
FERMI: Fair Empirical Risk Minimization Via Exponential Rényi Mutual Information

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Despite the success of large-scale empirical risk minimization (ERM) at achieving
2 high accuracy across a variety of machine learning tasks, fair ERM is hindered by
3 the incompatibility of fairness constraints with stochastic optimization. In this pa-
4 per, we propose the fair empirical risk minimization via exponential Rényi mutual
5 information (FERMI) framework. FERMI is built on a stochastic estimator for ex-
6 ponential Rényi mutual information (ERMI), an information divergence measuring
7 the degree of the dependence of predictions on sensitive attributes. Theoretically,
8 we show that ERMI upper bounds existing popular fairness violation metrics, thus
9 controlling ERMI provides guarantees on other commonly used violations, such as
10 L_∞ . We derive an unbiased estimator for ERMI, which we use to derive the FERMI
11 algorithm. We prove that FERMI converges for demographic parity, equalized
12 odds, and equal opportunity notions of fairness in stochastic optimization. Em-
13 pirically, we show that FERMI is amenable to large-scale problems with multiple
14 (non-binary) sensitive attributes and non-binary targets. Extensive experiments
15 show that FERMI achieves the most favorable tradeoffs between fairness violation
16 and test accuracy across all tested setups compared with state-of-the-art baselines
17 for demographic parity, equalized odds, equal opportunity. These benefits are
18 especially significant for non-binary classification with large sensitive sets and
19 small batch sizes, showcasing the effectiveness of the FERMI objective and the
20 developed stochastic algorithm for solving it.

21 1 Introduction

22 Ensuring that decisions made using machine learning algorithms are fair to different subgroups is
23 of utmost importance. Without any mitigation strategy, machine learning algorithms may result in
24 discrimination against certain subgroups based on sensitive attributes, such as gender or race, even if
25 such discrimination is absent in the training data (Datta et al., 2015; Sweeney, 2013; Bolukbasi et al.,
26 2016; Angwin et al., 2016; Calmon et al., 2017b; Feldman et al., 2015; Hardt et al., 2016; Fish et al.,
27 2016; Woodworth et al., 2017; Zafar et al., 2017; Bechavod & Ligett, 2017; Kearns et al., 2018).
28 Algorithmic fairness literature aims to remedy such discrimination issues.

29 A machine learning algorithm satisfies the *demographic parity* fairness notion, if the predicted target
30 is independent of the sensitive attributes (Dwork et al., 2012). Promoting demographic parity can
31 lead to poor performance, especially if the true outcome is not independent of the sensitive attributes.
32 To remedy this, Hardt et al. (2016) proposed *equalized odds* to ensure that the predicted target is
33 conditionally independent of the sensitive attributes given the true label. A further relaxed version of
34 this notion is *equal opportunity* which is satisfied if predicted target is conditionally independent of
35 sensitive attributes given that the true label is in an advantaged class (Hardt et al., 2016). The inherent
36 assumption in such conditional notions is that the true labels are fair. These notions suffer from a
37 potential amplification of the inherent discrimination that may exist in the training data. Tackling
38 such bias is beyond the scope of this work; cf. Kilbertus et al. (2020) and Bechavod et al. (2019).

Reference	NB target	NB attrib.	NB code	Fairness notion			Beyond logistic	Stoch. alg. (unbiased**)	Converg. (stoch.)
				dp	eod	eop			
FERMI (this work)	✓	✓	✓	✓	✓	✓	✓	✓ (✓)	✓ (✓)
(Cho et al., 2020b)	✓	✓	✓	✓	✓	✗	✓	✓ (✗)	✗
(Cho et al., 2020a)	✓	✓	✗	✓	✓	✗	✓	✓ (✓)	✗
(Baharlouei et al., 2020)	✓	✓	✓	✓	✓	✓	✓	✗	✓ (✗)
(Rezaei et al., 2020)	✗	✗	✗	✓	✓	✗	✗	✗	✗
(Jiang et al., 2020)*	✗	✓	✗	✓	✗	✗	✗	✗	✗
(Mary et al., 2019)	✓	✓	✓	✓	✓	✗	✓	✓ (✗)	✗
(Donini et al., 2018)	✗	✓	✗	✗	✓	✗	✓	✗	✗
(Zhang et al., 2018)	✓	✓	✗	✓	✓	✗	✓	✓ (✗)	✗

Table 1: Comparison of state-of-the-art in-processing methods. **NB = non-binary**, dp = demographic parity, eod = equalized odds, eop = equal opportunity. While satisfying eod guarantees satisfying eop, an eod algorithm does not necessarily achieve a favorable tradeoff between performance and fairness violation in eop; we only credit those works that provide/implement algorithms for a given fairness notion. FERMI is the only method compatible with stochastic optimization and guaranteed convergence. The only existing baselines for non-binary classification with non-binary sensitive attributes are (Mary et al., 2019; Baharlouei et al., 2020; Cho et al., 2020b) (NB code). *We refer to the in-processing method of (Jiang et al., 2020), not their post-processing method. **We use the term “unbiased” to refer to unbiased estimation in statistical sense; it is not to be confused with bias in the fairness sense, for which we use the term discrimination.

39 **Measuring fairness violation.** In practice, the learner only has access to finite samples and cannot
40 verify demographic parity, equalized odds, or equal opportunity. This has led the machine learning
41 community to define several fairness violation metrics that quantify the degree of (conditional)
42 independence between random variables, e.g., L_∞ distance (Dwork et al., 2012; Hardt et al., 2016),
43 mutual information (Kamishima et al., 2011; Rezaei et al., 2020; Steinberg et al., 2020; Zhang
44 et al., 2018; Cho et al., 2020a), Pearson correlation (Zafar et al., 2017), false positive/negative rates
45 (Bechavod & Ligett 2017), Hilbert Schmidt independence criterion (HSIC) (Pérez-Suay et al., 2017),
46 Rényi correlation (Mary et al., 2019; Baharlouei et al., 2020; Gari et al., 2019, 2020), and exponential
47 Rényi mutual information (ERMI) (Mary et al., 2019). In this paper, we focus on three variants of
48 ERMI specialized to demographic parity, equalized odds, and equal opportunity. We prove that ERMI
49 provides an upper bound on the rest of the above existing notions of fairness violation. Consequently,
50 a model trained to reduce ERMI will also provide guarantees on these other fairness violations.
51 We also develop a stochastic estimator for ERMI that is compatible with large-scale stochastic
52 optimization, and use it as a regularizer in within ERM, and call it FERMI. We theoretically show
53 that FERMI is convergent, and empirically demonstrate that it outperforms all other state-of-the-art
54 baselines, including (Mary et al., 2019) which solves the same objective as FERMI.

55 **Related work & contributions.** Fairness-promoting machine learning algorithms can be categorized
56 in three main classes: *pre-processing*, *post-processing*, and *in-processing* methods. Pre-processing
57 algorithms (Feldman et al., 2015; Zemel et al., 2013; Calmon et al., 2017b) transform the biased
58 data features to a new space in which the labels and sensitive attributes are statistically independent.
59 This transform is oblivious to the training procedure. Post-processing approaches (Hardt et al., 2016;
60 Pleiss et al., 2017) mitigate the discrimination of the classifier by altering the the final decision.
61 In-processing approaches focus on the training procedure and impose the notions of fairness as
62 constraints or regularization terms in the training procedure. Several regularization-based methods
63 are proposed in the literature to promote fairness in decision-trees (Kamiran et al., 2010; Raff et al.,
64 2018; Aghaei et al., 2019), support vector machines (Donini et al., 2018), neural networks (Gari
65 et al., 2020; Cho et al., 2020b), or (logistic) regression models (Zafar et al., 2017; Berk et al., 2017;
66 Taskesen et al., 2020; Chzhen & Schreuder 2020; Baharlouei et al., 2020; Jiang et al., 2020; Gari
67 et al., 2019). While in-processing approaches generally give rise to better tradeoffs between fairness
68 violation and performance, existing approaches are mostly incompatible with large-scale stochastic
69 optimization. This paper addresses this problem. See below for a summary of our contributions and
70 Table 1 for a summary of the main differences between FERMI and existing in-processing methods.

- 71 1. We analyze a notion of fairness violation called ERMI. We show that ERMI is a stronger notion of
72 fairness violation than all existing notions. Therefore, a model that ensures small ERMI violation
73 is guaranteed to have small fairness violation with respect to all other notions as well.
- 74 2. We formulate an empirical objective, called FERMI objective, for using ERMI as a regularizer
75 with empirical risk minimization. We propose a solver for FERMI, which is the first stochastic
76 in-processing fairness algorithm with guaranteed convergence. The existing stochastic fairness

77 algorithms by Zhang et al. (2018); Mary et al. (2019); Cho et al. (2020a,b) are not guaranteed to
 78 converge.

79 3. We demonstrate through extensive numerical experiments that FERMI achieves superior fair-
 80 ness-accuracy tradeoff curves against all comparable baselines, even when fairness violation is
 81 measured in terms of commonly used L_∞ (for demographic parity, equalized odds, and equal
 82 opportunity). In particular, the performance gap is very large when minibatch size is small (as is
 83 practically necessary for large-scale problems), and the number of sensitive attributes is large.

84 2 Fairness notions: demographic parity, equalized odds, equal opportunity

85 In this section, we state a notion of fairness that generalizes demographic parity, equalized odds,
 86 and equal opportunity fairness definitions (the three notions considered in this paper). This will be
 87 convenient for presenting our theoretical results. Consider a learner who trains a model to make
 88 a prediction, \hat{Y} , e.g., whether or not to extend a loan, supported on \mathcal{Y} which can be discrete or
 89 continuous. The prediction is made using a set of features, \mathbf{X} , e.g., financial history features. We
 90 assume that there is a set of discrete sensitive attributes, S , e.g., race and sex, supported on \mathcal{S} ,
 91 associated with each sample. Further, let $\mathcal{A} \subseteq \mathcal{Y}$ denote an advantaged outcome class, e.g., the
 92 outcome where a loan is extended.

93 **Definition 1** ((Z, \mathcal{Z}) -fairness). *Given a random variable Z , let \mathcal{Z} be a subset of values that Z can
 94 take. We say that a learning machine satisfies (Z, \mathcal{Z}) -fairness if for every $z \in \mathcal{Z}$, \hat{Y} is conditionally
 95 independent of S given $Z = z$, i.e. $\forall \hat{y} \in \mathcal{Y}, s \in \mathcal{S}, z \in \mathcal{Z}, p_{\hat{Y}, S|Z}(\hat{y}, s|z) = p_{\hat{Y}|Z}(\hat{y}|z)p_{S|Z}(s|z)$.*

96 (Z, \mathcal{Z}) -fairness includes the popular demographic parity, equalized odds, and equal opportunity
 97 notions of fairness as special cases:

- 98 1. (Z, \mathcal{Z}) -fairness recovers demographic parity (Dwork et al., 2012) if $Z = 0$ and $\mathcal{Z} = \{0\}$. In this
 99 case, conditioning on Z has no effect, and hence $(0, \{0\})$ fairness is equivalent to the independence
 100 between \hat{Y} and S (see Definition 6, Appendix A).
- 101 2. (Z, \mathcal{Z}) -fairness recovers equalized odds (Hardt et al., 2016) if $Z = Y$ and $\mathcal{Z} = \mathcal{Y}$. In this case,
 102 $Z \in \mathcal{Z}$ is trivially satisfied. Hence, conditioning on Z is equivalent to conditioning on Y , which
 103 recovers the equalized odds notion of fairness, i.e., conditional independence of \hat{Y} and S given Y
 104 (see Definition 7, Appendix A).
- 105 3. (Z, \mathcal{Z}) -fairness recovers equal opportunity (Hardt et al., 2016) if $Z = Y$ and $\mathcal{Z} = \mathcal{A}$. This is also
 106 similar to the previous case with \mathcal{Y} replaced with \mathcal{A} (see Definition 8, Appendix A).

107 Note that verifying (Z, \mathcal{Z}) -fairness requires having access to the joint distribution of random variables
 108 (Z, \hat{Y}, S) . This joint distribution is unavailable to the learner in the context of machine learning, and
 109 hence the learner would resort to empirical estimation of the amount of violation of independence,
 110 measured through some divergence. See (Williamson & Menon, 2019) for a related discussion.

111 3 Measuring fairness violation using exponential Rényi mutual information

112 Most existing fairness violations can be viewed as a (conditional) f -divergence between the joint
 113 distribution of sensitive attributes and predicted targets, $p_{\hat{Y}, S|Z}$, and the Kronecker product of the
 114 marginals, $p_{\hat{Y}|Z} \otimes p_{S|Z}$. In this section, we focus on ERMI and show that several existing fairness
 115 violations are upper bounded by ERMI. For brevity, we present all definitions and results (Z, \mathcal{Z}) .

116 **Definition 2** (ERMI – exponential Rényi mutual information). *We define the exponential Rényi
 117 mutual information between \hat{Y} and S given $Z \in \mathcal{Z}$ as*

$$D_R(\hat{Y}; S|Z \in \mathcal{Z}) := \mathbb{E}_{Z, \hat{Y}, S} \left\{ \frac{p_{\hat{Y}, S|Z}(\hat{Y}, S|Z)}{p_{\hat{Y}|Z}(\hat{Y}|Z)p_{S|Z}(S|Z)} \middle| Z \in \mathcal{Z} \right\} - 1. \quad (\text{ERMI})$$

118 In Appendix B we unravel the definition for the special cases of interest corresponding to demo-
 119 graphic parity, equalized odds, and equal opportunity. We also discuss that ERMI is the χ^2 -divergence
 120 (which is an f -divergence) between the joint distribution, $p_{\hat{Y}, S|Z}$, and the Kronecker product of
 121 marginals, $p_{\hat{Y}|Z} \otimes p_{S|Z}$ (Calmon et al., 2017a). In particular, ERMI is non-negative, and zero if
 122 and only if (Z, \mathcal{Z}) -fairness is satisfied. In the context of algorithmic fairness, ERMI was first used
 123 by Mary et al. (2019) as a regularizer. We will provide a new stochastic solver/estimator for ERMI,
 124 which theoretically converges and empirically outperforms the one by Mary et al. (2019).

125 **Definition 3** (Rényi mutual information (Rényi, 1961)). Let the Rényi mutual information of order
 126 $\alpha > 1$ between random variables \hat{Y} and S given $Z \in \mathcal{Z}$ be defined as:

$$I_\alpha(\hat{Y}; S|Z \in \mathcal{Z}) := \frac{1}{\alpha - 1} \log \left(\mathbb{E}_{Z, \hat{Y}, S} \left\{ \left(\frac{p_{\hat{Y}, S|Z}(\hat{Y}, S|Z)}{p_{\hat{Y}|Z}(\hat{Y}|Z)p_{S|Z}(S|Z)} \right)^{\alpha-1} \middle| Z \in \mathcal{Z} \right\} \right), \quad (\text{RMI})$$

127 which generalizes Shannon mutual information

$$I_1(\hat{Y}; S|Z \in \mathcal{Z}) := \mathbb{E}_{Z, \hat{Y}, S} \left\{ \log \left(\frac{p_{\hat{Y}, S|Z}(\hat{Y}, S|Z)}{p_{\hat{Y}|Z}(\hat{Y}|Z)p_{S|Z}(S|Z)} \right) \middle| Z \in \mathcal{Z} \right\}, \quad (\text{MI})$$

128 and recovers it as $\lim_{\alpha \rightarrow 1^+} I_\alpha(\hat{Y}; S|Z \in \mathcal{Z}) = I_1(\hat{Y}; S|Z \in \mathcal{Z})$.

129 Note that $I_\alpha(\hat{Y}; S|Z \in \mathcal{Z}) \geq 0$ with equality if and only if (Z, \mathcal{Z}) -fairness is satisfied.

130 **Theorem 1** (ERMI is stronger than Shannon mutual information). We have

$$0 \leq I_1(\hat{Y}; S|Z \in \mathcal{Z}) \leq I_2(\hat{Y}; S|Z \in \mathcal{Z}) \leq e^{I_2(\hat{Y}; S|Z \in \mathcal{Z})} - 1 = D_R(\hat{Y}; S|Z \in \mathcal{Z}). \quad (1)$$

131 All proofs are relegated to the appendix. Theorem 1 establishes that ERMI is a stronger measure of
 132 fairness violation in the sense that driving it to zero would also bound the Shannon mutual information,
 133 which is used for promoting fairness in recent literature (Cho et al., 2020a). It also shows that ERMI
 134 is exponentially related to the Rényi mutual information of order 2.

136 **Definition 4** (Rényi correlation (Hirschfeld, 1935; Gebelein, 1941; Rényi, 1959)). Let \mathcal{F} and \mathcal{G}
 137 be the set of measurable functions such that for random variables \hat{Y} and S , $\mathbb{E}_{\hat{Y}}\{f(\hat{Y}; z)\} =$
 138 $\mathbb{E}_S\{g(S; z)\} = 0$, $\mathbb{E}_{\hat{Y}}\{f(\hat{Y}; z)^2\} = \mathbb{E}_S\{g(S; z)^2\} = 1$, for all $z \in \mathcal{Z}$. Rényi correlation is:

$$\rho_R(\hat{Y}, S|Z \in \mathcal{Z}) := \sup_{f, g \in \mathcal{F} \times \mathcal{G}} \mathbb{E}_{Z, \hat{Y}, S} \left\{ f(\hat{Y}; Z)g(S; Z) \middle| Z \in \mathcal{Z} \right\}. \quad (\text{RC})$$

139

140 Rényi correlation generalizes Pearson correlation,

$$\rho(\hat{Y}, S|Z \in \mathcal{Z}) := \mathbb{E}_Z \left\{ \frac{\mathbb{E}_{\hat{Y}, S}\{\hat{Y}S|Z\}}{\sqrt{\mathbb{E}_{\hat{Y}}\{\hat{Y}^2|Z\}\mathbb{E}_S\{S^2|Z\}}} \middle| Z \in \mathcal{Z} \right\}, \quad (\text{PC})$$

141

142 to capture nonlinear dependencies between the random variables by finding functions of random
 143 variables that maximize the Pearson correlation coefficient between the random variables. In fact,
 144 it is true that $\rho_R(\hat{Y}, S|Z \in \mathcal{Z}) \geq 0$ with equality if and only if (Z, \mathcal{Z}) -fairness is satisfied. Rényi
 145 correlation has gained popularity as a measure of fairness violation (Mary et al., 2019; Baharlouei
 146 et al., 2020; Grari et al., 2020). Rényi correlation is also upper bounded by ERMI. The following
 147 result has already been shown by Mary et al. (2019) and we present it for completeness.

148 **Theorem 2** (ERMI is stronger than Rényi correlation). We have

$$0 \leq |\rho(\hat{Y}, S|Z \in \mathcal{Z})| \leq \rho_R(\hat{Y}, S|Z \in \mathcal{Z}) \leq D_R(\hat{Y}; S|Z \in \mathcal{Z}), \quad (2)$$

149 and if $|\mathcal{S}| = 2$, $D_R(\hat{Y}; S|Z \in \mathcal{Z}) = \rho_R(\hat{Y}, S|Z \in \mathcal{Z})$.

150 **Definition 5** (L_q fairness violation). We define the L_q fairness violation for $q \geq 1$ by:

$$L_q(\hat{Y}, S|Z \in \mathcal{Z}) := \mathbb{E}_Z \left\{ \left(\int_{\hat{y} \in \mathcal{Y}_0} \sum_{s \in \mathcal{S}_0} |p_{\hat{Y}, S|Z}(\hat{y}, s|Z) - p_{\hat{Y}|Z}(\hat{y}|Z)p_{S|Z}(s|Z)|^q dy \right)^{\frac{1}{q}} \middle| Z \in \mathcal{Z} \right\}. \quad (\text{Lq})$$

151 Note that $L_q(\hat{Y}, S|Z \in \mathcal{Z}) = 0$ if and only if (Z, \mathcal{Z}) -fairness is satisfied. In particular, L_∞ fairness
 152 violation recovers demographic parity violation (Kearns et al., 2018, Definition 2.1) if we let $\mathcal{Z} = \{0\}$
 153 and $Z = 0$. It also recovers equal opportunity violation (Hardt et al., 2016) if $\mathcal{Z} = \mathcal{A}$ and $Z = Y$.

154 **Theorem 3** (ERMI is stronger than L_∞ fairness violation). Let \hat{Y} be a discrete or continuous random
 155 variable, and S be a discrete random variable supported on a finite set. Then for any $q \geq 1$,

$$0 \leq L_q(\hat{Y}, S|Z \in \mathcal{Z}) \leq \sqrt{D_R(\hat{Y}, S|Z \in \mathcal{Z})}. \quad (3)$$

156

157 The above theorem says that if a method controls ERMI value for imposing fairness, then L_∞
 158 violation is controlled. In particular, the variant of ERMI that is specialized to demographic parity
 159 also controls L_∞ demographic parity violation (Kearns et al., 2018). The variant of ERMI that is
 160 specialized to equal opportunity also controls the L_∞ equal opportunity violation (Hardt et al., 2016).
 161 While our algorithm uses ERMI as a regularizer, in our experiments, we measure fairness violation
 162 through the more commonly used L_∞ violation. Despite this, we show that our approach leads to
 163 better tradeoff curves between fairness violation and performance.

164 **Remark.** The bounds in Theorems 1-3 are not tight in general, but this is not of practical concern.
 165 They show that bounding ERMI is sufficient because any model that achieves small ERMI is
 166 guaranteed to satisfy any other fairness violation. This makes ERMI an effective regularizer for
 167 promoting fairness. In fact, in Sec. 5, we see that the proposed algorithm, FERMI, achieves the best
 168 tradeoffs between fairness violation and performance across state-of-the-art baselines.

169 4 FERMI: fair empirical risk minimization through ERMI regularization

170 Our goal is to train a model that balances fairness and accuracy objectives. To this end, we introduce
 171 fair risk minimization through exponential Rényi mutual information framework defined below:¹

$$\min_{\theta} \left\{ \text{FRMI}(\theta) := \mathbb{E}_{\mathbf{X}, Y, S} \{ \ell(\mathbf{X}, Y; \theta) \} + \lambda D_R(\hat{Y}(\mathbf{X}; \theta); S) \right\}, \quad (\text{FRMI obj.})$$

172 where ℓ denotes the loss function, such as L_2 loss or cross entropy loss; $\lambda > 0$ is a scalar balancing
 173 the accuracy versus fairness objectives; $D_R(\hat{Y}(\mathbf{X}; \theta); S)$ is the notion of ERMI given in Eq. (ERMI)
 174 particularized to demographic parity (see Eq. (5)); and $\hat{Y}(\mathbf{X}; \theta)$ is the output of the learned model
 175 (e.g., the output of a classification or a regression task, or the cluster number in a clustering task).
 176 While $\hat{Y}(\mathbf{X}; \theta)$ inherently depends on \mathbf{X} and θ , in the rest of this paper, we sometimes leave the
 177 dependence of \hat{Y} on \mathbf{X} and/or θ implicit for brevity of notation. Notice that we have also left the
 178 dependence of the loss on the predicted outcome \hat{Y} implicit.

179 In practice, the true joint distribution of $(\mathbf{X}, S, Y, \hat{Y})$ is unknown and we only have N samples at
 180 our disposal, making it impossible to solve FRMI. Let $\{\mathbf{x}_i, s_i, y_i, \hat{y}_i(\mathbf{x}_i; \theta)\}_{i \in [N]}$ denote the features,
 181 sensitive attributes, targets, and the predictions of the model parameterized by θ for these samples.
 182 Mary et al. (2019) considered the same objective Eq. (FRMI obj.), and tried to empirically solve it
 183 through a kernel approximation. We propose a completely different approach to solving this problem:
 184 fair empirical risk minimization via exponential Rényi mutual information (FERMI). FERMI results
 185 in a provably convergent algorithm, and empirically outperforms the algorithm by Mary et al. (2019).
 186 It is straightforward to derive an unbiased estimate for $\mathbb{E}_{\mathbf{X}, Y, S} \{ \ell(\mathbf{X}, Y; \theta) \}$ through the empirical
 187 risk, e.g., $\frac{1}{|B|} \sum_{i \in B} \ell(\mathbf{x}_i, y_i; \theta)$ where $B \subseteq [N]$ is a random minibatch of data points. However,
 188 estimating $D_R(\hat{Y}, S)$ in the objective function in Eq. (FRMI obj.) is more difficult. In what follows,
 189 we present our approach to deriving an *unbiased stochastic estimator* of $D_R(\hat{Y}, S)$ given a random
 190 batch of data points B . The following theorem is the key tool we use to obtain an unbiased estimator:

191 **Theorem 4.** For discrete random variables $\hat{Y} = \hat{Y}(\mathbf{X}; \theta)$ and S where $\hat{Y} \in [m], S \in [k]$, we have

$$D_R(\hat{Y}; S) = \max_{W \in \mathbb{R}^{k \times m}} \left\{ -\text{Tr}(W P_{\hat{Y}} W^T) + 2 \text{Tr}(W P_{\hat{Y}, S} P_S^{-1/2}) - 1 \right\}, \quad (4)$$

192 where $P_{\hat{Y}} = \text{diag}(p_{\hat{Y}}(1), \dots, p_{\hat{Y}}(m))$, $P_S = \text{diag}(p_S(1), \dots, p_S(k))$, and

$$P_{\hat{Y}, S} = \begin{pmatrix} p_{\hat{Y}, S}(1, 1) & \dots & p_{\hat{Y}, S}(1, k) \\ \vdots & \ddots & \vdots \\ p_{\hat{Y}, S}(m, 1) & \dots & p_{\hat{Y}, S}(m, k) \end{pmatrix}.$$

193 Let $\hat{\mathbf{Y}}, \hat{\mathbf{y}}_i \in \{0, 1\}^m$ and $\mathbf{S}, \mathbf{s}_i \in \{0, 1\}^k$ be the one-hot encodings of \hat{Y}, \hat{y}_i and S, s_i , respectively.
 194 Then, the above theorem implies that we can compute an unbiased estimate of Eq. (FRMI obj.):

¹In this section, we present all results in the context of $Z = 0$ and $Z = \{0\}$ (demographic parity), leaving off all conditional expectations for clarity of presentation. The results are readily generalized for general (Z, \mathcal{Z}) by using $D_R(\hat{Y}, S|Z \in \mathcal{Z})$ in Eq. (FRMI obj.); we have used the resulting algorithms for empirical experiments.

195 **Lemma 1** (Unbiased estimator of ERMI). Let $(\mathbf{X}, S, Y, \hat{Y}(\mathbf{X}; \boldsymbol{\theta}))$ be a random draw from $P_{\mathbf{X}, S, Y, \hat{Y}}$.
 196 Further, let

$$\psi(\mathbf{X}, S, Y, \hat{Y}; \boldsymbol{\theta}, W) := -\text{Tr}(W \hat{Y}(\mathbf{X}; \boldsymbol{\theta}) \hat{Y}^T(\mathbf{X}; \boldsymbol{\theta}) W^T) + 2 \text{Tr}(W \hat{Y}(\mathbf{X}; \boldsymbol{\theta}) \mathbf{S}^T P_s^{-1/2}) - 1.$$

197
 198 Then, $\max_{W \in \mathbb{R}^{k \times m}} \psi(\mathbf{X}, S, Y, \hat{Y}; \boldsymbol{\theta}, W)$ is an unbiased estimator of ERMI in Eq. (FERMI obj.), i.e.,

$$\mathbb{E}_{\mathbf{X}, S, Y} \left\{ \max_{W \in \mathbb{R}^{k \times m}} \psi(\mathbf{X}, S, Y, \hat{Y}; \boldsymbol{\theta}, W) \right\} = D_R(\hat{Y}(\mathbf{X}; \boldsymbol{\theta}); S).$$

199 The stochastic estimator, $\psi(\mathbf{X}, S, Y, \hat{Y}; \boldsymbol{\theta}, W)$, in Lemma 1 requires the knowledge of P_s , and
 200 computation of $P_s^{-1/2}$. This can be estimated with high fidelity (for small to moderate sensitive set)
 201 through a single initial pass over the entire dataset in practice. Hence, we consider it to be known.
 202 Now, we are equipped to state the empirical objective function that we solve in this paper:
 203

$$\min_{\boldsymbol{\theta}} \max_{W \in \mathbb{R}^{k \times m}} \left\{ \text{FERMI}(\boldsymbol{\theta}, W) := \frac{1}{N} \sum_{i \in [N]} [\ell(\mathbf{x}_i, y_i; \boldsymbol{\theta}) + \lambda \psi_i(\boldsymbol{\theta}, W)] \right\}, \quad (\text{FERMI obj.})$$

204 where

$$\psi_i(\boldsymbol{\theta}, W) := -\text{Tr}(W \hat{y}_i(\mathbf{x}_i; \boldsymbol{\theta}) \hat{y}_i^T(\mathbf{x}_i; \boldsymbol{\theta}) W^T) + 2 \text{Tr}(W \hat{y}_i(\mathbf{x}_i; \boldsymbol{\theta}) \mathbf{s}_i^T P_s^{-1/2}) - 1.$$

205 In particular, Lemma 1 says that, for any N , Eq. (FERMI obj.) (and its gradients) is an unbiased and
 206 consistent estimator of the Eq. (FERMI obj.) objective function (and its gradients) by an empirical
 207 average over the minibatch. This is in contrast to the density estimation methods used by Mary et al.
 208 (2019) and Baharlouei et al. (2020), which are biased but consistent. We will see in the experiments
 209 that the unbiased estimator empirically offers large performance improvements.

210 This observations leads us to deriving a stochastic algorithm, presented in Algorithm 1, which is
 guaranteed to converge for any batch size $1 \leq |B| \leq N$ since the stochastic gradients are unbiased.

Algorithm 1 (FERMI Algorithm). Two-Time Scale SGDA for solving FERMI objective

- 1: **Input:** $\boldsymbol{\theta}^0 \in \mathbb{R}^{d_\theta}$, $W^0 \in \mathcal{W} \subset \mathbb{R}^{k \times m}$, step-sizes (η_θ, η_w) , mini-batch $B \subseteq [N]$, fairness parameter $\lambda \geq 0$, iteration number R .
 - 2: **for** $t = 0, 1, \dots, R$ **do**
 - 3: Draw a mini-batch B of data points $\{(\mathbf{x}_i, s_i, y_i)\}_{i \in B}$
 - 4: Set $\boldsymbol{\theta}^{t+1} \leftarrow \boldsymbol{\theta}^t - \frac{\eta_\theta}{|B|} \sum_{i \in B} [\nabla_{\boldsymbol{\theta}} \ell(\mathbf{x}_i, y_i; \boldsymbol{\theta}^t) + \lambda \nabla_{\boldsymbol{\theta}} \psi_i(\boldsymbol{\theta}^t, W^t)]$.
 - 5: Set $W^{t+1} \leftarrow \Pi_{\mathcal{W}} \left(W^t + \frac{2\lambda\eta_w}{|B|} \sum_{i \in B} \left[-W \hat{y}_i(\mathbf{x}_i; \boldsymbol{\theta}^t) \hat{y}_i^T(\mathbf{x}_i; \boldsymbol{\theta}^t) + P_s^{-1/2} \mathbf{s}_i \hat{y}_i^T(\mathbf{x}_i; \boldsymbol{\theta}^t) \right] \right)$
 - 6: **end for**
 - 7: Pick \hat{t} uniformly at random from $\{1, \dots, R\}$.
 - 8: **Return:** $\boldsymbol{\theta}^{\hat{t}}$.
-

211
 212 **Theorem 5.** (Informal statement) Algorithm 1 converges to the set of ϵ -first order stationary points
 213 of the Eq. (FERMI obj.) objective in $O(\frac{1}{\epsilon^4})$ iterations (stochastic gradient evaluations).

214 The formal statement of this theorem can be found in Theorem 10 in Appendix D. A faster convergence
 215 rate of $O(\frac{1}{\epsilon^3})$ could be obtained by using the (more complicated) SREDA method of Luo et al. (2020)
 216 instead of SGDA to solve FERMI objective. We omit the details here. In the next section, we
 217 numerically evaluate the performance FERMI algorithm in several numerical experiments.

218 5 Numerical experiments

219 5.1 Binary classification and binary sensitive attribute

220 For our first set of experiments, we evaluate the fairness-accuracy tradeoffs of FERMI in binary
 221 classification problems with a binary sensitive attribute. This is a common setup, so we are able to
 222 compare against many existing baseline methods (Zafar et al., 2017; Feldman et al., 2015; Kamishima
 223 et al., 2011; Jiang et al., 2020; Hardt et al., 2016; Baharlouei et al., 2020; Rezaei et al., 2020; Donini
 224 et al., 2018; Cho et al., 2020b). We run experiments on three data sets: Adult, German Credit, and
 225 COMPAS. To implement FERMI, we train a logistic regression model (same model for all baselines)
 226 with an ERMI regularizer. Details about the datasets and experiments can be found in Appendix E.

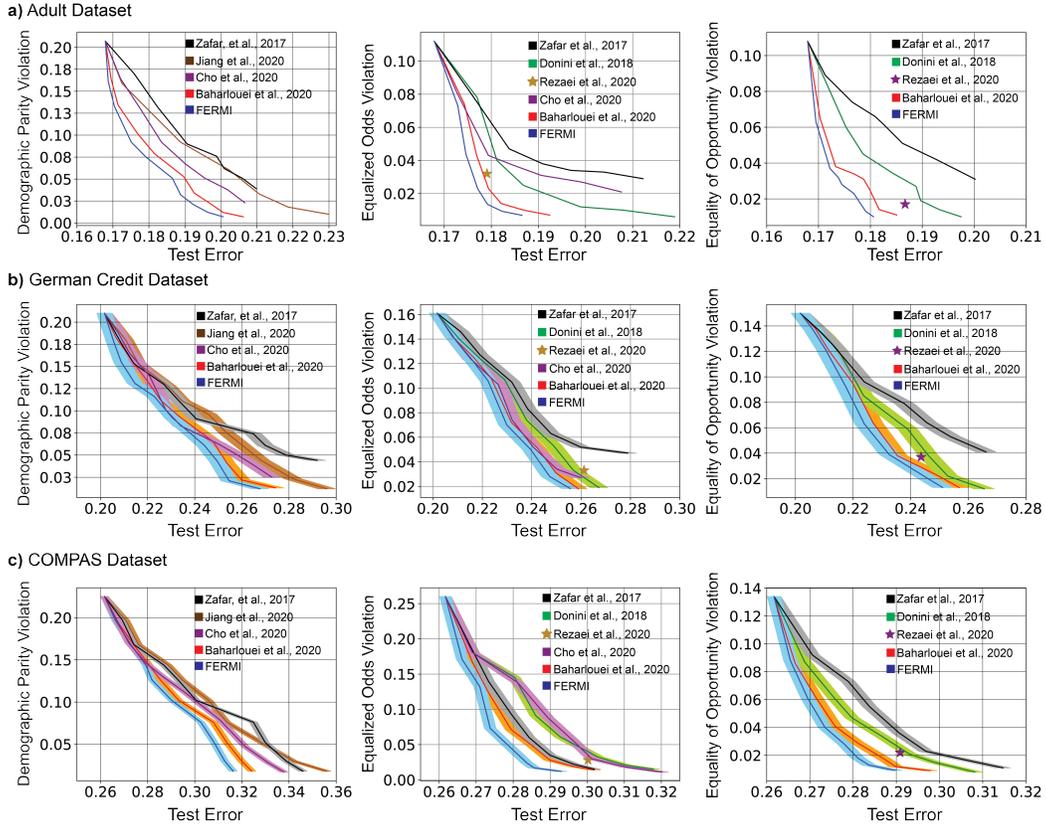


Figure 1: Binary classification with binary sensitive attribute using logistic regression. Tradeoff of fairness violation vs. test error for state-of-the-art fair classifiers on German Credit, Adult, and COMPAS datasets. FERMI offers the best fairness vs. accuracy tradeoff curve in all experiments against all baselines. Rezaei et al. (2020) only allow for a single output and do not yield a tradeoff curve. Further, the algorithms by Mary et al. (2019) and Baharlouei et al. (2020) are equivalent in this binary setting and shown by the red curve. FERMI, Mary et al. (2019) and Baharlouei et al. (2020) try to empirically solve the same risk function Eq. (FERMI obj.). However, the empirical formulation used by FERMI, Eq. (FERMI obj.) and its solver result in a better performance even-though we are using a full-batch for all baselines in this experiment.

227 In Fig. 1, we report the fairness violation vs. test error, for three notions of fairness: demographic
 228 parity, equalized odds, and equal opportunity. We have only included in-processing methods, which
 229 outperform pre-processing and post-processing methods. Complete experimental results are included
 230 in the appendix. We measure fairness violation through conditional demographic parity L_∞ violation
 231 (Definition 9), conditional equal opportunity L_∞ violation (Definition 10) and its generalization,
 232 conditional equalized odds violation. As can be seen, FERMI offers a fairness-accuracy tradeoff
 233 curve that dominates all existing state-of-the-art baselines in each experiment and with respect to
 234 each notion of fairness. This demonstrates the efficacy of having a strong regularizer such as ERMI:
 235 by enforcing small ERMI violation, our model simultaneously achieves small fairness violation with
 236 respect to these other notions which are upper bounded by ERMI.

237 It is noteworthy that the empirical objective function of Mary et al. (2019) and Baharlouei et al.
 238 (2020) is exactly the same in this setting, and their algorithms also coincide to the red curve in
 239 Fig. 1. Additionally, like FERMI, they are trying to empirically solve Eq. (FERMI obj.), albeit
 240 using different estimation techniques, i.e., their empirical objective is different from Eq. (FERMI
 241 obj.). This demonstrates the effectiveness of our empirical formulation (FERMI obj.) – which is
 242 both unbiased and consistent whereas theirs is biased. It also shows the effectiveness of our solver
 243 (Algorithm 1) even-though we are using all baselines in full batch mode in this experiment. In the
 244 following experiments, we will demonstrate that using smaller batch sizes results in much more
 245 pronounced advantages of FERMI over these baselines.

²Exponential Rényi mutual information is equal to Rényi correlation for binary targets and/or binary sensitive attributes (see Theorem 2), which is the setting of all experiments in Sec. 5.1

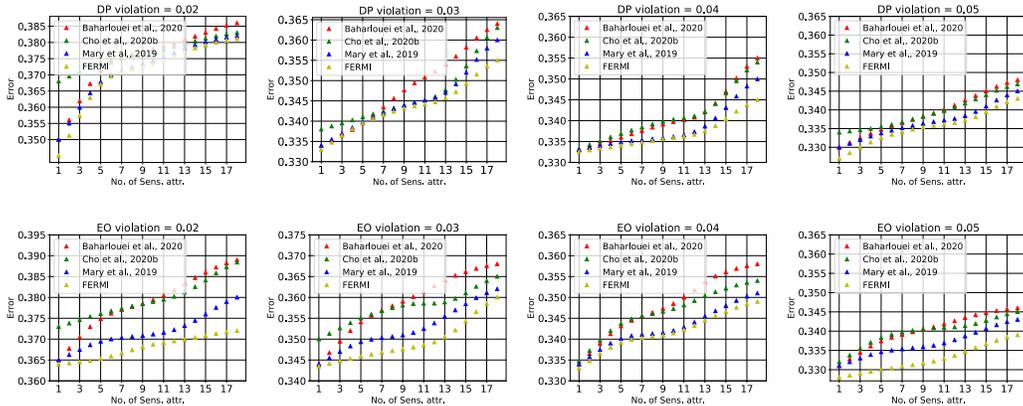


Figure 2: Comparison between FERMI, Baharlouei et al. (2020), and Cho et al. (2020b) on Communities dataset. (Mary et al., 2019) outperforms (Baharlouei et al., 2020; Cho et al., 2020b), which we believe could be attributed to the effectiveness of ERMI as a regularizer. FERMI outperforms Mary et al. (2019), which we attribute to our empirical formulation of ERMI and the effectiveness of its solver, given that we try to empirically solve the same risk function with different formulations.

246 5.2 Non-binary fair classification with a non-binary sensitive attribute

247 Next, we consider a non-binary classification problem with non-binary sensitive set. In this case, we
 248 consider the Communities and Crime dataset, which has 18 binary sensitive attributes in total, and we
 249 pick a subset of 1, 2, 3, ..., 18 sensitive attributes out of those for our experiments, which corresponds
 250 to $|\mathcal{S}| \in \{2, 4, 8, \dots, 2^{18}\}$. We discretize the target into three classes {high, medium, low}. The only
 251 baselines that we are aware of that can handle non-binary classification with non-binary sensitive
 252 attributes are (Mary et al., 2019), (Baharlouei et al., 2020), (Cho et al., 2020b), (Cho et al., 2020a),
 253 and (Zhang et al., 2018). We used the publicly available implementations of (Baharlouei et al., 2020)
 254 and (Cho et al., 2020b) and extended their binary classification algorithms to the non-binary setting.

255 The results are presented in Fig. 2, where we use conditional demographic parity L_∞ violation
 256 (Definition 9) and conditional equal opportunity L_∞ violation (Definition 10) as the fairness violation
 257 notions for the two experiments. For all baselines, test error increases as the number of sensitive
 258 attributes increases. As can be seen, compared to the baselines, FERMI offers the most favorable test
 259 error vs. fairness violation tradeoffs, particularly as the number of sensitive attributes increases and
 260 for the more stringent fairness violation levels, e.g., 0.02.

261 5.3 Domain generalization through FERMI

262 In our last experiment, our goal is to showcase the efficacy of FERMI in stochastic optimization with
 263 neural network approximation. For this experiment, we consider the Color MNIST dataset (Li &
 264 Vasconcelos, 2019), where all 60,000 training MNIST digits are colored with different colors drawn
 265 from a class conditional Gaussian distribution with variance σ around a certain average color for
 266 each digit, while the test set remains black and white. Li & Vasconcelos (2019) show that as $\sigma \rightarrow 0$,
 267 a convolutional network model overfits significantly to each digit’s color on the training set, and
 268 achieves vanishing training accuracy. However, the learned representation does not generalize to the
 269 regular black and white test set, in absence of the spurious correlation between digits and color.

270 Conceptually, the goal of the classifier in this problem is to achieve high classification accuracy with
 271 predictions that are independent of the color of the digit. We view color as the sensitive attribute
 272 in this experiment, and apply fairness baselines for the demographic parity notion of fairness. One
 273 would expect that by promoting such independence through a fairness regularizer generalization
 274 would improve (i.e. lower test error on the black and white test set), at the cost of increased training
 275 error (on the colored training set). We compare against Mary et al. (2019), Baharlouei et al. (2020),
 276 and Cho et al. (2020b) as baselines in this experiment.

277 The results of this experiment are as illustrated in Fig. 3. The details about the dataset and experimental
 278 setup is provided in Appendix E. In the left panel, we see that with no regularization ($\lambda = 0$); the
 279 test error is around 80%. As λ increases, all methods achieve smaller test error while training error
 280 increases. We also observe that FERMI offers the best test error in this setup. In the right panel,
 281 we observe that decreasing the batch size results in significantly worse generalization for all three

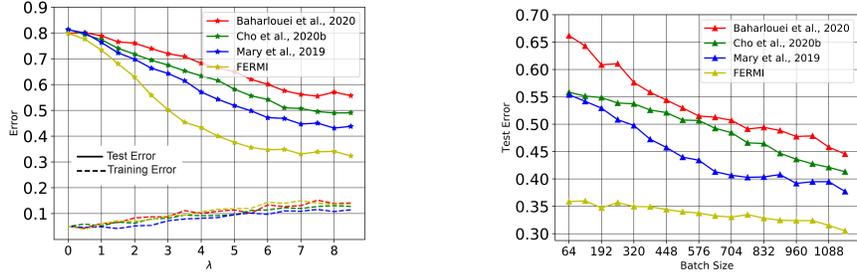


Figure 3: Domain generalization on Color MNIST (Li & Vasconcelos, 2019) using in-process fair algorithms for demographic parity. **Left panel:** The dashed line is the training error and the solid line is test error. As λ increases, fairness regularization results in a learned representation that is less dependent on color; hence training error increases while test error decreases (all algorithms reach a plateau around $\lambda = 8$). We use $|B| = 512$ for all baselines. **Right panel:** We plot test error vs. batch size using an optimized value of λ for each algorithm selected via a validation set. The performance of all baselines drops 10-20% as batch size becomes small whereas FERMI is relatively insensitive to batch size.

282 baselines considered (due to their biased estimators for the regularizer). However, the impact is much
 283 less on FERMI. In particular, the performance gap between FERMI and other baselines is more than
 284 20% for $|B| = 64$. Finally, FERMI with minibatch size $|B| = 64$ still outperforms all other baselines
 285 with $|B| > 1,000$. Finally, notice that the test error achieved by FERMI when $\sigma = 0$ is $\sim 30\%$, as
 286 compared to more than 50% obtained using REPAIR (Li & Vasconcelos, 2019) for $\sigma \leq 0.05$.

287 6 Discussion & concluding remarks

288 In this paper, we studied three variants of a notion of fairness violation, called exponential Rényi
 289 mutual information (ERMI), developed for demographic parity, equalized odds, and equal opportunity
 290 notions of fairness. We showed that ERMI is a strong fairness violation divergence providing upper
 291 bound guarantees on other popular violation divergences, namely Shannon mutual information,
 292 Rényi mutual information (Theorem 1), Pearson correlation, Rényi correlation (Theorem 2), and L_q
 293 distance violation (Theorem 3).

294 We derived an unbiased estimator for ERMI (Lemma 1), based on which we formulated an empirical
 295 objective (FERMI obj.) for solving fair empirical risk minimization with ERMI regularization
 296 to balance performance and fairness. We provided a stochastic algorithm for solving FERMI
 297 (Algorithm 1) and proved its convergence (Theorem 5); for non-binary sensitive attributes, non-binary
 298 target variables, regardless of the batch size. From an experimental perspective, we showed that
 299 FERMI leads to better fairness-accuracy tradeoffs than all of the state-of-the-art baselines on a wide
 300 variety of binary and non-binary classification tasks (for demographic parity, equalized odds, and
 301 equal opportunity). We also showed that these benefits are particularly significant when the number of
 302 sensitive attributes grows or the batch size is small. In particular, we observed that FERMI consistently
 303 outperforms Mary et al. (2019) (which tries to empirically solve the same objective Eq. (FRMI obj.))
 304 by up to 20% when the batch size is small, suggesting that the unbiasedness of the FERMI estimator
 305 is essential in achieving good empirical performance.

306 There are several possible explanations for the superior empirical performance of FERMI compared
 307 to baselines. One possible reason is that the objective function Eq. (FERMI obj.) is easier to optimize
 308 than the objectives of competing in-processing methods: ERMI is smooth; and in the discrete case, is
 309 equal to the trace of a matrix (see Theorem 7, appendix), which is easy to compute. Contrast this with
 310 the larger computational overhead of Rényi correlation used by Baharlouei et al. (2020), for example,
 311 which requires finding the second singular value of a matrix. Furthermore, the sample complexity of
 312 estimating Rényi mutual information of order 2 (and consequently that of ERMI) scales as $\Theta(\sqrt{|\mathcal{S}|})$
 313 as compared to Shannon mutual information which scales as $\Theta(|\mathcal{S}|/\log |\mathcal{S}|)$ (Acharya et al., 2014).
 314 Moreover, the fact that ERMI is a stronger fairness violation seems to imply that FERMI would
 315 generalize well to other fairness notions, a hypothesis that is supported by our experimental results.
 316 Together, these facts suggest that ERMI serves as an efficient and easily optimizable proxy for these
 317 other fairness notions, making Eq. (FERMI obj.) a good surrogate objective to optimize for all three
 318 notions of fairness considered (demographic parity, equalized odds, and equal opportunity). We leave
 319 it as future work to rigorously understand which of these (or other) factors are most responsible for
 320 the favorable performance tradeoffs observed from FERMI.

321 References

- 322 Acharya, J., Orlitsky, A., Suresh, A. T., and Tyagi, H. The complexity of estimating Rényi entropy.
323 *arXiv:1408.1000v1*, 2014.
- 324 Aghaei, S., Azizi, M. J., and Vayanos, P. Learning optimal and fair decision trees for non-
325 discriminative decision-making. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
326 volume 33, pp. 1418–1426, 2019.
- 327 Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias. *ProPublica*, 2016.
- 328 Baharlouei, S., Nouiehed, M., Beirami, A., and Razaviyayn, M. Rényi fair inference. In *ICLR*, 2020.
- 329 Bechavod, Y. and Ligett, K. Penalizing unfairness in binary classification. *arXiv preprint*
330 *arXiv:1707.00044*, 2017.
- 331 Bechavod, Y., Ligett, K., Roth, A., Waggoner, B., and Wu, Z. S. Equal opportunity in online
332 classification with partial feedback. *arXiv preprint arXiv:1902.02242*, 2019.
- 333 Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., and Roth, A. A
334 convex framework for fair regression. *arXiv preprint arXiv:1706.02409*, 2017.
- 335 Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. Man is to computer
336 programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural*
337 *information processing systems*, pp. 4349–4357, 2016.
- 338 Calmon, F., Makhdoumi, A., Médard, M., Varia, M., Christiansen, M., and Duffy, K. R. Principal
339 inertia components and applications. *IEEE Transactions on Information Theory*, 63(8):5011–5038,
340 2017a.
- 341 Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., and Varshney, K. R. Optimized pre-
342 processing for discrimination prevention. In *Advances in Neural Information Processing Systems*,
343 pp. 3992–4001, 2017b.
- 344 Cho, J., Hwang, G., and Suh, C. A fair classifier using mutual information. In *2020 IEEE International*
345 *Symposium on Information Theory (ISIT)*, pp. 2521–2526. IEEE, 2020a.
- 346 Cho, J., Hwang, G., and Suh, C. A fair classifier using kernel density estimation. In Larochelle,
347 H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information*
348 *Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*,
349 *NeurIPS 2020, December 6-12, 2020, virtual*, 2020b.
- 350 Chzhen, E. and Schreuder, N. A minimax framework for quantifying risk-fairness trade-off in
351 regression. *arXiv preprint arXiv:2007.14265*, 2020.
- 352 Cover, T. M. and Thomas, J. A. Information theory and statistics. *Elements of Information Theory*, 1
353 (1):279–335, 1991.
- 354 Csiszár, I. and Shields, P. C. *Information theory and statistics: A tutorial*. Now Publishers Inc, 2004.
- 355 Datta, A., Tschantz, M. C., and Datta, A. Automated experiments on ad privacy settings. *Proceedings*
356 *on privacy enhancing technologies*, 2015(1):92–112, 2015.
- 357 Dembo, A. and Zeitouni, O. *Large deviations techniques and applications*. Springer Science &
358 Business Media, 2009.
- 359 Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S., and Pontil, M. Empirical risk minimization
360 under fairness constraints. In *Advances in Neural Information Processing Systems*, pp. 2791–2801,
361 2018.
- 362 Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In
363 *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226, 2012.
- 364 Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying
365 and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference*
366 *on knowledge discovery and data mining*, pp. 259–268, 2015.

- 367 Fish, B., Kun, J., and Lelkes, Á. D. A confidence-based approach for balancing fairness and accuracy.
 368 In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pp. 144–152. SIAM,
 369 2016.
- 370 Gebelein, H. Das statistische problem der korrelation als variations-und eigenwertproblem und
 371 sein zusammenhang mit der ausgleichsrechnung. *ZAMM-Journal of Applied Mathematics and*
 372 *Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 21(6):364–379, 1941.
- 373 Gari, V., Ruf, B., Lamprier, S., and Detryniecki, M. Fairness-aware neural Rényi minimization for
 374 continuous features. *arXiv preprint arXiv:1911.04929*, 2019.
- 375 Gari, V., Hajouji, O. E., Lamprier, S., and Detryniecki, M. Learning unbiased representations via
 376 Rényi minimization. *arXiv preprint arXiv:2009.03183*, 2020.
- 377 Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in*
 378 *neural information processing systems*, pp. 3315–3323, 2016.
- 379 Hirschfeld, H. O. A connection between correlation and contingency. In *Proceedings of the*
 380 *Cambridge Philosophical Society*, volume 31, pp. 520–524, 1935.
- 381 Jiang, R., Pacchiano, A., Stepleton, T., Jiang, H., and Chiappa, S. Wasserstein fair classification. In
 382 *Uncertainty in Artificial Intelligence*, pp. 862–872. PMLR, 2020.
- 383 Kamiran, F., Calders, T., and Pechenizkiy, M. Discrimination aware decision tree learning. In *2010*
 384 *IEEE International Conference on Data Mining*, pp. 869–874. IEEE, 2010.
- 385 Kamishima, T., Akaho, S., and Sakuma, J. Fairness-aware learning through regularization approach.
 386 In *2011 IEEE 11th International Conference on Data Mining Workshops*, pp. 643–650. IEEE,
 387 2011.
- 388 Kearns, M., Neel, S., Roth, A., and Wu, Z. S. Preventing fairness gerrymandering: Auditing and
 389 learning for subgroup fairness. In *International Conference on Machine Learning*, pp. 2564–2572,
 390 2018.
- 391 Kilbertus, N., Rodriguez, M. G., Schölkopf, B., Muandet, K., and Valera, I. Fair decisions despite
 392 imperfect predictions. In Chiappa, S. and Calandra, R. (eds.), *Proceedings of the Twenty Third*
 393 *International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of*
 394 *Machine Learning Research*, pp. 277–287. PMLR, 26–28 Aug 2020.
- 395 Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document
 396 recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- 397 Li, Y. and Vasconcelos, N. Repair: Removing representation bias by dataset resampling. In
 398 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9572–9581,
 399 2019.
- 400 Lin, T., Jin, C., and Jordan, M. I. On gradient descent ascent for nonconvex-concave minimax
 401 problems. *arXiv: 1906.00331v6*, 2020.
- 402 Luo, L., Ye, H., and Zhang, T. Stochastic recursive gradient descent ascent for stochastic nonconvex-
 403 strongly-concave minimax problems. *arXiv: 2001.03724*, 2020.
- 404 Mary, J., Calauzenes, C., and El Karoui, N. Fairness-aware learning for continuous attributes and
 405 treatments. In *International Conference on Machine Learning*, pp. 4382–4391. PMLR, 2019.
- 406 Pérez-Suay, A., Laparra, V., Mateo-García, G., Muñoz-Marí, J., Gómez-Chova, L., and Camps-Valls,
 407 G. Fair kernel learning. In *Joint European Conference on Machine Learning and Knowledge*
 408 *Discovery in Databases*, pp. 339–355. Springer, 2017.
- 409 Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., and Weinberger, K. Q. On fairness and calibration.
 410 In *Advances in Neural Information Processing Systems*, pp. 5680–5689, 2017.
- 411 Raff, E., Sylvester, J., and Mills, S. Fair forests: Regularized tree induction to minimize model bias.
 412 In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 243–250, 2018.

- 413 Rényi, A. On measures of dependence. *Acta Mathematica Academiae Scientiarum Hungarica*, 10
414 (3-4):441–451, 1959.
- 415 Rényi, A. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium*
416 *on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*.
417 The Regents of the University of California, 1961.
- 418 Rezaei, A., Fathony, R., Memarrast, O., and Ziebart, B. D. Fairness for robust log loss classification.
419 In *AAAI*, pp. 5511–5518, 2020.
- 420 Steinberg, D., Reid, A., O’Callaghan, S., Lattimore, F., McCalman, L., and Caetano, T. Fast fair
421 regression via efficient approximations of mutual information. *arXiv preprint arXiv:2002.06200*,
422 2020.
- 423 Sweeney, L. Discrimination in online ad delivery. *arXiv preprint arXiv:1301.6822*, 2013.
- 424 Taskesen, B., Nguyen, V. A., Kuhn, D., and Blanchet, J. A distributionally robust approach to fair
425 classification. *arXiv preprint arXiv:2007.09530*, 2020.
- 426 Williamson, R. and Menon, A. Fairness risk measures. In *International Conference on Machine*
427 *Learning*, pp. 6786–6797. PMLR, 2019.
- 428 Witsenhausen, H. S. On sequences of pairs of dependent random variables. *SIAM Journal on Applied*
429 *Mathematics*, 28(1):100–113, 1975.
- 430 Woodworth, B., Gunasekar, S., Ohanessian, M. I., and Srebro, N. Learning non-discriminatory
431 predictors. *arXiv preprint arXiv:1702.06081*, 2017.
- 432 Zafar, M. B., Valera, I., Ródriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for
433 fair classification. In *Artificial Intelligence and Statistics*, pp. 962–970. PMLR, 2017.
- 434 Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In
435 *International Conference on Machine Learning*, pp. 325–333, 2013.
- 436 Zhang, B. H., Lemoine, B., and Mitchell, M. Mitigating unwanted biases with adversarial learning.
437 In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.

438 **Checklist**

- 439 1. For all authors...
- 440 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
441 contributions and scope? [Yes]
- 442 (b) Did you describe the limitations of your work? [Yes]
- 443 (c) Did you discuss any potential negative societal impacts of your work? [Yes]
- 444 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
445 them? [Yes]
- 446 2. If you are including theoretical results...
- 447 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- 448 (b) Did you include complete proofs of all theoretical results? [Yes]
- 449 3. If you ran experiments...
- 450 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
451 mental results (either in the supplemental material or as a URL)? [Yes]
- 452 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
453 were chosen)? [Yes]
- 454 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
455 ments multiple times)? [Yes]
- 456 (d) Did you include the total amount of compute and the type of resources used (e.g., type
457 of GPUs, internal cluster, or cloud provider)? [No]
- 458 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 459 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 460 (b) Did you mention the license of the assets? [N/A]
- 461 (c) Did you include any new assets either in the supplemental material or as a URL? [No]
- 462 (d) Did you discuss whether and how consent was obtained from people whose data you're
463 using/curating? [No]
- 464 (e) Did you discuss whether the data you are using/curating contains personally identifiable
465 information or offensive content? [No]
- 466 5. If you used crowdsourcing or conducted research with human subjects...
- 467 (a) Did you include the full text of instructions given to participants and screenshots, if
468 applicable? [N/A]
- 469 (b) Did you describe any potential participant risks, with links to Institutional Review
470 Board (IRB) approvals, if applicable? [N/A]
- 471 (c) Did you include the estimated hourly wage paid to participants and the total amount
472 spent on participant compensation? [N/A]