
Stable and Fair Classification

Lingxiao Huang¹ Nisheeth K. Vishnoi²

Abstract

In a recent study, Friedler et al. (Friedler et al., 2019) observed that fair classification algorithms may not be stable with respect to variations in the training dataset – a crucial consideration in several real-world applications. Motivated by their work, this paper initiates a study of designing classification algorithms that are both fair and stable. We propose an extended framework based on fair classification algorithms that are formulated as optimization problems, by introducing a stability-focused regularization term. Theoretically, we prove a stability guarantee, that was lacking in fair classification algorithms, and also provide an accuracy guarantee for our extended framework. Our accuracy guarantee can be used to inform the selection of the regularization parameter in our framework. We assess the benefits of our approach empirically by extending several fair classification algorithms that are shown to achieve a good balance between fairness and accuracy over the **Adult** dataset, and show that our framework improves the stability at only a slight sacrifice in accuracy.

1. Introduction

Fair classification has fast become a central problem in machine learning due to concerns of bias with respect to sensitive attributes in automated decision making, e.g., against African-Americans while predicting future criminals (Flores et al., 2016; Angwin et al., 2016; Berk, 2009), granting loans (Dedman et al., 1988), or NYPD stop-and-frisk (Goel et al., 2016). Consequently, a host of fair classification algorithms have been proposed; see (Bellamy et al., 2018).

In a recent study, (Friedler et al., 2019) pointed out that several existing fair classification algorithms are not “stable”. In particular, they considered the standard deviation of a fair-

ness metric (statistical rate, that measures the discrepancy between the positive proportions of two groups; see Eq. (4)) and accuracy over ten random training-testing splits with respect to race/sex attribute over the **Adult** dataset. They observed that the standard deviation of the fairness metric is 2.4% for the algorithm in (Kamishima et al., 2012) (**KAAS**) with respect to the race attribute, and is 4.1% for that in (Zafar et al., 2017b) (**ZVRG**) with respect to the sex attribute. These significant standard deviations imply that the classifier learnt from the respective fair classification algorithms might perform differently depending on the training dataset.

Stability is a crucial consideration in classification (Bousquet & Elisseeff, 2002; Mukherjee et al., 2006; Briand et al., 2009; Fawzi et al., 2018), and has been investigated in several real-world applications, e.g., advice-giving agents (Gershoff et al., 2003; Van Swol & Sniezek, 2005), recommendation systems (Adomavicius & Zhang, 2012; 2011; 2016), and judicial decision-making (Shapiro, 1965). Stable classification algorithms can also provide defense for data poisoning attacks, whereby adversaries want to corrupt the learned model by injecting false training data (Biggio et al., 2012; Mei & Zhu, 2015; Steinhardt et al., 2017).

There is a growing number of scenarios in which stable *and* fair classification algorithms are desired. One example is recommendation systems that rely on classification algorithms (Park et al., 2012; Portugal et al., 2018). Fairness is often desired in recommendation systems, e.g., to check gender inequality in recommending high-paying jobs (Farahat & Bailey, 2012; Datta et al., 2015; Sweeney, 2013). Moreover, stability is also important for the reliability and acceptability of recommendation systems (Adomavicius & Zhang, 2012; 2011; 2016). Another example is that of a judicial decision-making system, in which fair classification algorithms are being deployed to avoid human biases for specific sensitive attributes, e.g., against African-Americans (Flores et al., 2016; Angwin et al., 2016; Berk, 2009). The dataset, that incorporates collected personal information, may be noisy due to measurement errors, privacy issues, or even data poisoning attacks (Lam & Riedl, 2004; Mobasher et al., 2007; O’mahony et al., 2004; Barreno et al., 2010) and, hence, it is desirable that the fair classifier also be stable against perturbations in the dataset.

¹EPFL, Switzerland ²Yale University, USA. Correspondence to: <nisheeth.vishnoi@yale.edu>.

1.1. Our contributions

In this paper, we initiate a study of stable and fair classifiers in automated decision-making tasks. In particular, we consider the class of fair classification algorithms that are formulated as optimization problems that minimize the empirical risk while being constrained to being fair. The collection \mathcal{F} of possible classifiers is assumed to be a reproducing kernel Hilbert space (RKHS) (see Program (ConFair) for a definition); this includes many recent fair classifiers such as (Zafar et al., 2017a;b; Goel et al., 2018). Our main contribution is an algorithmic framework that incorporates the notion of uniform stability (Bousquet & Elisseeff, 2002) – the maximum l_∞ -distance between the risks of two classifiers learned from two training sets that differ in a single sample (see Definition 2.1). This allows us to address the stability issue observed by (Friedler et al., 2019). To achieve uniform stability, we introduce a stability-focused regularization term to the objective function of fair classifier (Program (Stable-Fair)), which is motivated by the work of (Bousquet & Elisseeff, 2002). Although some existing fair classification algorithms (Kamishima et al., 2012; Goel et al., 2018) use regularizers, they do not seem to realize that (and show how) the regularization term can also make the algorithm more stable. Under mild assumptions on the loss function (Definition 3.1), we prove that our extended framework indeed has an additional uniform stability guarantee $\tilde{O}(\frac{1}{\lambda N})$, where λ is the regularization parameter and N is the size of the training set (Theorems 3.2). Moreover, if \mathcal{F} is a linear model, we can achieve a slightly better stability guarantee (Theorem 3.7). Our stability guarantee also implies an empirical risk guarantee that can be used to inform the selection of the regularization parameter in our framework. By letting $\lambda = \Theta(\frac{1}{\sqrt{N}})$, the increase in the empirical risk by introducing the regularization term can be bounded by $\tilde{O}(\frac{1}{\sqrt{N}})$ (Theorems 3.2 and 3.7, Remark 3.3). As a consequence, our stability guarantee also implies a generalization bound – the expected difference between the expected risk and the empirical risk is $\tilde{O}(\frac{1}{\lambda N})$ (Corollaries 3.6 and 3.8).

Further, we conduct an empirical evaluation over the **Adult** dataset and apply our framework to several fair classification algorithms, including **KAAS** (Kamishima et al., 2012), **ZVRG** (Zafar et al., 2017a) and **GYF** (Goel et al., 2018) (Section 4). Similar to (Friedler et al., 2019), we evaluate the fairness metric and accuracy of these algorithms and our extended algorithms. Besides, we also compute the expected number of different predictions over the test dataset between classifiers learned from two random training sets as a stability measure stab (Eq. (3)). The empirical results show that our classification algorithms indeed achieve better stability guarantee, while being fair. For instance, with respect to the sex attribute, the standard deviation of the fairness metric of **ZVRG** improves from 4.1% ((Friedler et al., 2019)) to about 1% using our extended algorithm,

and the stability measure stab decreases from 70 ($\lambda = 0$) to 25 ($\lambda = 0.02$). Meanwhile, the loss in accuracy due to imposing stability-focused regularization term is small (at most 1.5%).

Overall, we provide the first extended framework for stable and fair classification, which makes it flexible and easy to use, slightly sacrifices accuracy, and performs well in practice.

1.2. Other related work

From a technical view, most relevant prior works formulated the fair classification problem as a constrained optimization problem, e.g., constrained to statistical parity (Zafar et al., 2017b; Menon & Williamson, 2018; Goel et al., 2018; Celis et al., 2019b), or equalized odds (Hardt et al., 2016a; Zafar et al., 2017a; Menon & Williamson, 2018; Celis et al., 2019b). Our extended framework can be applied to this type of fair classification. Another approach for fair classification is to shift the decision boundary of a baseline classifier, e.g., (Fish et al., 2016; Hardt et al., 2016a; Goh et al., 2016; Pleiss et al., 2017; Woodworth et al., 2017; Dwork et al., 2018). Finally, a different line of research pre-processes the training data with the goal of removing the bias for learning, e.g., (Kamiran & Calders, 2009; Luong et al., 2011; Kamiran & Calders, 2012; Zemel et al., 2013; Feldman et al., 2015; Krasanakis et al., 2018).

Several prior works (Bousquet & Elisseeff, 2002; Shalev-Shwartz et al., 2010; Maurer, 2017; Meng et al., 2017) study the stability property for empirical risk minimization. (Hardt et al., 2016b), (London, 2016) and (Kuzborskij & Lampert, 2018) showed that the stochastic gradient descent method is stable. Moreover, several recent works studied stability in deep neural networks (Raghu et al., 2017; Vidal et al., 2017). Stability has been investigated in other automated decision-making tasks, e.g., feature selection (Nogueira et al., 2018) and structured prediction (London et al., 2013; 2014; 2016).

There exists a related notion to stability, called differential privacy, where the prediction for any sample should not change with high probability if the training set varies a single element. By (Wang et al., 2016), differential privacy implies a certain stability guarantee. Hence, it is possible to achieve stable and fair classifiers by designing algorithms that satisfy differential privacy and fairness simultaneously. Recent studies (Hajian et al., 2016; 2015; Kashid et al., 2015; 2017; Ruggieri et al., 2014) have expanded the application of methods to achieve both goals; see a recent paper (Ekstrand et al., 2018) for more discussions. However, these methods are almost all heuristic and without theoretical guarantee. There also remains the open problem of characterizing under what circumstances and definitions, privacy and fairness are simultaneously achievable, and when they compete with each other.

2. Our model

2.1. Preliminaries

We consider the Bayesian model for classification, in which \mathfrak{S} denote a joint distribution over the domain $\mathcal{D} = \mathcal{X} \times [p] \times \{-1, 1\}$ where \mathcal{X} is the feature space. Each sample (X, Z, Y) is drawn from \mathfrak{S} where $Z \in [p]$ represents a sensitive attribute,¹ and $Y \in \{-1, 1\}$ is the label of (X, Z) that we want to predict.

Let \mathcal{F} denote the collection of all possible classifiers $f : \mathcal{X} \rightarrow \mathbb{R}$. Given a loss function $L(\cdot, \cdot)$ that takes a classifier f and a distribution \mathfrak{S} as arguments, the goal of fair classification is to learn a classifier $f \in \mathcal{F}$ that minimizes the expected risk $R(f) := \mathbb{E}_{s \sim \mathfrak{S}} [L(f, s)]$. However, since \mathfrak{S} is often unknown, we usually use the empirical risk to estimate the expected risk (Bousquet & Elisseeff, 2002; Shalev-Shwartz et al., 2010; Maurer, 2017), i.e., given a training set $S = \{s_i = (x_i, z_i, y_i)\}_{i \in [N]}$ (where $(x_i, z_i, y_i) \in \mathcal{D}$), the objective is to learn a classifier $f \in \mathcal{F}$ that minimizes the empirical risk $E(f) := \frac{1}{N} \sum_{i \in [N]} L(f, s_i)$. Denote by $\Pr_{\mathfrak{S}}[\cdot]$ the probability with respect to \mathfrak{S} . If \mathfrak{S} is clear from context, we simply denote $\Pr_{\mathfrak{S}}[\cdot]$ by $\Pr[\cdot]$. A fair classification algorithm \mathcal{A} can be considered as a mapping $\mathcal{A} : \mathcal{D}^* \rightarrow \mathcal{F}$, which learns a classifier $\mathcal{A}_S \in \mathcal{F}$ from a training set $S \in \mathcal{D}^*$.

2.2. Stability measure

In this paper, we focus on the following stability measure introduced by (Bousquet & Elisseeff, 2002), which was also used by (Shalev-Shwartz et al., 2010; Maurer, 2017; Meng et al., 2017). This notion of stability measures whether the risk of the learnt classifier is stable under replacing one sample in the training dataset.

Definition 2.1 (Uniform stability (Bousquet & Elisseeff, 2002)). *Given an integer N , a real-valued classification algorithm \mathcal{A} is β_N -uniformly stable with respect to the loss function $L(\cdot, \cdot)$ if the following holds: for all $i \in [N]$ and $S, S^i \in \mathcal{D}^N$, $\|L(\mathcal{A}_S, \cdot) - L(\mathcal{A}_{S^i}, \cdot)\|_{\infty} \leq \beta_N$, i.e., for any training set $S, S^i \in \mathcal{D}^N$, the l_{∞} -distance between the risks of \mathcal{A}_S and \mathcal{A}_{S^i} is at most β_N .*

By definition, algorithm \mathcal{A} is stable if β_N is small.

Since classification algorithms usually minimize the empirical risk, it is easier to bound to provide theoretical bounds on the risk difference. This is the reason we consider the notion of uniform stability. Moreover, uniform stability might imply that the prediction variation is small with a slight perturbation on the training set. Given an algorithm \mathcal{A} and a sample $x \in \mathcal{X}$, we predict the label to be +1 if $\mathcal{A}(x) \geq 0$ and to be -1 if $\mathcal{A}(x) < 0$. In the following, we summarize

¹Our results can be generalized to multiple sensitive attributes Z_1, \dots, Z_m where $Z_i \in [p_i]$. We omit the details.

the stability property considered in (Friedler et al., 2019).

Definition 2.2 (Prediction stability). *Given an integer N , a real-valued classification algorithm \mathcal{A} is β_N -prediction stable if the following holds: for all $i \in [N]$,*

$$\Pr_{S, S^i \in \mathcal{D}^N, X \sim \mathfrak{S}} [\mathbb{I}[\mathcal{A}_S(X) \geq 0] \neq \mathbb{I}[\mathcal{A}_{S^i}(X) \geq 0]] \leq \beta_N,$$

i.e., given two training sets $S, S^i \in \mathcal{D}^N$ that differ by a single sample, the probability that \mathcal{A}_S and \mathcal{A}_{S^i} predict differently is at most β_N .

The following lemma shows that uniform stability implies prediction stability.

Lemma 2.3 (Uniform stability implies prediction stability). *Given an integer N , if algorithm \mathcal{A} is β_N -uniformly stable with respect to the loss function $L(\cdot, \cdot)$ and the loss function satisfies that for any $f, f' \in \mathcal{F}$, $s = (x, z, y) \in \mathcal{D}$,*

$$|f(x) - f'(x)| \leq \tau \cdot |L(f, s) - L(f', s)|,$$

then the prediction stability \mathcal{A} is upper bounded by $\Pr_{S, \mathcal{A}} [|\mathcal{A}_S(X)| \leq \tau \beta_N]$.

Proof. For any $S, S^i \in \mathcal{D}^N$ and $s = (x, z, y)$, we have

$$|\mathcal{A}_S(x) - \mathcal{A}_{S^i}(x)| \leq \tau \cdot |L(\mathcal{A}_S, s) - L(\mathcal{A}_{S^i}, s)| \leq \tau \beta_N.$$

Hence, if $|\mathcal{A}_S(x)| > \tau \beta_N$, then we have $\mathbb{I}[\mathcal{A}_S(x) \geq 0] = \mathbb{I}[\mathcal{A}_{S^i}(x) \geq 0]$. By Definition 2.2, this implies the lemma. \square

2.3. The stable and fair optimization problem

Our goal is to design fair classification algorithms that have a uniform stability guarantee. We focus on extending fair classification algorithms that are formulated as constrained empirical risk minimization problem over the collection \mathcal{F} of classifiers (that is a reproducing kernel Hilbert space (RKHS), e.g., (Zafar et al., 2017a;b; Goel et al., 2018)); see the following program.

$$\begin{aligned} \min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i \in [N]} L(f, s_i) \quad s.t. \\ \Omega(f) \leq 0. \end{aligned} \quad (\text{ConFair})$$

Here, $\Omega(\cdot) : \mathcal{F} \rightarrow \mathbb{R}^a$ is a convex function given explicitly for a specific fairness requirement. For instance, if we consider the statistical rate $\gamma(f)$ (Eq. (4)) as the fairness metric, then the fairness requirement can be $0.8 - \gamma(f) \leq 0$. However, $0.8 - \gamma(f)$ is non-convex with respect to f . To address this problem, in the literature, one usually defines

²Here, $\mathbb{I}[\cdot]$ is the indicator function.

a convex function $\Omega(f)$ to estimate $0.8 - \gamma(f)$, e.g., $\Omega(f)$ is formulated as a covariance-type function which is the average signed distance from the feature vectors to the decision boundary in (Zafar et al., 2017b), and is formulated as the weighted sum of the logs of the empirical estimate of favorable bias in (Goel et al., 2018). In what follows, a fair classification algorithm is an algorithm that solves Program (ConFair).

Note that an empirical risk minimizer of Program (ConFair) might heavily depend on and even overfit the training set. Hence, replacing a sample from the training set might cause a significant change in the learnt fair classifier – the uniform stability guarantee might be large. To address this problem, a useful high-level idea is to introduce a regularization term to the objective function, which can penalize the “complexity” of the learned classifier. Intuitively, this can make the change in the learnt classifier smaller when a sample from the training set is replaced. This idea comes from (Bousquet & Elisseeff, 2002) who considered stability for unconstrained empirical risk minimization.

Motivated by the above intuition, we consider the following constrained optimization problem which is an extension of Program (ConFair) by introducing a stability-focused regularization term $\lambda \|f\|_k^2$. Here, $\lambda > 0$ is a regularization parameter and $\|f\|_k^2$ is the norm of f in RKHS \mathcal{F} where k is the kernel function (defined later in Definition 2.5). We consider such a regularization term since it satisfies a nice property that relates $|f(x)|$ and $\|f\|_k$ for any $x \in \mathcal{X}$ (Claim 2.6). This property is useful for proving making the intuition above concrete and providing a uniform stability guarantee.

$$\min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i \in [N]} L(f, s_i) + \lambda \|f\|_k^2 \quad \text{s.t.}$$

$$\Omega(f) \leq 0.$$

(Stable-Fair)

Our extended algorithm \mathcal{A} is to compute a minimizer \mathcal{A}_S of Program (Stable-Fair) by classic methods, e.g., stochastic gradient descent (Boyd & Mutapcic, 2008).

Remark 2.4. We discuss the motivation of considering fair classification algorithms formulated as Program (ConFair). The main reason is that such algorithms can achieve a good balance between fairness and accuracy, but might not be stable. For instance, (Friedler et al., 2019) observed that ZVRG (Zafar et al., 2017b) achieves better fairness than, and comparable accuracy to, other algorithms with respect to race/sex attribute over the **Adult** dataset. However, as mentioned in Section 1, ZVRG is not stable depending on a random training set. Hence, we would like to improve

the stability of ZVRG while keeping its balance between fairness and accuracy. Note that our extended framework can incorporate multiple sensitive attributes if the fairness constraint $\Omega(f) \leq 0$ deals with multiple sensitive attributes, e.g., (Zafar et al., 2017a;b; Goel et al., 2018).

It remains to define the regularization term $\|f\|_k$ in RKHS.

Definition 2.5 (Regularization in RKHS). We call $T(\cdot) : \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$ a regularization term in an RKHS \mathcal{F} if, for any $f \in \mathcal{F}$, $T(f) := \|f\|_k^2$, where k is a kernel function satisfying that 1) $\{k(x, \cdot) : x \in \mathcal{X}\}$ is a span of \mathcal{F} ; 2) for any $x \in \mathcal{X}$ and $f \in \mathcal{F}$, $f(x) = \langle f, k(x, \cdot) \rangle$.

Given a training set $S = \{s_i = (x_i, z_i, y_i)\}_{i \in [N]}$ and a kernel function $k : S \times S \rightarrow \mathbb{R}$, by definition, each classifier is a vector space by linear combinations of $k(x_i, \cdot)$, i.e., $f(\cdot) = \sum_{i \in [N]} \alpha_i k(x_i, \cdot)$. Then for any $x \in \mathcal{X}$,

$$f(x) = \left\langle \sum_{i \in [N]} \alpha_i k(x_i, \cdot), k(x, \cdot) \right\rangle = \sum_{i \in [N]} \alpha_i k(x_i, x). \quad (1)$$

For instance, if $k(x, y) = \langle x, y \rangle$, then each classifier f can be represented by

$$\begin{aligned} f(x) &\stackrel{(1)}{=} \sum_{i \in [N]} \alpha_i k(x_i, x) = \sum_{i \in [N]} \alpha_i \langle x_i, x \rangle \\ &= \left\langle \sum_{i \in [N]} \alpha_i x_i, x \right\rangle = \langle \beta, x \rangle, \end{aligned}$$

where $\beta = \sum_{i \in [N]} \alpha_i x_i$. Thus, by the Cauchy-Schwarz inequality, we have the following useful property.

Claim 2.6. ((Bousquet & Elisseeff, 2002)) Suppose \mathcal{F} is a RKHS with a kernel k . For any $f \in \mathcal{F}$ and any $x \in \mathcal{X}$, we have $|f(x)| \leq \|f\|_k \sqrt{k(x, x)}$.

Remark 2.7. There exists another class of fair classification algorithms, which introduce a fairness-focused regularization term $\mu \cdot \Omega(\cdot)$ to the objective function; see the following program.

$$\min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i \in [N]} L(f, s_i) + \mu \cdot \Omega(f). \quad (\text{RegFair})$$

This approach is applied in several prior work, e.g., (Kamishima et al., 2012; Corbett-Davies et al., 2017; Goel et al., 2018). We can also extend this program by introducing a stability-focused regularization term $\lambda \|f\|_k^2$.

$$\min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i \in [N]} L(f, s_i) + \mu \cdot \Omega(f) + \lambda \|f\|_k^2.$$

By Lagrangian principle, there exists a value $\mu \geq 0$ such that Program (RegFair) is equivalent to Program (ConFair). Thus, by solving the above program, we can obtain the same stability guarantee, empirical risk guarantee and generalization bound as for Program (Stable-Fair).

3. Theoretical results

In this section, we analyze the performance of algorithm \mathcal{A} that solves Program (Stable-Fair) (Theorem 3.2). Moreover, if \mathcal{F} is a linear model, we can achieve a slightly better stability guarantee (Theorem 3.7) based on different assumptions of both the model and the loss function.

Given a training set $S = \{s_i = (x_i, z_i, y_i)\}_{i \in [N]}$ by replacing the i -th element from S , we denote

$$S^i := \{s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_N\}.$$

Before analyzing the performance of algorithm \mathcal{A} , we give the following definition for a loss function.

Definition 3.1 (σ -admissible (Bousquet & Elisseeff, 2002)). *The loss function $L(\cdot, \cdot)$ is called σ -admissible with respect to \mathcal{F} if for any $f \in \mathcal{F}$, $x, x' \in \mathcal{X}$ and $y \in \{-1, 1\}$,*

$$|L(f(x), y) - L(f(x'), y)| \leq \sigma |f(x) - f(x')|.$$

By definition, $L(\cdot, \cdot)$ is σ -admissible if $L(f, y)$ is σ -Lipschitz with respect to f .

3.1. Main theorem for Program (Stable-Fair)

Now we can state our main theorem which indicates that under reasonable assumptions of the loss function and the kernel function, algorithm \mathcal{A} is uniformly stable.

Theorem 3.2 (Stability and empirical risk guarantee by solving Program (Stable-Fair)). *Let \mathcal{F} be a RKHS with kernel k such that $\forall x \in \mathcal{X}$, $k(x, x) \leq \kappa^2 < \infty$. Let $L(\cdot, \cdot)$ be a σ -admissible differentiable function with respect to \mathcal{F} . Suppose algorithm \mathcal{A} computes a minimizer \mathcal{A}_S of Program (Stable-Fair). Then \mathcal{A} is $\frac{\sigma^2 \kappa^2}{\lambda N}$ -uniformly stable.*

Moreover, denote f^* to be an optimal fair classifier that minimizes the expected risk and satisfies $\|f^*\|_k \leq B$, i.e., $f^* := \arg \min_{f \in \mathcal{F}: \Omega(f) \leq 0} \mathbb{E}_{s \in \mathfrak{S}} [L(f, s)]$. We have

$$\mathbb{E}_{S \sim \mathfrak{S}^N} [R(\mathcal{A}_S)] - \mathbb{E}_{s \sim \mathfrak{S}} [L(f^*, s)] \leq \frac{\sigma^2 \kappa^2}{\lambda N} + \lambda B^2.$$

Remark 3.3. *We show the assumptions of Theorem 3.2 are reasonable. We first give some examples of $L(\cdot, \cdot)$ in which σ is constant. In the main body, we directly give the constant due to the space limit. The details can be found in Appendix C. 1) Prediction error: Suppose $f(x) \in \{-1, 1\}$ for any pair (f, x) . Then $L(f(x), y) = \mathbb{I}[f(x) \neq y]$ is $\frac{1}{2}$ -admissible. 2) Soft margin SVM: $L(f, s) = (1 - yf(x))_+^3$*

³ $(a)_+ = a$ if $a \geq 0$ and otherwise $(a)_+ = 0$.

is 1-admissible. 3) Least Squares regression: Suppose $f(x) \in [-1, 1]$ for any $x \in \mathcal{X}$. Then we have that $L(f, s) = (f(x) - y)^2$ is 4-admissible. 4) Logistic regression: $L(f, s) = \ln(1 + e^{-yf(x)})$ is 1-admissible.

Then we give examples of kernel k in which κ^2 is constant. 1) Linear: $k(x, y) = \langle x, y \rangle$. Then $k(x, x) = \|x\|_2^2$ and we can let $\kappa^2 = \max_{x \in \mathcal{X}} \|x\|_2^2$. 2) Gaussian RBF: $k(x, y) = e^{-\|x-y\|^2}$. Then we can let $\kappa^2 = k(x, x) = 1$. 3) Multiquadric: $k(x, y) = (\|x - y\|^2 + c^2)^{1/2}$ for some constant $c > 0$. Then we can let $\kappa^2 = k(x, x) = c$. 4) Inverse Multiquadric: $k(x, y) = (\|x - y\|^2 + c^2)^{-1/2}$ for some constant $c > 0$. Then we can let $\kappa^2 = k(x, x) = 1/c$.

Remark 3.4. *The statement of Theorem 3.2 seems similar to Lemma 4.1 of (Bousquet & Elisseeff, 2002), while the analysis should be different due to the additional fairness constraints. The critical difference is that the gradient of the objective function of Program 2 might not be 0 at the optimal point any more. Thus, we need to develop a new analysis by applying the convexity of $\Omega(f)$.*

Theorem 3.2 can be used to inform the selection of the regularization parameter λ . On the one hand, the stability guarantee is tighter as λ increases. On the other hand, the bound for the increase of the empirical risk contains a term λB^2 and hence λ should not increase to infinity. Hence, there exists a balance between achieving stability guarantee and utility guarantee. For instance, to minimize the increase of the empirical risk, we can set $\lambda = \frac{\sigma \kappa}{B \sqrt{N}}$. Then the stability guarantee is upper bounded by $\frac{\sigma \kappa B}{\sqrt{N}}$ and the increase of the empirical risk is upper bounded by $\frac{2\sigma \kappa B}{\sqrt{N}}$.

Now we are ready to prove Theorem 3.2. For convenience, we define $g = \mathcal{A}_S$ and $g^i = \mathcal{A}_{S^i}$. We first give Lemma 3.5 for preparation. This lemma is the one of the places different from the argument in (Bousquet & Elisseeff, 2002) since our framework includes a fairness constraint. To prove the lemma, we need to use the fact that $\Omega(f)$ is a convex function of $f \in \mathcal{F}$.

Lemma 3.5. *For any $i \in [N]$, we have*

$$\|g - g^i\|_k^2 \leq \frac{\sigma}{2\lambda N} (|g(x_i) - g^i(x_i)| + |g(x'_i) - g^i(x'_i)|).$$

Due to the space limit, we defer the proof of Lemma 3.5 to Appendix A. Roughly speaking, we use the fact that $|g - g^i|_k^2$ is equivalent to the Bregman divergence between g and g^i . Then by the fact that $\Omega(f)$ is convex, we can upper bound the Bregman divergence by the right side of the inequality. Combining Lemma 3.5 and Claim 2.6, we can upper bound $|g(x) - g^i(x)|$ for any $x \in \mathcal{X}$. This implies a uniform stability guarantee by the assumption that $L(\cdot, \cdot)$ is σ -admissible.

Proof of Theorem 3.2. By Claim 2.6, we have

$$|g(x_i) - g^i(x_i)| \leq \|g - g^i\|_k \sqrt{k(x_i, x_i)} \leq \kappa \|g - g^i\|_k,$$

$$|g(x'_i) - g^i(x'_i)| \leq \|g - g^i\|_k \sqrt{k(x'_i, x'_i)} \leq \kappa \|g - g^i\|_k.$$

Combining the above inequalities with Lemma 3.5, we have $\|g - g^i\|_k \leq \frac{\sigma\kappa}{\lambda N}$. Hence, for any sample $s = (x, z, y) \in \mathcal{D}$, we have $|g(x) - g^i(x)| \leq \kappa \|g - g^i\|_k \leq \frac{\sigma\kappa^2}{\lambda N}$. Moreover, since $L(\cdot, \cdot)$ is σ -admissible, we have

$$|L(g, s) - L(g^i, s)| \leq \sigma |g(x) - g^i(x)| \leq \frac{\sigma^2 \kappa^2}{\lambda N}.$$

By definitions of g and g^i , the above inequality completes the proof of stability guarantee.

For the increase of the empirical risk, let $F(f) := \frac{1}{N} \sum_{i \in [N]} L(f, s_i) + \lambda \|f\|_k^2$ for any $f \in \mathcal{F}$. By Theorem 8 of (Shalev-Shwartz et al., 2010), we have the following claim: for any classifier $h \in \mathcal{F}$ satisfying that $\Omega(h) \leq 0$, $\mathbb{E}_{S \sim \mathfrak{S}^N} [F(g) - F(h)]$ is consistent with the uniform stability guarantee of \mathcal{A} , i.e.,

$$\mathbb{E}_{S \sim \mathfrak{S}^N} [F(g) - F(h)] \leq \frac{\sigma^2 \kappa^2}{\lambda N}. \quad (2)$$

Let $h = f^*$, we have

$$\begin{aligned} & \mathbb{E}_{S \sim \mathfrak{S}^N} [R(\mathcal{A}_S)] - \mathbb{E}_{s \sim \mathfrak{S}} [L(f^*, s)] \\ &= \mathbb{E}_{S \sim \mathfrak{S}^N} \left[R(\mathcal{A}_S) - \frac{1}{N} \sum_{i \in [N]} L(f^*, s_i) \right] \\ &= \mathbb{E}_{S \sim \mathfrak{S}^N} [F(g) - \lambda \|g\|_k^2 - F(f^*) + \lambda \|f^*\|_k^2] \\ & \quad (\text{Defns. of } g \text{ and } F(\cdot)) \\ & \leq \mathbb{E}_{S \sim \mathfrak{S}^N} [F(g) - F(f^*)] + \lambda \|f^*\|_k^2 \quad (\|g\|_k^2 \geq 0) \\ & \leq \frac{\sigma^2 \kappa^2}{\lambda N} + \lambda B^2 \quad (\text{Ineq. (2) and } \|f^*\|_k \leq B). \end{aligned}$$

This completes the proof. \square

The generalization bound, i.e., the quality of the estimation $|R(\mathcal{A}_S) - E(\mathcal{A}_S)|$, depends on the number of samples N and algorithm \mathcal{A} , and has been well studied in the literature (Adomavicius & Zhang, 2011; Bousquet & Elisseeff, 2002; Wainwright, 2006; London et al., 2016). Existing literature (Bousquet & Elisseeff, 2002; Feldman & Vondrak, 2018) mainly claimed that uniform stability implies a generalization bound. We have the following corollary.

Corollary 3.6 (Generalization bound from Theorem 3.2). *Let \mathcal{A} denote the $\frac{\sigma^2 \kappa^2}{\lambda N}$ -uniformly stable algorithm as stated in Theorem 3.2. We have*

$$1. \mathbb{E}_{S \sim \mathfrak{S}^N} [R(\mathcal{A}_S) - E(\mathcal{A}_S)] \leq \frac{\sigma^2 \kappa^2}{\lambda N}.$$

2. *Suppose S is a random draw of size N from \mathfrak{S} . With probability at least $1 - \delta$,*

$$R(\mathcal{A}_S) \leq E(\mathcal{A}_S) + 8 \sqrt{\left(\frac{2\sigma^2 \kappa^2}{\lambda N} + \frac{1}{N} \right) \cdot \ln(8/\delta)}.$$

Proof. The first generalization bound is directly implied by Lemma 7 of (Bousquet & Elisseeff, 2002). The second generalization bound is a direct corollary of Theorem 1.2 of (Feldman & Vondrak, 2018). \square

3.2. Better stability guarantee for linear models

In this section, we consider the case that $k(x, y) = \langle \phi(x), \phi(y) \rangle$ where $\phi: \mathcal{X} \rightarrow \mathbb{R}^d$ is a feature map. It implies that $f(x) = \alpha^\top \phi(x)$ for some $\alpha \in \mathbb{R}^d$, i.e., \mathcal{F} is the family of all linear functions. In this case, we provide a stronger stability guarantee by the following theorem.

Theorem 3.7 (Stability and utility guarantee for linear models). *Let \mathcal{F} be the family of all linear classifiers $\{f = \alpha^\top \phi(\cdot) \mid \alpha \in \mathbb{R}^d\}$. Let $G = \sup_{f = \alpha^\top \phi(\cdot) \in \mathcal{F}: \Omega(f) \leq 0} \sup_{s \in \mathcal{D}} \|\nabla_\alpha L(f, s)\|_2$. Suppose algorithm \mathcal{A} computes a minimizer \mathcal{A}_S of Program (Stable-Fair). Then \mathcal{A} is $\frac{G^2}{\lambda N}$ -uniformly stable.*

Moreover, denote f^ to be an optimal fair classifier that minimizes the expected risk and satisfies $\|f^*\|_k \leq B$, i.e., $f^* := \arg \min_{f \in \mathcal{F}: \Omega(f) \leq 0} \mathbb{E}_{s \in \mathfrak{S}} [L(f, s)]$. We have*

$$\mathbb{E}_{S \sim \mathfrak{S}^N} [R(\mathcal{A}_S)] - \mathbb{E}_{s \sim \mathfrak{S}} [L(f^*, s)] \leq \frac{G^2}{\lambda N} + \lambda B^2.$$

Note that we only have an assumption for the gradient of the loss function. Given a sample $s = (x, z, y) \in \mathcal{D}$ such that $G = \sup_{f \in \mathcal{F}: \Omega(f) \leq 0} \|\nabla_\alpha L(f, s)\|_2$, we have

$$G = \|\nabla_\alpha L(f, s)\|_2 = \|\nabla_f L(f, s) \cdot \phi(x)\|_2.$$

Under the assumption of Theorem 3.2, we have 1) $\|\nabla_f L(f, s)\| \leq \sigma$ since $L(\cdot, \cdot)$ is σ -admissible with respect to \mathcal{F} ; 2) $\|\phi(x)\|_2 = \sqrt{k(x, x)} \leq \kappa$. Hence, $G \leq \sigma\kappa$ which implies that Theorem 3.2 is stronger than Theorem 3.7 for linear models. The proof idea is similar to that of Theorem 3.2 and hence we defer the proof to Appendix B. Moreover, we directly have the following corollary similar to Corollary 3.6.

Corollary 3.8 (Generalization bound by Theorem 3.7).

Let \mathcal{A} denote the $\frac{G^2}{\lambda N}$ -uniformly stable algorithm as stated in Theorem 3.7. We have

$$1. \mathbb{E}_{S \sim \mathfrak{S}^N} [R(\mathcal{A}_S) - E(\mathcal{A}_S)] \leq \frac{G^2}{\lambda N}.$$

2. *Suppose S is a random draw of size N from \mathfrak{S} . With probability at least $1 - \delta$,*

$$R(\mathcal{A}_S) \leq E(\mathcal{A}_S) + 8 \sqrt{\left(\frac{2G^2}{\lambda N} + \frac{1}{N} \right) \cdot \ln(8/\delta)}.$$

Table 1: The performance (mean and standard deviation in parenthesis), of **KAAS-St** and **ZVRG-St** with respect to accuracy and the fairness metrics γ on the **Adult** dataset with race/sex attribute.

			λ					
			0	0.01	0.02	0.03	0.04	0.05
ZVRG-St	Race	Acc.	0.844(0.001)	0.842(0.001)	0.841(0.001)	0.840(0.001)	0.838(0.001)	0.838(0.001)
		γ	0.577(0.031)	0.667(0.020)	0.686(0.015)	0.711(0.016)	0.743(0.013)	0.761(0.012)
	Sex	Acc.	0.844(0.001)	0.840(0.001)	0.838(0.001)	0.838(0.001)	0.837(0.001)	0.836(0.001)
		γ	0.331(0.041)	0.501(0.011)	0.495(0.009)	0.478(0.009)	0.463(0.009)	0.469(0.009)
KAAS-St	Race	Acc.	0.850(0.001)	0.844(0.001)	0.843(0.001)	0.839(0.001)	0.837(0.001)	0.835(0.001)
		γ	0.571(0.019)	0.359(0.024)	0.302(0.011)	0.301(0.011)	0.300(0.015)	0.298(0.015)
	Sex	Acc.	0.850(0.002)	0.848(0.001)	0.844(0.001)	0.839(0.001)	0.837(0.001)	0.835(0.001)
		γ	0.266(0.011)	0.226(0.011)	0.165(0.008)	0.136(0.007)	0.128(0.006)	0.128(0.005)
GYF-St	Race	Acc.	0.849(0.001)	0.845(0.001)	0.844(0.001)	0.842(0.001)	0.840(0.001)	0.835(0.001)
		γ	0.558(0.020)	0.679(0.013)	0.690(0.017)	0.710(0.018)	0.740(0.014)	0.753(0.013)
	Sex	Acc.	0.850(0.002)	0.845(0.001)	0.844(0.001)	0.842(0.001)	0.840(0.001)	0.839(0.001)
		γ	0.275(0.010)	0.245(0.004)	0.242(0.004)	0.241(0.005)	0.245(0.005)	0.234(0.008)

4. Empirical results

4.1. Empirical setting

Algorithms and baselines. We select three fair classification algorithms designed to ensure statistical parity that can be formulated in the convex optimization framework of Program (ConFair). We choose **ZVRG** (Zafar et al., 2017a) since it is reported to achieve the better fairness than, and comparable accuracy to, other algorithms (Friedler et al., 2019). We also select **KAAS** (Kamishima et al., 2012) and **GYF** (Goel et al., 2018) as representatives of algorithms that are formulated as Program (RegFair). Specifically, (Goel et al., 2018) showed that the performance of **GYF** is comparable to **ZVRG** over the **Adult** dataset. We extend them by introducing a stability-focused regularization term.

- **ZVRG** (Zafar et al., 2017b). Zafar et al. re-express fairness constraints (which can be nonconvex) via a convex relaxation. This allows them to maximize accuracy subject to fairness constraints.⁴ We denote the extended, stability included, algorithm by **ZVRG-St**.
- **KAAS** (Kamishima et al., 2012). Kamishima et al. introduce a fairness-focused regularization term and apply it to a logistic regression classifier. Their approach requires numerical input and a binary sensitive attribute. Let the extended algorithm be **KAAS-St**.
- **GYF** (Goel et al., 2018). Goel et al. introduce *negative weighted sum of logs* as fairness-focused regularization term and apply it to a logistic regression classifier. Let the extended algorithm be **GYF-St**.

Dataset. Our simulations are over an income dataset **Adult** (Dheeru & Karra Taniskidou, 2017), that records the demographics of 45222 individuals, along with a binary label indicating whether the income of an individual is

⁴There exists a threshold parameter in the constraints. In this paper, we set the parameter to be default 0.1.

greater than 50k USD or not. We use the pre-processed dataset as in (Friedler et al., 2019). We take race and sex to be the sensitive attributes, that are binary in the dataset.

Stability metrics. The following stability metric that measures the prediction difference between classifiers learnt from two random training sets. Given an integer N , a testing set T and algorithm \mathcal{A} , we define $\text{stab}_T(\mathcal{A})$ as $|T| \cdot \Pr_{S, S' \sim \mathfrak{S}^N, X \sim T, \mathcal{A}} [\mathbb{I}[\mathcal{A}_S(X) \geq 0] \neq \mathbb{I}[\mathcal{A}_{S'}(X) \geq 0]]$.

$\text{stab}_T(\mathcal{A})$ indicates the expected number of different predictions of \mathcal{A}_S and $\mathcal{A}_{S'}$ over the testing set T . Note that this metric is considered in (Friedler et al., 2019), but is slightly different from prediction stability since S and S' may differ by more than one training sample. We investigate $\text{stab}_T(\mathcal{A})$ instead of prediction stability so that we can distinguish the performances of prediction difference under different regularization parameters. Since \mathfrak{S} is unknown, we generate n training sets S_1, \dots, S_n and use the following metric to estimate $\text{stab}_T(\mathcal{A})$:



$$\text{stab}_{T,n}(\mathcal{A}) := \frac{1}{n(n-1)} \sum_{i,j \in [n]: i \neq j} \sum_{s=(x,z,y) \in T} |\mathbb{I}[\mathcal{A}_{S_i}(x) \geq 0] - \mathbb{I}[\mathcal{A}_{S_j}(x) \geq 0]|. \quad (3)$$

Note that we have $\mathbb{E}_{S_1, \dots, S_n} [\text{stab}_{T,n}(\mathcal{A})] = \text{stab}_T(\mathcal{A})$.

Fairness metric. Let D denote the empirical distribution over the testing set. Given a classifier f , we consider a fairness metric for statistical rate, which has been applied in (Menon & Williamson, 2018; Agarwal et al., 2018). Suppose the sensitive attribute is binary, i.e., $Z \in \{0, 1\}$.

$$\gamma(f) := \min \left\{ \frac{\Pr_D[f=1 | Z=0]}{\Pr_D[f=1 | Z=1]}, \frac{\Pr_D[f=1 | Z=1]}{\Pr_D[f=1 | Z=0]} \right\}. \quad (4)$$

Our framework can be easily extended to other fairness metrics; see a summary in Table 1 of (Celis et al., 2019b).

Figure 1: stab vs. λ for race attribute.Figure 2: stab vs. λ for sex attribute.

Implementation details. We first generate a common testing set (20%). Then we perform 50 repetitions, in which we uniformly sample a training set (75%) from the remaining data. For all three algorithms, we set the regularization parameter λ to be 0, 0.01, 0.02, 0.03, 0.04, 0.05 and compute the resulting stability metric stab, average accuracy and average fairness. Note that $\lambda = 0$ is equivalent to the case without stability-focused regularization term.

4.2. Results

Our simulations indicate that introducing a stability-focused regularization term can make the algorithm more stable by slightly sacrificing accuracy. Table 1 summarizes the accuracy and fairness metric under different regularization parameters λ . As λ increases, the average accuracy slightly decreases, by at most 1.5%, for all algorithms including **ZVRG-St**, **KAAS-St** and **GYR-St**. As for the fairness metric, as λ increases, the mean of γ decreases for **KAAS-St** and increases for **ZVRG-St** for both race and sex attribute. For **GYF-St**, the performance of fairness metric depends on the sensitive attribute: as λ increases, the mean of γ decreases for the sex attribute and increases for the race attribute. Note that the fairness metric γ of **KAAS-St** and **GYF-St** is usually smaller than that of **ZVRG-St** with the same λ . The results indicate that **ZVRG-St** achieves the better fairness than, and comparable accuracy to, other algorithms. Another observation is that the standard deviation of γ decreases by introducing the regularization term. Specifically, considering the sex attribute, the mean of γ is 4.1% when $\lambda = 0$ and decreases to about 1% by introducing a stability-focused regularization term. This observation implies that our extended framework improves the stability.

Figures 1 and 2 summarize the stability metrics stab under different regularization parameters λ . By introducing stability-focused regularization term, stab indeed decreases for both race and sex attributes. Observe that stab can decrease by a half by introducing the regularization term for

all three algorithms. Note that stab of **KAAS-St** is always larger than that of **ZVRG-St** and **GYF-St** with the same λ . The stability of **ZVRG-St** and **GYF-St** is comparable. Interestingly, stab does not monotonically decrease as λ increases due to the fairness requirements. The reason might be as follows: as λ increases, the model parameters of the learned classifiers should decrease monotonically. However, it is possible that a classifier with smaller model parameters is more sensitive to random training sets. In this case, if the effect of λ to stab is less when compared to the effect of model parameters, stab might not decrease monotonically with λ . Hence, selecting a suitable regularization parameter λ is valuable in practice, e.g., considering **ZVRG-St** for sex attribute, letting $\lambda = 0.03$ achieves better performance of accuracy, fairness and stability than letting $\lambda = 0.05$.

5. Conclusion and future directions

We propose an extended framework for fair classification algorithms that are formulated as optimization problems. Our framework comes with a stability guarantee and we also provide an analysis of the resulting accuracy. The analysis can be used to inform the selection of the regularization parameter. The empirical results show that our framework indeed improves stability by slightly sacrificing the accuracy.

There exist other fair classification algorithms that are not formulated as optimization problems, e.g., shifting the decision boundary of a baseline classifier (Fish et al., 2016; Hardt et al., 2016a) or pre-processing the training data (Feldman et al., 2015; Krasanakis et al., 2018). It is interesting to investigate and improve the stability guarantee of those algorithms. Another potential direction is to combine stability and fairness for other automated decision-making tasks, e.g., ranking (Celis et al., 2018c; Yang & Stoyanovich, 2017), summarization (Celis et al., 2018b), personalization (Celis & Vishnoi, 2017; Celis et al., 2019c), multiwinner voting (Celis et al., 2018a), and online advertising (Celis et al., 2019a).

References

- Adomavicius, G. and Zhang, J. Maximizing stability of recommendation algorithms: A collective inference approach. In *21st Workshop on Information Technologies and Systems (WITS)*, pp. 151–56. Citeseer, 2011.
- Adomavicius, G. and Zhang, J. Stability of recommendation algorithms. *ACM Transactions on Information Systems (TOIS)*, 30(4):23, 2012.
- Adomavicius, G. and Zhang, J. Classification, ranking, and top-k stability of recommendation algorithms. *INFORMS Journal on Computing*, 28(1):129–147, 2016.
- Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. M. A reductions approach to fair classification. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, pp. 60–69, 2018.
- Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. *ProPublica*, May, 2016.
- Barreno, M., Nelson, B., Joseph, A. D., and Tygar, J. D. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.
- Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilovic, A., et al. Ai fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943*, 2018.
- Berk, R. The role of race in forecasts of violent crime. *Race and social problems*, 2009.
- Biggio, B., Nelson, B., and Laskov, P. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 2012.
- Bousquet, O. and Elisseeff, A. Stability and generalization. In *Journal of machine learning research*, volume 2, pp. 499–526, 2002.
- Boyd, S. and Mutapcic, A. Stochastic subgradient methods. *Lecture Notes for EE364b, Stanford University*, 2008.
- Briand, B., Ducharme, G. R., Parache, V., and Mercat-Rommens, C. A similarity measure to assess the stability of classification trees. *Computational Statistics & Data Analysis*, 53(4):1208–1217, 2009.
- Celis, E., Mehrotra, A., and Vishnoi, N. Towards controlling discrimination in online ad auctions. In *International Conference on Machine Learning*, 2019a.
- Celis, L. E. and Vishnoi, N. K. Fair personalization. In *Fairness, Accountability, and Transparency in Machine Learning*, 2017.
- Celis, L. E., Huang, L., and Vishnoi, N. K. Multiwinner voting with fairness constraints. In *Proceedings of the Twenty-seventh International Joint Conference on Artificial Intelligence and the Twenty-third European Conference on Artificial Intelligence, IJCAI-ECAI*, 2018a.
- Celis, L. E., Keswani, V., Straszak, D., Deshpande, A., Kathuria, T., and Vishnoi, N. K. Fair and diverse dpp-based data summarization. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, pp. 715–724, 2018b.
- Celis, L. E., Straszak, D., and Vishnoi, N. K. Ranking with fairness constraints. In *Proceedings of the forty-fifth International Colloquium on Automata, Languages, and Programming ICALP*, 2018c.
- Celis, L. E., Huang, L., Keswani, V., and Vishnoi, N. K. Classification with fairness constraints: A meta-algorithm with provable guarantees. *The second annual ACM FAT* Conference*, 2019b.
- Celis, L. E., Kapoor, S., Salehi, F., and Vishnoi, N. K. Controlling polarization in personalization: An algorithmic framework. In *Fairness, Accountability, and Transparency in Machine Learning*, 2019c.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 797–806, 2017.
- Datta, A., Tschantz, M. C., and Datta, A. Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies*, 2015.
- Dedman, B. et al. The color of money. *Atlanta Journal-Constitution*, 1988.
- Dheeru, D. and Karra Taniskidou, E. UCI machine learning repository. <http://archive.ics.uci.edu/ml>, 2017.
- Dwork, C., Immorlica, N., Kalai, A. T., and Leiserson, M. D. M. Decoupled classifiers for group-fair and efficient machine learning. In *Fairness, Accountability, and Transparency in Machine Learning*, pp. 119–133, 2018.
- Ekstrand, M. D., Joshaghani, R., and Mehrpouyan, H. Privacy for all: Ensuring fair and equitable privacy protections. In *Conference on Fairness, Accountability and Transparency*, pp. 35–47, 2018.

- Farahat, A. and Bailey, M. C. How effective is targeted advertising? In *Proceedings of the 21st international conference on World Wide Web*, pp. 111–120. ACM, 2012.
- Fawzi, A., Fawzi, O., and Frossard, P. Analysis of classifiers’ robustness to adversarial perturbations. In *Machine Learning*, volume 107, pp. 481–508. Springer, 2018.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 259–268. ACM, 2015.
- Feldman, V. and Vondrak, J. Generalization bounds for uniformly stable algorithms. In *Advances in Neural Information Processing Systems*, pp. 9769–9779, 2018.
- Fish, B., Kun, J., and Lelkes, Á. D. A confidence-based approach for balancing fairness and accuracy. In *Proceedings of the 2016 SIAM International Conference on Data Mining, 2016*, pp. 144–152. SIAM, 2016.
- Flores, A. W., Bechtel, K., and Lowenkamp, C. T. False positives, false negatives, and false analyses: A rejoinder to machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. *Fed. Probation*, 80:38, 2016.
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. A comparative study of fairness-enhancing interventions in machine learning. In *Fairness, Accountability, and Transparency in Machine Learning*, 2019.
- Gershoff, A. D., Mukherjee, A., and Mukhopadhyay, A. Consumer acceptance of online agent advice: Extremity and positivity effects. *Journal of Consumer Psychology*, 13:161–170, 2003.
- Goel, N., Yaghini, M., and Faltings, B. Non-discriminatory machine learning through convex fairness criteria. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Goel, S., Rao, J. M., Shroff, R., et al. Precinct or prejudice? understanding racial disparities in new york city’s stop-and-frisk policy. *The Annals of Applied Statistics*, 10(1): 365–394, 2016.
- Goh, G., Cotter, A., Gupta, M. R., and Friedlander, M. P. Satisfying real-world goals with dataset constraints. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems*, pp. 2415–2423, 2016.
- Hajian, S., Domingo-Ferrer, J., Monreale, A., Pedreschi, D., and Giannotti, F. Discrimination-and privacy-aware patterns. *Data Mining and Knowledge Discovery*, 29(6): 1733–1782, 2015.
- Hajian, S., Bonchi, F., and Castillo, C. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 2125–2126. ACM, 2016.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing*, pp. 3315–3323, 2016a.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016*, pp. 1225–1234, 2016b.
- Kamiran, F. and Calders, T. Classifying without discriminating. In *Computer, Control and Communication, 2009. IC4 2009. 2nd International Conference on*, pp. 1–6. IEEE, 2009.
- Kamiran, F. and Calders, T. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- Kamishima, T., Akaho, S., Asoh, H., and Sakuma, J. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 35–50. Springer, 2012.
- Kashid, A., Kulkarni, V., and Patankar, R. Discrimination prevention using privacy preserving techniques. *International Journal of Computer Applications*, 120(1), 2015.
- Kashid, A., Kulkarni, V., and Patankar, R. Discrimination-aware data mining: a survey. *International Journal of Data Science*, 2(1):70–84, 2017.
- Krasanakis, E., Spyromitros-Xioufis, E., Papadopoulos, S., and Kompatsiaris, Y. Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018*. International World Wide Web Conferences Steering Committee, 2018.
- Kuzborskij, I. and Lampert, C. H. Data-dependent stability of stochastic gradient descent. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, pp. 2820–2829, 2018.
- Lam, S. K. and Riedl, J. Shilling recommender systems for fun and profit. In *Proceedings of the 13th international conference on World Wide Web*, pp. 393–402. ACM, 2004.

- London, B. Generalization bounds for randomized learning with application to stochastic gradient descent. In *NIPS Workshop on Optimizing the Optimizers*, 2016.
- London, B., Huang, B., Taskar, B., and Getoor, L. Collective stability in structured prediction: Generalization from one example. In *International Conference on Machine Learning*, pp. 828–836, 2013.
- London, B., Huang, B., Taskar, B., and Getoor, L. Bayesian collective stability. In *Artificial Intelligence and Statistics*, pp. 585–594, 2014.
- London, B., Huang, B., and Getoor, L. Stability and generalization in structured prediction. *The Journal of Machine Learning Research*, 17(1):7808–7859, 2016.
- Luong, B. T., Ruggieri, S., and Turini, F. k-nn as an implementation of situation testing for discrimination discovery and prevention. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 502–510. ACM, 2011.
- Maurer, A. A second-order look at stability and generalization. In *Conference on Learning Theory*, pp. 1461–1475, 2017.
- Mei, S. and Zhu, X. Using machine teaching to identify optimal training-set attacks on machine learners. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 2871–2877, 2015.
- Meng, Q., Wang, Y., Chen, W., Wang, T., Ma, Z., and Liu, T. Generalization error bounds for optimization algorithms via stability. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pp. 2336–2342, 2017.
- Menon, A. K. and Williamson, R. C. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency, FAT*, pp. 107–118, 2018.
- Mobasher, B., Burke, R., Bhaumik, R., and Williams, C. Toward trustworthy recommender systems: An analysis of attack models and algorithm robustness. *ACM Transactions on Internet Technology (TOIT)*, 7(4):23, 2007.
- Mukherjee, S., Niyogi, P., Poggio, T. A., and Rifkin, R. M. Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Adv. Comput. Math.*, 25(1-3):161–193, 2006.
- Nogueira, S., Sechidis, K., and Brown, G. On the stability of feature selection algorithms. *Journal of Machine Learning Research*, 18(174):1–54, 2018.
- O’mahony, M. P., Hurley, N. J., and Silvestre, G. C. An evaluation of neighbourhood formation on the performance of collaborative filtering. *Artificial Intelligence Review*, 21(3-4):215–228, 2004.
- Park, D. H., Kim, H. K., Choi, I. Y., and Kim, J. K. A literature review and classification of recommender systems research. *Expert Syst. Appl.*, 39(11):10059–10072, 2012.
- Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J. M., and Weinberger, K. Q. On fairness and calibration. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pp. 5684–5693, 2017.
- Portugal, I., Alencar, P. S. C., and Cowan, D. D. The use of machine learning algorithms in recommender systems: A systematic review. *Expert Syst. Appl.*, 97:205–227, 2018.
- Raghu, M., Poole, B., Kleinberg, J. M., Ganguli, S., and Sohl-Dickstein, J. On the expressive power of deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, pp. 2847–2854, 2017.
- Ruggieri, S., Hajian, S., Kamiran, F., and Zhang, X. Anti-discrimination analysis using privacy attack strategies. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 694–710. Springer, 2014.
- Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11(Oct):2635–2670, 2010.
- Shapiro, M. Stability and change in judicial decision-making: Incrementalism or stare decisis. In *Law in Transition Quarterly*, volume 2, pp. 134–157, 1965.
- Steinhardt, J., Koh, P. W., and Liang, P. S. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pp. 3520–3532, 2017.
- Sweeney, L. Discrimination in online ad delivery. *Commun. ACM*, 56(5):44–54, 2013.
- Van Swol, L. M. and Sniezek, J. A. Factors affecting the acceptance of expert advice. *British Journal of Social Psychology*, 44(3):443–461, 2005.
- Vidal, R., Bruna, J., Gyries, R., and Soatto, S. Mathematics of deep learning. *arXiv preprint arXiv:1712.04741*, 2017.
- Wainwright, M. J. Estimating the “wrong” graphical model: Benefits in the computation-limited setting. *Journal of Machine Learning Research*, 7(Sep):1829–1859, 2006.

- Wang, Y., Lei, J., and Fienberg, S. E. Learning with differential privacy: Stability, learnability and the sufficiency and necessity of ERM principle. *Journal of Machine Learning Research*, 17:183:1–183:40, 2016.
- Woodworth, B. E., Gunasekar, S., Ohannessian, M. I., and Srebro, N. Learning non-discriminatory predictors. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017*, pp. 1920–1953, 2017.
- Yang, K. and Stoyanovich, J. Measuring fairness in ranked outputs. In *SSDBM*, 2017.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017*, pp. 1171–1180, 2017a.
- Zafar, M. B., Valera, I., Gomez-Rodriguez, M., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, pp. 962–970, 2017b.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013*, pp. 325–333, 2013.