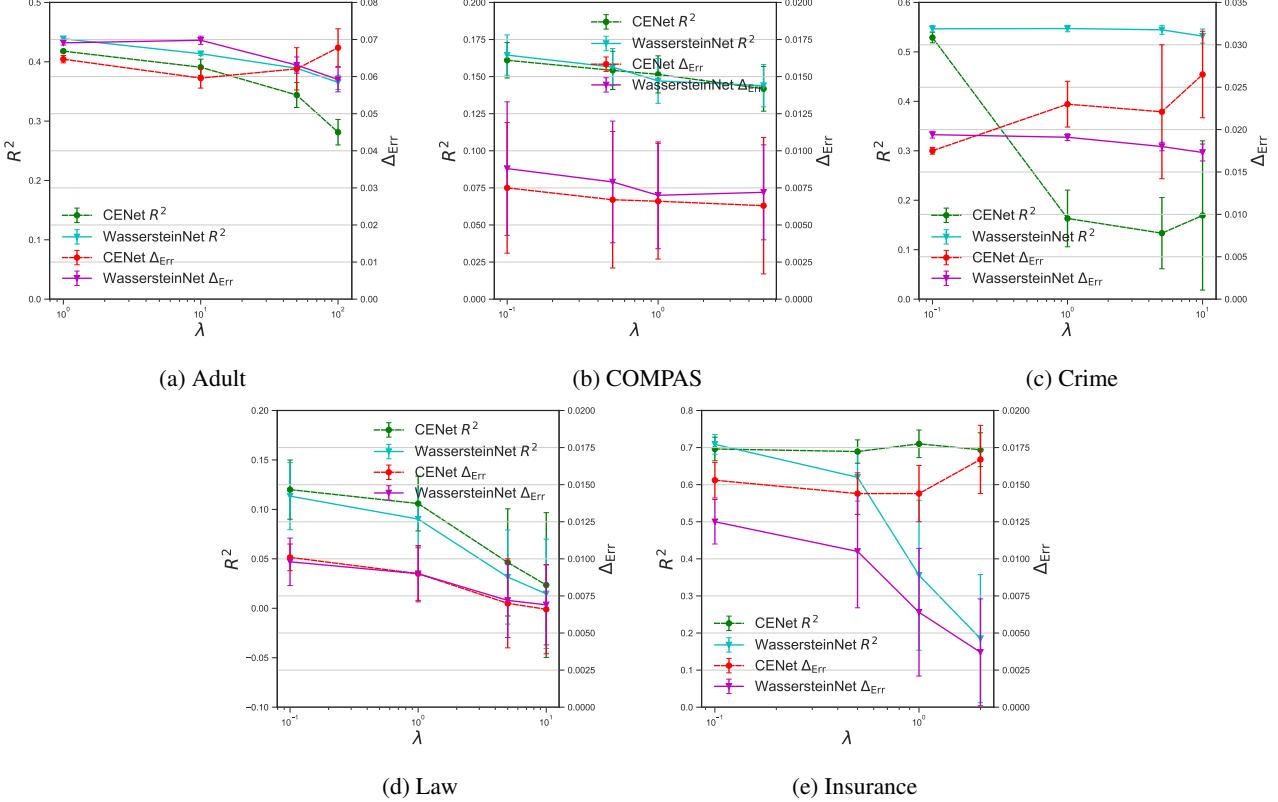


Figure 4: R^2 regression scores and error gaps when λ changes in CENET and WASSERSTEINNET. The general trend is that with the increase of λ , the error gap values and R^2 scores gradually decrease, except the cases where λ increases in CENET in Adult, Crime and Insurance dataset. The exceptions are caused by the instability of the training processes of CENET (Arjovsky and Bottou 2017).

fixing the regression model architectures to be the same. We train the regression models via minimizing mean squared loss. Among all methods, we vary the trade-off parameter (*i.e.*, λ in CENET and WASSERSTEINNET and ϵ in BGL and COD) and report and the corresponding R^2 scores and the error gap values. For each experiment, we average the results for ten different random seeds. Note that COD cannot be implemented on the Adult dataset since the size of the Adult dataset is large and the QCQP optimization algorithm to solve COD needs a quadratic memory usage of the dataset size. We refer readers to Appendix B for detailed hyper-parameter settings in our experiments and Appendix C for additional experimental results.

4.2. Results and Analysis

The overall results are visualized in Figure 3. The following summarizes our observations and analyses: (1) Our proposed methods WASSERSTEINNET and CENET are most effective in reducing the error gap values in all datasets compared to the baselines. Our proposed methods also achieve the best trade-offs in Adult, COMPAS, Crime and Insurance

datasets: with the similar error gap values (R^2 scores), our methods achieve the highest R^2 scores (lowest error gap values). In the Law dataset, the error gap values decrease with high utility losses in our proposed methods due the significant trade-offs between the predictive power of the regressors and accuracy parity. We suspect this is because the feature noise distribution in one group differs significantly than the others in the Law dataset. (2) Among our proposed methods, WASSERSTEINNET are more effective in reducing the error gap values while CENET might fail to decrease the error gaps in Adult, Crime and Insurance datasets and might even cause non-negligible reductions in the predictive performance of the regressors in Adult and Crime datasets. The reason behind it is that the minimax optimization in the training of CENET could lead to an unstable training process under the presence of a noisy approximation to the optimal discriminator (Arjovsky and Bottou 2017). We will provide more analysis in Figure 4 next. (3) Compared to our proposed methods, BGL and COD can also decrease error gaps to a certain extent. This is because: (i) BGL aims to keep errors remaining relatively low in each group, which helps to reduce accuracy disparity; (ii) CoD aims to reduce

the correlation between the sensitive attributes and the predictions (or the inputs) in the feature space, which might somehow reduce the dependency between the distributions of these two variables.

We further analyze how the trade-off parameter λ in the objective functions affect the performance of our methods. Figure 4 shows R^2 regression scores and error gaps when λ changes in CENET and WASSERSTEINNET. We see the general trend is that with the increase of the trade-off parameter λ , the error gap values and R^2 scores gradually decrease. Plus, the increase of λ generally leads to the instability of training processes with larger variances of both R^2 scores and error gap values. In Adult, Crime and Insurance datasets, WASSERSTEINNET is more effective in mitigating accuracy disparity when λ increases, while CENET fails to decrease the error gap values and might suffer from significant accuracy loss. The failure to decrease the error gap values with significant accuracy loss and variance indicates the estimation of total variation in minimax optimization for CENET could lead to a highly unstable training process (Arjovsky and Bottou, 2017).

5. Related Work

Algorithmic Fairness In the literature, two main notions of fairness, i.e., *group fairness* and *individual fairness*, has been widely studied (Dwork et al., 2012; Zemel et al., 2013; Feldman et al., 2015; Zafar et al., 2017a; Hardt et al., 2016; Zafar et al., 2017b; Hashimoto et al., 2018; Madras et al., 2019). In particular, Chen et al. (2018) analyzed the impact of data collection on discrimination (e.g., false positive rate, false negative rate, and zero-one loss) from the perspectives of bias-variance-noise decomposition, and they suggested collecting more training examples and collect additional variables to reduce discrimination. Khani and Liang (2019) argued that the loss difference among different groups is determined by the amount of latent (unobservable) feature noise and the difference between means, variances, and sizes of the groups with an assumption that there are a latent random feature and a noise feature that are involved in the generation of the observable features. Khani and Liang (2020) further found out that spurious features from inputs can hurt accuracy and affect groups disproportionately. Zhao and Gordon (2019) proposed an error decomposition theorem which upper bounds accuracy disparity in the classification setting by three terms: the sum of group-wise noise, the distance of marginal input distributions across groups and the discrepancy of group-wise optimal decision functions. However, their error decomposition theorem does not lead to any mitigation approaches in classification: minimizing the distance of marginal input distributions across groups does not necessarily mitigate accuracy disparity since it could possibly exacerbate the noise

term and the discrepancy of group-wise optimal decision functions in the meantime. Besides, the optimal group-wise decision functions are unknown and intractable to approximate in the feature spaces, which also adds to the difficulty of applying their upper bound directly. In comparison, our work only assumes that there is a joint distribution where all variables are sampled and precisely characterizes disparate predictive accuracy in regression in terms of the distance between marginal label distributions and the distance between conditional representations. Inspired by our theoretical results, we also propose practical algorithms to mitigate the problem when collecting more data becomes infeasible.

Fair Regression A series of works focus on fairness under the regression problems (Calders et al., 2013; Johnson et al., 2016; Berk et al., 2018; Komiyama et al., 2018; Chzhen et al., 2020b; Bigot, 2020). To the best of our knowledge, no previous study aimed to minimize accuracy disparity in regression from representation learning. However, there are different fairness notions and techniques proposed for regression: Agarwal et al. (2019) proposed fair regression with bounded group loss (i.e., it asks that the prediction error for any protected group remains below some pre-defined level) and used exponentiated-gradient approach to satisfy BGL. Komiyama et al. (2018) aimed to reduce the coefficient of determination between the sensitive attributes between the predictions to some pre-defined level and used an off-the-shelf convex optimizer to solve the problem. Mary et al. (2019) used the Hirschfeld-Gebelein-Rényi Maximum Correlation Coefficient to generalize fairness measurement to continuous variables and ensured equalized odds (demographic parity) constraint by minimizing the χ^2 divergence between the predicted variable and the sensitive variable (conditioned on target variable). Zink and Rose (2020) considered regression problems in health care spending and proposed five fairness criteria (e.g., covariance constraint, net compensation penalization, etc.) in the healthcare domain. Narasimhan et al. (2020) proposed pairwise fairness notions (e.g., pairwise equal opportunity requires each pair from two arbitrary different groups to be equally-likely to be ranked correctly) for ranking and regression models. Chzhen et al. (2020a) studied the regression problem with demographic parity constraint and showed the optimal fair predictor is achieved in the Wasserstein barycenter of group distributions. In contrast, we source out the root of accuracy disparity in regression through the lens of information theory and reduce it via distributional alignment using TV distance and Wasserstein distance in the minimax games.

Fair Representation A line of works focus on building algorithmic fair decision making systems using adversarial techniques to learn fair representations (Edwards and Storkey, 2015; Beutel et al., 2017; Zhao et al., 2019). The main idea behind is to learn a good representation of the

data so that the data owner can maximize the accuracy while removing the information related to the sensitive attribute. [Madras et al. \(2018\)](#) proposed a generalized framework to learn adversarially fair and transferable representations and suggests using the label information in the adversary to learn equalized odds or equal opportunity representations in the classification setting. Apart from adversarial representation, recent work also proposed to use distance metrics, *e.g.*, the maximum mean discrepancy [Louizos et al. \(2015\)](#) and the Wasserstein distance [Jiang et al. \(2019\)](#) to remove group-related information. Prior to this work, it is not clear aligning conditional distributions via adversarial representation learning could lead to (approximate) accuracy parity. Our analysis is the first work to connect accuracy parity and (conditional) distributional alignment in regression and we also provide algorithm interventions to mitigate the problem where it is challenging to align conditional distributions in regression problems.

6. Conclusion

In this paper, we theoretically and empirically study accuracy disparity in regression problems. Specifically, we prove an information-theoretic lower bound on the joint error and a complementary upper bound on the error gap across groups to depict the feasible region of group-wise errors. Our theoretical results indicate that accuracy disparity occurs inevitably due to the marginal label distributions differ across groups. To reduce such disparity, we further propose to achieve accuracy parity by learning conditional group-invariant representations using statistical distances. The game-theoretic optima of the objective functions in our proposed methods are achieved when the accuracy disparity is minimized. Our empirical results on five benchmark datasets demonstrate that our proposed algorithms help to reduce accuracy disparity effectively. We believe our results take an important step towards better understanding accuracy disparity in machine learning models.

Acknowledgements

We thank anonymous reviewers for their insightful feedback and suggestions. JC and YT would like to acknowledge support from NSF CNS 1823325, NSF CNS 1850479, and NSF OAC 2002985. HZ thanks the DARPA XAI project, contract #FA87501720152, for support. GG thanks Microsoft Research for support.

References

Alekh Agarwal, Miroslav Dudik, and Zhiwei Steven Wu. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, pages 120–129, 2019.

Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *arxiv e-prints*, art. *arXiv preprint arXiv:1701.04862*, 2017.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.

Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. Differential privacy has disparate impact on model accuracy. In *Advances in Neural Information Processing Systems*, pages 15453–15462, 2019.

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

Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, page 0049124118782533, 2018.

Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017.

Jérémie Bigot. Statistical data analysis in the wasserstein space. *ESAIM: Proceedings and Surveys*, 68:1–19, 2020.

Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. Fairlearn: A toolkit for assessing and improving fairness in AI. Technical Report MSR-TR-2020-32, Microsoft, May 2020.

Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91, 2018.

Toon Calders, Asim Karim, Faisal Kamiran, Wasif Ali, and Xiangliang Zhang. Controlling attribute effect in linear regression. In *2013 IEEE 13th international conference on data mining*, pages 71–80. IEEE, 2013.

Irene Chen, Fredrik D Johansson, and David Sontag. Why is my classifier discriminatory? In *Advances in Neural Information Processing Systems*, pages 3539–3550, 2018.

Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. Fair regression with wasserstein barycenters. *arXiv preprint arXiv:2006.07286*, 2020a.

Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. Fair regression via plug-in estimator and recalibration with statistical guarantees.

Advances in Neural Information Processing Systems, 33, 2020b.

Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems*, pages 9236–9246, 2018.

William Dieterich, Christina Mendoza, and Tim Brennan. Compas risk scales: Demonstrating accuracy equity and predictive parity. *Northpointe Inc*, 2016.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.

Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.

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

Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268. ACM, 2015.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.

Moritz Hardt, Eric Price, Nati Srebro, et al. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pages 3315–3323, 2016.

Tatsunori B Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. *arXiv preprint arXiv:1806.08010*, 2018.

Markus Hittmeir, Andreas Ekelhart, and Rudolf Mayer. Utility and privacy assessments of synthetic data for regression tasks. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 5763–5772. IEEE, 2019.

Like Hui and Mikhail Belkin. {EVALUATION} {of} {neural} {architectures} {trained} {with} {square} {loss} {vs} {cross}-{entropy} {in} {classification} {tasks}. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=hsFN92eQE1a>.

Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. Wasserstein fair classification. *arXiv preprint arXiv:1907.12059*, 2019.

Kory D Johnson, Dean P Foster, and Robert A Stine. Impartial predictive modeling: Ensuring fairness in arbitrary models. *arXiv preprint arXiv:1608.00528*, 2016.

Fereshte Khani and Percy Liang. Noise induces loss discrepancy across groups for linear regression. *arXiv preprint arXiv:1911.09876*, 2019.

Fereshte Khani and Percy Liang. Removing spurious features can hurt accuracy and affect groups disproportionately. *arXiv preprint arXiv:2012.04104*, 2020.

Pauline T Kim. Data-driven discrimination at work. *Wm. & Mary L. Rev.*, 58:857, 2016.

Junpei Komiyama, Akiko Takeda, Junya Honda, and Haimie Shimao. Nonconvex optimization for regression with fairness constraints. In *International conference on machine learning*, pages 2737–2746, 2018.

Brett Lantz. *Machine learning with R*. Packt publishing ltd, 2013.

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

David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. *arXiv preprint arXiv:1802.06309*, 2018.

David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Fairness through causal awareness: Learning causal latent-variable models for biased data. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 349–358, 2019.

Jérémie Mary, Clément Calauzènes, and Noureddine El Karoui. Fairness-aware learning for continuous attributes and treatments. In *International Conference on Machine Learning*, pages 4382–4391, 2019.

Vidya Muthukumar, Adhyyan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, and Anant Sahai. Classification vs regression in overparameterized regimes: Does the loss function matter? *arXiv preprint arXiv:2005.08054*, 2020.

Harikrishna Narasimhan, Andrew Cotter, Maya R Gupta, and Serena Wang. Pairwise fairness for ranking and regression. In *AAAI*, pages 5248–5255, 2020.

Arvind Narayanan. Translation tutorial: 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp.*, New York, USA, 2018.

Yangchen Pan, Ehsan Imani, Amir-massoud Farahmand, and Martha White. An implicit function learning approach for parametric modal regression. *Advances in Neural Information Processing Systems*, 33, 2020.

Belisario Panay, Nelson Baloian, José A Pino, Sergio Peñafiel, Horacio Sanson, and Nicolas Bersano. Predicting health care costs using evidence regression. In *Multidisciplinary Digital Publishing Institute Proceedings*, volume 31, page 74, 2019.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037, 2019.

Qichao Que and Mikhail Belkin. Back to the future: Radial basis function networks revisited. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 1375–1383, Cadiz, Spain, 09–11 May 2016. PMLR. URL <http://proceedings.mlr.press/v51/que16.html>.

Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

Linda F Wightman and Henry Ramsey. *LSAC national longitudinal bar passage study*. Law School Admission Council, 1998.

Mohammad Yaghini, Bogdan Kulynych, and Carmela Troncoso. Disparate vulnerability: On the unfairness of privacy attacks against machine learning. *arXiv preprint arXiv:1906.00389*, 2019.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1171–1180. International World Wide Web Conferences Steering Committee, 2017a.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970, 2017b.

Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International Conference on Machine Learning*, pages 325–333, 2013.

Han Zhao and Geoffrey J Gordon. Inherent tradeoffs in learning fair representations. In *Advances in neural information processing systems*, 2019.

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

Han Zhao, Jianfeng Chi, Yuan Tian, and Geoffrey J Gordon. Trade-offs and guarantees of adversarial representation learning for information obfuscation. *Advances in Neural Information Processing Systems*, 33, 2020.

Anna Zink and Sherri Rose. Fair regression for health care spending. *Biometrics*, 76(3):973–982, 2020.

Appendix

In the appendix, we give the proofs of the theorems and claims in our paper, the experimental details and more experimental results.

A. Missing Proofs

Proposition 3.1. Assume both $\mathbb{E}_{\mathcal{D}_a}[Y]$ and $\mathbb{E}_{\mathcal{D}_a}[Y^2]$ are equivalent for any $A = a$, then using a constant predictor ensures accuracy parity in regression.

Proof. For $a \in \{0, 1\}$, we have

$$\begin{aligned} & \text{Err}_{\mathcal{D}_a}(h) \\ &= \mathbb{E}_{\mathcal{D}_a}[(h(X) - Y)^2] \\ &= \mathbb{E}_{\mathcal{D}_a}[(h(X) - \mathbb{E}_{\mathcal{D}_a}(Y) + \mathbb{E}_{\mathcal{D}_a}(Y) - Y)^2] \\ &= \mathbb{E}_{\mathcal{D}_a}[(h(X) - \mathbb{E}_{\mathcal{D}_a}(Y))^2] + \mathbb{E}_{\mathcal{D}_a}[(Y - \mathbb{E}_{\mathcal{D}_a}(Y))^2] - 2\mathbb{E}_{\mathcal{D}_a}[(h(X) - \mathbb{E}_{\mathcal{D}_a}(Y))(Y - \mathbb{E}_{\mathcal{D}_a}(Y))]. \end{aligned}$$

It is easy to see the first two terms are equal across different groups since $\mathbb{E}_{\mathcal{D}_a}[Y]$, $\mathbb{E}_{\mathcal{D}_a}[Y^2]$ and $h(X)$ are the same across different groups. For the third term, we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}_a}[(h(X) - \mathbb{E}_{\mathcal{D}_a}(Y))(Y - \mathbb{E}_{\mathcal{D}_a}(Y))] \\ &= \mathbb{E}_{\mathcal{D}_a(X)}[\mathbb{E}_{\mathcal{D}_a(Y|X)}[(h(X) - \mathbb{E}_{\mathcal{D}_a}(Y))(Y - \mathbb{E}_{\mathcal{D}_a}(Y)) | X]] \\ &= \mathbb{E}_{\mathcal{D}_a(X)}[(h(X) - \mathbb{E}_{\mathcal{D}_a}[Y | X])(\mathbb{E}_{\mathcal{D}_a}[Y | X]) - \mathbb{E}_{\mathcal{D}_a}[Y | X])] \\ &= 0. \end{aligned}$$

Thus, the errors across different groups made by the constant predictor are the same if $\mathbb{E}_{\mathcal{D}_a}[Y]$ and $\mathbb{E}_{\mathcal{D}_a}[Y^2]$ are equivalent across different groups. \blacksquare

Lemma 3.1. Let $\hat{Y} = h(X)$, then for $a \in \{0, 1\}$, $W_1(\mathcal{D}_a(Y), h_{\sharp}\mathcal{D}_a) \leq \sqrt{\text{Err}_{\mathcal{D}_a}(h)}$.

Proof. The prediction error conditioned on $a \in \{0, 1\}$ is

$$\begin{aligned} \text{Err}_{\mathcal{D}_a}(h) &= \mathbb{E}[(Y - h(X))^2 | A = a] \\ &\geq \mathbb{E}^2[|Y - h(X)|| |A = a|] \\ &\geq \left(\inf_{\Gamma(\mathcal{D}_a(Y), \mathcal{D}_a(h(X)))} \mathbb{E}[|Y - h(X)|] \right)^2 \\ &= W_1^2(\mathcal{D}_a(Y), h_{\sharp}\mathcal{D}_a). \end{aligned}$$

Taking square root at both sides then completes the proof. \blacksquare

Theorem 3.1. Let $\hat{Y} = h(X)$ be the predicted variable, then $\text{Err}_{\mathcal{D}_0}(h) + \text{Err}_{\mathcal{D}_1}(h) \geq \frac{1}{2}[(W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1))_+]^2$.

Proof. Since $W_1(\cdot, \cdot)$ is a distance metric, the result follows immediately the triangle inequality and Lemma 3.1:

$$W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) \leq \sqrt{\text{Err}_{\mathcal{D}_0}(h)} + W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1) + \sqrt{\text{Err}_{\mathcal{D}_1}(h)}.$$

Rearrange the equation above and by AM-GM inequality, we have

$$W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1) \leq \sqrt{\text{Err}_{\mathcal{D}_0}(h)} + \sqrt{\text{Err}_{\mathcal{D}_1}(h)} \leq \sqrt{2(\text{Err}_{\mathcal{D}_0}(h) + \text{Err}_{\mathcal{D}_1}(h))}.$$

Taking square at both sides then completes the proof. \blacksquare

Corollary 3.1. Let $\hat{Y} = h(X)$ and $\alpha = \mathcal{D}(A = 0) \in [0, 1]$, we have $\text{Err}_{\mathcal{D}}(h) \geq \frac{1}{2} \min\{\alpha, 1 - \alpha\} \cdot [(W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1))_+]^2$.

Proof. The joint error is

$$\begin{aligned}
 & \text{Err}_{\mathcal{D}}(h) \\
 &= \alpha \text{Err}_{\mathcal{D}_0}(h) + (1 - \alpha) \text{Err}_{\mathcal{D}_1}(h) \\
 &\geq \min\{\alpha, 1 - \alpha\} (\text{Err}_{\mathcal{D}_0}(h) + \text{Err}_{\mathcal{D}_1}(h)) \\
 &\geq \frac{1}{2} \min\{\alpha, 1 - \alpha\} [(W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) - W_1(h_{\sharp}\mathcal{D}_0, h_{\sharp}\mathcal{D}_1))_+]^2. \quad (\text{Theorem 3.1})
 \end{aligned}$$

■

Theorem 3.2. If Assumption 2.1 holds, then for $\forall h \in \mathcal{H}$, let $\hat{Y} = h(X)$, the following inequality holds:

$$\begin{aligned}
 \Delta_{\text{Err}}(h) &\leq 8M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) \\
 &\quad + 3M \min\{\mathbb{E}_{\mathcal{D}_0}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|], \\
 &\quad \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|]\}.
 \end{aligned}$$

Proof. First, we show that for $a \in \{0, 1\}$:

$$\text{Err}_{\mathcal{D}_a}(h) = \mathbb{E}_{\mathcal{D}_a}[(h(X) - Y)^2] = \mathbb{E}_{\mathcal{D}_a}[h^2(X) - 2Yh(X) + Y^2] = \mathbb{E}_{\mathcal{D}_a}[h^2(X) - 2Yh(X)] + \mathbb{E}_{\mathcal{D}_a}[Y^2].$$

Next, we bound the error gap:

$$\begin{aligned}
 & |\text{Err}_{\mathcal{D}_0}(h) - \text{Err}_{\mathcal{D}_1}(h)| \\
 &= |\mathbb{E}_{\mathcal{D}_0}[h^2(X) - 2Yh(X)] + \mathbb{E}_{\mathcal{D}_0}[Y^2] - \mathbb{E}_{\mathcal{D}_1}[h^2(X) - 2Yh(X)] - \mathbb{E}_{\mathcal{D}_1}[Y^2]| \\
 &\leq |\mathbb{E}_{\mathcal{D}_0}[h^2(X) - 2Yh(X)] - \mathbb{E}_{\mathcal{D}_1}[h^2(X) - 2Yh(X)]| + |\mathbb{E}_{\mathcal{D}_0}[Y^2] - \mathbb{E}_{\mathcal{D}_1}[Y^2]|. \quad (\text{Triangle inequality})
 \end{aligned}$$

For the second term, we can easily prove that

$$|\mathbb{E}_{\mathcal{D}_0}[Y^2] - \mathbb{E}_{\mathcal{D}_1}[Y^2]| = |\langle Y^2, d\mathcal{D}_0 - d\mathcal{D}_1 \rangle| \leq \|Y\|_{\infty}^2 \|d\mathcal{D}_0 - d\mathcal{D}_1\|_1 \leq 2M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)),$$

where the second equation follows Hölder's inequality and the last equation follows the definition of total variation distance. Now it suffices to bound the remaining term:

$$\begin{aligned}
 & |\mathbb{E}_{\mathcal{D}_0}[h^2(X) - 2Yh(X)] - \mathbb{E}_{\mathcal{D}_1}[h^2(X) - 2Yh(X)]| \\
 &= \left| \int h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}, y) - \int h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_1(\mathbf{x}, y) \right| \\
 &\leq \left| \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_0(y) - \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_1(y) \right| \quad (\text{Triangle inequality}) \\
 &\quad + \left| \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_1(\mathbf{x}|y) d\mu_1(y) - \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_1(y) \right|.
 \end{aligned}$$

We upper bound the first term:

$$\begin{aligned}
 & \left| \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_0(y) - \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_1(y) \right| \\
 &\leq \iint |h(\mathbf{x})(h(\mathbf{x}) - 2y)(d\mu_0(y) - d\mu_1(y))| d\mu_0(\mathbf{x}|y) \\
 &\leq \int |d\mu_0(y) - d\mu_1(y)| \int |\sup_{\mathbf{x}} h(\mathbf{x})| |h(\mathbf{x}) - 2y| d\mu_0(\mathbf{x}|y) \\
 &\leq M \int \mathbb{E}_{\mathcal{D}_0}[|h(X) - 2Y| |Y = y|] |d\mu_0(y) - d\mu_1(y)| \quad (\text{Assumption 2.1}) \\
 &\leq 3M^2 \int |d\mu_0(y) - d\mu_1(y)| \quad (\text{Assumption 2.1}) \\
 &\leq 6M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)).
 \end{aligned}$$

Note that the last equation follows the definition of total variation distance. For the second term, we have:

$$\begin{aligned}
 & \left| \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_1(\mathbf{x}|y) d\mu_1(y) - \iint h(\mathbf{x})(h(\mathbf{x}) - 2y) d\mu_0(\mathbf{x}|y) d\mu_1(y) \right| \\
 & \leq \left| \iint h^2(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| + \left| \iint 2y h(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| \quad (\text{Triangle inequality}) \\
 & \leq 3M \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|]. \quad (\text{Assumption 2.1})
 \end{aligned}$$

To prove the last equation, we first see that:

$$\begin{aligned}
 & \left| \iint h^2(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| \\
 & \leq \left| \iint (\sup_{\mathbf{x}} h(\mathbf{x})) h(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| \\
 & \leq M \int |\mathbb{E}_{\mathcal{D}_0}[h(X)|Y=y] - \mathbb{E}_{\mathcal{D}_1}[h(X)|Y=y]| d\mu_1(y) \quad (\text{Assumption 2.1}) \\
 & = M \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|].
 \end{aligned}$$

Similarly, we also have:

$$\begin{aligned}
 & \left| \iint 2y h(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| \\
 & \leq 2 \left| \iint (\sup y) h(\mathbf{x})(d\mu_1(\mathbf{x}|y) - d\mu_0(\mathbf{x}|y)) d\mu_1(y) \right| \\
 & \leq 2M \int |\mathbb{E}_{\mathcal{D}_0}[h(X)|Y=y] - \mathbb{E}_{\mathcal{D}_1}[h(X)|Y=y]| d\mu_1(y) \quad (\text{Assumption 2.1}) \\
 & = 2M \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|].
 \end{aligned}$$

By symmetry, we can also see that:

$$|\mathbb{E}_{\mathcal{D}_0}[h^2(X) - 2Yh(X)] - \mathbb{E}_{\mathcal{D}_1}[h^2(X) - 2Yh(X)]| \leq 6M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) + 3M \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|].$$

Combine the above two equations yielding:

$$\begin{aligned}
 & |\mathbb{E}_{\mathcal{D}_0}[h^2(X) - 2Yh(X)] - \mathbb{E}_{\mathcal{D}_1}[h^2(X) - 2Yh(X)]| \\
 & \leq 6M^2 d_{\text{TV}}(\mathcal{D}_0(Y), \mathcal{D}_1(Y)) + 3M \min\{\mathbb{E}_{\mathcal{D}_0}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|], \mathbb{E}_{\mathcal{D}_1}[|\mathbb{E}_{\mathcal{D}_0^y}[\hat{Y}] - \mathbb{E}_{\mathcal{D}_1^y}[\hat{Y}]|]\}.
 \end{aligned}$$

Incorporating the terms back to the upper bound of the error gap then completes the proof. \blacksquare

Theorem 3.3. Consider the minimax game in (I). The equilibrium (g^*, f^*) of the game is attained when 1). $Z = g^*(X)$ is independent of A conditioned on Y ; 2). $f^*(Z, Y) = \mathcal{D}(A = 1 | Y, Z)$.

Proof. To prove Theorem 3.3, we first give Proposition A.1.

Proposition A.1. For any feature map $g : \mathcal{X} \rightarrow \mathcal{Z}$, assume that \mathcal{F} contains all the randomized binary classifiers and $\mathcal{F} \ni f : \mathcal{Z} \times \mathcal{Y} \rightarrow \mathcal{A}$, then $\min_{f \in \mathcal{F}} \text{CE}_{\mathcal{D}}(A \parallel f(g(X), Y)) = H(A | Z, Y)$.

Proof. By the definition of cross-entropy loss, we have:

$$\begin{aligned}
 \text{CE}_{\mathcal{D}}(A \parallel f) &= -\mathbb{E}_{\mathcal{D}}[\mathbb{I}(A=0) \log(1 - f(g(X), Y)) + \mathbb{I}(A=1) \log(f(g(X), Y))] \\
 &= -\mathbb{E}_{g \in \mathcal{D}}[\mathbb{I}(A=0) \log(1 - f(Z, Y)) + \mathbb{I}(A=1) \log(f(Z, Y))] \\
 &= -\mathbb{E}_{Z,Y} \mathbb{E}_{A|Z,Y}[\mathbb{I}(A=0) \log(1 - f(Z, Y)) + \mathbb{I}(A=1) \log(f(Z, Y))] \\
 &= -\mathbb{E}_{Z,Y}[\mathcal{D}(A=0 | Z, Y) \log(1 - f(Z, Y)) + \mathcal{D}(A=1 | Z, Y) \log(f(Z, Y))] \\
 &= \mathbb{E}_{Z,Y}[D_{\text{KL}}(\mathcal{D}(A | Z, Y) \parallel f(Z, Y))] + H(A | Z, Y) \\
 &\geq H(A | Z, Y),
 \end{aligned}$$

where $D_{\text{KL}}(\cdot \parallel \cdot)$ denotes the KL divergence between two distributions. From the above inequality, it is also clear that the minimum value of the cross-entropy loss is achieved when $f(Z, Y)$ equals the conditional probability $\mathcal{D}(A = 1 \mid Z, Y)$, i.e., $f^*(Z, Y) = \mathcal{D}(A = 1 \mid Z = g(X), Y)$. \blacksquare

Proposition A.1 states that the minimum cross-entropy loss that the discriminator can achieve is $H(A \mid Z, Y)$ when f is the conditional distribution $\mathcal{D}(A = 1 \mid Z = g(X), Y)$. By the basic property of conditional entropy, we have:

$$\min_{f \in \mathcal{F}} \text{CE}_{\mathcal{D}}(A \parallel f(g(X), Y)) = H(A \mid Z, Y) = H(A \mid Y) - I(A; Z \mid Y).$$

Note that $H(A \mid Y)$ is a constant given the distribution \mathcal{D} , so the maximization of g is equivalent to the minimization of $\min_{Z=g(X)} I(A; Z \mid Y)$, and it follows that the optimal strategy for the transformation g is the one that induces conditionally invariant features, e.g., $I(A; Z \mid Y) = 0$. On the other hand, if g^* plays optimally, then the optimal response of the discriminator f is given by

$$f^*(Z, Y) = \mathcal{D}(A = 1 \mid Z = g^*(X), Y) = \mathcal{D}(A = 1 \mid Y).$$

\blacksquare

Theorem 3.4. Let the optimal feature transformation $g^* := \arg \min_g W_1(\mathcal{D}_0(g(X), Y), \mathcal{D}_1(g(X), Y))$, then $\mathcal{D}_0^Y(Z = g^*(X)) = \mathcal{D}_1^Y(Z = g^*(X))$ almost surely.

Proof. By the definition of Wasserstein distance, we have:

$$\begin{aligned} W_1(\mathcal{D}_0(Z, Y), \mathcal{D}_1(Z, Y)) &= \inf_{\gamma \in \Gamma(\mathcal{D}_0, \mathcal{D}_1)} \int d((\mathbf{z}_0, y_0), (\mathbf{z}_1, y_1)) d\gamma((\mathbf{z}_0, y_0), (\mathbf{z}_1, y_1)) \\ &= \inf_{\gamma \in \Gamma(\mathcal{D}_0, \mathcal{D}_1)} \iint d((\mathbf{z}_0, y_0), (\mathbf{z}_1, y_1)) d\gamma(\mathbf{z}_0, \mathbf{z}_1 \mid y_0, y_1) d\gamma(y_0, y_1) \\ &= \inf_{\gamma \in \Gamma(\mathcal{D}_0, \mathcal{D}_1)} \iint \|\mathbf{z}_0 - \mathbf{z}_1\|_1 + |y_0 - y_1| d\gamma(\mathbf{z}_0, \mathbf{z}_1 \mid y_0, y_1) d\gamma(y_0, y_1) \\ &\geq \inf_{\gamma \in \Gamma(\mathcal{D}_0, \mathcal{D}_1)} \iint |y_0 - y_1| d\gamma(y_0, y_1) d\gamma(\mathbf{z}_0, \mathbf{z}_1 \mid y_0, y_1) \\ &= \inf_{\gamma \in \Gamma(\mathcal{D}_0(Y), \mathcal{D}_1(Y))} \int |y_0 - y_1| d\gamma(y_0, y_1) \\ &= W_1(\mathcal{D}_0(Y), \mathcal{D}_1(Y)). \end{aligned}$$

To finish the proof, next we prove the lower bound is achieved when $\mathcal{D}_0^Y(Z = g^*(X)) = \mathcal{D}_1^Y(Z = g^*(X))$: it is easy to see $W_1(\mathcal{D}_0^Y(Z), \mathcal{D}_1^Y(Z)) = \int \|\mathbf{z}_0 - \mathbf{z}_1\|_1 d\gamma(\mathbf{z}_0, \mathbf{z}_1 \mid y_0, y_1) = 0$ when the conditional distributions are equal. In this case, when the Wasserstein distance is minimized, then Z is conditionally independent of A given Y almost surely. \blacksquare

B. Experimental Details

Adult The Adult dataset contains 48,842 examples for income prediction. The task is to predict whether the annual income of an individual is greater or less than 50K/year based on the attributes of the individual, such as education level, age, occupation, etc. In our experiment, we use gender (binary) as the sensitive attribute. The target variable (income) is an ordinal binary variable: 0 if < 50K/year otherwise 1. After data pre-processing, the dataset contains 30,162/15,060 training/test instances where the input dimension of each instance is 113. We show the data distributions for different demographic subgroups in Table I.

To preprocess the dataset, we first filter out the data records that contain the missing values. We then remove the sensitive attribute from the input features and normalize the input features with its means and standard deviations. Note that we use one-hot encoding for the categorical attributes.

For our proposed methods, we use a three-layer neural network with ReLU as the activation function of the hidden layers and the sigmoid function as the output function for the prediction task (we take the first two layers as the feature mapping).

The number of neurons in the hidden layers is 60. We train the neural networks with the ADADELTA algorithm with the learning rate 0.1 and a batch size of 512. The models are trained in 50 epochs. For the adversary networks in CENET and WASSERSTEINNET, we use a two-layer neural network with ReLU as the activation function. The number of neurons in the hidden layers of the adversary networks is 60. The adversary network in CENET also uses sigmoid function as the output function. The weight clipping norm in the adversary network of WASSERSTEINNET is 0.005. We use the gradient reversal layer (Ganin et al., 2016) to implement the gradient descent ascent (GDA) algorithm for optimization of the minimax problem since it makes the training process more stable (Daskalakis and Panageas, 2018). For the rest of the datasets we used in our experiments, we also use a gradient reversal layer to implement our algorithms.

We use the Fairlearn toolkit (Bird et al., 2020) to implement BGL: we use the exponentiated-gradient algorithm with the default setting as the mitigator and vary the upper bound $\epsilon \in \{0.1, 0.2, 0.3, 0.5\}$ of the bounded group loss constraint. For each value of ϵ , we average the results of ten different random seeds.

COMPAS The COMPAS dataset contains 6,172 instances to predict whether a criminal defendant will recidivate within two years or not. It contains attributes such as age, race, etc. In our experiment, we use race (white or non-white) as the sensitive attribute and recidivism as the target variable. We split the dataset into train and test sets with the ratio 7/3. We show the data distributions for different demographic subgroups in Table 2.

For all methods, we use a two-layer neural network with ReLU as the activation function of the hidden layers and the sigmoid function as the output function for the prediction task (we take the first layer as the feature mapping). The number of neurons in the hidden layers is 60. We train the neural networks with the ADADELTA algorithm with the learning rate 1.0 and a batch size of 512. The models are trained in 50 epochs. For the adversary networks in CENET and WASSERSTEINNET, we use a two-layer neural network with ReLU as the activation function. The number of neurons in the hidden layers of the adversary networks is 10. The adversary network in CENET also uses sigmoid function as the output function. The weight clipping norm in the adversary network of WASSERSTEINNET is 0.05.

We use the Fairlearn toolkit to implement BGL: we use the exponentiated-gradient algorithm with the default setting as the mitigator and vary the upper bound $\epsilon \in \{0.1, 0.2, 0.3, 0.5\}$ of the bounded group loss constraint. For each value of ϵ , we average the results of ten different random seeds.

As for COD, we follow the source implementation.³ We use the same hyper-parameter settings as (Komiyama et al., 2018): We use the kernelized optimization with the random Fourier features and the RBF kernel (we vary hyper-parameter of the RBF kernel $\gamma \in \{0.1, 1.0, 10, 100\}$) and report the best results with minimal MSE loss for each time we change the fairness budget ϵ . We also vary $\epsilon \in \{0.01, 0.1, 0.5, 1.0\}$ and average the results of ten different random seeds.

Table 1: Data distribution of Y and A in Adult dataset.

	$Y = 0$	$Y = 1$
$A = 0$	20988	9539
$A = 1$	13026	1669

Table 2: Data distribution of Y and A in COMPAS dataset.

	$Y = 0$	$Y = 1$
$A = 0$	1849	1148
$A = 1$	1514	1661

Communities and Crime The Communities and Crime dataset contains 1,994 examples of socio-economic, law enforcement, and crime data about communities in the United States. The task is to predict the number of violent crimes per 100K population. All attributes in the dataset have been curated and normalized to $[0, 1]$. In our experiment, we use race (binary) as the sensitive attribute: 1 if the population percentage of the white is greater or equal to 80% otherwise 0. After data pre-processing, the dataset contains 1,595/399 training/test instances where the input dimension of each instance is 96. We visualize the data distributions for different demographic subgroups in Figure 5a.

To preprocess the dataset, we first remove the non-predictive attributes and sensitive attributes from the input features. Note that all features in the dataset have already been normalized in $[0, 1]$ so that we do not perform additional normalization to the features. We then replace the missing values with the mean values of the corresponding attributes.

³<https://github.com/jkomiyama/fairregression>

For all methods, we use a two-layer neural network with ReLU as the activation function of the hidden layers and the sigmoid function as the output function for the prediction task (we take the first layer as the feature mapping). The number of neurons in the hidden layers is 50. We train the neural networks with the ADADELTA algorithm with the learning rate 0.1 and a batch size of 256. The models are trained in 100 epochs. For the adversary networks in CENET and WASSERSTEINNET, we use a two-layer neural network with ReLU as the activation function. The number of neurons in the hidden layers of the adversary networks is 100. The adversary network in CENET also uses sigmoid function as the output function. The weight clipping norm in the adversary network of WASSERSTEINNET is 0.002.

We use the Fairlearn toolkit to implement BGL: we use the exponentiated-gradient algorithm with the default setting as the mitigator and vary the upper bound $\epsilon \in \{0.01, 0.02, 0.03, 0.05\}$ of the bounded group loss constraint. For each value of ϵ , we average the results of ten different random seeds. Note that our experiment setup is different from (Agarwal et al., 2019), so our results cannot be directly compared to theirs.

As for COD, we follow the same hyper-parameter settings as (Komiyama et al., 2018): We use the kernelized optimization with the random Fourier features and the RBF kernel (we vary hyper-parameter of the RBF kernel $\gamma \in \{0.1, 1.0, 10, 100\}$) and report the best results with minimal MSE loss for each time we change the fairness budget ϵ . The hyper-parameter settings follow from (Komiyama et al., 2018). We also vary $\epsilon \in \{0.01, 0.1, 0.5, 1.0\}$ and average the results of ten different random seeds. Note that our experiment setup is different from (Komiyama et al., 2018), so our results cannot be directly compared to theirs.

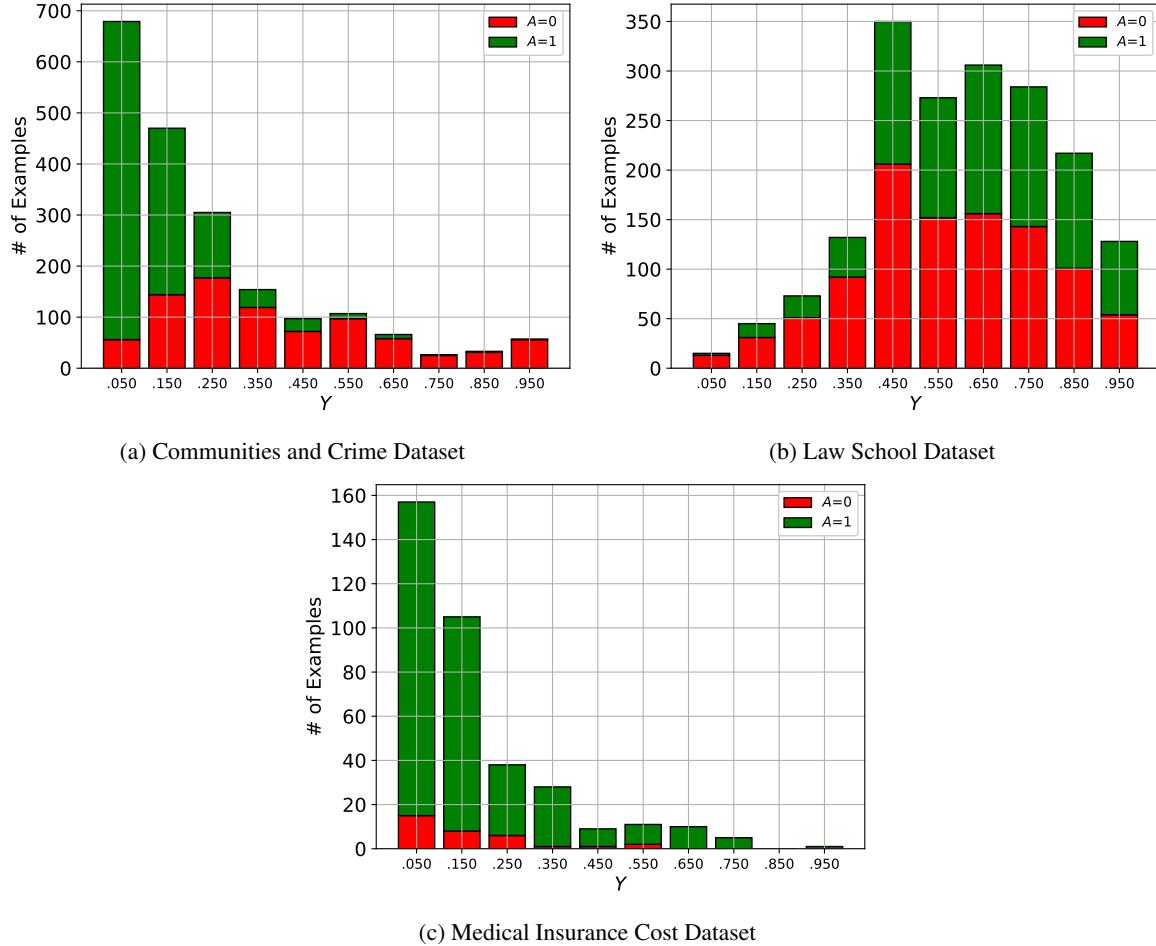


Figure 5: Data distributions for different demographic subgroups in three datasets.

Law School The Law School dataset contains 1,823 records for law students who took the bar passage study for Law School Admission⁴. The features in the dataset include variables such as undergraduate GPA, LSAT score, full-time status, family income, gender, etc. In our experiment, we use gender as the sensitive attribute and undergraduate GPA as the target variable. We split the dataset into train and test sets with the ratio 8/2. We show the data distributions for different demographic subgroups in Figure 5b.

For all methods, we use a two-layer neural network with ReLU as the activation function of the hidden layers and the sigmoid function as the output function for the prediction task (we take the first layer as the feature mapping). The number of neurons in the hidden layers is 10. We train the neural networks with the ADADELTA algorithm with the learning rate 0.1 and a batch size of 256. The models are trained in 100 epochs. For the adversary networks in CENET and WASSERSTEINNET, we use a two-layer neural network with ReLU as the activation function. The number of neurons in the hidden layers of the adversary networks is 10. The adversary network in CENET also uses sigmoid function as the output function. The weight clipping norm in the adversary network of WASSERSTEINNET is 0.2.

We use the Fairlearn toolkit to implement BGL: we use the exponentiated-gradient algorithm with the default setting as the mitigator and vary the upper bound $\epsilon \in \{0.01, 0.02, 0.03, 0.05\}$ of the bounded group loss constraint. For each value of ϵ , we average the results of ten different random seeds. Note that our experiment setup is different from (Agarwal et al., 2019), so our results cannot be directly compared to theirs.

As for COD, we follow the same hyper-parameter settings as (Komiyama et al., 2018): We use the kernelized optimization with the random Fourier features and the RBF kernel (we vary hyper-parameter of the RBF kernel $\gamma \in \{0.1, 1.0, 10, 100\}$) and report the best results with minimal MSE loss for each time we change the fairness budget ϵ . The hyper-parameter settings follow from (Komiyama et al., 2018). We also vary $\epsilon \in \{0.01, 0.1, 0.5, 1.0\}$ and average the results of ten different random seeds. Note that our experiment setup is different from (Komiyama et al., 2018), so our results cannot be directly compared to theirs.

Medical Insurance Cost The medical insurance cost dataset (Lantz, 2013) is a simulated dataset which was created using real-world demographic statistics from the U.S. Census Bureau⁵. The dataset reflect approximately reflect real-world conditions and has been used in the research of regression (Panay et al., 2019; Hittmeir et al., 2019; Pan et al., 2020). It contains 1,338 medical expense examples for patients in the United States, with features such as gender, age, BMI, etc., indicating characteristics of the patient and total annual medical expenses charged to the patients. In our experiment, we use gender as the sensitive attribute and the charged medical expenses as the target variable. In order to reflect the real-world scenarios where the accuracy disparity is significant due to the small and imbalanced dataset, we sub-sample the dataset: we randomly subsample 5% of examples with gender as male and 50% of examples with gender as female. After sub-sampling, we get 364 examples in total (33 male examples and 331 female examples). We split the dataset into train and test sets with the ratio 7/3. We visualize the data distributions for different demographic subgroups in Figure 5c.

For all methods, we use a two-layer neural network with ReLU as the activation function of the hidden layers and the sigmoid function as the output function for the prediction task (we take the first layer as the feature mapping). The number of neurons in the hidden layers is 7. We train the neural networks with the SGD algorithm with the learning rate 0.1 and a batch size of 64. The models are trained in 750 epochs. For the adversary networks in CENET and WASSERSTEINNET, we use a two-layer neural network with ReLU as the activation function. The number of neurons in the hidden layers of the adversary networks is 7. The adversary network in CENET also uses sigmoid function as the output function. The weight clipping norm in the adversary network of WASSERSTEINNET is 0.2.

We use the Fairlearn toolkit to implement BGL: we use the exponentiated-gradient algorithm with the default setting as the mitigator and vary the upper bound $\epsilon \in \{0.01, 0.1, 0.5, 1.0\}$ of the bounded group loss constraint. For each value of ϵ , we average the results of ten different random seeds.

As for COD, we follow the same hyper-parameter settings as (Komiyama et al., 2018): We use the kernelized optimization with the random Fourier features and the RBF kernel (we vary hyper-parameter of the RBF kernel $\gamma \in \{0.1, 1.0, 10, 100\}$) and report the best results with minimal MSE loss for each time we change the fairness budget ϵ . The hyper-parameter settings follow from (Komiyama et al., 2018). We also vary $\epsilon \in \{0.01, 0.1, 0.5, 1.0\}$ and average the results of ten different

⁴We use the edited public version of the dataset which can be download here: <https://github.com/algowatchpenn/GerryFair/blob/master/dataset/lawschool.csv>

⁵We download the public version of data here: <https://www.kaggle.com/mirichoi0218/insurance>

random seeds.

C. Additional Experimental Results and Analysis

In this section, we provide additional experimental results and analysis.

C.1. Classification Accuracy vs. Error Gaps in Adult and COMPAS Datasets

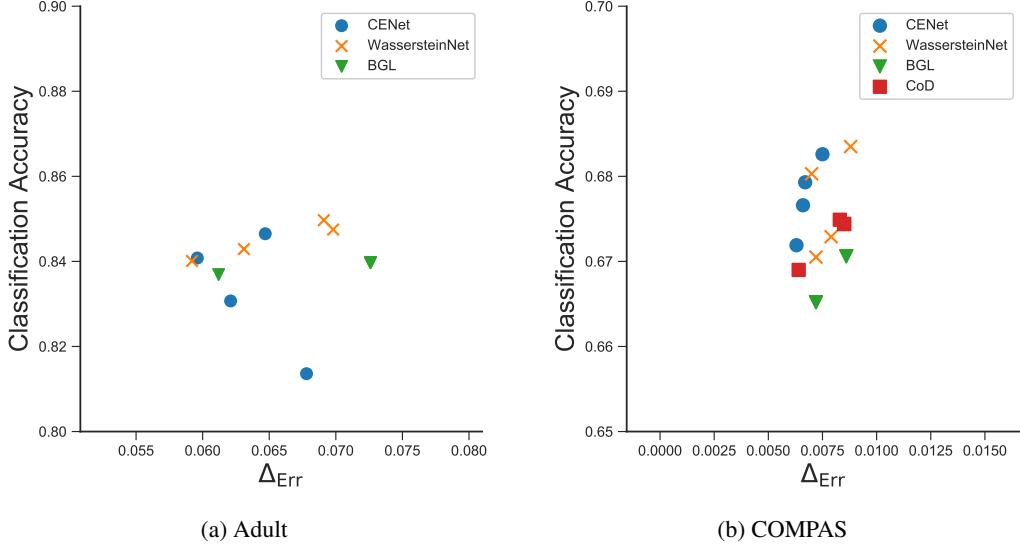


Figure 6: Classification accuracy and error gaps of different methods in Adult and COMPAS datasets.

We also report the corresponding classification accuracy for Adult and COMPAS datasets here. In Figure 6, we can see that our proposed methods achieve the best trade-offs in terms of classification accuracies and error gap values.

C.2. Impact of Fairness Trade-off in the Baseline Methods

We present additional experimental results and analyses to gain more insights into how the fairness trade-off parameters (*e.g.*, ϵ) affect the performance of the model predictive performance and accuracy disparity in baseline methods.

Table 3: R^2 regression scores and error gaps when ϵ changes in BGL.

	ϵ	0.1	0.2	0.3	0.5
Adult	R^2	0.3508	0.3696	0.3696	0.3696
	Δ_{Err}	0.0612	0.0726	0.0726	0.0726
COMPAS	ϵ	0.1	0.2	0.3	0.5
	R^2	0.1478	0.1478	0.1507	0.1507
Crime	Δ_{Err}	0.0072	0.0072	0.0086	0.0086
	ϵ	0.01	0.02	0.03	0.05
Law	R^2	0.3922	0.3922	0.5380	0.5380
	Δ_{Err}	0.0189	0.0189	0.0238	0.0238
Insurance	ϵ	0.01	0.02	0.03	0.05
	R^2	0.1407	0.1407	0.1407	0.1412
	Δ_{Err}	0.0094	0.0094	0.0094	0.0101
	ϵ	0.0001	0.01	0.05	0.1
	R^2	0.6804	0.6855	0.6855	0.6855
	Δ_{Err}	0.0145	0.0144	0.0144	0.0144

Table 3 shows R^2 regression scores and error gaps when ϵ changes in BGL. We see that with the decrease of the trade-off parameter ϵ , both the values of R^2 and error gaps decrease. This is because when the upper bound of ϵ in BGL is small, the accuracy disparity is also mitigated. When ϵ is above/below a certain threshold, R^2 scores and error gap values then increase/decrease.

Table 4: R^2 regression scores and error gaps when ϵ changes in COD.

	ϵ	0.01	0.1	0.5	1.0
COMPAS	R^2	0.1033	0.1144	0.1146	0.1146
	Δ_{Err}	0.0064	0.0083	0.0085	0.0085
	ϵ	0.01	0.1	0.5	1.0
Crime	R^2	0.1262	0.3284	0.3603	0.3603
	Δ_{Err}	0.0312	0.0307	0.0343	0.0343
	ϵ	0.01	0.1	0.5	1.0
Law	R^2	0.1262	0.3284	0.3606	0.3603
	Δ_{Err}	0.0312	0.0307	0.0343	0.0343
	ϵ	0.01	0.1	0.5	1.0
Insurance	R^2	0.2711	0.2691	0.2689	0.2689
	Δ_{Err}	0.0203	0.0210	0.0211	0.0211
	ϵ	0.01	0.1	0.5	1.0

Table 4 shows R^2 regression scores and error gaps when ϵ changes in COD. We see that with the decrease of the trade-off parameter ϵ , both the values of R^2 and error gaps decrease in general.

C.3. Visualization of Training Processes

We visualize the training processes of our proposed methods CENET and WASSERSTEINNET in the Adult dataset and COMPAS dataset in Figure 7 and Figure 8, respectively. We also compare their training dynamics with the model performance when we solely minimize the MSE loss (*i.e.*, $\lambda = 0$) and we term it as NO DEBIAS.

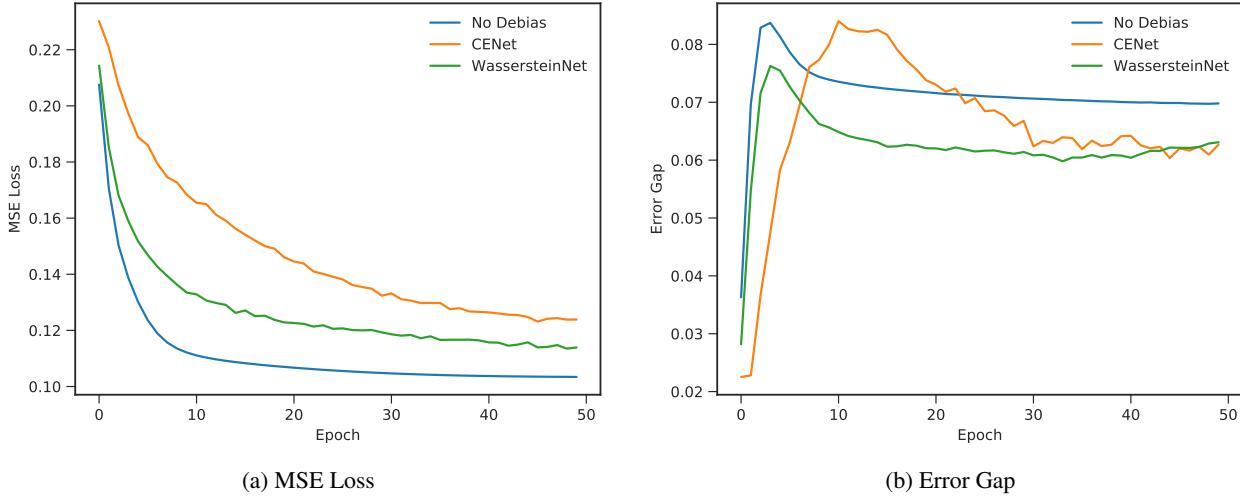


Figure 7: Training visualization of CENET, WASSERSTEINNET ($\lambda = 50$) and NO DEBIAS ($\lambda = 0$) in the Adult dataset.

In Figure 7 and Figure 8, we can see that as the training progresses go on, the MSE losses in both datasets are decreasing and finally converge. However, the training dynamics of error gaps are much more complex even in the NO DEBIAS case. Before convergence, the training dynamics of error gaps differs among different datasets. Our methods enforce the models to converge to the points where error gap are smaller while preserving the models' predictive performance. It is also worth to note that minimax optimization makes the training processes somehow unstable, especially when training CENET.

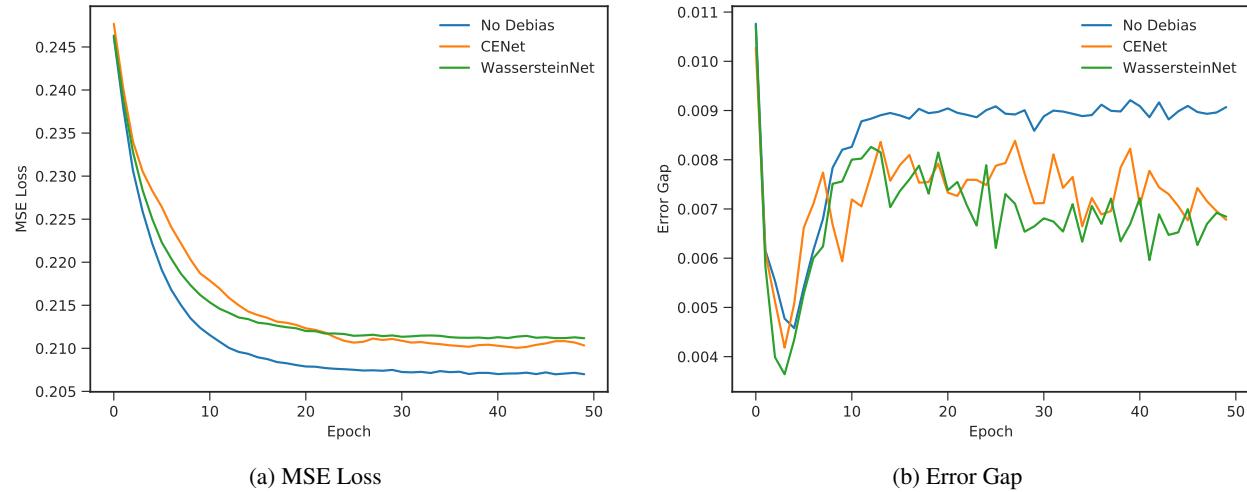


Figure 8: Training visualization of CENET, WASSERSTEINNET ($\lambda = 5$) and NO DEBIAS ($\lambda = 0$) in the COMPAS dataset.